

Oracle® Fusion Middleware

Configuring and Managing JDBC Data Sources for Oracle
WebLogic Server

11g Release 1 (10.3.4)

E13737-03

January 2011

This book contains JDBC data source configuration and
administration information.

Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server, 11g Release 1 (10.3.4)

E13737-03

Copyright © 2007, 2011, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this software or related documentation is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of this software. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software in dangerous applications.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

This software and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

Preface	xi
Documentation Accessibility	xi
Conventions	xi
1 Introduction and Roadmap	
1.1 Document Scope and Audience.....	1-1
1.2 Guide to this Document	1-1
1.3 Related Documentation.....	1-2
1.4 JDBC Samples and Tutorials	1-2
1.4.1 Avitek Medical Records Application (MedRec) and Tutorials.....	1-2
1.4.2 JDBC Examples in the WebLogic Server Distribution	1-3
1.5 New and Changed JDBC Data Source Features in This Release.....	1-3
1.5.1 Enhanced Oracle RAC Support	1-3
1.5.2 Reorganization of JDBC Resources	1-3
1.5.3 Data Source Factories	1-4
1.5.4 Performance Enhancements.....	1-4
1.5.5 Connection-based System Properties	1-4
1.5.6 Keep Connection After Local Transaction	1-4
2 Configuring WebLogic JDBC Resources	
2.1 Understanding JDBC Resources in WebLogic Server	2-1
2.2 Ownership of Configured JDBC Resources	2-1
2.3 Data Source Configuration Files.....	2-2
2.3.1 JDBC System Modules	2-2
2.3.2 JDBC Application Modules	2-4
2.3.3 JDBC Module File Naming Requirements.....	2-5
2.3.4 JDBC Modules in Versioned Applications.....	2-5
2.3.5 JDBC Schema	2-6
2.4 JMX and WLST Access for JDBC Resources	2-6
2.4.1 JDBC MBeans for System Resources.....	2-6
2.4.2 JDBC Management Objects in the Java EE Management Model (JSR-77 Support) ...	2-7
2.4.3 Using WLST to Create JDBC System Resources	2-8
2.4.4 How to Modify and Monitor JDBC Resources.....	2-9
2.4.5 Best Practices when Using WLST to Configure JDBC Resources.....	2-9
2.5 Overview of Clustered JDBC Resources.....	2-10

3 Configuring JDBC Data Sources

3.1	Understanding JDBC Data Sources.....	3-1
3.2	Types of WebLogic Server JDBC Data Sources	3-1
3.3	Creating a JDBC Data Source	3-2
3.3.1	JDBC Data Source Properties	3-2
3.3.1.1	Data Source Names	3-2
3.3.1.2	JNDI Names	3-2
3.3.1.3	Selecting a Database Type	3-2
3.3.1.4	Selecting a JDBC Driver.....	3-3
3.3.2	Configure Transaction Options	3-4
3.3.3	Configure Connection Properties.....	3-5
3.3.4	Test Connections.....	3-5
3.3.5	Target the Data Source.....	3-5
3.4	Configuring Connection Pool Features	3-5
3.4.1	Enabling JDBC Driver-Level Features	3-6
3.4.2	Enabling Connection-based System Properties.....	3-6
3.4.3	Initializing Database Connections with SQL Code	3-6
3.5	Configuring Oracle Parameters	3-7
3.6	Configuring an ONS Client	3-7
3.7	Tuning Generic Data Source Connection Pools	3-7
3.8	Setting Database Security Credentials	3-7
3.8.1	Types of Data Source Pools	3-7
3.8.2	Using a User Name/Password.....	3-7
3.8.3	Set Client ID On Connection	3-8
3.8.4	Identity-based Connection Pooling.....	3-8
3.8.4.1	How Heterogeneous Connections are Created.....	3-9
3.8.4.2	Using Identity-based Pooling with Global Transactions.....	3-10
3.8.4.3	Using Identity-based Pooling with LLR.....	3-10
3.9	JDBC Data Source Factories (Deprecated).....	3-10

4 Using GridLink Data Sources

4.1	What is a GridLink Data Source	4-1
4.1.1	Fast Connection Failover	4-2
4.1.2	Runtime Connection Load Balancing	4-3
4.1.3	Graceful Handling for Oracle RAC Outages	4-4
4.1.3.1	Handling for Oracle RAC Outages Prior to Oracle RAC 11.2	4-4
4.1.4	XA Affinity.....	4-4
4.1.5	SCAN Addresses	4-5
4.1.6	Secure Communication using Oracle Wallet.....	4-5
4.2	Creating a GridLink Data Source	4-5
4.2.1	JDBC Data Source Properties	4-6
4.2.1.1	Data Source Names	4-6
4.2.1.2	JNDI Names	4-6
4.2.1.3	Select an XA or Non-XA Driver	4-6
4.2.2	Configure Transaction Options	4-6
4.2.3	Configure Connection Properties.....	4-7
4.2.3.1	Enter Connection Properties.....	4-7

4.2.3.2	Enter a Complete URL	4-7
4.2.4	Test Connections	4-8
4.2.5	Configure an ONS Client Configuration.....	4-8
4.2.5.1	Secure ONS Client Communication	4-8
4.2.6	Test ONS Client Configuration.....	4-8
4.2.7	Target the Data Source	4-9
4.3	Enable Socket Direct Protocol for a Grid Link Data Source	4-9
4.4	Configuring Connection Pool Features	4-9
4.4.1	Enabling JDBC Driver-Level Features	4-10
4.4.2	Enabling Connection-based System Properties.....	4-10
4.4.3	Initializing Database Connections with SQL Code	4-10
4.5	Configuring Oracle Parameters	4-11
4.6	Configuring an ONS Client	4-11
4.6.1	Enabling FAN Events.....	4-11
4.6.2	Configuring Generic Data Source Connection Testing	4-11
4.6.3	Using a Wallet File.....	4-11
4.7	Configuring Oracle Parameters	4-11
4.8	Tuning GridLink Data Source Connection Pools	4-12
4.9	Setting Database Security Credentials	4-12
4.9.1	Using a User Name/Password	4-12
4.9.2	Set Client ID On Connection	4-12
4.10	Monitoring GridLink JDBC Resources	4-13
4.10.1	Viewing Run-Time Statistics	4-13
4.10.1.1	JDBCOracleDataSourceInstanceRuntimeMBean	4-13
4.10.1.2	JDBCDataSourceRuntimeMBean	4-13
4.10.1.3	ONSDaemonRuntimeMBean.....	4-13
4.10.2	Debug GridLink Data Sources	4-14
4.10.2.1	JDBC Debugging Scopes	4-14
4.10.2.2	UCP JDK Logging.....	4-14
4.10.2.3	Enable Debugging Using the Command Line	4-14

5 Configuring JDBC Multi Data Sources

5.1	Multi Data Source Features	5-1
5.1.1	Removing a Database Node.....	5-2
5.1.2	Adding a Database Node	5-2
5.2	Creating and Configuring Multi Data Sources.....	5-2
5.3	Choosing the Multi Data Source Algorithm	5-3
5.3.1	Failover.....	5-3
5.3.2	Load Balancing.....	5-3
5.4	Multi Data Source Fail-Over Limitations and Requirements.....	5-3
5.4.1	Test Connections on Reserve to Enable Fail-Over	5-3
5.4.2	No Fail-Over for In-Use Connections	5-4
5.5	Multi Data Source Failover Enhancements	5-4
5.5.1	Connection Request Routing Enhancements When a Data Source Fails	5-4
5.5.2	Automatic Re-enablement on Recovery of a Failed Data Source within a Multi Data Source	5-4
5.5.3	Enabling Failover for Busy Data Sources in a Multi Data Source	5-5

5.5.4	Controlling Multi Data Source Failover with a Callback.....	5-5
5.5.4.1	Callback Handler Requirements	5-5
5.5.4.2	Callback Handler Configuration.....	5-6
5.5.4.3	How It Works—Failover	5-6
5.5.5	Controlling Multi Data Source Failback with a Callback	5-7
5.5.5.1	How It Works—Failback	5-7
5.6	Deploying JDBC Multi Data Sources on Servers and Clusters	5-8

6 Advanced Configuration for Oracle Drivers

7 JDBC Data Source Transaction Options

7.1	Enabling Support for Global Transactions with a Non-XA JDBC Driver	7-2
7.2	Understanding the Logging Last Resource Transaction Option	7-2
7.2.1	Advantages to Using the Logging Last Resource Optimization	7-3
7.2.2	Enabling the Logging Last Resource Transaction Optimization	7-4
7.2.3	Programming Considerations and Limitations for LLR Data Sources.....	7-4
7.2.4	Administrative Considerations and Limitations for LLR Data Sources.....	7-5
7.3	Understanding the Emulate Two-Phase Commit Transaction Option.....	7-6
7.3.1	Limitations and Risks When Emulating Two-Phase Commit Using a Non-XA Driver.....	7-7
7.3.1.1	Heuristic Completions and Data Inconsistency	7-7
7.3.1.2	Cannot Recover Pending Transactions	7-8
7.3.1.3	Possible Performance Loss with Non-XA Resources in Multi-Server Configurations 7-8	
7.3.1.4	Multiple Non-XA Participants.....	7-8

8 Using Roles and Policies to Secure JDBC Data Sources

8.1	Setting Security Policies for JDBC Resources.....	8-1
8.2	Security Roles for JDBC MBeans	8-1
8.2.1	JDBC Domain Configuration MBeans	8-1
8.2.2	JDBC System Module MBeans.....	8-2

9 Deploying Data Sources on Servers and Clusters

9.1	Deploying Data Sources on Servers and Clusters.....	9-1
9.2	Minimizing Server Startup Hang Caused By an Unresponsive Database	9-1

10 Tuning Data Source Connection Pools

10.1	Increasing Performance with the Statement Cache	10-1
10.1.1	Statement Cache Algorithms	10-2
10.1.1.1	LRU (Least Recently Used)	10-2
10.1.1.2	Fixed	10-2
10.1.2	Statement Cache Size.....	10-2
10.1.3	Usage Restrictions for the Statement Cache	10-3
10.1.3.1	Calling a Stored Statement After a Database Change May Cause Errors	10-3
10.1.3.2	Using setNull In a Prepared Statement.....	10-3
10.1.3.3	Statements in the Cache May Reserve Database Cursors.....	10-3

10.2	Connection Testing Options for a Data Source	10-4
10.2.1	Database Connection Testing Semantics.....	10-4
10.2.1.1	Connection Testing When Database Connections are Created	10-5
10.2.1.2	Periodic Connection Testing	10-5
10.2.1.3	Testing Reserved Connections.....	10-6
10.2.1.4	Minimized Connection Test Delay After Database Connectivity Loss	10-6
10.2.1.5	Minimized Connection Request Delay After Connection Test Failures.....	10-6
10.2.1.6	Minimized Connection Request Delays After Loss of DBMS Connectivity	10-7
10.2.1.7	Minimizing Connection Request Delay with Seconds to Trust an Idle Pool Connection	10-7
10.2.2	Database Connection Testing Configuration Recommendations.....	10-7
10.2.3	Default Test Table Name	10-8
10.3	Enabling Connection Creation Retries.....	10-9
10.4	Enabling Connection Requests to Wait for a Connection	10-9
10.4.1	Connection Reserve Timeout	10-9
10.4.2	Limiting the Number of Waiting Connection Requests.....	10-10
10.5	Automatically Recovering Leaked Connections	10-10
10.6	Avoiding Server Lockup with the Correct Number of Connections	10-10
10.7	Limiting Statement Processing Time with Statement Timeout.....	10-11
10.8	Using Pinned-To-Thread Property to Increase Performance	10-11
10.8.1	Changes to Connection Pool Administration Operations When PinnedToThread is Enabled	10-11
10.8.2	Additional Database Resource Costs When PinnedToThread is Enabled	10-12
10.9	Using Unwrapped Data Type Objects	10-12
10.9.1	How to Disable Wrapping.....	10-13
10.9.1.1	Disable Wrapping using the Administration Console.....	10-13
10.9.1.2	Disable Wrapping using WLST	10-13

11 Using WebLogic Server with Oracle RAC

11.1	Overview of Oracle Real Application Clusters	11-1
11.2	Software Requirements	11-2
11.3	JDBC Driver Requirements	11-2
11.4	Hardware Requirements.....	11-2
11.4.1	WebLogic Server Cluster	11-2
11.4.2	Oracle RAC Cluster	11-2
11.4.3	Shared Storage.....	11-3
11.5	Configuration Options in WebLogic Server with Oracle RAC	11-3
11.5.1	Choosing a WebLogic Server Configuration for Use with Oracle RAC.....	11-3
11.5.2	Validating Connections when using WebLogic Server with Oracle RAC	11-4
11.5.3	Additional Considerations When Using WebLogic Server with Oracle RAC	11-4

12 Using JDBC Drivers with WebLogic Server

12.1	JDBC Drivers Installed with WebLogic Server.....	12-1
12.2	Setting the Environment for a Type-4 Third-Party JDBC Driver.....	12-2
12.3	Globalization Support for the Oracle 11g Thin Driver	12-2
12.4	Using the Oracle Thin Driver in Debug Mode	12-2

12.5	Using the Oracle Thin Driver Over SSL	12-2
------	---	------

13 Monitoring WebLogic JDBC Resources

13.1	Viewing Run-Time Statistics	13-1
13.1.1	Data Source Statistics	13-1
13.1.2	Prepared Statement Cache Statistics	13-2
13.2	Collecting Profile Information	13-2
13.2.1	Profile Types.....	13-2
13.2.1.1	Connection Usage (PROFILE_TYPE_CONN_USAGE_STR).....	13-2
13.2.1.2	Connection Reservation Wait (PROFILE_TYPE_CONN_RESV_WAIT_STR).....	13-3
13.2.1.3	Connection Reservation Failed (PROFILE_TYPE_CONN_RESV_FAIL_STR).....	13-3
13.2.1.4	Connection Leak (PROFILE_TYPE_CONN_LEAK_STR)	13-3
13.2.1.5	Connection Last Usage (PROFILE_TYPE_CONN_LAST_USAGE_STR)	13-3
13.2.1.6	Connection Multithreaded Usage (PROFILE_TYPE_CONN_MT_USAGE_STR).....	13-4
13.2.1.7	Statement Cache Entry (PROFILE_TYPE_STMT_CACHE_ENTRY_STR)	13-4
13.2.1.8	Statements Usage (PROFILE_TYPE_STMT_USAGE_STR).....	13-4
13.2.2	Accessing Diagnostic Data	13-4
13.2.3	Callbacks for Monitoring Driver-Level Statistics.....	13-5
13.3	Debugging JDBC Data Sources.....	13-5
13.3.1	Enabling Debugging.....	13-5
13.3.1.1	Enable Debugging Using the Command Line	13-5
13.3.1.2	Enable Debugging Using the WebLogic Server Administration Console	13-6
13.3.1.3	Enable Debugging Using the WebLogic Scripting Tool	13-6
13.3.1.4	Changes to the config.xml File	13-7
13.3.2	JDBC Debugging Scopes.....	13-8
13.3.3	Request Dyeing	13-8

14 Managing WebLogic JDBC Resources

14.1	Testing Data Sources and Database Connections	14-1
14.2	Managing the Statement Cache for a Data Source.....	14-2
14.2.1	Clearing the Statement Cache for a Data Source	14-2
14.2.2	Clearing the Statement Cache for a Single Connection.....	14-2
14.3	Shrinking a Connection Pool.....	14-3
14.4	Resetting a Connection Pool.....	14-3
14.5	Suspending a Connection Pool	14-3
14.6	Resuming a Connection Pool	14-3
14.7	Shutting Down a Data Source	14-4
14.8	Starting a Data Source	14-4
14.9	Managing DBMS Network Failures	14-4

A Configuring JDBC Application Modules for Deployment

A.1	Packaging a JDBC Module with an Enterprise Application: Main Steps	A-1
A.2	Creating Packaged JDBC Modules.....	A-1
A.2.1	Creating a JDBC Data Source Module Using the Administration Console	A-2

A.2.2	JDBC Packaged Module Requirements	A-2
A.2.3	JDBC Application Module Limitations	A-3
A.2.4	Creating a JDBC Data Source Module	A-3
A.2.5	Creating a JDBC Multi Data Source Module	A-4
A.2.6	Encrypting Database Passwords in a JDBC Module	A-5
A.2.6.1	Deploying JDBC Modules to New Domains	A-5
A.2.7	Application Scoping for a Packaged JDBC Module	A-5
A.3	Referencing a JDBC Module in Java EE Descriptor Files	A-6
A.3.1	Packaged JDBC Module References in weblogic-application.xml	A-6
A.3.2	Packaged JDBC Module References in Other Descriptors	A-7
A.4	Packaging an Enterprise Application with a JDBC Module	A-7
A.5	Deploying an Enterprise Application with a JDBC Module	A-8
A.6	Getting a Database Connection from a Packaged JDBC Module	A-8

B Using Multi Data Sources with Oracle RAC

B.1	Overview of Oracle Real Application Clusters	B-1
B.1.1	Oracle RAC Scalability with WebLogic Server Multi Data Sources	B-2
B.1.2	Oracle RAC Availability with WebLogic Server Multi Data Sources	B-3
B.1.3	Oracle RAC Load Balancing with WebLogic Server Multi Data Sources	B-3
B.2	Software Requirements	B-3
B.3	JDBC Driver Requirements	B-3
B.4	Hardware Requirements	B-3
B.4.1	WebLogic Server Cluster	B-4
B.4.2	Oracle RAC Cluster	B-4
B.4.3	Shared Storage	B-4
B.5	Configuring Multi Data Sources with Oracle RAC	B-4
B.5.1	Choosing a Multi Data Source Configuration for Use with Oracle RAC	B-4
B.5.2	Configuring Multi Data Sources for use with Oracle RAC	B-5
B.5.2.1	Attributes of a Multi Data Source	B-6
B.5.3	Configuration Considerations for Failover	B-6
B.5.3.1	Multi Data Source-Managed Failover and Load Balancing	B-6
B.5.3.2	Delays During Failover	B-7
B.5.3.3	Failure Handling Walkthrough for Global Transactions	B-8
B.5.4	Configuring the Listener Process for Each Oracle RAC Instance	B-8
B.5.5	Configuring Multi Data Sources When Remote Listeners are Enabled or Disabled	B-10
B.5.6	Additional Configuration Considerations	B-10
B.6	Using Multi Data Sources with Global Transactions	B-11
B.6.1	Rules for Data Sources within a Multi Data Source Using Global Transactions	B-11
B.6.2	Required Attributes of Data Sources within a Multi Data Source Using Global Transactions	B-12
B.6.3	Sample Configuration Code	B-12
B.7	Using Multi Data Sources without Global Transactions	B-14
B.7.1	Attributes of Data Sources within a Multi Data Source Not Using Global Transactions	B-14
B.7.2	Sample Configuration Code	B-14
B.8	Configuring Connections to Services on Oracle RAC Nodes	B-16

B.8.1	Configuring a Data Source to Connect to a Service.....	B-16
B.8.2	Service Connection Configurations	B-17
B.8.2.1	Workload Management.....	B-17
B.8.2.2	Load Balancing.....	B-18
B.8.3	Connection Pool Capacity Planning	B-19
B.9	XA Considerations and Limitations when using multi Data Sources with Oracle RAC	B-22
B.9.1	Oracle RAC XA Requirements when using multi Data Sources	B-23
B.9.1.1	Use Multi Data Sources	B-23
B.9.1.2A	Global Transaction Must Be Initiated, Prepared, and Concluded in the Same Instance of the Oracle RAC Cluster	B-23
B.9.1.3	Transaction IDs Must Be Unique Within the Oracle RAC Cluster	B-23
B.9.2	Known Limitations When Using Oracle RAC with multi Data Sources.....	B-23
B.9.2.1	Potential for Data Deadlocks in Some Failure Scenarios.....	B-23
B.9.2.2	Potential for Transactions Completed Out of Sequence	B-24
B.9.3	Known Issue Occurring After Database Server Crash.....	B-24
B.10	JDBC Store Recovery with Oracle RAC.....	B-24
B.10.1	Configuring a JDBC Store for Use with Oracle RAC.....	B-24
B.10.2	Automatic Retry	B-25

C Using Connect-Time Failover with Oracle RAC

C.1	Using Connect-Time Failover without Global Transactions	C-1
C.2	Attributes of a Connect-Time Failover Configuration without Global Transactions	C-3
C.3	Sample Configuration Code.....	C-3

D Using Fast Connection Failover with Oracle RAC

D.1	JDBC Driver Configuration for use with Oracle Fast Connection Failover	D-1
-----	--	-----

Preface

This preface describes the document accessibility features and conventions used in this guide—*Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server*.

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at <http://www.oracle.com/accessibility/>.

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/support/contact.html> or visit <http://www.oracle.com/accessibility/support.html> if you are hearing impaired.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
<code>monospace</code>	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Introduction and Roadmap

This section describes the contents and organization of this guide—*Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server*.

- [Section 1.1, "Document Scope and Audience"](#)
- [Section 1.2, "Guide to this Document"](#)
- [Section 1.3, "Related Documentation"](#)
- [Section 1.4, "JDBC Samples and Tutorials"](#)
- [Section 1.5, "New and Changed JDBC Data Source Features in This Release"](#)

1.1 Document Scope and Audience

This document is a resource for software developers and system administrators who develop and support applications that use the Java Database Connectivity (JDBC) API. It also contains information that is useful for business analysts and system architects who are evaluating WebLogic Server. The topics in this document are relevant during the evaluation, design, development, pre-production, and production phases of a software project.

This document does not address specific JDBC programming topics. For links to WebLogic Server documentation and resources for this topic, see [Section 1.3, "Related Documentation."](#)

It is assumed that the reader is familiar with Java EE and JDBC concepts. This document emphasizes the value-added features provided by WebLogic Server.

1.2 Guide to this Document

- This chapter, [Chapter 1, "Introduction and Roadmap,"](#) introduces the organization of this guide and lists new features in the current release.
- [Chapter 2, "Configuring WebLogic JDBC Resources,"](#) provides an overview of WebLogic JDBC resources.
- [Chapter 3, "Configuring JDBC Data Sources,"](#) describes WebLogic JDBC data source configuration.
- [Chapter 4, "Using GridLink Data Sources,"](#) describes WebLogic JDBC GridLink data source configuration.
- [Chapter 5, "Configuring JDBC Multi Data Sources,"](#) describes WebLogic JDBC multi data source configuration.

- [Chapter 11, "Using WebLogic Server with Oracle RAC,"](#) describes how to configure WebLogic Server for use with Oracle Real Application Clusters.
- [Chapter 12, "Using JDBC Drivers with WebLogic Server,"](#) describes how to use JDBC drivers from other sources in your WebLogic JDBC data source configuration.
- [Chapter 13, "Monitoring WebLogic JDBC Resources,"](#) describes how to monitor JDBC resources, gather profile information about database connection usage, and enable JDBC debugging.
- [Chapter 14, "Managing WebLogic JDBC Resources,"](#) describes how to administer data sources.
- [Appendix A, "Configuring JDBC Application Modules for Deployment,"](#) describes how to package a WebLogic JDBC module with your enterprise application.
- [Appendix B, "Using Multi Data Sources with Oracle RAC,"](#) describes how to configure generic data sources for use with Oracle Real Application Clusters.

1.3 Related Documentation

This document contains JDBC data source configuration and administration information.

For comprehensive guidelines for developing, deploying, and monitoring WebLogic Server applications, see the following documents:

- *Programming JDBC for Oracle WebLogic Server* is a guide to JDBC API programming with WebLogic Server.
- *Developing Applications for Oracle WebLogic Server* is a guide to developing WebLogic Server applications.
- *Deploying Applications to Oracle WebLogic Server* is the primary source of information about deploying WebLogic Server applications in development and production environments.

1.4 JDBC Samples and Tutorials

In addition to this document, Oracle provides a variety of JDBC code samples and tutorials that show configuration and API use, and provide practical instructions on how to perform key JDBC development tasks.

1.4.1 Avitek Medical Records Application (MedRec) and Tutorials

MedRec is an end-to-end sample Java EE application shipped with WebLogic Server that simulates an independent, centralized medical record management system. The MedRec application provides a framework for patients, doctors, and administrators to manage patient data using a variety of different clients.

MedRec demonstrates WebLogic Server and Java EE features, and highlights Oracle-recommended best practices. MedRec is included in the WebLogic Server distribution, and can be accessed from the Start menu on Windows machines. For Linux and other platforms, you can start MedRec from the `WL_HOME\samples\domains\medrec` directory, where `WL_HOME` is the top-level installation directory for WebLogic Platform.

1.4.2 JDBC Examples in the WebLogic Server Distribution

WebLogic Server optionally installs API code examples in `WL_HOME\samples\server\examples\src\examples`, where `WL_HOME` is the top-level directory of your WebLogic Server installation. You can start the examples server, and obtain information about the samples and how to run them from the WebLogic Server Start menu.

1.5 New and Changed JDBC Data Source Features in This Release

This release includes the following new and changed features:

- [Section 1.5.1, "Enhanced Oracle RAC Support"](#)
- [Section 1.5.2, "Reorganization of JDBC Resources"](#)
- [Section 1.5.3, "Data Source Factories"](#)
- [Section 1.5.4, "Performance Enhancements"](#)
- [Section 1.5.5, "Connection-based System Properties"](#)
- [Section 1.5.6, "Keep Connection After Local Transaction"](#)

For a comprehensive listing of the new WebLogic Server features introduced in this release, see *What's New in Oracle WebLogic Server*.

1.5.1 Enhanced Oracle RAC Support

This release provides a new data source type, a GridLink Data Source, to provide enhanced support for Oracle RAC. GridLink Data Sources provide:

- A simplified and consolidated method to configure Oracle RAC connectivity between a single WebLogic Server data source and a service targeted to an Oracle RAC cluster.
- The ability to adaptively respond to state changes in the Oracle RAC Cluster, including support for:
 - Fast Connection Failover (FCF)
 - Runtime Connection Load Balancing (RCLB)
 - Graceful data base shutdown
- Additional performance features, such as XA Affinity for global transactions, that ensures all the data base operations performed on an Oracle RAC cluster are directed to the same Oracle RAC instance.
- Supports for Oracle Single Client Access Name (SCAN) addresses. This allows multiple GridLink data source database listener addresses and ONS daemon addresses to register to SCAN listeners and appear as a single name used to access an Oracle RAC instance. A GridLink data source containing SCAN addresses does not need to change if you add or remove Oracle RAC nodes.

For more information, see [Section 4, "Using GridLink Data Sources"](#)

1.5.2 Reorganization of JDBC Resources

JDBC resources have been reorganized within the WebLogic Server Administration Console domain tree. Generic Data Sources, GridLink Data Sources, and Multi Data Sources are all located under the **Services > Data Sources** branch of the Domain Tree in the Administration Console.

1.5.3 Data Source Factories

Data Source Factories (deprecated) are still supported in this release but do not appear in the domain tree unless you have upgraded an existing domain that contains Data Source Factories. If your upgraded domain contains Data Source Factories, they will be located under **Services > Data Source Factories** branch of the Domain Tree in the Administration Console.

1.5.4 Performance Enhancements

This release supports provides several Data Source performance enhancements when using a supported Oracle driver, including:

- Oracle Net Sdp
- Oracle Dynamic Buffering
- Oracle Reduced Data Copies

These enhancements may require specific hardware, software, and configuration. For more information, see [Section 3.5, "Configuring Oracle Parameters."](#)

1.5.5 Connection-based System Properties

This release supports setting driver properties using the value of system properties see [Section 3.4.2, "Enabling Connection-based System Properties."](#)

1.5.6 Keep Connection After Local Transaction

This release provides a new attribute, `Keep Connection After Local Transaction`, that enables WebLogic Server to keep a physical database connection associated with a logical connection when committing or rolling back a local transaction. See `KeepConnAfterLocalTx` in the *Oracle WebLogic Server MBean Reference*.

Configuring WebLogic JDBC Resources

The following sections describe WebLogic JDBC resources, how they are configured, and how those resources apply to a WebLogic domain:

- [Section 2.1, "Understanding JDBC Resources in WebLogic Server"](#)
- [Section 2.2, "Ownership of Configured JDBC Resources"](#)
- [Section 2.3, "Data Source Configuration Files"](#)
- [Section 2.4, "JMX and WLST Access for JDBC Resources"](#)
- [Section 2.5, "Overview of Clustered JDBC Resources"](#)

2.1 Understanding JDBC Resources in WebLogic Server

In WebLogic Server, you can configure database connectivity by configuring JDBC data sources and multi data sources and then targeting or deploying the JDBC resources to servers or clusters in your WebLogic domain.

Each data source that you configure contains a pool of database connections that are created when the data source instance is created—when it is deployed or targeted, or at server startup. Applications lookup a data source on the JNDI tree or in the local application context (`java:comp/env`), depending on how you configure and deploy the object, and then request a database connection. When finished with the connection, the application calls `connection.close()`, which returns the connection to the connection pool in the data source. For more information about data sources in WebLogic Server, see [Chapter 3, "Configuring JDBC Data Sources"](#) and

A multi data source is an abstraction around a selected list of data sources that provides load balancing or failover processing between the data sources associated with the multi data source. Multi data sources are bound to the JNDI tree or local application context just like data sources are bound to the JNDI tree. Applications lookup a multi data source on the JNDI tree or in the local application context (`java:comp/env`) just like they do for data sources, and then request a database connection. The multi data source determines which data source to use to satisfy the request depending on the algorithm selected in the multi data source configuration: load balancing or failover. For more information about multi data sources, see [Chapter 5, "Configuring JDBC Multi Data Sources."](#)

2.2 Ownership of Configured JDBC Resources

A key to understanding WebLogic JDBC data source configuration and management is that *who* creates a JDBC resource or *how* a JDBC resource is created determines how a

resource is deployed and modified. Both WebLogic Administrators and programmers can create JDBC resources:

- WebLogic Administrators typically use the Administration Console or the WebLogic Scripting Tool (WLST) to create and deploy (target) JDBC modules. These JDBC modules are considered *system modules*. See [Section 2.3.1, "JDBC System Modules"](#) for more details.
- Programmers create modules in a development tool that supports creating an XML descriptor file, then package the JDBC modules with an application and pass the application to a WebLogic Administrator to deploy. These JDBC modules are considered *application modules*. See [Section 2.3.2, "JDBC Application Modules"](#) for more details.

[Table 2–1](#) lists the JDBC module types and how they can be configured and modified.

Table 2–1 JDBC Module Types and Configuration and Management Options

Module Type	Created with	Add/Remove Modules with Administration Console	Modify with JMX (remotely)	Modify with JSR-88 (non-remotely)	Modify with Administration Console
System	Administration Console or WLST	Yes	Yes	No	Yes—via JMX
Application	Oracle Enterprise Pack for Eclipse (OEPE), Oracle JDeveloper, another IDE, or an XML editor	No	No	Yes—via a deployment plan	Yes—via a deployment plan

2.3 Data Source Configuration Files

WebLogic JDBC data source configurations are stored in XML documents that conform to the `jdbc-data-source.xsd` schema (available at <http://xmlns.oracle.com/weblogic/jdbc-data-source/1.0/jdbc-data-source.xsd>).

You create and manage JDBC resources either as system modules or as application modules. JDBC application modules are a WebLogic-specific extension of Java EE modules and can be configured either within a Java EE application or as stand-alone modules.

Regardless of whether you are using JDBC system modules or JDBC application modules, each JDBC data source or multi data source is represented by an XML file (a module).

2.3.1 JDBC System Modules


When you create a JDBC resource (data source or multi data source) using the Administration Console or using the WebLogic Scripting Tool (WLST), WebLogic Server creates a JDBC module in the `config/jdbc` subdirectory of the domain directory, and adds a reference to the module in the domain's `config.xml` file. The JDBC module conforms to the `jdbc-data-source.xsd` schema (available at <http://xmlns.oracle.com/weblogic/jdbc-data-source/1.0/jdbc-data-source.xsd>).

JDBC resources that you configure this way are considered *system modules*. System modules are owned by an Administrator, who can delete, modify, or add similar resources at any time. System modules are globally available for targeting to servers and clusters configured in the domain, and therefore are available to all applications

deployed on the same targets and to client applications. System modules are also accessible through JMX as `JDBCSystemResourceMBeans`.

Data source system modules are included in the domain's `config.xml` file as a `JDBCSystemResource` element, which includes the name of the JDBC module file and the list of target servers and clusters on which the module is deployed. [Figure 2-1](#) shows an example of a data source listing in a `config.xml` file and the module that it maps to.

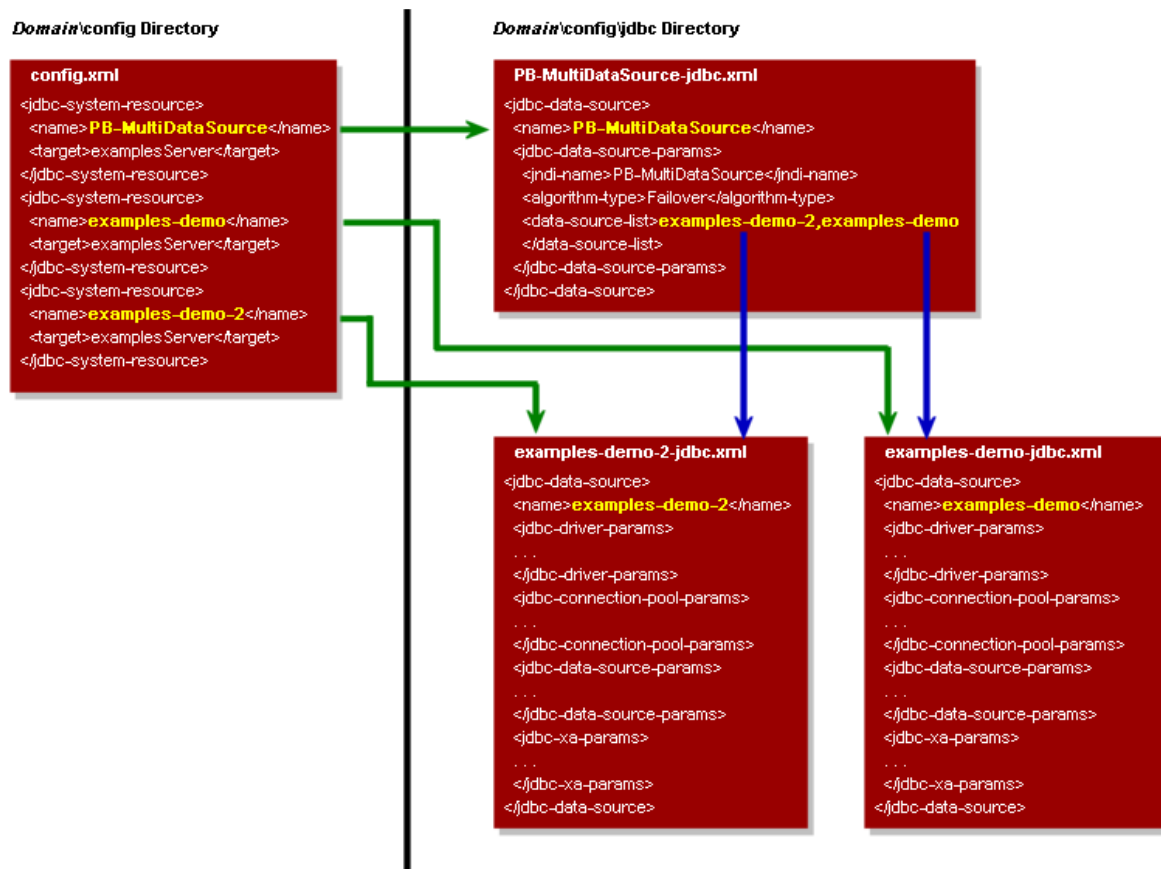
Figure 2-1 Reference from config.xml to a Data Source System Module



In this illustration, the `config.xml` file lists the `examples-demo` data source as a `jdbc-system-resource` element, which maps to the `examples-demo-jdbc.xml` module in the `domain\config\jdbc` folder.

Similarly, multi data source system modules are included in the domain's `config.xml` file as a `jdbc-system-resource` element. The multi data source module includes a `data-source-list` parameter that maps to the data source modules used by the multi data source. The individual data source modules are also included in the `config.xml` file. [Figure 2-2](#) shows the relationship between elements in the `config.xml` file and the system modules in the `config/jdbc` directory.

Figure 2–2 Reference from config.xml to Multi Data Source and Data Source System Modules



In this illustration, the `config.xml` file lists three JDBC modules—one multi data source and the two data sources used by the multi data source, which are also listed within the multi data source module. Your application can look up any of these modules on the JNDI tree and request a database connection. If you look up the multi data source, the multi data source determines which of the other data sources to use to supply the database connection, depending on the data sources in the `data-source-list` parameter, the order in which the data sources are listed, and the algorithm specified in the `algorithm-type` parameter. For more information about multi data sources, see [Chapter 5, "Configuring JDBC Multi Data Sources."](#)

2.3.2 JDBC Application Modules

JDBC resources can also be managed as application modules, similar to standard Java EE modules. A JDBC application module is simply an XML file that conforms to the `jdbc-data-source.xsd` schema (available at <http://xmlns.oracle.com/weblogic/jdbc-data-source/1.0/jdbc-data-source.xsd>) and represents a data source or a multi data source.

JDBC modules can be included as part of an Enterprise Application as a *packaged module*. Packaged modules are bundled with an EAR or exploded EAR directory, and are referenced in all appropriate deployment descriptors, such as the `weblogic-application.xml` and `ejb-jar.xml` deployment descriptors. The JDBC module is deployed along with the enterprise application, and can be configured to be available only to the enclosing application or to all applications. Using packaged modules ensures that an application always has access to required resources and simplifies the process of moving the application into new environments. With

packaged JDBC modules, you can migrate your application and the required JDBC configuration from environment to environment, such as from a testing environment to a production environment, without opening an EAR file and without extensive manual data source reconfiguration.

In contrast to system resource modules, JDBC modules that are packaged with an application are owned by the developer who created and packaged the module, rather than the Administrator who deploys the module. This means that the Administrator has more limited control over packaged modules. When deploying a resource module, an Administrator can change resource properties that were specified in the module, but the Administrator cannot add or delete modules. (As with other Java EE modules, deployment configuration changes for a resource module are stored in a deployment plan for the module, leaving the original module untouched.)

By definition, packaged JDBC modules are included in an enterprise application, and therefore are deployed when you deploy the enterprise application. For more information about deploying applications with packaged JDBC modules, see *Deploying Applications to Oracle WebLogic Server*.

A JDBC application module can also be deployed as a stand-alone resource using the `weblogic.Deployer` utility or the Administration Console, in which case the resource is typically available to the server or cluster targeted during the deployment process. JDBC resources deployed in this manner are called *stand-alone modules* and can be reconfigured using the Administration Console or a JSR-88 compliant tool, but are unavailable through JMX or WLST.

Stand-alone JDBC modules promote sharing and portability of JDBC resources. You can create a data source configuration and distribute it to other developers. Stand-alone JDBC modules can also be used to move data source configuration between domains, such as between the development domain and the staging domain.

Note: When deploying JDBC modules as standalone modules, a multi data source needs to have a deployment order that is greater than the deployment orders of its member data sources.

For more information about JDBC application modules, see [Appendix A, "Configuring JDBC Application Modules for Deployment."](#)

For information about deploying stand-alone JDBC modules, see "Deploying JDBC, JMS, WLDF Application Modules" in *Deploying Applications to Oracle WebLogic Server*."

2.3.3 JDBC Module File Naming Requirements

All WebLogic JDBC module files must end with the `-jdbc.xml` suffix, such as `examples-demo-jdbc.xml`. WebLogic Server checks the file name when you deploy the module. If the file does not end in `-jdbc.xml`, the deployment will fail and the server will not boot.

2.3.4 JDBC Modules in Versioned Applications

When you use production redeployment (versioning) to deploy a version of an application that includes a packaged JDBC module, WebLogic Server identifies the data source defined in the JDBC module with a name in the following format:

```
application_id#version_id@module_name@data_source_name
```

This name is used for data source run-time MBeans and for registering the data source instance with the WebLogic Server transaction manager.

If transactions in a retiring version of an application time out and the version of the application is then undeployed, you may have to manually resolve any pending or incomplete transactions on the data source in the retired version of the application. After a data source is undeployed (in this case, with the retired version of the application), the WebLogic Server transaction manager cannot recover pending or incomplete transactions.

For more information about production redeployment, see:

- "Developing Applications for Production Redeployment" in *Developing Applications for Oracle WebLogic Server*
- "Using Production Redeployment to Update Applications" in *Deploying Applications to Oracle WebLogic Server*

2.3.5 JDBC Schema

In support of the modular deployment model for JDBC resources in WebLogic Server, Oracle provides a schema for WebLogic JDBC objects: `weblogic-jdbc.xsd`. When you create JDBC resource modules (descriptors), the modules must conform to the schema. IDEs and other tools can validate JDBC resource modules based on the schema.

The schema is available at

<http://xmlns.oracle.com/weblogic/jdbc-data-source/1.0/jdbc-data-source.xsd>.

2.4 JMX and WLST Access for JDBC Resources

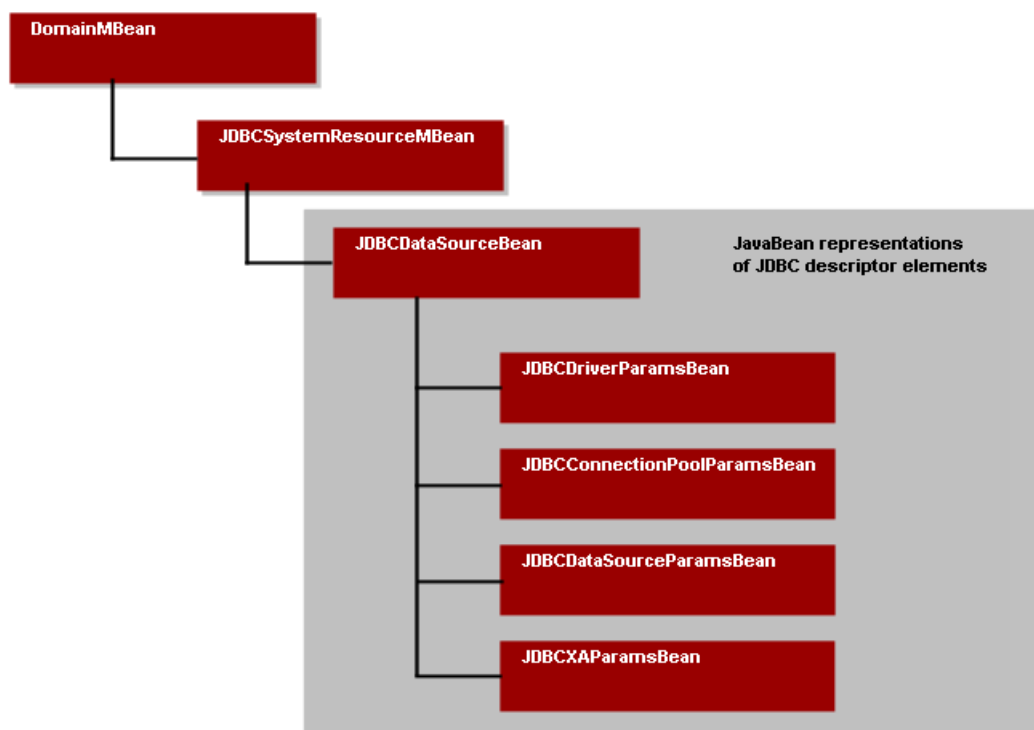
When you create JDBC resources using the Administration Console or WLST, WebLogic Server creates MBeans (Managed Beans) for each of the resources. You can then access these MBeans using JMX or the WebLogic Scripting Tool (WLST). See *Developing Custom Management Utilities With JMX for Oracle WebLogic Server* and *Oracle WebLogic Scripting Tool* for more information.

- [Section 2.4.1, "JDBC MBeans for System Resources"](#)
- [Section 2.4.2, "JDBC Management Objects in the Java EE Management Model \(JSR-77 Support\)"](#)
- [Section 2.4.3, "Using WLST to Create JDBC System Resources"](#)
- [Section 2.4.4, "How to Modify and Monitor JDBC Resources"](#)
- [Section 2.4.5, "Best Practices when Using WLST to Configure JDBC Resources"](#)

2.4.1 JDBC MBeans for System Resources

[Figure 2–3](#) shows the hierarchy of the MBeans for JDBC objects in a WebLogic domain.

Figure 2–3 JDBC Bean Tree



The `JDBCSystemResourceMBean` is a container for the JavaBeans created from a data source module. However, all JMX access for a JDBC data source is through the `JDBCSystemResourceMBean`. You cannot directly access the individual JavaBeans created from the data source module.

2.4.2 JDBC Management Objects in the Java EE Management Model (JSR-77 Support)

The WebLogic Server JDBC subsystem supports JSR-77, which defines the Java EE Management Model. The Java EE Management Model is used for monitoring the run-time state of a Java EE Web application server and its resources. You can access the Java EE Management Model to monitor resources, including the WebLogic JDBC subsystem as a whole, JDBC drivers loaded into memory, and JDBC data sources.

To comply with the specification, Oracle added the following run-time MBean types for the WebLogic JDBC subsystem:

- `JDBCServiceRuntimeMBean`—Which represents the JDBC subsystem and provides methods to access the list of `JDBCDriverRuntimeMBeans` and `JDBCDataSourceRuntimeMBeans` currently available in the system.
- `JDBCDriverRuntimeMBean`—Which represents a JDBC driver that the server loaded into memory.
- `JDBCDataSourceRuntimeMBeans`—Which represents a JDBC data source deployed on a server or cluster.

Note: WebLogic JDBC run-time MBeans do not implement the optional Statistics Provider interfaces specified by JSR-77.

For more information about using the Java EE management model with WebLogic Server, see *Monitoring and Managing With the Java EE Management APIs for Oracle WebLogic Server*.

2.4.3 Using WLST to Create JDBC System Resources

Basic tasks you need to perform when creating JDBC resources with the WLST are:

- Start an edit session.
- Create a JDBC system module that includes JDBC system resources, such as pools, data sources, multi data sources, and JDBC drivers.
- Target your JDBC system module.

Example 2-1 WLST Script to Create JDBC Resources

```
#-----
# Create JDBC Resources
#-----

import sys
from java.lang import System

print "### Starting the script ..."
global props

url = sys.argv[1]
usr = sys.argv[2]
password = sys.argv[3]

connect(usr,password, url)
edit()
startEdit()

servermb=getMBean("Servers/examplesServer")
if servermb is None:
    print '### No server MBean found'
else:
    def addJDBC(prefix):

        print("")
        print("**** Creating JDBC resources with property prefix " + prefix)

# Create the Connection Pool. The system resource will have
# generated name of <PoolName>+"-jdbc"

myResourceName = props.getProperty(prefix+"PoolName")
print("Here is the Resource Name: " + myResourceName)

jdbcSystemResource = wl.create(myResourceName, "JDBCSystemResource")
myFile = jdbcSystemResource.getDescriptorFileName()
print ("HERE IS THE JDBC FILE NAME: " + myFile)

jdbcResource = jdbcSystemResource.getJDBCResource()
jdbcResource.setName(props.getProperty(prefix+"PoolName"))

# Create the DataSource Params
dpBean = jdbcResource.getJDBCDataSourceParams()
myName=props.getProperty(prefix+"JNDIName")
dpBean.setJNDINames([myName])
```



```

# Create the Driver Params
drBean = jdbcResource.getJDBCDriverParams()
drBean.setPassword(props.getProperty(prefix+"Password"))
drBean.setUrl(props.getProperty(prefix+"URLName"))
drBean.setDriverName(props.getProperty(prefix+"DriverName"))

propBean = drBean.getProperties()
driverProps = Properties()
driverProps.setProperty("user",props.getProperty(prefix+"UserName"))

e = driverProps.propertyNames()
while e.hasMoreElements() :
    propName = e.nextElement()
    myBean = propBean.createProperty(propName)
    myBean.setValue(driverProps.getProperty(propName))

# Create the ConnectionPool Params
ppBean = jdbcResource.getJDBCConnectionPoolParams()
ppBean.setInitialCapacity(int(props.getProperty(prefix+"InitialCapacity")))
ppBean.setMaxCapacity(int(props.getProperty(prefix+"MaxCapacity")))
ppBean.setCapacityIncrement(int(props.getProperty(prefix+"CapacityIncrement")))

if not props.getProperty(prefix+"ShrinkPeriodMinutes") == None:
    ppBean.setShrinkFrequencySeconds(int(props.getProperty(prefix+"ShrinkPeriodMinutes")))
if not props.getProperty(prefix+"TestTableName") == None:
    ppBean.setTestTableName(props.getProperty(prefix+"TestTableName"))

if not props.getProperty(prefix+"LoginDelaySeconds") == None:
    ppBean.setLoginDelaySeconds(int(props.getProperty(prefix+"LoginDelaySeconds")))

# Adding KeepXaConnTillTxComplete to help with in-doubt transactions.
xaParams = jdbcResource.getJDBCXAParams()
xaParams.setKeepXaConnTillTxComplete(1)

# Add Target
jdbcSystemResource.addTarget(wl.getMBean("/Servers/examplesServer"))
.
.
.

```

2.4.4 How to Modify and Monitor JDBC Resources

You can modify or monitor JDBC objects and attributes by using the appropriate method available from the MBean.

- You can modify JDBC objects and attributes using the `set`, `target`, `untarget`, and `delete` methods.
- You can monitor JDBC run-time objects using `get` methods.

For more information, see "Navigating MBeans (WLST Online)" in *Oracle WebLogic Scripting Tool*.

2.4.5 Best Practices when Using WLST to Configure JDBC Resources

This section provides best practices information when using WLST to configure JDBC resources:

- Trap for Null MBean objects (such as pools, data sources, drivers) before trying to manipulate the MBean object.

2.5 Overview of Clustered JDBC Resources

You can target or deploy JDBC resources to a cluster to improve the availability of cluster-hosted applications. For information about JDBC objects in a clustered environment, see "JDBC Connections" in *Using Clusters for Oracle WebLogic Server*.

Multi data sources are supported for use in clusters. However, note that multi data sources can only use data sources in the same JVM. Multi data sources cannot use data sources from other cluster members.

Configuring JDBC Data Sources

This section includes the following information:

- [Section 3.1, "Understanding JDBC Data Sources"](#)
- [Section 3.2, "Types of WebLogic Server JDBC Data Sources"](#)
- [Section 3.3, "Creating a JDBC Data Source"](#)
- [Section 3.4, "Configuring Connection Pool Features"](#)
- [Section 3.5, "Configuring Oracle Parameters"](#)
- [Section 3.6, "Configuring an ONS Client"](#)
- [Section 3.7, "Tuning Generic Data Source Connection Pools"](#)
- [Section 3.8, "Setting Database Security Credentials"](#)
- [Section 3.9, "JDBC Data Source Factories \(Deprecated\)"](#)

3.1 Understanding JDBC Data Sources

In WebLogic Server, you configure database connectivity by adding data sources to your WebLogic domain. WebLogic JDBC data sources provide database access and database connection management. Each data source contains a pool of database connections that are created when the data source is created and at server startup. Applications reserve a database connection from the data source by looking up the data source on the JNDI tree or in the local application context and then calling `getConnection()`. When finished with the connection, the application should call `connection.close()` as early as possible, which returns the database connection to the pool for other applications to use.

3.2 Types of WebLogic Server JDBC Data Sources

WebLogic Server provides three types of data sources:

- **Generic Data Sources**—Generic data sources and their connection pools provide connection management processes that help keep your system running efficiently. You can set options in the data source to suit your applications and your environment.
- **GridLink Data Sources**—An event-based data source that adaptively responds to state changes in an Oracle RAC instance. See [Section 4, "Using GridLink Data Sources."](#)

- Multi data sources—A *multi data source* is an abstraction around a group of generic data sources that provides load balancing or failover processing. See [Section 5, "Configuring JDBC Multi Data Sources."](#)

3.3 Creating a JDBC Data Source

You can create JDBC data sources in your WebLogic domain using the Administration Console or the WebLogic Scripting Tool (WLST):

- "Create a JDBC Data Source" in the *Oracle WebLogic Server Administration Console Help*.
- The sample WLST script `SAMPLES_HOME\server\examples\src\examples\wlst\online\jdbc_data_source_creation.py`, where `SAMPLES_HOME` refers to the main examples directory of your WebLogic Server installation. See "WLST Online Sample Scripts" in *Oracle WebLogic Scripting Tool*

The following sections provide an overview of the basics steps used in the data source configuration wizard to create a data source using the Administration console:

- [Section 3.3.1, "JDBC Data Source Properties"](#)
- [Section 3.3.2, "Configure Transaction Options"](#)
- [Section 3.3.3, "Configure Connection Properties"](#)
- [Section 3.3.4, "Test Connections"](#)
- [Section 3.3.5, "Target the Data Source"](#)

3.3.1 JDBC Data Source Properties

JDBC Data Source Properties include options that determine the identity of the data source and the way the data is handled on a database connection.

3.3.1.1 Data Source Names

JDBC data source names are used to identify the data source within the WebLogic domain. For system resource data sources, names must be unique among all other JDBC system resources, including data sources and multi data sources. To avoid naming conflicts, data source names should also be unique among other configuration object names, such as servers, clusters, and JMS queues, topics, and servers. For JDBC application modules scoped to an application, data source names must be unique among JDBC data sources and multi data sources that are similarly scoped.

3.3.1.2 JNDI Names

You can configure a data source so that it binds to the JNDI tree with a single or multiple names. You can use a multi-JNDI-named data source in place of legacy configurations that included multiple data sources that pointed to a single JDBC connection pool. For more information, see `rss`

3.3.1.3 Selecting a Database Type

Select a DBMS. For information about supported databases, see "Supported Database Configurations" in *System Requirements and Supported Platforms for Oracle WebLogic Server* at http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html

3.3.1.4 Selecting a JDBC Driver

When creating a JDBC data source using the Administration Console, you are prompted to select a JDBC driver class. The Administration Console provides most of the more common driver class names and in most cases tries to help you construct the URL as required by the driver. You should verify, however, that the URL is as you want it before asking the console to test it. The driver you select must be in the classpath on all servers on which you intend to deploy the data source. Some but not all JDBC drivers listed in the Administration Console are shipped (and/or are already in the classpath) with WebLogic Server:

- Oracle Thin Driver
 - Oracle Thin Driver XA
 - Oracle Thin Driver non-XA
- Third-party JDBC drivers (see [Chapter 12, "Using JDBC Drivers with WebLogic Server"](#)):
 - MySQL (non-XA)
- Oracle Type 4 JDBC Drivers from DataDirect for the following database management systems (see in *Type 4 JDBC Drivers for Oracle WebLogic Server*):
 - DB2
 - Informix
 - Microsoft SQL Server
 - Sybase

All of these drivers are referenced by the `weblogic.jar` manifest file and do not need to be explicitly defined in a server's classpath.

When deciding which JDBC driver to use to connect to a database, you should try drivers from various vendors in your environment. In general, JDBC driver performance is dependent on many factors, especially the SQL code used in applications and the JDBC driver implementation.

For information about supported JDBC drivers, see "Supported Database Configurations" in *System Requirements and Supported Platforms for Oracle WebLogic Server* at

http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html.

Note: JDBC drivers listed in the Administration Console when creating a data source are not necessarily certified for use with WebLogic Server. JDBC drivers are listed as a convenience to help you create connections to many of the database management systems available.

You must install JDBC drivers in order to use them to create database connections in a data source on each server on which the data source is deployed. Drivers are listed in the Administration Console with known required configuration options to help you configure a data source. The JDBC drivers in the list are not necessarily installed. Driver installation can include setting system Path, Classpath, and other environment variables. See [Section 12.2, "Setting the Environment for a Type-4 Third-Party JDBC Driver."](#) When a JDBC driver is updated, configuration requirements may change. The Administration Console uses known configuration requirements at the time the WebLogic Server software was released. If configuration options for your JDBC driver have changed, you may need to manually override the configuration options when creating the data source or in the property pages for the data source after it is created.

3.3.2 Configure Transaction Options

When you configure a JDBC data source using the Administration Console, WebLogic Server automatically selects specific transaction options based on the type of JDBC driver:

- **For XA drivers**, the system automatically selects the **Two-Phase Commit** protocol for global transaction processing.
- **For non-XA drivers**, local transactions are supported by definition, and WebLogic Server offers the following options

Supports Global Transactions: (selected by default) Select this option if you want to use connections from the data source in global transactions, even though you have not selected an XA driver. See [Section 7.1, "Enabling Support for Global Transactions with a Non-XA JDBC Driver"](#) for more information.

When you select Supports Global Transactions, you must also select the protocol for WebLogic Server to use for the transaction branch when processing a global transaction:

- **Logging Last Resource:** With this option, the transaction branch in which the connection is used is processed as the last resource in the transaction and is processed as a local transaction. Commit records for two-phase commit (2PC) transactions are inserted in a table on the resource itself, and the result determines the success or failure of the prepare phase of the global transaction. This option offers some performance benefits and greater data safety than Emulate Two-Phase Commit, but it has some limitations. See [Section 7.2, "Understanding the Logging Last Resource Transaction Option."](#)

Note: Logging Last Resource is not supported for data sources used by a multi data source except when used with Oracle RAC version 10G Release 2 (10GR2) and greater versions as described in [Section 7.2.4, "Administrative Considerations and Limitations for LLR Data Sources."](#)

- **Emulate Two-Phase Commit:** With this option, the transaction branch in which the connection is used always returns success for the prepare phase of the transaction. It offers performance benefits, but also has risks to data in some failure conditions. Select this option only if your application can tolerate heuristic conditions. See [Section 7.3, "Understanding the Emulate Two-Phase Commit Transaction Option."](#)
- **One-Phase Commit:** (selected by default) With this option, a connection from the data source can be the only participant in the global transaction and the transaction is completed using a one-phase commit optimization. If more than one resource participates in the transaction, an exception is thrown when the transaction manager calls `XAResource.prepare` on the 1PC resource.

For more information on configuring transaction support for a data source, see [Section 7, "JDBC Data Source Transaction Options."](#)

3.3.3 Configure Connection Properties

Connection Properties are used to configure the connection between the data source and the DBMS. Typical attributes are the database name, host name, port number, user name, and password.

Note: You can use a Single Client Access Name (SCAN) address to represent the host name.

3.3.4 Test Connections

Test Database Connection allows you to test a database connection before the data source configuration is finalized using a table name or SQL statement. If necessary, you can test additional configuration information using the `Properties` and `System Properties` attributes.

3.3.5 Target the Data Source

You can select one or more targets to deploy your new JDBC data source. If you don't select a target, the data source will be created but not deployed. You will need to deploy the data source at a later time.

3.4 Configuring Connection Pool Features

Each JDBC data source has a pool of JDBC connections that are created when the data source is deployed or at server startup. Applications use a connection from the pool then return it when finished using the connection. Connection pooling enhances performance by eliminating the costly task of creating database connections for the application.

The following sections include information about connection pool options for a JDBC data source.

- [Section 3.4.1, "Enabling JDBC Driver-Level Features"](#)
- [Section 3.4.2, "Enabling Connection-based System Properties."](#)
- [Section 3.4.3, "Initializing Database Connections with SQL Code"](#)

You can see more information and set these and other related options through the:

- **JDBC Data Source: Configuration: Connection Pool** page in the Administration Console. See "JDBC Data Source: Configuration: Connection Pool" in the *Oracle WebLogic Server Administration Console Help*
- `JDBCConnectionPoolParamsBean`, which is a child MBean of the `JDBCDataSourceBean`

3.4.1 Enabling JDBC Driver-Level Features

WebLogic JDBC data sources support the `javax.sql.ConnectionPoolDataSource` interface implemented by JDBC drivers. You can enable driver-level features by adding the property and its value to the `Properties` attribute in a JDBC data source. Driver-level properties in the `Properties` attribute are set on the driver's `ConnectionPoolDataSource` object.

3.4.2 Enabling Connection-based System Properties

WebLogic JDBC data sources support setting driver properties using the value of system properties. The value of each property is derived at runtime from the named system property. You can configure connection-based system properties using the Administration Console by editing the `System Properties` attribute of your data source configuration.

3.4.3 Initializing Database Connections with SQL Code

When WebLogic Server creates database connections in a data source, the server can automatically run SQL code to initialize the database connection. To enable this feature, enter `SQL` followed by a space and the SQL code you want to run in the `Init SQL` attribute on the **JDBC Data Source: Configuration: Connection Pool** page in the Administration Console. If you leave this attribute blank (the default), WebLogic Server does not run any code to initialize database connections.

WebLogic Server runs this code whenever it creates a database connection for the data source, which includes at server startup, when expanding the connection pool, and when refreshing a connection.

You can use this feature to set DBMS-specific operational settings that are connection-specific or to ensure that a connection has memory or permissions to perform required actions.

Start the code with `SQL` followed by a space. For example:

```
SQL alter session set NLS_DATE_FORMAT='YYYY-MM-DD HH24:MI:SS'
```

or

```
SQL SET LOCK MODE TO WAIT
```

Options that you can set using `InitSQL` vary by DBMS.

Note: `Init SQL` is not a dynamic attribute. When you change the value for `Init SQL`, you must either undeploy and redeploy the data source or restart the server.

3.5 Configuring Oracle Parameters

WebLogic Server provides several attributes that provide improved Data Source performance when using Oracle drivers, for more information, see [Section 6, "Advanced Configuration for Oracle Drivers."](#)

3.6 Configuring an ONS Client

Configuring an ONS client changes a generic data source to a GridLink data source. For more detailed configuration information and additional environment requirements, see [Section 4, "Using GridLink Data Sources."](#)

3.7 Tuning Generic Data Source Connection Pools

By properly configuring the connection pool attributes in JDBC data sources in your WebLogic Server domain, you can improve application and system performance. For more information, see [Section 10, "Tuning Data Source Connection Pools."](#)

3.8 Setting Database Security Credentials

The following sections provide information on how to pass security credentials to a DBMS:

- [Section 3.8.1, "Types of Data Source Pools"](#)
- [Section 3.8.2, "Using a User Name/Password"](#)
- [Section 3.8.3, "Set Client ID On Connection"](#)
- [Section 3.8.4, "Identity-based Connection Pooling"](#)

3.8.1 Types of Data Source Pools

Weblogic Server provides two types of data source pools based on security privileges:

- Homogeneous—Regardless of the end user of the application, all connections in the pool use the same security credentials to access the DBMS.
- Heterogeneous—Allows applications to use a JDBC connection with a specific DBMS credential by pooling physical connections with different DBMS credentials.

This section compares methods of passing security credentials to a DBMS.

Table 3–1 Comparing Methods of Passing Security Credentials

Method	Type of Connection Pool
Using a User Name/Password	Homogeneous pool of connections.
Set Client ID On Connection	Homogeneous pool of connections.
Identity-based Connection Pooling	Heterogeneous pool of connections.

3.8.2 Using a User Name/Password

The simplest type of credential is to provide the connection pool a user account name and password for the DBMS. All the connections in the pool then use the same credentials to access a DBMS. See "Create JDBC data sources" in *Oracle WebLogic Server Administration Console Help*.

Note: You can enter the password as a name-value pair in the Properties field (not permitted for production environments) or you can enter it in the Password field. The value in the Password field overrides any password value defined in the Properties passed to the JDBC Driver when creating physical database connections. Oracle recommends that you use the Password attribute in place of the password property in the properties string because the Password value is encrypted in the configuration file (stored as the password-encrypted attribute in the `jdbc-driver-params` tag in the module file) and is hidden in the administration console.

3.8.3 Set Client ID On Connection

If the `Set Client ID On Connection` attribute is enabled on the data source, when an application requests a database connection from the data source, the WebLogic Server instance determines the current WebLogic user ID and then sets the mapped database ID as a light-weight client ID. All the connections in the pool have the same credentials to access a DBMS. Basic configuration steps are:

1. Select `Set Client ID On Connection`, see "Enable Set Client ID On Connection for a JDBC data source" in *Oracle WebLogic Server Administration Console Help*.

Note: Credential mapping to map the WebLogic user ID and the database ID is only supported on the Oracle database with the Oracle Thin driver. This feature is not supported with the Oracle DMS driver.

2. Map the WebLogic user ID and the database ID. See "Configure credential mapping for a JDBC data source" in the *Oracle WebLogic Server Administration Console Help*.

This feature relies on features in the JDBC driver and DBMS. It is only supported for use with Oracle and DB2 databases using a vendor extension method:

- `oracle.jdbc.OracleConnection.setClientIdentifier(String id)`
- `com.ibm.db2.jcc.DB2Connection.setDB2ClientUser(String user)`

Note: `Set Client ID On Connection` and `Enable Identity Based Connection Pooling` are mutually exclusive. If you think you need both mechanisms to pass security credentials in your application environment, create separate data sources—one for with `Set Client ID On Connection` and one with `Enable Identity Based Connection Pooling`.

3.8.4 Identity-based Connection Pooling

Identity-based connection pooling allows applications to use a JDBC connection with a specific DBMS credential by pooling physical connections with different DBMS credentials.

If the `Enable Identity Based Connection Pooling` attribute is enabled on the data source, when an application requests a database connection, the WebLogic Server instance selects an existing physical connection or creates a new physical connection with requested DBMS identity based on a map of WebLogic user credentials and DBMS credentials. Basic configuration steps are:

1. Select `Enable Identity Based Connection Pooling`, see "Enable identity-based connection pooling for a JDBC data source" in *Oracle WebLogic Server Administration Console Help*.
2. Map WebLogic user credentials and DBMS credentials. See "Configure credential mapping for a JDBC data source" in the *Oracle WebLogic Server Administration Console Help*.

Note: Set Client ID On Connection and Enable Identity Based Connection Pooling are mutually exclusive. If you think you need both mechanisms to pass security credentials in your application environment, create separate data sources—one for with Set Client ID On Connection and one with Enable Identity Based Connection Pooling.

3.8.4.1 How Heterogeneous Connections are Created

The following section provides information on how heterogeneous connections are created:

1. At connection pool initialization, the physical JDBC connections are created with the default DBMS credential of the data source.
2. An application tries to get a connection from a data source.
3. The current server instance credential is mapped to a DBMS credential. See "Configure credential mapping for a JDBC data source" in the *Oracle WebLogic Server Administration Console Help*.
 - If no match is found, the default DBMS credential is used.

Note: The default DBMS credential should have minimum DBMS privileges, such as the ability to execute XA transactions and perform connection test operations.

- If a match is found, it is used to find physical connections matching the DBMS credential.
 - If a match is found, the connection is reserved and returned to the application.
 - If no match is found, a connection is created or reused based on the maximum capacity of the pool:
 - * If the maximum capacity has not been reached, a new connection is created with the DBMS credential, reserved, and returned to the application.
 - * If the pool has reached maximum capacity, based on the least recently used (LRU) algorithm, a physical connection is selected from the pool and destroyed. A new connection is created with the DBMS credential, reserved, and returned to the application.

Regardless of how physical connections are created, each physical connection in the pool has its own DBMS credential information maintained by the pool. Once a physical connection is reserved by the pool, it does not change its DBMS credential even if the current thread changes its WebLogic user credential and continues to use the same connection.

3.8.4.2 Using Identity-based Pooling with Global Transactions

When executing inside a global transaction, an application may change the credential on the current thread and get multiple JDBC connections under different credentials. However, the Identity-based Pooling feature maps multiple logical JDBC connections of a WebLogic JDBC data source inside of a global transaction into a single physical JDBC connection. This means that only one DBMS credential per WebLogic JDBC data source per WebLogic server instance is honored for a global transaction.

3.8.4.3 Using Identity-based Pooling with LLR

You must make the following changes to use Logging Last Resource (LLR) transaction optimization with Identity-based Pooling:

- You must configure a custom schema for LLR using a fully qualified LLR table name. All LLR connections will then use the named schema rather than the default schema when accessing the LLR transaction table.
- Use database specific administration tools to grant permission to access the named LLR table to all users. By default, the LLR table is created during boot by the user configured for the connection in the data source. In most cases, the database will only allow access to this user and not allow access to mapped users.

3.9 JDBC Data Source Factories (Deprecated)

In previous releases of WebLogic Server, application-scoped JDBC connection pools relied on JDBC data source factories to provide default connection pool values. JDBC data source factories are deprecated in WebLogic Server 9.2 and are included in the release for backward compatibility only. Application-scoped JDBC connection pools are replaced by JDBC application modules. For more information, see [Section A.2.7, "Application Scoping for a Packaged JDBC Module."](#)

Using GridLink Data Sources


This section includes the following information:

- [Section 4.1, "What is a GridLink Data Source"](#)
- [Section 4.2, "Creating a GridLink Data Source"](#)
- [Section 4.3, "Enable Socket Direct Protocol for a Grid Link Data Source"](#)
- [Section 4.4, "Configuring Connection Pool Features"](#)
- [Section 4.5, "Configuring Oracle Parameters"](#)
- [Section 4.6, "Configuring an ONS Client"](#)
- [Section 4.7, "Configuring Oracle Parameters"](#)
- [Section 4.8, "Tuning GridLink Data Source Connection Pools"](#)
- [Section 4.9, "Setting Database Security Credentials"](#)
- [Section 4.10, "Monitoring GridLink JDBC Resources"](#)

4.1 What is a GridLink Data Source

A single GridLink data source provides connectivity between WebLogic Server and an Oracle Database service targeted to an Oracle RAC cluster. It uses the Oracle Notification Service (ONS) to adaptively respond to state changes in an Oracle RAC instance. An Oracle Database service represents a workload with common attributes that enables administrators to manage the workload as a single entity. You scale the number of GridLink data sources as the number of services increases in the data base, independent of the number of nodes in the cluster.

Figure 4–1 GridLink Data Source Connectivity




A GridLink data source includes the features of generic data sources plus the following support for Oracle RAC:

- [Section 4.1.1, "Fast Connection Failover"](#)
- [Section 4.1.2, "Runtime Connection Load Balancing"](#)
- [Section 4.1.3, "Graceful Handling for Oracle RAC Outages"](#)
- [Section 4.1.4, "XA Affinity"](#)
- [Section 4.1.5, "SCAN Addresses"](#)
- [Section 4.1.6, "Secure Communication using Oracle Wallet."](#)

4.1.1 Fast Connection Failover

A GridLink data source uses Fast Connection Failover and responds to Oracle RAC events using ONS. This ensures that the connection pool in the GridLink data source contains valid connections (including reserved connections) without the need to poll and test connections.

Figure 4–2 Fast Connection Failover

A GridLink data source uses Fast Connection Failover to:

- Provide rapid failure detection.
- Abort and remove invalid connections from the connection pool.
- Perform graceful shutdown for planned and unplanned Oracle RAC node outages. See [Section 4.1.3, "Graceful Handling for Oracle RAC Outages."](#)
- Adapt to changes in topology, such as adding or removing a node.
- Distribute runtime work requests to all active Oracle RAC instances, including those rejoining a cluster.

See Fast Connection Failover in the *Oracle Database JDBC Developer's Guide and Reference*.


4.1.2 Runtime Connection Load Balancing

GridLink data sources provide load balancing in XA and non-XA environments. GridLink data sources use runtime connection load balancing to distribute connections to Oracle RAC instances based on Oracle FAN events issued by the database. This simplifies data source configuration and improves performance as the database drives load balancing of connections through the GridLink data source, independent of the database topology.

Runtime Connection Load Balancing allows WebLogic Server to:

- Adjust the distribution of work based on back end node capacities such as CPU, availability, and response time.
- React to changes in Oracle RAC topology.
- Manage pooled connections for high performance and scalability.

Figure 4–3 Runtime Connection Load Balancing



If FAN is not enabled, GridLink data sources use a round-robin load balancing algorithm to allocate connections to Oracle RAC nodes.

4.1.3 Graceful Handling for Oracle RAC Outages

A GridLink data source provides graceful handling for the planned and unplanned shutdown of an Oracle RAC service:

- For planned shutdowns, the data source allows in-progress transactions to complete before closing connections. New Requests are load balanced to active Oracle RAC instances.
- For unplanned shutdowns, the data source rolls back in-progress transactions and closes the connections. New Requests are load balanced to active Oracle RAC instances.


4.1.3.1 Handling for Oracle RAC Outages Prior to Oracle RAC 11.2

In Releases prior to Oracle RAC 11.2, manually shutting down an Oracle RAC instance without first shutting down the corresponding services results in an unplanned shutdown.

4.1.4 XA Affinity

XA Affinity for global transactions ensures all the data base operations for a global transaction performed on an Oracle RAC cluster are directed to the same Oracle RAC instance. The first connection request for an XA transaction is load balanced using RCLB and is assigned an Affinity context. All subsequent connection requests are routed to the same Oracle RAC instance using the Affinity context of the first connection.

Figure 4–4 XA Affinity



4.1.5 SCAN Addresses

Oracle Single Client Access Name (SCAN) addresses can be used to specify the host and port for both the TNS listener and the ONS listener in the WebLogic console. A GridLink data source containing SCAN addresses does not need to change if you add or remove Oracle RAC nodes. Contact your network administrator for appropriately configured SCAN urls for your environment. For more information, see <http://www.oracle.com/technetwork/database/clustering/overview/s-can-129069.pdf>.

4.1.6 Secure Communication using Oracle Wallet

Allows you to configure secure communication with the ONS listener using Oracle Wallet. See [Section 4.2.5.1, "Secure ONS Client Communication."](#)

4.2 Creating a GridLink Data Source

To create a GridLink data source in your WebLogic domain, you can use the Administration Console or the WebLogic Scripting Tool (WLST).

See the following for more information:

- "Create a JDBC GridLink Data Source" in the *Oracle WebLogic Server Administration Console Help*.
- The sample WLST script `SAMPLES_HOME\server\examples\src\examples\wlst\online\jdbc_data_source_creation.py`, where `SAMPLES_HOME` refers to the main examples directory of your WebLogic Server installation. This example creates a generic data source. See "WLST Online Sample Scripts" in *Oracle WebLogic Scripting Tool*.

The following sections provide an overview of the basics steps used in the data source configuration wizard to create a data source using the Administration console:

- [Section 4.2.1, "JDBC Data Source Properties"](#)
- [Section 4.2.2, "Configure Transaction Options"](#)
- [Section 4.2.3, "Configure Connection Properties"](#)

- [Section 4.2.4, "Test Connections"](#)
- [Section 4.2.5, "Configure an ONS Client Configuration"](#)
- [Section 4.2.6, "Test ONS Client Configuration"](#)
- [Section 4.2.7, "Target the Data Source"](#)

4.2.1 JDBC Data Source Properties

JDBC Data Source Properties include options that determine the identity of the data source and the way the data is handled on a database connection.

4.2.1.1 Data Source Names

JDBC data source names are used to identify the data source within the WebLogic domain. For system resource data sources, names must be unique among all other JDBC system resources, including data sources and multi data sources. To avoid naming conflicts, data source names should also be unique among other configuration object names, such as servers, clusters, and JMS queues, topics, and servers. For JDBC application modules scoped to an application, data source names must be unique among JDBC data sources and multi data sources that are similarly scoped.

4.2.1.2 JNDI Names

You can configure a data source so that it binds to the JNDI tree with a single or multiple names. You can use a multi-JNDI-named data source in place of legacy configurations that included multiple data sources that pointed to a single JDBC connection pool. For more information, see [rss](#)

4.2.1.3 Select an XA or Non-XA Driver

Specify how your data source handles global transactions.

4.2.2 Configure Transaction Options

When you configure a JDBC data source using the Administration Console, WebLogic Server automatically selects specific transaction options based on the type of JDBC driver:

- **For XA drivers**, the system automatically selects the **Two-Phase Commit** protocol for global transaction processing.
- **For non-XA drivers**, local transactions are supported by definition, and WebLogic Server offers the following options

Supports Global Transactions: (selected by default) Select this option if you want to use connections from the data source in global transactions, even though you have not selected an XA driver. See [Section 7.1, "Enabling Support for Global Transactions with a Non-XA JDBC Driver"](#) for more information.

When you select Supports Global Transactions, you must also select the protocol for WebLogic Server to use for the transaction branch when processing a global transaction:

- **Logging Last Resource:** With this option, the transaction branch in which the connection is used is processed as the last resource in the transaction and is processed as a local transaction. Commit records for two-phase commit (2PC) transactions are inserted in a table on the resource itself, and the result determines the success or failure of the prepare phase of the global transaction. This option offers some performance benefits and greater data

safety than Emulate Two-Phase Commit, but it has some limitations. See [Section 7.2, "Understanding the Logging Last Resource Transaction Option."](#)

- **Emulate Two-Phase Commit:** With this option, the transaction branch in which the connection is used always returns success for the prepare phase of the transaction. It offers performance benefits, but also has risks to data in some failure conditions. Select this option only if your application can tolerate heuristic conditions. See [Section 7.3, "Understanding the Emulate Two-Phase Commit Transaction Option."](#)
- **One-Phase Commit:** (selected by default) With this option, a connection from the data source can be the only participant in the global transaction and the transaction is completed using a one-phase commit optimization. If more than one resource participates in the transaction, an exception is thrown when the transaction manager calls `XAResource.prepare` on the 1PC resource.

For more information on configuring transaction support for a data source, see [Section 7, "JDBC Data Source Transaction Options."](#)

4.2.3 Configure Connection Properties

Connection Properties are used to configure the connection between the data source and the DBMS. Typical attributes are the service name, database name, host name, port number, user name, and password. The console allows you to enter connection properties in one of the following ways:

- [Section 4.2.3.1, "Enter Connection Properties"](#)
- [Section 4.2.3.2, "Enter a Complete URL"](#)

4.2.3.1 Enter Connection Properties

On the **One the GridLink data source connection Properties Options** page, select **Enter individual listener information** and click **Next**. Enter the connection properties. For example:

- Enter **myService** in `Service Name`.
- Enter **left:1234, center:1234, right:1234** in the `Host and Port`:. Separate the host and port of each listener with colon.
- Enter **myDataBase** in `Database User Name`.
- Enter **myPassword1** in `Password`.

The console automatically generates the complete JDBC URL. For example:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=left)(PORT=1234))(ADDRESS=(PRO
TOCOL=TCP)(HOST=right)(PORT=1234))(ADDRESS=(PROTOCOL=TCP)(HOST=c
enter)(PORT=1234)))(CONNECT_DATA=(SERVICE_NAME=myService)))
```

4.2.3.2 Enter a Complete URL

On the **One the GridLink data source connection Properties Options** page, select **Enter complete JDBC URL** and click **Next**. Enter the connection properties. For example:

- In **Complete JDBC URL**, enter the JDBC URL. For example:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=left)(PORT=1234))(ADDRESS=(
PROTOCOL=TCP)(HOST=right)(PORT=1234))(ADDRESS=(PROTOCOL=TCP)(
```

```
HOST=center) (PORT=1234)) (CONNECT_DATA=(SERVICE_
NAME=myService))
```

You can also use a SCAN address. For example:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=MyScanAddr-scn.myCompany.co
m) (PORT=1234)) (CONNECT_DATA=(SERVICE_NAME=myService)))
```

- Enter **myDataBase** in Database User Name.
- Enter **myPassword1** in Password.

4.2.4 Test Connections

Test Database Connection allows you to test a database connection before the data source configuration is finalized using a table name or SQL statement. If necessary, you can test additional configuration information using the `Properties` and `System Properties` attributes.

4.2.5 Configure an ONS Client Configuration

ONS client configuration allows the data source to subscribe to and process Oracle FAN events. To configure an ONS client:

- Select `Fan Enabled`.
- In **ONS host and port**, enter a comma-separated list of ONS daemon listen addresses and ports for receiving ONS-based FAN events. You can use Single Client Access Name (SCAN) addresses to access FAN notifications.
- Optionally, configure secure ONS client communication using SSL. See [Section 4.2.5.1, "Secure ONS Client Communication."](#)

4.2.5.1 Secure ONS Client Communication

To use an Oracle Wallet file with WebLogic Server, you must:

- Add the following files to the WebLogic Server classpath:
 - `$MW_HOME/modules/com.oracle.osdt_cert_1.0.0.0.jar`
 - `$MW_HOME/modules/com.oracle.osdt_core_1.0.0.0.jar`
 - `$MW_HOME/modules/com.oracle.oraclepki_1.0.0.0.jar`
- Update your GridLink Data Source configuration to include the directory of the Oracle wallet file in which the SSL certificates are stored and the ONS Wallet password. See *Secure ONS Listener using Oracle Wallet* in *Oracle WebLogic Server Administration Console Help*.
- For more information on Oracle Wallet, see the *Database Advanced Security Administrator's Guide*.

4.2.6 Test ONS Client Configuration

Test ONS client configuration allows you to test a database connection before the data source configuration is finalized.

4.2.7 Target the Data Source

You can select one or more targets to deploy your new JDBC GridLink data source. If you don't select a target, the data source will be created but not deployed. You will need to deploy the data source at a later time.

4.3 Enable Socket Direct Protocol for a Grid Link Data Source

Note: To use Socket Direct Protocol (SDP), your database network must be configured to use Infiniband. See *Configuring SDP Support for InfiniBand Connections* in the *Oracle Database Net Services Administrator's Guide*.

To enable SDP for an existing Grid Link data source, use the following steps:

- Add system property `-Djava.net.preferIPv4Stack=true` to the server startup command line.
- To enable SPD for a Grid Link data source using the Administration Console:
 1. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
 2. In the **Domain Structure** tree, expand **Services**, then select **Data Sources**.
 3. On the Summary of Data Sources page, click the data source name.
 4. Select the **Configuration: Connection Pool** tab.
 5. In **Url**, edit the url, replacing instances of `PROTOCOL=TCP` with `PROTOCOL=SDP`.
 6. Click **Save**.
 7. To activate these changes, in the Change Center of the Administration Console, click **Activate Changes**.

This change does not take effect immediately—it requires that the data source be redeployed (untargeted and retargeted) or the server be restarted.

For example:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=SDP)(HOST=left)(PORT=1234))(ADDRESS=(PRO
TOCOL=SDP)(HOST=right)(PORT=1234))(ADDRESS=(PROTOCOL=SDP)(HOST=c
enter)(PORT=1234)))(CONNECT_DATA=(SERVICE_NAME=myService)))
```

4.4 Configuring Connection Pool Features

Each JDBC data source has a pool of JDBC connections that are created when the data source is deployed or at server startup. Applications use a connection from the pool then return it when finished using the connection. Connection pooling enhances performance by eliminating the costly task of creating database connections for the application.

The following sections include information about connection pool options for a JDBC data source.

- [Section 4.4.1, "Enabling JDBC Driver-Level Features"](#)
- [Section 4.4.2, "Enabling Connection-based System Properties."](#)

- [Section 4.4.3, "Initializing Database Connections with SQL Code"](#)

You can see more information and set these and other related options through the:

- **JDBC Data Source: Configuration: Connection Pool** page in the Administration Console. See "JDBC Data Source: Configuration: Connection Pool" in the *Oracle WebLogic Server Administration Console Help*
- `JDBCConnectionPoolParamsBean`, which is a child MBean of the `JDBCDataSourceBean`

4.4.1 Enabling JDBC Driver-Level Features

WebLogic JDBC data sources support the `javax.sql.ConnectionPoolDataSource` interface implemented by JDBC drivers. You can enable driver-level features by adding the property and its value to the `Properties` attribute in a JDBC data source. Driver-level properties in the `Properties` attribute are set on the driver's `ConnectionPoolDataSource` object.

4.4.2 Enabling Connection-based System Properties

WebLogic JDBC data sources support setting driver properties using the value of system properties. The value of each property is derived at runtime from the named system property. You can configure connection-based system properties using the Administration Console by editing the `System Properties` attribute of your data source configuration.

4.4.3 Initializing Database Connections with SQL Code

When WebLogic Server creates database connections in a data source, the server can automatically run SQL code to initialize the database connection. To enable this feature, enter `SQL` followed by a space and the SQL code you want to run in the `Init SQL` attribute on the JDBC Data Source: Configuration: Connection Pool page in the Administration Console. If you leave this attribute blank (the default), WebLogic Server does not run any code to initialize database connections.

WebLogic Server runs this code whenever it creates a database connection for the data source, which includes at server startup, when expanding the connection pool, and when refreshing a connection.

You can use this feature to set DBMS-specific operational settings that are connection-specific or to ensure that a connection has memory or permissions to perform required actions.

Start the code with `SQL` followed by a space. For example:

```
SQL alter session set NLS_DATE_FORMAT='YYYY-MM-DD HH24:MI:SS'
```

or

```
SQL SET LOCK MODE TO WAIT
```

Options that you can set using `InitSQL` vary by DBMS.

Note: `Init SQL` is not a dynamic attribute. When you change the value for `Init SQL`, you must either undeploy and redeploy the data source or restart the server.

4.5 Configuring Oracle Parameters

WebLogic Server provides several attributes that provide improved Data Source performance when using Oracle drivers, for more information, see [Section 6, "Advanced Configuration for Oracle Drivers."](#)

4.6 Configuring an ONS Client

The following section provides information on how to configure an ONS client.

- [Section 4.6.1, "Enabling FAN Events"](#)
- [Section 4.6.2, "Configuring Generic Data Source Connection Testing."](#)
- [Section 4.6.3, "Using a Wallet File"](#)

4.6.1 Enabling FAN Events

Enabling a data source to subscribe to and process Oracle Fast Application Notification (FAN) events.

1. Select `Fan Enabled`
2. Provide a comma-separated list of ONS daemon listen addresses and ports for receiving ONS-based FAN events. You can use Single Client Access Name (SCAN) addresses to access FAN notifications.

See Configure ONS client parameters in *Oracle WebLogic Server Administration Console Help*.

4.6.2 Configuring Generic Data Source Connection Testing

Enabling FAN events automatically disables generic data source connection testing. However, if the ONS server is down or in unhealthy state, you can disable `Fan Enabled` and use generic data source connection testing to maintain the health state of database connections. For more information, see [Section 10.2, "Connection Testing Options for a Data Source."](#)

4.6.3 Using a Wallet File

To communicate with ONS daemons using SSL, you must use a wallet file. See [Section 4.2.5.1, "Secure ONS Client Communication."](#)

4.7 Configuring Oracle Parameters

WebLogic Server provides several attributes that provide improved Data Source performance when using Oracle drivers, for more information, see [Section 6, "Advanced Configuration for Oracle Drivers."](#)

4.8 Tuning GridLink Data Source Connection Pools

By properly configuring the connection pool attributes in JDBC data sources in your WebLogic Server domain, you can improve application and system performance. For more information, see [Section 10, "Tuning Data Source Connection Pools."](#)

4.9 Setting Database Security Credentials

The following sections provide information on how to pass security credentials to a DBMS:

- [Section 4.9.1, "Using a User Name/Password"](#)
- [Section 4.9.2, "Set Client ID On Connection"](#)

Note: GridLink Data Sources do not support identity-based connection pooling.

4.9.1 Using a User Name/Password

The simplest type of credential is to provide the connection pool a user account name and password for the DBMS. All the connections in the pool then use the same credentials to access a DBMS. See "Create JDBC data sources" in *Oracle WebLogic Server Administration Console Help*.

Note: You can enter the password as a name-value pair in the Properties field (not permitted for production environments) or you can enter it in the Password field. The value in the Password field overrides any password value defined in the Properties passed to the JDBC Driver when creating physical database connections. Oracle recommends that you use the Password attribute in place of the password property in the properties string because the Password value is encrypted in the configuration file (stored as the password-encrypted attribute in the `jdbc-driver-params` tag in the module file) and is hidden in the administration console.

4.9.2 Set Client ID On Connection

If the `Set Client ID On Connection` attribute is enabled on the data source, when an application requests a database connection from the data source, the WebLogic Server instance determines the current WebLogic user ID and then sets the mapped database ID as a light-weight client ID. All the connections in the pool have the same credentials to access a DBMS. Basic configuration steps are:

1. Select `Set Client ID On Connection`, see "Enable Set Client ID On Connection for a JDBC data source" in *Oracle WebLogic Server Administration Console Help*.

Note: Credential mapping to map the WebLogic user ID and the database ID is only supported on the Oracle database with the Oracle Thin driver. This feature is not supported with the Oracle DMS driver.

2. Map the WebLogic user ID and the database ID. See "Configure credential mapping for a JDBC data source" in the *Oracle WebLogic Server Administration Console Help*.

This feature relies on features in the JDBC driver and DBMS. It is only supported for use with Oracle and DB2 databases using a vendor extension method:

- `oracle.jdbc.OracleConnection.setClientIdentifier(String id)`
- `com.ibm.db2.jcc.DB2Connection.setDB2ClientUser(String user)`

Note: Set Client ID On Connection and Enable Identity Based Connection Pooling are mutually exclusive. If you think you need both mechanisms to pass security credentials in your application environment, create separate data sources—one for with Set Client ID On Connection and one with Enable Identity Based Connection Pooling.

4.10 Monitoring GridLink JDBC Resources

The following sections include details about monitoring GridLink JDBC objects:

- [Section 4.10.1, "Viewing Run-Time Statistics"](#)
- [Section 4.10.2, "Debug GridLink Data Sources."](#)

For more information on JDBC monitoring, see [Section 13, "Monitoring WebLogic JDBC Resources."](#)

4.10.1 Viewing Run-Time Statistics

You can view run-time statistics for a GridLink data source via the Administration Console or through the associated runtime MBeans.

4.10.1.1 JDBCOracleDataSourceInstanceRuntimeMBean

The `JDBCOracleDataSourceInstanceRuntimeMBean` provides methods for getting the current state of the data source instance. For more information, see "JDBCOracleDataSourceInstanceRuntimeMBean" in the *Oracle WebLogic Server MBean Reference*.

4.10.1.2 JDBCDataSourceRuntimeMBean

The `JDBCDataSourceRuntimeMBean` provides methods for getting the current state of the data source instance. The `JDBCDataSourceRuntimeMBean` provides methods for getting the current state of the data source and for getting statistics about the data source, such as the average number of active connections, the current number of active connections, and the highest number of active connections. For more information, see "JDBCDataSourceRuntimeMBean" in the *Oracle WebLogic Server MBean Reference*.

4.10.1.3 ONSDaemonRuntimeMBean

The `ONSDaemonRuntimeMBean` provides methods for monitoring the ONS client configuration that is associated with a GridLink data source. For more information, see "ONSDaemonRuntimeMBean" in the *Oracle WebLogic Server MBean Reference*.

4.10.2 Debug GridLink Data Sources

You can activate WebLogic Server's debugging features to track down the specific problem within the application.

4.10.2.1 JDBC Debugging Scopes

The following are registered debugging scopes for JDBC:

- `DebugJDBCRCAC` (scope `weblogic.jdbc.racl`) - prints information about GridLink data source lifecycle, UCP callback, and connection information.
- `DebugJDBC CONS` (scope `weblogic.jdbc.connection`) - traces ONS client information, including the LBA event body.

4.10.2.2 UCP JDK Logging

You can enable UPC JDK logging by following the instructions at http://download.oracle.com/docs/cd/B28359_01/java.111/e10788/get_started.htm#sthref67.

4.10.2.3 Enable Debugging Using the Command Line

Set the appropriate GridLink data source debugging properties on the command line. For example,

```
-Dweblogic.debug.DebugJDBCRCAC=true  
-Dweblogic.debug.DebugJDBCONS=true  
-Dweblogic.debug.DebugJDBCUCP=true
```

This method is static and can only be used at server startup.

Configuring JDBC Multi Data Sources

A *multi data source* is an abstraction around a group of data sources that provides load balancing or failover processing at the time of connection requests, between the data sources associated with the multi data source. Multi data sources are bound to the JNDI tree or local application context just like data sources are bound to the JNDI tree. Applications lookup a multi data source on the JNDI tree or in the local application context (`java:comp/env`) just as they do for data sources, and then request a database connection. The multi data source determines which data source to use to satisfy the request depending on the algorithm selected in the multi data source configuration: load balancing or failover.

This section includes the following information:

- [Section 5.1, "Multi Data Source Features"](#)
- [Section 5.2, "Creating and Configuring Multi Data Sources"](#)
- [Section 5.3, "Choosing the Multi Data Source Algorithm"](#)
- [Section 5.4, "Multi Data Source Fail-Over Limitations and Requirements"](#)
- [Section 5.5, "Multi Data Source Failover Enhancements"](#)
- [Section 5.6, "Deploying JDBC Multi Data Sources on Servers and Clusters"](#)

5.1 Multi Data Source Features

A multi data source can be thought of as a pool of data sources. Multi data sources are best used for failover or load balancing between nodes of a highly available database system, such as redundant databases or Oracle Real Application Clusters (Oracle RAC).

The data source member list for a Multi data source supports dynamic updates. This allows environments, such as those using Oracle RAC, to add and remove database nodes and corresponding data sources without redeployment and provide the ability to:

- Grow and shrink Oracle RAC clusters in response to throughput. See [Section 5.1.2, "Adding a Database Node."](#)
- Shutdown Oracle RAC nodes for maintenance. See [Section 5.1.1, "Removing a Database Node."](#)

See [Section B, "Using Multi Data Sources with Oracle RAC."](#)

Note: Multi data sources do not provide any synchronization between databases. It is assumed that database synchronization is handled properly outside of WebLogic Server so that data integrity is maintained.

5.1.1 Removing a Database Node

You can remove a database node and corresponding data sources without redeployment. This capability provides you the ability to shutdown a node for maintenance or shrink a cluster. Use the following high-level steps to shutdown a database node:

Note: Failure to follow these step may cause transaction roll-backs.

1. Remove the data source from the multi data source. See "Add or remove data sources in a JDBC multi data source" in *Oracle WebLogic Server Administration Console Help*
2. When all transactions have completed, suspend the data source. See "Suspend JDBC data sources" in *Oracle WebLogic Server Administration Console Help*
3. When all transactions have completed, shut down the data source. See "Shut down JDBC data sources" in *Oracle WebLogic Server Administration Console Help*
4. Shut down the database node.

5.1.2 Adding a Database Node

You can add a database node and corresponding data sources without redeployment. This capability provides you the ability to start a node after maintenance or grow a cluster. Use the following high-level steps to add a database node:

1. Restart the database node.
2. Restart the data source. See "Start JDBC data sources" in *Oracle WebLogic Server Administration Console Help*
3. Add the data source back to the multi data source. See "Add or remove data sources in a JDBC multi data source" in *Oracle WebLogic Server Administration Console Help*

5.2 Creating and Configuring Multi Data Sources

You create a multi data source by first creating data sources, then creating the multi data source using the Administration Console or the WebLogic Scripting Tool and then assigning the data sources to the multi data source.

For instructions to create a multi data source, see "Configure JDBC multi data sources" in the *Oracle WebLogic Server Administration Console Help*.

For information about the configuration files created when configuring a multi data source, see [Section 2.1, "Understanding JDBC Resources in WebLogic Server."](#) Also see [Section A.2.5, "Creating a JDBC Multi Data Source Module."](#)

5.3 Choosing the Multi Data Source Algorithm

Before you set up a multi data source, you need to determine the primary purpose of the multi data source—failover or load balancing. You can choose the algorithm that corresponds with your requirements.

5.3.1 Failover

The Failover algorithm provides an ordered list of data sources to use to satisfy connection requests. Normally, every connection request to this kind of multi data source is served by the first data source in the list. If a database connection test fails and the connection cannot be replaced, or if the data source is suspended, a connection is sought sequentially from the next data source on the list.

Note: This algorithm requires that Test Reserved Connections (`TestConnectionsOnReserve`) on the data source is enabled. If enabled, a connection in the first data source is tested to verify if the data source is healthy. If the connection fails the test, the multi data source uses a connection from the next data source listed in the multi data source. See [Section 10.2, "Connection Testing Options for a Data Source"](#) for information about configuring `TestConnectionsOnReserve`.

JDBC is a highly stateful client-DBMS protocol, in which the DBMS connection and transactional state are tied directly to the socket between the DBMS process and the client (driver). For this reason, failover of a connection while it is in use is not supported.

5.3.2 Load Balancing

Connection requests to a load-balancing multi data source are served from any data source in the list. The multi data source selects data sources to use to satisfy connection requests using a round-robin scheme. When the multi data source provides a connection, it selects a connection from the data source listed just after the last data source that was used to provide a connection. Multi data sources that use the Load Balancing algorithm also fail over to the next data source in the list if a database connection test fails and the connection cannot be replaced, or if the data source is suspended.

5.4 Multi Data Source Fail-Over Limitations and Requirements

WebLogic Server provides the Failover algorithm for multi data sources so that if a data source fails (for example, if the database management system crashes), your system can continue to operate. However, you must consider the following limitations and requirements when configuring your system.

5.4.1 Test Connections on Reserve to Enable Fail-Over

Data sources rely on the Test Reserved Connections (`TestConnectionsOnReserve`) feature on the data source to know when database connectivity is lost. Testing reserved connections must be enabled for the data sources within the multi data source. WebLogic Server will test each connection before giving it to an application. With the Failover algorithm, the multi data source uses the results from connection test to determine when to fail over to the next data source in the multi data source. After a test failure, the data source attempts to recreate the connection. If that attempt fails, the multi data source fails over to the next data source.

5.4.2 No Fail-Over for In-Use Connections

It is possible for a connection to fail after being reserved, in which case your application must handle the failure. WebLogic Server cannot provide fail-over for connections that fail while being used by an application. Any failure while using a connection requires that the application code close the failed connection, and the transaction must be restarted from the beginning with a new connection.

5.5 Multi Data Source Failover Enhancements

The following enhancements improve failover processing for multi data sources:

- Connection request routing enhancements to avoid requesting a connection from an automatically disabled (dead) data source within a multi data source. See [Section 5.5.1, "Connection Request Routing Enhancements When a Data Source Fails."](#)
- Automatic failback on recovery of a failed data source within a multi data source. See [Section 5.5.2, "Automatic Re-enablement on Recovery of a Failed Data Source within a Multi Data Source."](#)
- Failover for busy data sources within a multi data sources. See [Section 5.5.3, "Enabling Failover for Busy Data Sources in a Multi Data Source."](#)
- Failover callbacks for multi data sources with the Failover algorithm. See [Section 5.5.4, "Controlling Multi Data Source Failover with a Callback."](#)
- Failback callbacks for multi data sources with either algorithm. See [Section 5.5.5, "Controlling Multi Data Source Failback with a Callback."](#)

5.5.1 Connection Request Routing Enhancements When a Data Source Fails

To improve performance when a data source within a multi data source fails, WebLogic Server automatically disables the data source when a pooled connection fails a connection test. After a data source is disabled, WebLogic Server does not route connection requests from applications to the data source. Instead, it routes connection requests to the next available data source listed in the multi data source.

This feature requires that data source testing options are configured for all data sources in a multi data source, specifically Test Table Name and Test Reserved Connections. See [Section 10.2, "Connection Testing Options for a Data Source."](#)

If a callback handler is registered for the multi data source, WebLogic Server calls the callback handler before failing over to the next data source in the list. See [Section 5.5.4, "Controlling Multi Data Source Failover with a Callback"](#) for more details.

5.5.2 Automatic Re-enablement on Recovery of a Failed Data Source within a Multi Data Source

After a data source is automatically disabled because a connection failed a connection test, the multi data source periodically tests a connection from the disabled data source to determine when the data source (or underlying database) is available again. When the data source becomes available, the multi data source automatically re-enables the data source and resumes routing connection requests to the data source, depending on the multi data source algorithm and the position of the data source in the list of included data sources. Frequency of these tests is controlled by the Test Frequency Seconds attribute of the multi data source. The default value for Test Frequency is 120 seconds, so if you do not specifically set a value for the option, the multi data source

will test disabled data sources every 120 seconds. See "JDBC Multi Data Source: Configuration: General" in the *Oracle WebLogic Server Administration Console Help*.

WebLogic Server does not test and automatically re-enable data sources that you manually disable. It only tests data sources that are automatically disabled.

If a callback handler is registered for the multi data source, WebLogic Server calls the callback handler before re-enabling the data source. See [Section 5.5.5, "Controlling Multi Data Source Failback with a Callback"](#) for more details.

5.5.3 Enabling Failover for Busy Data Sources in a Multi Data Source

By default, for multi data sources with the Failover algorithm, when the number of requests for a database connection exceeds the number of available connections in the current data source in the multi data source, subsequent connection requests fail.

To enable the multi data source to failover when all connections in the current data source are in use, you can enable the Failover Request if Busy option on the JDBC Multi Data Source: Configuration: General page in the Administration Console. (Also available as the `FailoverRequestIfBusy` attribute in the `JDBCDataSourceParamsBean`). If enabled (set to `true`), when all connections in the current data source are in use, application requests for connections will be routed to the next available data source within the multi data source. When disabled (set to `false`, the default), connection requests do not failover.

If a `ConnectionPoolFailoverCallbackHandler` is included in the multi data source configuration, WebLogic Server calls the callback handler before failing over. See [Section 5.5.4, "Controlling Multi Data Source Failover with a Callback"](#) for more details.

5.5.4 Controlling Multi Data Source Failover with a Callback

You can register a callback handler with WebLogic Server that controls when a multi data source with the Failover algorithm fails over connection requests from one JDBC data source in the multi data source to the next data source in the list.

You can use callback handlers to control if or when the failover occurs so that you can make any other system preparations before the failover, such as priming a database or communicating with a high-availability framework.

Callback handlers are registered via the Failover Callback Handler attribute of the multi data source and are registered per multi data source. You must register the callback handler for each multi data source to which you want the callback handler to apply. And you can register different callback handlers for each multi data source in your domain.

5.5.4.1 Callback Handler Requirements

A callback handler used to control the failover and failback within a multi data source must include an implementation of the `weblogic.jdbc.extensions.ConnectionPoolFailoverCallback` interface. When the multi data source needs to failover to the next data source in the list or when a previously disabled data source becomes available, WebLogic Server calls the `allowPoolFailover()` method in the `ConnectionPoolFailoverCallback` interface, and passes a value for the three parameters, `currPool`, `nextPool`, and `opcode`, as defined below. WebLogic Server then waits for the return from the callback handler before completing the task.

Your application must return `OK`, `RETRY_CURRENT`, or `DONOT_FAILOVER` as defined below. The application should handle failover and fallback cases.

See the Javadoc for the `weblogic.jdbc.extensions.ConnectionPoolFailoverCallback` interface for more details.

Note: Failover callback handlers are optional. If no callback handler is specified in the multi data source configuration, WebLogic Server proceeds with the operation (failing over or re-enabling the disabled data source).

5.5.4.2 Callback Handler Configuration

There are two multi data source configuration attributes associated with the failover and fallback functionality:

- **Failover Callback Handler** (`ConnectionPoolFailoverCallbackHandler`)—To register a failover callback handler for a multi data source, you add a value for this attribute to the multi data source configuration. The value must be an absolute name, such as `com.bea.samples.wls.jdbc.MultiDataSourceFailoverCallbackApplication`. You can set the Failover Callback Handler using the Administration Console (see "Register a failover callback handler" in the *Oracle WebLogic Server Administration Console Help*) or on the `JDBCDataSourceParamsBean` for the multi data source using WLST.
- **Test Frequency** (`TestFrequencySeconds`)—To control how often the multi data source checks disabled (dead) data sources to see if they are now available. See [Section 5.5.2, "Automatic Re-enablement on Recovery of a Failed Data Source within a Multi Data Source"](#) for more details.

5.5.4.3 How It Works—Failover

WebLogic Server attempts to failover connection requests to the next data source in the list when the current data source fails a connection test or, if you enabled `FailoverRequestIfBusy`, when all connections in the current data source are busy.

To enable the callback feature, you register the callback handler with WebLogic Server using Failover Callback Handler in the multi data source configuration.

With the Failover algorithm, connection requests are served from the first data source in the list. If a connection from that data source fails a connection test, WebLogic Server marks the data source as dead and disables it. If a callback handler is registered, WebLogic Server calls the callback handler, passing the following information, and waits for a return:

- `currPool`—For failover, this is the name of data source currently being used to supply database connections. This is the "failover from" data source.
- `nextPool`—The name of next available data source listed in the multi data source. For failover, this is the "failover to" data source.
- `opcode`—A code that indicates the reason for the call:
 - `OPCODE_CURR_POOL_DEAD`—WebLogic Server determined that the current data source is dead and has disabled it.
 - `OPCODE_CURR_POOL_BUSY`—All database connections in the data source are in use. (Requires `FailoverIfBusy=true` in the multi data source configuration. See [Section 5.5.3, "Enabling Failover for Busy Data Sources in a Multi Data Source."](#))

Failover is synchronous with the connection request: Failover occurs only when WebLogic Server is attempting to satisfy a connection request.

The return from the callback handler can indicate one of three options:

- OK—proceed with the operation. In this case, that means to failover to the next data source in the list.
- `RETRY_CURRENT`—Retry the connection request with the current data source.
- `DONOT_FAILOVER`—Do not retry the current connection request and do not failover. WebLogic Server will throw a `weblogic.jdbc.extensions.PoolUnavailableSQLException`.

WebLogic Server acts according to the value returned by the callback handler.

If the secondary data sources fails, WebLogic Server calls the callback handler again, as in the previous failover, in an attempt to failover to the next available data source in the multi data source, if there is one.

Note: WebLogic Server does not call the callback handler when you manually disable a data source.

For multi data sources with the Load-Balancing algorithm, WebLogic Server does not call the callback handler when a data source is disabled. However, it does call the callback handler when attempting to re-enable a disabled data source. See the following section for more details.

5.5.5 Controlling Multi Data Source Failback with a Callback

If you register a failover callback handler for a multi data source, WebLogic Server calls the same callback handler when re-enabling a data source that was automatically disabled. You can use the callback to control if or when the disabled data source is re-enabled so that you can make any other system preparations before the data source is re-enabled, such as priming a database or communicating with a high-availability framework.

See the following sections for more details about the callback handler:

- ["Callback Handler Requirements"](#)
- ["Callback Handler Configuration"](#)

5.5.5.1 How It Works—Failback

WebLogic Server periodically checks the status of data sources in a multi data source that were automatically disabled. (See [Section 5.5.2, "Automatic Re-enablement on Recovery of a Failed Data Source within a Multi Data Source."](#)) If a disabled data source becomes available and if a failover callback handler is registered, WebLogic Server calls the callback handler with the following information and waits for a return:

- `currPool`—For failback, this is the name of the data source that was previously disabled and is now available to be re-enabled.
- `nextPool`—For failback, this is null.
- `opcode`—A code that indicates the reason for the call. For failback, the code is always `OPCODE_REENABLE_CURR_POOL`, which indicates that the data source named in `currPool` is now available.

Failback, or automatically re-enabling a disabled data source, differs from failover in that failover is *synchronous* with the connection request, but failback is *asynchronous* with the connection request.

The callback handler can return one of the following values:

- `OK`—proceed with the operation. In this case, that means to re-enable the indicated data source. WebLogic Server resumes routing connection requests to the data source, depending on the multi data source algorithm and the position of the data source in the list of included data sources.
- `DONOT_FAILOVER`—Do not re-enable the `currPool` data source. Continue to serve connection requests from the data source(s) in use.

WebLogic Server acts according to the value returned by the callback handler.

If the callback handler returns `DONOT_FAILOVER`, WebLogic Server will attempt to re-enable the data source during the next testing cycle as determined by the `TestFrequencySeconds` attribute in the multi data source configuration, and will call the callback handler as part of that process.

The order in which data sources are listed in a multi data source is very important. A multi data source with the Failover algorithm will always attempt to serve connection requests from the first available data source in the list of data sources in the multi data source. Consider the following scenario:

1. `MultiDataSource_1` uses the Failover algorithm, has a registered `ConnectionPoolFailoverCallbackHandler`, and includes three data sources: `DS1`, `DS2`, and `DS3`, listed in that order.
2. `DS1` becomes disabled, so `MultiDataSource_1` fails over connection requests to `DS2`.
3. `DS2` then becomes disabled, so `MultiDataSource_1` fails over connection requests to `DS3`.
4. After some time, `DS1` becomes available again and the callback handler allows WebLogic Server to re-enable the data source. Future connection requests will be served by `DS1` because `DS1` is the first data source listed in the multi data source.
5. If `DS2` subsequently becomes available and the callback handler allows WebLogic Server to re-enable the data source, connection requests will continue to be served by `DS1` because `DS1` is listed before `DS2` in the list of data sources.

5.6 Deploying JDBC Multi Data Sources on Servers and Clusters

All data sources used by a multi data source to satisfy connection requests must be deployed on the same servers and clusters as the multi data source. A multi data source always uses a data source deployed on the same server to satisfy connection requests. Multi data sources do not route connection requests to other servers in a cluster or in a domain.

To deploy a multi data source to a cluster or server, you select the server or cluster as a deployment target. When a multi data source is deployed on a server, WebLogic Server creates an instance of the multi data source on the server. When you deploy a multi data source to a cluster, WebLogic Server creates an instance of the multi data source on each server in the cluster.

For instructions, see "Target and deploy JDBC multi data sources" in the *Oracle WebLogic Server Administration Console Help*.

Advanced Configuration for Oracle Drivers

WebLogic Server provides several attributes that provide improved Data Source performance when using Oracle drivers, including:

- Oracle Enable JavaNet Fastpath: Enables the Oracle JDBC JavaNet Fastpath to reduce data copies and fragmentation.
- Oracle Optimize UTF8 Conversion: Enables the Oracle JDBC optimize UTF-8 conversion option.

If your environment supports these attributes, see Configure Oracle parameters in the *Oracle WebLogic Server Administration Console Help*.

JDBC Data Source Transaction Options

When you configure a JDBC data source using the Administration Console, WebLogic Server automatically selects specific transaction options based on the type of JDBC driver:

- **For XA drivers**, the system automatically selects the **Two-Phase Commit** protocol for global transaction processing.
- **For non-XA drivers**, local transactions are supported by definition, and WebLogic Server offers the following options

Supports Global Transactions: (selected by default) Select this option if you want to use connections from the data source in global transactions, even though you have not selected an XA driver. See [Section 7.1, "Enabling Support for Global Transactions with a Non-XA JDBC Driver"](#) for more information.

When you select Supports Global Transactions, you must also select the protocol for WebLogic Server to use for the transaction branch when processing a global transaction:

- **Logging Last Resource:** With this option, the transaction branch in which the connection is used is processed as the last resource in the transaction and is processed as a local transaction. Commit records for two-phase commit (2PC) transactions are inserted in a table on the resource itself, and the result determines the success or failure of the prepare phase of the global transaction. This option offers some performance benefits and greater data safety than Emulate Two-Phase Commit, but it has some limitations. See [Section 7.2, "Understanding the Logging Last Resource Transaction Option."](#)

Note: Logging Last Resource is not supported for data sources used by a multi data source except when used with Oracle RAC version 10G Release 2 (10GR2) and greater versions as described in [Section 7.2.4, "Administrative Considerations and Limitations for LLR Data Sources."](#)

- **Emulate Two-Phase Commit:** With this option, the transaction branch in which the connection is used always returns success for the prepare phase of the transaction. It offers performance benefits, but also has risks to data in some failure conditions. Select this option only if your application can tolerate heuristic conditions. See [Section 7.3, "Understanding the Emulate Two-Phase Commit Transaction Option."](#)
- **One-Phase Commit:** (selected by default) With this option, a connection from the data source can be the only participant in the global transaction and the transaction is completed using a one-phase commit optimization. If more than

one resource participates in the transaction, an exception is thrown when the transaction manager calls `XAResource.prepare` on the 1PC resource.

This section includes the following information:

- [Section 7.1, "Enabling Support for Global Transactions with a Non-XA JDBC Driver"](#)
- [Section 7.2, "Understanding the Logging Last Resource Transaction Option"](#)
- [Section 7.3, "Understanding the Emulate Two-Phase Commit Transaction Option"](#)

7.1 Enabling Support for Global Transactions with a Non-XA JDBC Driver

If you use global transactions in your applications, you should use an XA JDBC driver to create database connections in the JDBC data source. If an XA driver is unavailable for your database, or you prefer not to use an XA driver, you should enable support for global transactions in the data source. You should also enable support for global transaction if your applications meet any of the following criteria:

- Use the EJB container in WebLogic Server to manage transactions
- Include multiple database updates within a single transaction
- Access multiple resources, such as a database and the Java Messaging Service (JMS), during a transaction
- Use the same data source on multiple servers (clustered or non-clustered)

With an EJB architecture, it is common for multiple EJBs that are doing database work to be invoked as part of a single transaction. Without XA, the only way for this to work is if all transaction participants use the exact same database connection. When you enable global transactions and select either Logging Last Resource or Emulate Two-Phase Commit, WebLogic Server internally uses the JTS driver to make sure all EJBs use the same database connection within the same transaction context without requiring you to explicitly pass the connection from EJB to EJB.

If multiple EJBs are participating in a transaction and you do not use an XA JDBC driver for database connections, configure a Data Source with the following options:

- Supports Global Transactions selected
- Logging Last Resource or Emulate Two-Phase Commit selected

This configuration will force the JTS driver to internally use the same database connection for all database work within the same transaction.

With XA (requires an XA driver), EJBs can use a different database connection for each part of the transaction. WebLogic Server coordinates the transaction using the two-phase commit protocol, which guarantees that all or none of the transaction will be completed.

7.2 Understanding the Logging Last Resource Transaction Option

WebLogic Server supports the Logging Last Resource (LLR) transaction optimization through JDBC data sources. LLR is a performance enhancement option that enables one non-XA resource to participate in a global transaction with the same ACID guarantee as XA. LLR is a refinement of the "Last Agent Optimization." It differs from Last Agent Optimization in that it is transactionally safe. The LLR resource uses a local transaction for its transaction work. The WebLogic Server transaction manager prepares all other resources in the transaction and then determines the commit

decision for the global transaction based on the outcome of the LLR resource's local transaction.

The LLR optimization improves performance by:

- Removing the need for an XA JDBC driver to connect to the database. XA JDBC drivers are typically inefficient compared to non-XA JDBC drivers.
- Reducing the number of processing steps to complete the transaction, which also reduces network traffic and the number of disk I/Os.
- Removing the need for XA processing at the database level

When a connection from a data source configured for LLR participates in a two-phase commit (2PC) global transaction, the WebLogic Server transaction manager completes the transaction by:

- Calling prepare on all other (XA-compliant) transaction participants.
- Inserting a commit record to a table on the LLR participant (rather than to the file-based transaction log).
- Committing the LLR participant's local transaction (which includes both the transaction commit record insert and the application's SQL work).
- Calling commit on all other transaction participants.

For a one-phase commit (1PC) global transaction, LLR eliminates the XA overhead by using a local transaction to complete the database operations, but no 2PC transaction record is written to the database.

The Logging Last Resource optimization maintains data integrity by writing the commit record on the LLR participant. If the transaction fails during the local transaction commit, the WebLogic Server transaction manager rolls back the transaction on all other transaction participants. For failure recovery, the WebLogic Server transaction manager reads the transaction log on the LLR resource along with other transaction log files in the default store and completes any transaction processing as necessary. Work associated with XA participants is committed if a commit record exists, otherwise their work is rolled back.

For instructions on how to create an LLR-enabled JDBC data source, see "Create LLR-enabled JDBC data sources" in the *Oracle WebLogic Server Administration Console Help*. For more details about the Logging Last Resource transaction processing, see "Logging Last Resource Transaction Optimization" in *Programming JTA for Oracle WebLogic Server*.

7.2.1 Advantages to Using the Logging Last Resource Optimization

Depending on your environment, you may want to consider the LLR transaction protocol in place of the two-phase commit protocol for transaction processing because of its performance benefits. The LLR transaction protocol offers the following advantages:

- Allows non-XA JDBC drivers and even non-XA-capable databases to safely participate in two-phase commit transactions.
- Eliminates the database's use of the XA protocol.
- Performs better than JDBC XA connections.
- Reduces the length of time that database row locks are held.
- Always commits database work prior to other XA work. In XA transactions, these operations are committed in parallel, so, for example, when a JMS send

participates in the transaction, the JMS message may be delivered before database work commits. With LLR, the database work in the local transaction is completed before all other transaction work.

- Has no increased risk of heuristic hazards, unlike the Emulate Two-Phase Commit option for a JDBC data source.

Note: The LLR optimization provides a significant increase in performance for insert, update, and delete operations. However, for read operations with LLR, performance is somewhat slower than read operations with XA.

For more information about performance tuning with LLR, see "Optimizing Performance with LLR" in *Programming JTA for Oracle WebLogic Server*.

7.2.2 Enabling the Logging Last Resource Transaction Optimization

To enable the LLR transaction optimization, you create a JDBC data source with the Logging Last Resource transaction protocol, then use database connections from the data source in your applications. WebLogic Server automatically creates the required table on the database.

See "Create LLR-enabled JDBC data sources" in the *Oracle WebLogic Server Administration Console Help*.

7.2.3 Programming Considerations and Limitations for LLR Data Sources

You use JDBC connections from an LLR-enabled data source in an application as you would use JDBC connections from any other data source: *after* beginning a transaction, you look up the data source on the JNDI tree, then request a connection from the data source. However, with the LLR optimization, WebLogic Server internally manages the connection request and handles the transaction processing differently than in an XA transaction. For more information about how Logging Last Resource works, see "Logging Last Resource Transaction Optimization" in *Programming JTA for Oracle WebLogic Server*.

Note the following:

- When programming with an LLR data source, you must start the global transaction before calling `getConnection` on the LLR data source. If you call `getConnection` before starting the global transaction, the connection will be independent, and will not be associated with any subsequently started global transaction. The connection will operate in the `autoCommit(true)` mode. In this mode, every update will commit automatically on its own, and there will be no way to roll back any update unless application code has explicitly set the `autoCommit` state to `false` and is explicitly managing its own local transaction.
- Only one internal JDBC LLR connection is reserved per transaction. And that connection is used throughout the transaction processing.
- The reserved connection is always hosted on the transaction's coordinator server. Make sure that the data source is targeted to the coordinating server or to the cluster. Also see "Optimizing Performance with LLR" in *Programming JTA for Oracle WebLogic Server*.
- For additional JDBC connection requests within the transaction from a same-named data source, operations are routed to the reserved connection from the original connection request, even if the subsequent connection request is made on a different instance of the data source (i.e., a data source deployed on a

different server than the original data source that supplied the connection for the first request). Note the following:

- Routed LLR connections may be less capable and less performant than locally hosted XA connections. (See [Section 7.3.1.3, "Possible Performance Loss with Non-XA Resources in Multi-Server Configurations."](#))
- Connection request routing limits the number of concurrent transactions. The maximum number of concurrent LLR transactions is equal to the configured size (`MaxCapacity`) of the coordinator's JDBC LLR data source.
- Routed connections have less capability than local connections, and may fail as a result. Specifically, non-serializable "custom" data types within a query `ResultSet` may fail.
- Only instances of a single LLR data source may participate in a particular transaction. A single LLR data source may have instances on multiple WebLogic servers, and two data sources are considered to be the same if they have the same configured name. If more than one LLR data source instance is detected and they are not instances of the same data source, the transaction manager will roll back the transaction.
- Resource adapters (connectors) that implement the `weblogic.transaction.nonxa.NonXAResource` interface cannot participate in global transaction in which an LLR resource also participates because both must be the last resource in the transaction. If both resource types participate in the same transaction, the `transaction.commit()` method throws a `javax.transaction.RollbackException` when this conflict is detected.
- Because the LLR connection uses a separate *local* transaction for database processing, any changes made (and locks held) to the same database using an XA connection are not visible during the LLR processing even though all of the processing occurs in the same *global* transaction. In some cases, this can cause deadlocks in the database. You should not combine XA and LLR processing in the same database in a single global transaction.
- Connections from an LLR data source cannot participate in transactions coordinated by foreign transaction managers, such as a transaction started by a remote object request broker or by Tuxedo.
- Global transactions cannot span to another legacy domain that includes a data source with the same name as an LLR data source.
- For JDBC LLR 2PC transactions, if the transaction data is too large to fit in the LLR table, the transaction will fail with a rollback exception thrown during commit. This can occur if your application adds many transaction properties during transaction processing. (See "Oracle WebLogic Extensions to JTA" in *Programming JTA for Oracle WebLogic Server*) Your database administrator can manually create a table with larger columns if this occurs.

7.2.4 Administrative Considerations and Limitations for LLR Data Sources

Consider the following requirements and limitations when configuring an LLR-enabled JDBC data source. For more information about how Logging Last Resource works, see "Logging Last Resource Transaction Optimization" in *Programming JTA for Oracle WebLogic Server*.

- There is one LLR table per server:
 - Multiple LLR data sources may share a table.

- WebLogic Server automatically creates the table if it is not found.
- Default name is `WL_LLRL_SERVERNAME`. You can configure the table name in the Administration Console on the Server > Configuration > General tab under Advanced options. See "Servers: Configuration: General" in *Oracle WebLogic Server Administration Console Help*.
- A server **will not boot** if the database is down or the LLR table is unreachable during boot.
- Multiple servers must *not* share the same LLR table. Boot checks to ensure domain and server name match the domain and server name stored in the table when the table is created. If WebLogic Server detects that more than one server is sharing the same LLR table, WebLogic Server will shut down one or more of the servers.
- LLR supports server migration and transaction recovery service migration. To use the transaction recovery service migration, ensure that each LLR resource be targeted to either the cluster or the set of candidate servers in the cluster. See "Recovering Transactions For a Failed Clustered Server" in *Programming JTA for Oracle WebLogic Server*.
- The LLR transaction option is not permitted for use in JDBC application modules.
- When using multi data sources, the LLR transaction option can only be used with Oracle RAC version 10G Release 2 (10GR2) and greater versions with the following settings:
 - All WebLogic application database JDBC interactions must use the `READ_COMMITTED` transaction isolation level (the default).
 - The Oracle RAC setting `MAX_COMMIT_PROPAGATION_DELAY` must be set to a value of 0 (zero, the default).

The use of LLR with multi data sources is supported only with Oracle RAC. LLR is not supported with any other multi data source configuration.

- If you use credential mapping on an LLR data source, all mapped users must have write permissions on the LLR table.
- You cannot use a JDBC XA driver to create database connections in a JDBC LLR data source. If the JDBC driver used in a JDBC LLR data source supports XA, a warning message is logged, and the data source participates in transactions as a full XA resource rather than as an LLR resource.
- Transaction statistics for LLR resources are tracked under "NonXAResource." See "View transaction statistics for non-XA resources" in the *Oracle WebLogic Server Administration Console Help*.

7.3 Understanding the Emulate Two-Phase Commit Transaction Option

If you need to support distributed transactions with a JDBC data source, but there is no available XA-compliant driver for your DBMS, you can select the Emulate Two-Phase Commit for non-XA Driver option for a data source to emulate two-phase commit for the transactions in which connections from the data source participate. This option is an advanced option on the Configuration > General tab of a data source configuration.

When the Emulate Two-Phase Commit for non-XA Driver option is selected (`EnableTwoPhaseCommit` is set to `true`), the non-XA JDBC resource always returns `XA_OK` during the `XAResource.prepare()` method call. The resource attempts to commit or roll back its local transaction in response to subsequent

`XAResource.commit()` or `XAResource.rollback()` calls. If the resource commit or rollback fails, a heuristic error results. Application data may be left in an inconsistent state as a result of a heuristic failure.

When the Emulate Two-Phase Commit for non-XA Driver option is not selected in the Console (`EnableTwoPhaseCommit` is set to `false`), the non-XA JDBC resource causes `XAResource.prepare()` to fail. When there is only one resource participating in a transaction, the one phase optimization bypasses `XAResource.prepare()`, and the transaction commits successfully in most instances.

Note: There are risks to data integrity when using the Emulate Two-Phase Commit for non-XA Driver option. Oracle recommends that you use an XA-compliant JDBC driver or the Logging Last Resource option rather than use the Emulate Two-Phase Commit option. Make sure you consider the risks below before enabling this option.

This non-XA JDBC driver support is often referred to as the "JTS driver" because WebLogic Server uses the WebLogic JTS Driver internally to support the feature. For more information about the WebLogic JTS Driver, see "Using the WebLogic JTS Driver" in *Programming JDBC for Oracle WebLogic Server*.

7.3.1 Limitations and Risks When Emulating Two-Phase Commit Using a Non-XA Driver

WebLogic Server supports the participation of non-XA JDBC resources in global transactions with the Emulate Two-Phase Commit data source transaction option, but there are limitations that you must consider when designing applications to use such resources. Because a non-XA driver does not adhere to the XA/2PC contracts and only supports one-phase commit and rollback operations, WebLogic Server (through the JTS driver) has to make compromises to allow the resource to participate in a transaction controlled by the Transaction Manager.

Consider the following limitations and risks before using the Emulate Two-Phase Commit for non-XA Driver option.

7.3.1.1 Heuristic Completions and Data Inconsistency

When Emulate Two-Phase Commit is selected for a non-XA resource, (`enableTwoPhaseCommit = true`), the prepare phase of the transaction for the non-XA resource always succeeds. Therefore, the non-XA resource does not truly participate in the two-phase commit (2PC) protocol and is susceptible to failures. If a failure occurs in the non-XA resource after the prepare phase, the non-XA resource is likely to roll back the transaction while XA transaction participants will commit the transaction, resulting in a heuristic completion and data inconsistencies.

Because of the data integrity risks, the Emulate Two-Phase Commit option should only be used in applications that can tolerate heuristic conditions.

7.3.1.2 Cannot Recover Pending Transactions

Because a non-XA driver manipulates local database transactions only, there is no concept of a transaction pending state in the database with regard to an external transaction manager. When `XAResource.recover()` is called on the non-XA resource, it always returns an empty set of Xids (transaction IDs), even though there may be transactions that need to be committed or rolled back. Therefore, applications

that use a non-XA resource in a global transaction cannot recover from a system failure and maintain data integrity.

7.3.1.3 Possible Performance Loss with Non-XA Resources in Multi-Server Configurations

Because WebLogic Server relies on the database local transaction associated with a particular JDBC connection to support non-XA resource participation in a global transaction, when the same JDBC data source is accessed by an application with a global transaction context on multiple WebLogic Server instances, the JTS driver will always route to the first connection established by the application in the transaction. For example, if an application starts a transaction on one server, accesses a non-XA JDBC resource, then makes a remote method invocation (RMI) call to another server and accesses a data source that uses the same underlying JDBC driver, the JTS driver recognizes that the resource has a connection associated with the transaction on another server and sets up an RMI redirection to the actual connection on the first server. All operations on the connection are made on the one connection that was established on the first server. This behavior can result in a performance loss due to the overhead associated with setting up these remote connections and making the RMI calls to the one physical connection.

7.3.1.4 Multiple Non-XA Participants

If you use more than one non-XA resource in a global transaction, it is possible to see JTA `SystemExceptions` in the event of a non-atomic outcome. The chance for non-atomic outcomes and `SystemExceptions` tends to increase with the number of two-phase-emulated datasource participants.

Note: The use of a two-phase-emulated datasource in a JTA transaction across domains of different versions is not supported.

Using Roles and Policies to Secure JDBC Data Sources

The following sections provide information on how WebLogic Server uses roles and policies to secure JDBC data sources:

- [Section 8.1, "Setting Security Policies for JDBC Resources"](#)
- [Section 8.2, "Security Roles for JDBC MBeans"](#)

8.1 Setting Security Policies for JDBC Resources

You can optionally restrict access to JDBC data sources. In WebLogic Server, security policies answer the question "who has access" to a WebLogic resource. A security policy is created when you define an association between a WebLogic resource and a user, group, or role. A WebLogic resource has no protection until you assign it a security policy. For instructions on how to set up security for all WebLogic Server resources, see "Use roles and policies to secure resources" in *Oracle WebLogic Server Administration Console Help*. For more information about securing server resources, see *Securing Resources Using Roles and Policies for Oracle WebLogic Server*.

You can protect JDBC resource operations by assigning Administrator methods which can limit the actions that an administrator may take upon a JDBC data source. See "Java DataBase Connectivity (JDBC) Resources" in *Securing Resources Using Roles and Policies for Oracle WebLogic Server*.

8.2 Security Roles for JDBC MBeans

JDBC MBeans allow only the Admin and Deployer roles. The following sections provide information on the security roles defined for JDBC MBeans:

- ["JDBC Domain Configuration MBeans"](#)
- ["JDBC System Module MBeans"](#)

See "Default Security Policies for MBeans" in *Oracle WebLogic Server MBean Reference* for information on default security settings for WebLogic Server.

8.2.1 JDBC Domain Configuration MBeans

The following domain configuration JDBC MBeans have settings that override the default security settings.

- JDBCConnectionPoolMBean (deprecated)
- JDBCDataSourceFactoryMBean (deprecated)

- JDBCDataSourceMBean (deprecated)
- JDBCMultiPoolMBean (deprecated)
- JDBCSystemResourceMBean
- JDBCTxDataSourceMBean (deprecated)

See "Domain Configuration MBeans" in *Oracle WebLogic Server MBean Reference*.

8.2.2 JDBC System Module MBeans

The following system module JDBC MBeans that have settings that override the default security settings.

- JDBCConnectionPoolParamsBean
- JDBCDataSourceBean
- JDBCDataSourceParamsBean
- JDBCDriverParamsBean
- JDBCPropertiesBean
- JDBCPropertyBean
- JDBCXAParamsBean

See "System Module MBeans" in *Oracle WebLogic Server MBean Reference*.

Deploying Data Sources on Servers and Clusters

This section includes the following information:

- [Section 9.1, "Deploying Data Sources on Servers and Clusters"](#)
- [Section 9.2, "Minimizing Server Startup Hang Caused By an Unresponsive Database"](#)

9.1 Deploying Data Sources on Servers and Clusters

To deploy a data source to a cluster or server, you select the server or cluster as a deployment target. When a data source is deployed on a server, WebLogic Server creates an instance of the data source on the server, including the pool of database connections in the data source. When you deploy a data source to a cluster, WebLogic Server creates an instance of the data source on each server in the cluster.

For instructions, see "Target JDBC data sources" in the *Oracle WebLogic Server Administration Console Help*.

9.2 Minimizing Server Startup Hang Caused By an Unresponsive Database

On server startup, WebLogic Server attempts to create database connections in the data sources deployed on the server. If a database is unreachable, server startup may hang in the *STANDBY* state for a long period of time. This is due to WebLogic Server threads that hang inside the JDBC driver code waiting for a reply from the database server. The duration of the hang depends on the JDBC driver and the TCP/IP timeout setting on the WebLogic Server machine.

To work around this issue, WebLogic Server includes the `JDBCLoginTimeoutSeconds` attribute on the `ServerMBean`. When you set a value for this attribute, the value is passed into `java.sql.DriverManager.setLoginTimeout()`. If the JDBC driver being used to create database connections implements the `setLoginTimeout` method, attempts to create database connections will wait only as long as the timeout specified.

Tuning Data Source Connection Pools

By properly configuring the connection pool attributes in JDBC data sources in your WebLogic Server domain, you can improve application and system performance. The following sections include information about tuning options for the connection pool in a JDBC data source:

- [Section 10.1, "Increasing Performance with the Statement Cache"](#)
- [Section 10.2, "Connection Testing Options for a Data Source"](#)
- [Section 10.3, "Enabling Connection Creation Retries"](#)
- [Section 10.4, "Enabling Connection Requests to Wait for a Connection"](#)
- [Section 10.5, "Automatically Recovering Leaked Connections"](#)
- [Section 10.6, "Avoiding Server Lockup with the Correct Number of Connections"](#)
- [Section 10.7, "Limiting Statement Processing Time with Statement Timeout"](#)
- [Section 10.8, "Using Pinned-To-Thread Property to Increase Performance"](#)
- [Section 10.9, "Using Unwrapped Data Type Objects"](#)

10.1 Increasing Performance with the Statement Cache

When you use a prepared statement or callable statement in an application or EJB, there is considerable processing overhead for the communication between the application server and the database server and on the database server itself. To minimize the processing costs, WebLogic Server can cache prepared and callable statements used in your applications. When an application or EJB calls any of the statements stored in the cache, WebLogic Server reuses the statement stored in the cache. Reusing prepared and callable statements reduces CPU usage on the database server, improving performance for the current statement and leaving CPU cycles for other tasks.

Each connection in a data source has its own individual cache of prepared and callable statements used on the connection. However, you configure statement cache options per data source. That is, the statement cache for each connection in a data source uses the statement cache options specified for the data source, but each connection caches its own statements. Statement cache configuration options include:

- **Statement Cache Type**—The algorithm that determines which statements to store in the statement cache. See [Section 10.1.1, "Statement Cache Algorithms."](#)
- **Statement Cache Size**—The number of statements to store in the cache for each connection. The default value is 10. See [Section 10.1.2, "Statement Cache Size."](#)

You can use the Administration Console to set statement cache options for a data source. See "Configure the statement cache for a JDBC data source" in the *Oracle WebLogic Server Administration Console Help*.

10.1.1 Statement Cache Algorithms

The Statement Cache Type (or algorithm) determines which prepared and callable statements to store in the cache for each connection in a data source. You can choose from the following options:

- "LRU (Least Recently Used)"
- "Fixed"

10.1.1.1 LRU (Least Recently Used)

When you select LRU (Least Recently Used, the default) as the Statement Cache Type, WebLogic Server caches prepared and callable statements used on the connection until the statement cache size is reached. When an application calls `Connection.prepareStatement()`, WebLogic Server checks to see if the statement is stored in the statement cache. If so, WebLogic Server returns the cached statement (if it is not already being used). If the statement is not in the cache, and the cache is full (number of statements in the cache = statement cache size), WebLogic Server determines which existing statement in the cache was the least recently used and replaces that statement in the cache with the new statement.

The LRU statement cache algorithm in WebLogic Server uses an approximate LRU scheme.

10.1.1.2 Fixed

When you select FIXED as the Statement Cache Type, WebLogic Server caches prepared and callable statements used on the connection until the statement cache size is reached. When additional statements are used, they are not cached.

With this statement cache algorithm, you can inadvertently cache statements that are rarely used. In many cases, the LRU algorithm is preferred because rarely used statements will eventually be replaced in the cache with frequently used statements.

10.1.2 Statement Cache Size

The Statement Cache Size attribute determines the total number of prepared and callable statements to cache for each connection in each instance of the data source. By caching statements, you can increase your system performance. However, you must consider how your DBMS handles open prepared and callable statements. In many cases, the DBMS will maintain a cursor for each open statement. This applies to prepared and callable statements in the statement cache. If you cache too many statements, you may exceed the limit of open cursors on your database server.

For example, if you have a data source with 10 connections deployed on 2 servers, if you set the Statement Cache Size to 10 (the default), you may open 200 (10 x 2 x 10) cursors on your database server for the cached statements.

10.1.3 Usage Restrictions for the Statement Cache

Using the statement cache can dramatically increase performance, but you must consider its limitations before you decide to use it. Please note the following restrictions when using the statement cache.

There may be other issues related to caching statements that are not listed here. If you see errors in your system related to prepared or callable statements, you should set the statement cache size to 0, which turns off statement caching, to test if the problem is caused by caching prepared statements.

10.1.3.1 Calling a Stored Statement After a Database Change May Cause Errors

Prepared statements stored in the cache refer to specific database objects at the time the prepared statement is cached. If you perform any DDL (data definition language) operations on database objects referenced in prepared statements stored in the cache, the statements may fail the next time you run them. For example, if you cache a statement such as `select * from emp` and then drop and recreate the `emp` table, the next time you run the cached statement, the statement may fail because the exact `emp` table that existed when the statement was prepared, no longer exists.

Likewise, prepared statements are bound to the data type for each column in a table in the database at the time the prepared statement is cached. If you add, delete, or rearrange columns in a table, prepared statements stored in the cache are likely to fail when run again.

These limitations depend on the behavior of your DBMS.

10.1.3.2 Using `setNull` In a Prepared Statement

If you cache a prepared statement that uses a `setNull` bind variable, you must set the variable to the proper data type. If you use a generic data type, as in the following example, data may be truncated or the statement may fail when it runs with a value other than null.

```
java.sql.Types.Long sal=null
.
.
.
if (sal == null)
    setNull(2,int)//This is incorrect
else
    setLong(2,sal)
```

Instead, use the following:

```
if (sal == null)
    setNull(2,long)//This is correct
else
    setLong(2,sal)
```

10.1.3.3 Statements in the Cache May Reserve Database Cursors

When WebLogic Server caches a prepared or callable statement, the statement may open a cursor in the database. If you cache too many statements, you may exceed the limit of open cursors for a connection. To avoid exceeding the limit of open cursors for a connection, you can change the limit in your database management system or you can reduce the statement cache size for the data source.

10.2 Connection Testing Options for a Data Source

To make sure that the database connections in a data source remain healthy, you should periodically test the connections. WebLogic Server includes two basic types of testing:

- Automatic testing that you configure with options on the data source so that WebLogic Server makes sure that database connections remain healthy.
- Manual testing that you can do to trouble-shoot a data source.

The following section discusses automatic connection testing options. For more information about manual connection testing, see [Section 14.1, "Testing Data Sources and Database Connections."](#)

To configure automatic testing options for a data source, you set the following options either through the Administration Console or through WLST using the `JDBCConnectionPoolParamsBean`:

- **Test Frequency**—(`TestFrequencySeconds` in the `JDBCConnectionPoolParamsBean`) Use this attribute to specify the number of seconds between tests of unused connections. WebLogic Server tests unused connections, and closes and replaces any faulty connections. You must also set the Test Table Name.
- **Test Reserved Connections**—(`TestConnectionsOnReserve` in the `JDBCConnectionPoolParamsBean`) Enable this option to test each connection before giving to a client. This may add a slight delay to the request, but it guarantees that the connection is healthy. You must also set a Test Table Name.
- **Test Table Name**—(`TestTableName` in the `JDBCConnectionPoolParamsBean`) Use this attribute to specify a table name to use in a connection test. You can also specify SQL code to run in place of the standard test by entering SQL followed by a space and the SQL code you want to run as a test. Test Table Name is required to enable any database connection testing.
- **Seconds to Trust an Idle Pool Connection**—(`SecondsToTrustAnIdlePoolConnection` in the `JDBCConnectionPoolParamsBean`) Use this option to specify the number of seconds after a connection has been proven to be OK that WebLogic Server trusts the connection is still viable and will skip the connection test, either before delivering it to an application or during the periodic connection testing process. This option is an optimization that minimizes the performance impact of connection testing, especially during heavy traffic. See [Section 10.2.1.7, "Minimizing Connection Request Delay with Seconds to Trust an Idle Pool Connection."](#)

See the [JDBC Data Source: Configuration: Connection Pool](#) page in the Administration Console or see "`JDBCConnectionPoolParamsBean`" in the *Oracle WebLogic Server MBean Reference* for more details about these options.

For instructions to set connection testing options, see "Configure testing options for a JDBC data source" in the *Oracle WebLogic Server Administration Console Help*.

10.2.1 Database Connection Testing Semantics

When WebLogic Server tests database connections in a data source, it reserves a connection from the data source, runs a small query on the connection, then returns the connection to the pool in the data source. The server instance tracks statistics on the pool status, including the amount of time a required to complete a connection test, the number of connections waiting for a connection, and the number of connections being tested. The history of recent test connection behavior is used to calculate the amount of time the server instance waits until a connection test is determined to have failed.

If a thread appears to be taking longer than normal to connect, the server instance may delay testing on other threads until the abnormally long-running test completes. If that thread hangs too long in connection testing (10 seconds by default), a pool may declare a DBMS connectivity failure, disable itself, and kill all connections, whether unreserved or in application hands.

This is very rare, and is intended to relieve the otherwise interminable hangs that can be caused by network cable disconnects and other problems that can lock any JVM thread which is doing a call in a socket read that the JVM will be unable to break until the OS TCP limit is hit (typically 10 minutes). If a pool disables itself in this manner, it will periodically try to reconnect to the DBMS (every 5 seconds by default). Once a new connection can be made, the pool will re-enable itself, and subsequent connection requests will be served as normal (with the pool repopulating itself as load requires).

The query used in testing is determined by the value in Test Table Name. If the value is a table name, the query is `select 1 from table_name`. If Test Table Name includes a full query starting with `SQL` followed by space and the query, WebLogic Server uses that query when testing database connections.

If a connection fails the test, WebLogic Server closes and recreates the connection, and then tests the new connection.

Details about the semantics of connection testing is discussed in the following sections:

- ["Connection Testing When Database Connections are Created"](#)
- ["Periodic Connection Testing"](#)
- ["Testing Reserved Connections"](#)
- ["Minimized Connection Test Delay After Database Connectivity Loss"](#)
- ["Minimized Connection Request Delay After Connection Test Failures"](#)
- ["Minimizing Connection Request Delay with Seconds to Trust an Idle Pool Connection"](#)

10.2.1.1 Connection Testing When Database Connections are Created

When connections are created in a data source, WebLogic Server tests each connection using the query defined by the value in Test Table Name. Connections are created when a data source is deployed, either at server startup or when creating a data source, when increasing capacity to meet demand for connections, or when recreating a connection that failed a connection test.

The purpose of this testing is to ensure that new connections are viable and ready for use when an application requests a connection.

10.2.1.2 Periodic Connection Testing

If `Test Frequency` is greater than 0, WebLogic Server periodically tests the pooled connections that are not currently reserved by applications. The test is based on the query defined in Test Table Name. If a connection fails the test, WebLogic Server closes the connection, recreates the connection, and tests the new connection before returning it to the pool.

10.2.1.3 Testing Reserved Connections

When `Test Connections On Reserve` is enabled, when your application requests a connection from the data source, WebLogic Server tests the connection using the query specified in Test Table Name before giving the connection to the application. The default value is not enabled.

Testing reserved connections can cause a delay in satisfying connection requests, but it makes sure that the connection is viable when the application gets the connection. You can minimize the impact of testing reserved connections by tuning Seconds to Trust an Idle Pool Connection. See [Section 10.2.1.7, "Minimizing Connection Request Delay with Seconds to Trust an Idle Pool Connection."](#)

10.2.1.4 Minimized Connection Test Delay After Database Connectivity Loss

When connectivity to the DBMS is lost, even if only momentarily, some or all of the JDBC connections in a data source typically become defunct. If the data source is configured to test connections on reserve, when an application requests a database connection, WebLogic Server tests the connection, discovers that the connection is dead, and tries to replace it with a new connection to satisfy the request. Ordinarily, when the DBMS comes back online, the refresh process succeeds. However, in some cases and for some modes of failure, testing a dead connection can impose a long delay.

To minimize this delay, WebLogic data sources include logic that considers *all* connections in the data source as dead after a number of consecutive test failures, and closes all connections in the data source. After all connections are closed, when an application requests a connection, the data source creates a connection without first having to test a dead connection. This behavior minimizes the delay for connection requests following the data source's connection pool flush.

WebLogic Server determines the number of test failures before closing all connections based on the Test Frequency setting for the data source:

- If Test Frequency is greater than 0, the number of test failures before closing all connections is set to 2.
- If Test Frequency is set to 0 (periodic testing is disabled), the number of test failures before closing all connections is set to 25% of the Maximum Capacity for the data source.

10.2.1.5 Minimized Connection Request Delay After Connection Test Failures

If your DBMS becomes and remains unavailable, the data source will persistently test and try to replace dead connections while trying to satisfy connection requests. This behavior is beneficial because it enables the data source to react immediately when the database becomes available. However, testing a dead database connection can take as long as the network timeout, and can cause a long delay for clients.

To minimize this delay, the WebLogic data sources include logic that disables the data source after 2 consecutive failures to replace a dead connection. When an application requests a connection from a disabled data source, WebLogic Server throws a `PoolDisabledSQLException` immediately to notify the client that a connection is not available.

For data sources that are disabled in this manner, WebLogic Server periodically runs a refresh process. The refresh process does the following:

- The server instance executes a health check on the database server every 5 seconds. This setting is not configurable.
- If the server instance recognizes that the database was recovered, it creates a new database connection and enables the data source.

You can also manually enable the data source using the Administration Console.

10.2.1.6 Minimized Connection Request Delays After Loss of DBMS Connectivity

If your DBMS becomes and remains unavailable, the data source will persistently test and try to replace dead connections while trying to satisfy connection requests. This behavior is beneficial because it enables the data source to react immediately when the database becomes available. However, testing a dead database connection and trying futilely to replace it can in some cases take as long as the OS network timeout (minutes), and can cause long delays to clients before getting the expected failure message.

Note: If a data source is added to a multi data source, the multi data source takes over the responsibility of disabling and re-enabling its data sources. By default, a multi data source will check every two minutes (configurable) and re-enable any of its data sources that can re-establish connections.

10.2.1.7 Minimizing Connection Request Delay with Seconds to Trust an Idle Pool Connection

For some applications that use DBMS connections in a lot of very short cycles (such as reserve-do_one_query-close), the data source's testing of the connection can contribute a significant amount of overhead to each use cycle. To minimize the impact of connection testing, you can set the Seconds To Trust An Idle Pool Connection attribute in the JDBC data source configuration to trust recently-used or recently-tested database connections and skip the connection test.

If Test Reserved Connections is enabled on your data source, when an application requests a database connection, WebLogic Server tests the database connection before giving it to the application. If the request is made within the time specified for Seconds to Trust an Idle Pool Connection, since the connection was tested or successfully used by an application, WebLogic Server skips the connection test before delivering it to an application.

If Test Frequency is greater than 0 for your data source (periodic testing is enabled), WebLogic Server also skips the connection test if the connection was successfully used and returned to the data source within the time specified for Seconds to Trust an Idle Pool Connection.

For instructions to set Seconds to Trust an Idle Pool Connection, see "Configure testing options for a JDBC data source" in the *Oracle WebLogic Server Administration Console Help*.

Seconds to Trust an Idle Pool Connection is a tuning feature that can improve application performance by minimizing the delay caused by database connection testing, especially during heavy traffic. However, it can reduce the effectiveness of connection testing, especially if the value is set too high. The appropriate value depends on your environment and the likelihood that a connection will become defunct.

10.2.2 Database Connection Testing Configuration Recommendations

You should set connection testing attributes so that they best fit your environment. For example, if your application cannot tolerate database connection failures, you should set Seconds to Trust an Idle Pool Connection to 0 and make sure Test Reserved Connections is enabled so that WebLogic Server will test every connection before giving it to an application. If your application is more sensitive to delays in getting a connection from the data source and can tolerate a possible application failure due to

using a dead connection, you should set Seconds to Trust an Idle Pool Connection to a higher number, set Test Frequency to a lower number, and enable Test Reserved Connections.

With these settings, your application will rely more on the data source testing connections in the pool when they are not in use, rather than when an application requests a connection.

Note: Ultimately, even if WebLogic does its best, a connection may fail in the instant after WebLogic successfully tested it, and just before the application uses it. Therefore, every application should be written to respond appropriately in the case of unexpected exceptions from a dead connection.

10.2.3 Default Test Table Name

When you create a data source using the Administration Console, the Administration Console automatically sets the Test Table Name attribute for a data source based on the DBMS that you select. The Test Table Name attribute is used in connection testing which is optionally performed periodically or when you create or reserve a connection, depending on how you configure the testing options. For database tests to succeed, the database user used to create database connections in the data source must have access to the database table. If not, you should either grant access to the user (make this change in the DBMS) or change the Test Table Name attribute to the name of a table to which the user does have access (make this change in the WebLogic Server Administration Console).

The Test Table Name is an overloaded parameter. Its simplest form is to name a table that WLS will query to test a connection. Setting it to any table, such as "DUAL" for Oracle, will cause the data source to run the query `select count(*) from DUAL`. If used in this mode, Oracle recommends that you choose a small, infrequently updated table (preferably a pseudo-table such as DUAL).

The second manner in which you can define this parameter is to allow any specific SQL string to be executed to test the connection. To use this option, set the parameter to "SQL " plus the desired sql string. For example `SQL select 1` works for SQLServer, which does not need a table in queries to select constants. This option is useful for adding DBMS-side control of WLS pool connection testing, and to make the test as fast as possible.

Table 10–1 Default Test Table Name by DBMS

DBMS	Default Test Table Name (Query)
Adabas for z/OS	SQL call shadow_adabas('select * from employees')
Cloudscape	SQL SELECT 1
DB2	SQL SELECT COUNT(*) FROM SYSIBM.SYSTABLES
FirstSQL	SQL SELECT 1
IMS/TM for z/OS	SQL call shadow_ims('otm','/dis','ctl')
Informix	SQL SELECT COUNT(*) FROM SYSTABLES
Microsoft SQL Server	SQL SELECT 1
MySQL	SQL SELECT 1
Oracle	SQL SELECT 1 FROM DUAL

Table 10–1 (Cont.) Default Test Table Name by DBMS

DBMS	Default Test Table Name (Query)
PostgreSQL	SQL SELECT 1
Progress	SQL SELECT COUNT(*) FROM SYSTABLES
Sybase	SQL SELECT 1

10.3 Enabling Connection Creation Retries

WebLogic JDBC data sources include the Connection Creation Retry Frequency option (`ConnectionCreationRetryFrequencySeconds` in the `JDBCConnectionPoolParamsBean`) that you can use to specify the number of seconds between attempts to establish connections to the database. If set and if the database is unavailable when the data source is created, WebLogic Server attempts to create connections in the pool again after the number of seconds you specify, and will continue to attempt to create the connections until it succeeds. This option applies to connections created when the data source is created at server startup or when the data source is deployed or if the initial capacity is increased. It does *not* apply to connections created for pool expansion or to replace a defunct connection in the pool.

By default, Connection Creation Retry Frequency is 0 seconds. When the value is set to 0, connection creation retries is disabled and data source creation fails if the database is unavailable.

See the JDBC Data Source: Configuration: Connection Pool page in the Administration Console or see "JDBCConnectionPoolParamsBean" in the *Oracle WebLogic Server MBean Reference* for more details about this option.

10.4 Enabling Connection Requests to Wait for a Connection

JDBC data sources have two attributes that you can set to enable connection requests to wait for a connection from a data source: Connection Reserve Timeout (`ConnectionReserveTimeoutSeconds`) and Maximum Waiting for Connection (`HighestNumWaiters`). You use these two attributes together to enable connection requests to wait for a connection without disabling your system by blocking too many threads.

See the JDBC Data Source: Configuration: Connection Pool page in the Administration Console or see "JDBCConnectionPoolParamsBean" in the *Oracle WebLogic Server MBean Reference* for more details about these options.

Also see "Enable connection requests to wait for a connection" in the *Administration Console Online Help*.

10.4.1 Connection Reserve Timeout

When an application requests a connection from a data source, if all connections in the data source are in use and if the data source has expanded to its maximum capacity, the application will get a Connection Unavailable SQL Exception. To avoid this, you can configure the Connection Reserve Timeout value (in seconds) so that connection requests will wait for a connection to become available. After the Connection Reserve Timeout has expired, if no connection becomes available, the request will fail and the application will get a `PoolLimitsSQLException` exception.

If you set Connection Reserve Timeout to `-1`, a connection request will timeout immediately if there is no connection available. If you set Connection Reserve Timeout to `0`, a connection request will wait indefinitely. The default value is 10 seconds.

See "Enable connection requests to wait for a connection" in the *Oracle WebLogic Server Administration Console Help*.

10.4.2 Limiting the Number of Waiting Connection Requests

Connection requests that wait for a connection block a thread. If too many connection requests concurrently wait for a connection and block threads, your system performance can degrade. To avoid this, you can set the Maximum Waiting for Connection (`HighestNumWaiters`) attribute, which limits the number connection requests that can concurrently wait for a connection.

If you set Maximum Waiting for Connection (`HighestNumWaiters`) to `MAX-INT` (the default), there is effectively no bound on how many connection requests can wait for a connection. If you set Maximum Waiting for Connection to `0`, connection requests cannot wait for a connection. If the maximum number of requests has been met, a `SQLException` is thrown when an application requests a connection.

10.5 Automatically Recovering Leaked Connections

A leaked connection is a connection that was not properly returned to the connection pool in the data source. To automatically recover leaked connections, you can specify a value for Inactive Connection Timeout on the JDBC Data Source:

Configuration: Connection Pool page in the Administration Console. When you set a value for Inactive Connection Timeout, WebLogic Server forcibly returns a connection to the data source when there is no activity on a reserved connection for the number of seconds that you specify. When set to `0` (the default value), this feature is turned off.

See the JDBC Data Source: Configuration: Connection Pool page in the Administration Console or see "JDBCConnectionPoolParamsBean" in the *Oracle WebLogic Server MBean Reference* for more details about this option.

Note that the actual timeout could exceed the configured value for Inactive Connection Timeout. The internal data source maintenance thread runs every 5 seconds. When it reaches the Inactive Connection Timeout (for example 30 seconds), it checks for inactive connections. To avoid timing out a connection that was reserved just before the current check or just after the previous check, the server gives an inactive connection a "second chance." On the next check, if the connection is still inactive, the server times it out and forcibly returns it to the data source. On average, there could be a delay of 50% more than the configured value.

10.6 Avoiding Server Lockup with the Correct Number of Connections

When your applications attempt to get a connection from a data source in which there are no available connections, the data source throws an exception stating that a connection is not available in the data source. To avoid this error, make sure your data source can expand to the size required to accommodate your peak load of connection requests. To increase the maximum number of connections available in the data source, increase the value for Maximum Capacity for the data source on the JDBC Data Source: Configuration: Connection Pool page in the Administration Console.

10.7 Limiting Statement Processing Time with Statement Timeout

With the Statement Timeout option on a JDBC data source, you can limit the amount of time that a statement takes to execute on a database connection reserved from the data source. When you set a value for Statement Timeout, WebLogic Server passes the time specified to the JDBC driver using the

`java.sql.Statement.setQueryTimeout()` method. WebLogic Server will make the call, and if the driver throws an exception, the value will be ignored. In some cases, the driver may silently not support the call, or may document limited support. Oracle recommends that you check the driver documentation to verify the expected behavior.

When Statement Timeout is set to -1, (the default) statements do not timeout.

See the JDBC Data Source: Configuration: Connection Pool page in the Administration Console or see "JDBCConnectionPoolParamsBean" in the *Oracle WebLogic Server MBean Reference* for more details about this option.

10.8 Using Pinned-To-Thread Property to Increase Performance

To minimize the time it takes for an application to reserve a database connection from a data source and to eliminate contention between threads for a database connection, you can add the Pinned-To-Thread property in the connection Properties list for the data source, and set its value to `true`.

When Pinned-To-Thread is enabled, WebLogic Server pins a database connection from the data source to an execution thread the first time an application uses the thread to reserve a connection. When the application finishes using the connection and calls `connection.close()`, which otherwise returns the connection to the data source, WebLogic Server keeps the connection with the execute thread and does not return it to the data source. When an application subsequently requests a connection using the same execute thread, WebLogic Server provides the connection already reserved by the thread. There is no locking contention on the data source that occurs when multiple threads attempt to reserve a connection at the same time and there is no contention for threads that attempt to reserve the same connection from a limited number of database connections.

Note: In this release, the Pinned-To-Thread feature does not work with multi data sources, Oracle RAC, and IdentityPool. These features rely on the ability to return a connection to the connection pool and reacquire it if there is a connection failure or connection identity does not match.

See "JDBC Data Source: Configuration: Connection Pool" in the *Oracle WebLogic Server Administration Console Help*.

10.8.1 Changes to Connection Pool Administration Operations When PinnedToThread is Enabled

Because the nature of connection pooling behavior is changed when PinnedToThread is enabled, some connection pool attributes or features behave differently or are disabled to suit the behavior change:

- **Maximum Capacity** is ignored. The number of connections in a connection pool equals the greater of either the initial capacity or the number of connections reserved from the connection pool.

- **Shrinking** does not apply to connection pools with `PinnedToThread` enabled because connections are never returned to the connection pool. Effectively, they are always reserved.
- When you **Reset** a connection pool, the reset connections from the connection pool are marked as `Test Needed`. The next time each connection is reserved, WebLogic Server tests the connection and recreates it if necessary. Connections are not tested synchronously when you reset the connection pool. This feature requires that `Test Connections on Reserve` is enabled and a `Test Table Name` or query is specified.

10.8.2 Additional Database Resource Costs When `PinnedToThread` is Enabled

When `PinnedToThread` is enabled, the maximum capacity of the connection pool (maximum number of database connections created in the connection pool) becomes the number of execute threads used to request a connection multiplied by the number of concurrent connections each thread reserves. This may exceed the `Maximum Capacity` specified for the connection pool. You may need to consider this larger number of connections in your system design and ensure that your database allows for additional associated resources, such as open cursors.

Also note that connections are never returned to the connection pool, which means that the connection pool can never shrink to reduce the number of connections and associated resources in use. You can minimize this cost by setting an additional driver parameter `onePinnedConnectionOnly`. When `onePinnedConnectionOnly=true`, only the first connection requested is pinned to the thread. Any additional connections required by the thread are taken from and returned to the connection pool as needed. Set `onePinnedConnectionOnly` using the `Properties` attribute, for example:

```
Properties="PinnedToThread=true;onePinnedConnectionOnly=true;user=examples"
```

If your system can handle the additional resource requirements, Oracle recommends that you use the `PinnedToThread` option to increase performance.

If your system cannot handle the additional resource requirements or if you see database resource errors after enabling `PinnedToThread`, Oracle recommends *not* using `PinnedToThread`.

10.9 Using Unwrapped Data Type Objects

Some JDBC objects from a driver that are returned from WebLogic Server are wrapped by default. Wrapping data source objects provides WebLogic Server the ability to:

- Generate debugging output from all method calls.
- Track connection utilization so that connections can be timed out appropriately.
- Provide transparent automatic transaction enlistment and security authorization.

WebLogic Server provides the ability to disable wrapping which provides the following benefits:

- Although WebLogic Server generates a dynamic proxy for vendor methods that implement an interface to show through the wrapper, some data types do not implement an interface. For example, Oracle data types `Array`, `Blob`, `Clob`, `NClob`, `Ref`, `SQLXML`, and `Struct` are classes that do not implement interfaces. Disabling wrapping allows applications to use native driver objects directly.
- Eliminating wrapping overhead can provide a significant performance improvement.

When wrapping is disabled (the `wrap-types` element is `false`), the following data types are not wrapped:

- Array
- Blob
- Clob
- NClob
- Ref
- SQLXML
- Struct
- ParameterMetaData
 - No connection testing performed.
- ResultSetMetaData
 - No connection testing performed.
 - No result set testing performed.
 - No JDBC MT profiling performed.

10.9.1 How to Disable Wrapping

You can use the Administration Console and WLST to disable data type wrapping.

10.9.1.1 Disable Wrapping using the Administration Console

To disable wrapping of JDBC data type objects:

1. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
2. In the **Domain Structure** tree, expand **Services**, then select **Data Sources**.
3. On the Summary of Data Sources page, click the data source name.
4. Select the **Configuration: Connection Pool** tab.
5. Scroll down and click **Advanced** to show the advanced connection pool options.
6. In **Wrap Data Types**, deselect the checkbox to disable wrapping.
7. Click Save.
8. To activate these changes, in the Change Center of the Administration Console, click **Activate Changes**.

This change does not take effect immediately—it requires that the data source be redeployed or the server be restarted.

10.9.1.2 Disable Wrapping using WLST

The following is a WLST code snippet to disable data type wrapping:

```
. . .
jdbcSR = create(dsname, "JDBCSystemResource");
theJDBCResource = jdbcSR.getJDBCResource();
poolParams = theJDBCResource.getJDBCConnectionPoolParams();
poolParams.setWrapTypes(false);
. . .
```

This change does not take effect immediately—it requires that the data source be redeployed or the server be restarted.

Using WebLogic Server with Oracle RAC

Oracle WebLogic Server provides strong support for Oracle Real Application Clusters (RAC) features in Oracle Database 11g, minimizing database access time while allowing transparent access to rich pooling management functions that maximizes both connection performance and availability.

The following sections describe the requirements and configuration tasks for using Oracle Real Application Clusters with WebLogic Server:

- [Section 11.1, "Overview of Oracle Real Application Clusters"](#)
- [Section 11.2, "Software Requirements"](#)
- [Section 11.3, "JDBC Driver Requirements"](#)
- [Section 11.4, "Hardware Requirements"](#)
- [Section 11.5, "Configuration Options in WebLogic Server with Oracle RAC"](#)

Both Oracle RAC and WebLogic Server are complex systems. To use them together requires specific configuration on both systems, as well as clustering software and a shared storage solution. This document describes the configuration required at a high level. For more details about configuring Oracle RAC, your clustering software, your operating system, and your storage solution, see the documentation from the respective vendors.

11.1 Overview of Oracle Real Application Clusters

Oracle Real Application Clusters (Oracle RAC) is a software component you can add to a high-availability solution that enables users on multiple machines to access a single database with increased performance. Oracle RAC comprises two or more Oracle database instances running on two or more clustered machines and accessing a shared storage device via cluster technology. To support this architecture, the machines that host the database instances are linked by a high-speed interconnect to form the cluster. The interconnect is a physical network used as a means of communication between the nodes of the cluster. Cluster functionality is provided by the operating system or compatible third party clustering software.

An Oracle RAC installation appears like a single standard Oracle database and is maintained using the same tools and practices. All the nodes in the cluster execute transactions against the same database and Oracle RAC coordinates each node's access to the shared data to maintain consistency and ensure integrity. You can add nodes to the cluster easily and there is no need to partition data when you add them. This means that you can horizontally scale the database tier as usage and demand grows by adding Oracle RAC nodes, storage, or both.

11.2 Software Requirements

To use WebLogic Server with Oracle RAC, you must install the following software on each Oracle RAC node:

- Operating system patches required to support Oracle RAC. See the release notes from Oracle for details.
- Oracle 10g or 11g database management system. See *Oracle® Fusion Middleware Licensing Information*.
- Clustering software for your operating system. See the Oracle documentation for supported clustering software and cluster configurations.
- Shared storage software, such as Veritas Cluster File System. Note that some clustering software includes a file storage solution, in which case additional shared storage software is not required.

Note: See the *Oracle Fusion Middleware Supported System Configurations* page at http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html for the latest WebLogic Server hardware platform and operating system support, and for the Oracle RAC versions supported by WebLogic Server versions and service packs. See the Oracle documentation for hardware and software requirements required for running the Oracle RAC software.

11.3 JDBC Driver Requirements

To use WebLogic Server with Oracle RAC, your WebLogic JDBC data sources must use the Oracle JDBC Thin driver 11g to create database connections.

11.4 Hardware Requirements

A typical WebLogic Server/Oracle RAC system includes a WebLogic Server cluster, an Oracle RAC cluster, and hardware for shared storage.

11.4.1 WebLogic Server Cluster

The WebLogic Server cluster can be configured in many ways and with various hardware options. See *Using Clusters for Oracle WebLogic Server* for more details about configuring a WebLogic Server cluster.

11.4.2 Oracle RAC Cluster

For the latest hardware requirements for Oracle RAC, see the Oracle RAC documentation. However, to use Oracle RAC with WebLogic Server, you must run Oracle RAC instances on robust, production-quality hardware. The Oracle RAC configuration must deliver database processing performance appropriate for reasonably-anticipated application load requirements. Unusual database response delays can lead to unexpected behavior during database failover scenarios.

11.4.3 Shared Storage

In an Oracle RAC configuration, all data files, control files, and parameter files are shared for use by all Oracle RAC instances. An HA storage solution that uses one of the following architectures is recommended:

- Direct Attached Storage (DAS), such as a dual ported disk array or a Storage Area Network (SAN)
- Network Attached Storage (NAS)

For a complete list of supported storage solutions, see your Oracle documentation.

11.5 Configuration Options in WebLogic Server with Oracle RAC

When using WebLogic Server with Oracle RAC, you must configure your WebLogic Domain so that it can interact with Oracle RAC instances and so that it performs as expected. The following sections describe configuration options and requirements:

- [Section 11.5.1, "Choosing a WebLogic Server Configuration for Use with Oracle RAC"](#)
- [Section 11.5.2, "Validating Connections when using WebLogic Server with Oracle RAC"](#)
- [Section 11.5.3, "Additional Considerations When Using WebLogic Server with Oracle RAC"](#)

11.5.1 Choosing a WebLogic Server Configuration for Use with Oracle RAC

Consider the following alternatives:

- Using GridLink data sources, see *Oracle® Fusion Middleware Licensing Information*.
- To connect to multiple Oracle RAC 10g or Oracle RAC 11g instances when using global transactions (XA), Oracle recommends the use of transaction-aware WebLogic JDBC multi data sources, which support failover and load balancing, to connect to the Oracle RAC nodes. For more information see [Section B.6, "Using Multi Data Sources with Global Transactions."](#)
- To connect to multiple Oracle RAC 10g or Oracle RAC 11g instances when not using XA, Oracle recommends the use of (non-transaction-aware) multi data sources to connect to the Oracle RAC nodes. Use the standard multi data source configuration, which supports failover and load balancing. For more information see [Section B.7, "Using Multi Data Sources without Global Transactions."](#)

The following table may help you as you try to determine which configuration is right for your particular application:

Table 11–1 Choosing Configurations to Use with Oracle RAC

Requires Load Balancing?	Requires Failover?	Requires Global Transactions (XA)?	Uses Oracle RAC Services	See...
Yes	Yes	Yes	Yes	Section 4, "Using GridLink Data Sources"
Yes	Yes	Yes	No	Section B.6, "Using Multi Data Sources with Global Transactions"

Table 11-1 (Cont.) Choosing Configurations to Use with Oracle RAC

Requires Load Balancing?	Requires Failover?	Requires Global Transactions (XA)?	Uses Oracle RAC Services	See...
Yes	Yes	Yes	Yes	Section B.8, "Configuring Connections to Services on Oracle RAC Nodes."
Yes	Yes	No	Yes	Section B.8, "Configuring Connections to Services on Oracle RAC Nodes."
Yes	Yes	No	No	Section B.7, "Using Multi Data Sources without Global Transactions"

11.5.2 Validating Connections when using WebLogic Server with Oracle RAC

Applications can use the JDBC 4.0 `Connection.isValid` API to verify connection viability.

Note: WebLogic Server does not support `oracle.ucp.jdbc.ValidConnection.isValid` or `oracle.ucp.jdbc.ValidConnection.setInvalid`.

11.5.3 Additional Considerations When Using WebLogic Server with Oracle RAC

Dynamic Transaction Processing (DTP) is not supported to coordinate transactions when using GridLink Data Sources or multi data sources with Oracle RAC.

Using JDBC Drivers with WebLogic Server

WebLogic Server works with JDBC drivers that offer the following functionality:

- Are thread-safe
- Can implement transactions using standard JDBC statements

JDBC drivers that do not implement `Serializable` or `Remote` interfaces cannot pass objects to a remote client application.

This section describes how to set up and use JDBC drivers with WebLogic Server. It includes the following sections:

- [Section 12.1, "JDBC Drivers Installed with WebLogic Server"](#)
- [Section 12.2, "Setting the Environment for a Type-4 Third-Party JDBC Driver"](#)
- [Section 12.3, "Globalization Support for the Oracle 11g Thin Driver"](#)
- [Section 12.4, "Using the Oracle Thin Driver in Debug Mode"](#)
- [Section 12.5, "Using the Oracle Thin Driver Over SSL"](#)

12.1 JDBC Drivers Installed with WebLogic Server

The 11g version of the Oracle Thin driver (`ojdbc6.jar` for JDK 6) is installed with Oracle WebLogic Server.

Note: The Oracle Type 4 JDBC drivers from DataDirect are also installed with WebLogic Server. See *Type 4 JDBC Drivers for Oracle WebLogic Server* for more information.

In addition to the Oracle Thin Driver, the `mysql-connector-java-commercial-5.0.x` (`mysql-connector-java-commercial-5.0.x-bin.jar`) JDBC driver is installed with WebLogic Server.

This driver is installed in the `WL_HOME\server\lib` folder (where `WL_HOME` is the folder where WebLogic Server is installed) with `weblogic.jar`. The manifest in `weblogic.jar` lists this file so that it is loaded when `weblogic.jar` is loaded (when the server starts). Therefore, you do not need to add this JDBC driver to your `CLASSPATH`. If you plan to use a third-party JDBC driver that is not installed with WebLogic Server, you must install the drivers, which includes updating your `CLASSPATH` with the path to the driver files, and may include updating your `PATH` with the path to database client files. See "Supported Database Configurations" on the *Oracle Fusion Middleware Supported System Configurations* page at

http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html.

If you plan to use a different version of any of the drivers installed with WebLogic Server, you can replace the driver file in `WL_HOME\server\lib` with an updated version of the file or add the new file to the front of your `CLASSPATH`.

Copies of the MySQL and Oracle Thin drivers installed with WebLogic Server and other supporting files are installed in `WL_HOME\server\ext\jdbc\`. There is a subdirectory in this folder for each DBMS. If you need to revert to the version of the driver installed with WebLogic Server, you can copy the file from `WL_HOME\server\ext\jdbc\DBMS` to `WL_HOME\server\lib`.

Note: WebLogic Server also includes a version of the Derby DBMS installed with the WebLogic Server examples in the `WL_HOME\common\derby` directory. Derby is an all-Java DBMS product included in the WebLogic Server distribution solely in support of demonstrating the WebLogic Server examples. For more information about Derby, see <http://db.apache.org/derby>.

12.2 Setting the Environment for a Type-4 Third-Party JDBC Driver

If you plan to use a third-party JDBC driver that is not installed with WebLogic Server, you need to update the WebLogic Server's classpath to include the location of the JDBC driver classes. Edit the `commEnv.cmd/sh` script in `WL_HOME/common/bin` and prepend your classes as described in "Modifying the Classpath" in the *Command Reference for Oracle WebLogic Server*.

12.3 Globalization Support for the Oracle 11g Thin Driver

For Globalization Support with the 11g version of the Oracle Thin driver, Oracle supplies the `orai18n.jar` file, which replaces `nls_charset.zip`. If you use character sets other than US7ASCII, WE8DEC, WE8ISO8859P1 and UTF8 with CHAR and NCHAR data in Oracle object types and collections, you must include `orai18n.jar` and `orai18n-mapping.jar` in your `CLASSPATH`.

The `orai18n.jar` and `orai18n-mapping.jar` are included with the WebLogic Server installation in the `WL_HOME\server\ext\jdbc\oracle\11g` folder. These files are *not* referenced by the `weblogic.jar` manifest file, so you must add them to your `CLASSPATH` before they can be used.

12.4 Using the Oracle Thin Driver in Debug Mode

The `WL_HOME\server\ext\jdbc\oracle\11g` folder includes the `ojdbc6_g.jar` (for JDK 6) and `ojdbc5_g.jar` (for JDK 5) files, which are the versions of the Oracle Thin driver with classes to support debugging and tracing. To use the Oracle Thin driver in debug mode, add the path to these files at the beginning of your `CLASSPATH`.

12.5 Using the Oracle Thin Driver Over SSL

This section describes how to configure the data source in WebLogic Server when using the Oracle Thin driver over SSL.

To configure the data source in WebLogic Server, perform the following steps.

1. In the Administration Console navigation pane, click the plus sign to expand Services.
2. Click the plus sign to expand JDBC.
3. Select Data Sources.
4. On the Summary of JDBC Data Sources page, select New.
5. On the Create a New JDBC Data Source page, provide a JDBC data source name and, optionally, the JNDI name.
Select Oracle as the database type, and the Oracle database driver (Thin XA or Thin) you want to use.
6. Click Next.
7. Select the transaction options, if any. Click Next.
8. On the Connection Properties page, enter the Database Name, the Host Name, the Port, and the Database User Name and password. These values are used to construct a URL for the data source. You will modify this URL on the next page, but you need to enter values now to proceed. Click Next.
9. On the Test Configuration page, make the following changes:
 - a. Verify that the driver Class Name is *oracle.jdbc.xa.client.OracleXADataSource* or *oracle.jdbc.OracleDriver*.
 - b. Modify the URL to use the following syntax and format:
 URL : jdbc:oracle:thin:@(DESCRIPTION = (ADDRESS_LIST = (ADDRESS = (PROTOCOL = tcps) (HOST = <hostname>) (PORT = <sslport>))) (CONNECT_DATA = (SERVICE_NAME = <service_name>)) (SECURITY=(SSL_SERVER_CERT_DN=\"CN=<.....>\")))

 (*SECURITY=(SSL_SERVER_CERT_DN=\"CN=<.....>\")*) is optional and is used only for server DN matching.
 - c. Enter the following properties:


```
User =<username> (Should already be present.)
javax.net.ssl.trustStore=<absolute path location of jks file>
javax.net.ssl.trustStoreType=JKS
javax.net.ssl.keyStore=<absolute path location of jks file>
javax.net.ssl.keyStoreType=JKS
javax.net.ssl.keyStorePassword=<password>
oracle.net.ssl_server_dn_match=true
```


javax.net.ssl.keyStore, *javax.net.ssl.keyStoreType*, and *javax.net.ssl.keyStorePassword* are needed only if the *listener.ora* and *sqlnet.ora* in the database are set to have *SSL_CLIENT_AUTHENTICATION=TRUE* for two-way SSL.

oracle.net.ssl_server_dn_match needs to be set only if server DN matching is needed.
10. Enter (or accept) the Test Table Name.
11. Optionally, click Test Configuration to test the configuration.
12. Click Finish.

Monitoring WebLogic JDBC Resources

This release of WebLogic Server includes the WebLogic Diagnostic Framework, which is a monitoring and diagnostic service that runs within the WebLogic Server process and participates in the standard server life cycle. This service enables you to create, collect, analyze, archive, and access diagnostic data generated by a running server and the applications deployed within its containers. This data provides insight into the run-time performance of servers and applications and enables you to isolate and diagnose faults when they occur. WebLogic JDBC takes advantage of this service to provide enhanced run-time statistics, profile information over a period of time, logging, and debugging to help you keep your WebLogic domain running smoothly.

You can use the run-time statistics to monitor the data sources in your WebLogic domain to see if there is a problem. If there is a problem, you can use profiling to determine which application is the source of the problem. Once you've narrowed it down to the application, you can then use JDBC debugging features to find the problem within the application.

The following sections include details about monitoring JDBC objects:

- [Section 13.1, "Viewing Run-Time Statistics"](#)
- [Section 13.2, "Collecting Profile Information"](#)
- [Section 13.3, "Debugging JDBC Data Sources"](#)

For more information about the WebLogic Diagnostic Framework, see *Configuring and Using the Diagnostics Framework for Oracle WebLogic Server*.

13.1 Viewing Run-Time Statistics

Viewing run-time statistics allows you to monitor the data sources in your WebLogic domain.

13.1.1 Data Source Statistics

You can view run-time statistics for a data source via the Administration Console or through the `JBCDataSourceRuntimeMBean`. The `JDBCDataSourceRuntimeMBean` provides methods for getting the current state of the data source and for getting statistics about the data source, such as the average number of active connections, the current number of active connections, the highest number of active connections, and so forth. For more information, see "`JDBCDataSourceRuntimeMBean`" in the *Oracle WebLogic Server MBean Reference*.

13.1.2 Prepared Statement Cache Statistics

You can view run-time statistics for a prepared statement cache via the Administration Console or through the `JDBCDataSourceRuntimeMBean`. For more information, see "JDBCDataSourceRuntimeMBean" in the *Oracle WebLogic Server MBean Reference*.

13.2 Collecting Profile Information

If the statistics that you are seeing indicate that there is a problem in your WebLogic Server domain, you can configure any data source to collect profile information to help you pinpoint the source of the problem. The collected profile information is stored in records in the WLDF Archive. The fields contain different information for different profile types, as described in the sections that follow.

When configuring your data source for profiling, you must specify the interval at which profile data is harvested (`Harvest Frequency Seconds`); if the interval is set to 0, harvesting of data is disabled. For more information, see "Configuring the Harvester for Metric Collection" in *Configuring and Using the Diagnostics Framework for Oracle WebLogic Server*.

13.2.1 Profile Types

You can choose to profile the following information about data sources and the prepared statement cache, as described in the next sections of this document:

- ["Connection Usage \(PROFILE_TYPE_CONN_USAGE_STR\)"](#)
- ["Connection Reservation Wait \(PROFILE_TYPE_CONN_RESV_WAIT_STR\)"](#)
- ["Connection Reservation Failed \(PROFILE_TYPE_CONN_RESV_FAIL_STR\)"](#)
- ["Connection Leak \(PROFILE_TYPE_CONN_LEAK_STR\)"](#)
- ["Connection Last Usage \(PROFILE_TYPE_CONN_LAST_USAGE_STR\)"](#)
- ["Connection Multithreaded Usage \(PROFILE_TYPE_CONN_MT_USAGE_STR\)"](#)
- ["Statement Cache Entry \(PROFILE_TYPE_STMT_CACHE_ENTRY_STR\)"](#)
- ["Statements Usage \(PROFILE_TYPE_STMT_USAGE_STR\)"](#)

13.2.1.1 Connection Usage (PROFILE_TYPE_CONN_USAGE_STR)

Enable connection usage profiling to collect information about threads currently using connections from the pool of connections in the data source. This profile information can help determine why applications are unable to get connections from the data source.

Note: By default, enabling connection usage profiling on its own will not provide a stack trace of the threads using the connections. To obtain this information you must enable profiling of connection leaks in addition to enabling connection. For more information about profiling connection leaks see [Section 13.2.1.4, "Connection Leak \(PROFILE_TYPE_CONN_LEAK_STR\)."](#)

The record contains the following information:

- `PoolName` - name of the data source to which this connection belongs
- `ID` - connection ID
- `User` - stack trace of the thread using the connection

- Timestamp - time stamp showing when the connection was given to the thread

13.2.1.2 Connection Reservation Wait (PROFILE_TYPE_CONN_RESV_WAIT_STR)

Enable connection reservation wait profiling to collect information about threads currently waiting to reserve a connection from the data source. This profile information can help determine why applications are unable to get connections from the data source or to wait for connections. The record contains the following information:

- PoolName - name of the data source to which this connection belongs
- ID - thread ID
- User - stack trace of the thread waiting for the connection
- Timestamp - time stamp showing when the thread started waiting for a connection

13.2.1.3 Connection Reservation Failed (PROFILE_TYPE_CONN_RESV_FAIL_STR)

Enable connection reservation failure profiling to collect information about threads that attempt to reserve a connection from the data source but fail to get that connection. This profile information can help determine why applications are unable to get connections from the data source even after reserving them. The record contains the following information:

- PoolName - name of the data source to which this connection belongs
- ID - thread ID
- User - stack trace of the thread waiting for the connection plus the exception received when the reservation request failed
- Timestamp - time stamp showing when the reservation request failed

13.2.1.4 Connection Leak (PROFILE_TYPE_CONN_LEAK_STR)

Enable connection leak profiling to collect information about threads that have reserved a connection from the data source and the connection leaked (was not properly returned to the pool of connections). This profile information can help determine which applications are not properly closing JDBC connections. The record contains the following information:

- PoolName - name of the data source to which this connection belongs
- ID - connection ID
- User - stack trace of the thread waiting for the connection
- Timestamp - time stamp showing when the connection leak was detected

13.2.1.5 Connection Last Usage (PROFILE_TYPE_CONN_LAST_USAGE_STR)

Enable connection last usage profiling to collect information about the previous thread that last used the connection. This information is useful when you are debugging problems with connections infected in pending transactions that cause subsequent XA operations on the connections to fail. The record contains the following information:

- PoolName - name of the data source to which this connection belongs
- ID - stack trace of the XA exception thrown
- User - stack trace of the thread that last used the connection

- Timestamp - timestamp showing when the exception was thrown

13.2.1.6 Connection Multithreaded Usage (PROFILE_TYPE_CONN_MT_USAGE_STR)

Enable connection multithreaded usage profiling to collect information about threads that erroneously use a connection that was previously obtained by a different thread. This information is useful when an application reports a problem that you suspect may have been caused by the simultaneous use of a connection by more than one thread. The record contains the following information:

- PoolName - name of the data source to which this connection belongs
- ID - stack trace of the other thread that was found using the connection
- User - stack trace of the thread that reserved the connection
- Timestamp - time stamp showing when usage of the connection by multiple threads was detected

13.2.1.7 Statement Cache Entry (PROFILE_TYPE_STMT_CACHE_ENTRY_STR)

Enable statement cache entry profiling to collect information for prepared and callable statements added to the statement cache, and for the threads that originated the cached statements. This information can help you determine how the cache is being used. The record contains the following information:

- PoolName - name of the data source to which this connection belongs
- ID - string representation of the statement
- User - stack trace of the thread using the statement
- Timestamp - time stamp showing when the statement was added to the cache

13.2.1.8 Statements Usage (PROFILE_TYPE_STMT_USAGE_STR)

Enable statements usage profiling to collect information about threads currently executing SQL statements from the statement cache. This information can help you determine how statements are being used. The record contains the following information:

- PoolName - name of the data source to which this connection belongs
- ID - SQL statement being executed via the statement
- User - stack trace of the thread using the statement
- Timestamp - duration of statement execution

13.2.2 Accessing Diagnostic Data

Once profile data has been collected, you can access it using code similar to that shown in the sample below, substituting the appropriate profile type for the one shown in boldface type. The first part of the code defines a vector variable (`dataV`) in which the collected profile data (`jdbcProfData`) is stored. The second part of the code gets the stored data from the vector variable and prints it.

```
}
jdbcProfData = getData("EventsDataArchive", "TYPE LIKE '%" +
JDBCLegalHelper.PROFILE_TYPE_CONN_USAGE_STR + "%'");
Vector dataV = new Vector();
for (int i = 0; i < jdbcProfData.length; i++) {
```

```

if (jdbcProfData[i].getPoolName().equalsIgnoreCase
(Data_Source_1))
dataV.add(jdbcProfData[i]);
}

System.out.println("records found for PROFILE_TYPE_CONN_USAGE_STR : " +
dataV.size());
for (int a = 0; a < dataV.size(); a++) {
System.out.println("ID : " + ((ProfileDataRecord)
dataV.get(a)).getId());
System.out.println("PoolName : " + ((ProfileDataRecord)
dataV.get(a)).getPoolName());
System.out.println("Time : " + ((ProfileDataRecord)
dataV.get(a)).getTimestamp());
System.out.println("User : " + ((ProfileDataRecord)
dataV.get(a)).getUser());
}

```

Another method for accessing diagnostic data is to use the Data Accessor component of the WebLogic Diagnostic Framework (WLDF); for more information see "Accessing Diagnostic Data With the Data Accessor" in *Configuring and Using the Diagnostics Framework for Oracle WebLogic Server*

13.2.3 Callbacks for Monitoring Driver-Level Statistics

WebLogic Server provides callbacks for methods called on a JDBC driver. You can use these callbacks to monitor and profile JDBC driver usage, including methods being executed, any exceptions thrown, and the time spent executing driver methods.

To enable the callback feature, you specify the fully qualified path of the callback handler for the driver-interceptor element in the JDBC data source descriptor (module). Your callback handler must implement the `weblogic.jdbc.extensions.DriverInterceptor` interface. When you enable JDBC driver callbacks, WebLogic Server calls the `preInvokeCallback()`, `postInvokeExceptionCallback()`, and `postInvokeCallback()` methods of the registered callback handler before and after invoking any method inside the JDBC driver.

Any time an application calls the JDBC driver, a callback is sent to the class that implemented the driver.

13.3 Debugging JDBC Data Sources

Once you have narrowed the problem down to a specific application, you can activate WebLogic Server's debugging features to track down the specific problem within the application.

13.3.1 Enabling Debugging

You can enable debugging by setting the appropriate `ServerDebug` configuration attribute to "true." Optionally, you can also set the server `StdoutSeverity` to "Debug".

You can modify the configuration attribute in any of the following ways.

13.3.1.1 Enable Debugging Using the Command Line

Set the appropriate properties on the command line. For example,

```
-Dweblogic.debug.DebugJDBCSQL=true  
-Dweblogic.log.StdoutSeverity="Debug"
```

This method is static and can only be used at server startup.

13.3.1.2 Enable Debugging Using the WebLogic Server Administration Console

Use the WebLogic Server Administration Console to set the debugging values:

1. If you have not already done so, in the Change Center of the Administration Console, click Lock & Edit (see "Using the Change Center" in *Introduction to Oracle WebLogic Server*).
2. In the left pane of the console, expand Environment and select Servers.
3. On the Summary of Servers page, click the server on which you want to enable or disable debugging to open the settings page for that server.
4. Click Debug.
5. Expand default.
6. Select the check box for the debug scopes or attributes you want to modify.
7. Select Enable to enable (or Disable to disable) the debug scopes or attributes you have checked.
8. To activate these changes, in the Change Center of the Administration Console, click Activate Changes.
9. Not all changes take effect immediately—some require a restart (see "Using the Change Center" in *Introduction to Oracle WebLogic Server*).

This method is dynamic and can be used to enable debugging while the server is running.

13.3.1.3 Enable Debugging Using the WebLogic Scripting Tool

Use the WebLogic Scripting Tool (WLST) to set the debugging values. For example, the following command runs a program for setting debugging values called `debug.py`:

```
java weblogic.WLST debug.py
```

The `debug.py` program contains the following code:

```
user='user1'  
password='password'  
url='t3://localhost:7001'  
connect(user, password, url)  
edit()  
cd('Servers/myserver/ServerDebug/myserver')  
startEdit()  
set('DebugJDBCSQL','true')  
save()  
activate()
```

Note that you can also use WLST from Java. The following example shows a Java file used to set debugging values:

```
import weblogic.management.scripting.utils.WLSTInterpreter;  
import java.io.*;  
import weblogic.jndi.Environment;  
import javax.naming.Context;
```

```

import javax.naming.InitialContext;
import javax.naming.NamingException;

public class test {
    public static void main(String args[]) {
        try {
            WLSTInterpreter interpreter = null;
            String user="user1";
            String pass="pw12ab";
            String url = "t3://localhost:7001";
            Environment env = new Environment();
            env.setProviderUrl(url);
            env.setSecurityPrincipal(user);
            env.setSecurityCredentials(pass);
            Context ctx = env.getInitialContext();

            interpreter = new WLSTInterpreter();
            interpreter.exec
            ("connect('"+user+"','"+pass+"','"+url+"')");
            interpreter.exec("edit()");
            interpreter.exec("startEdit()");
            interpreter.exec
            ("cd('Servers/myserver/ServerDebug/myserver')");
            interpreter.exec("set('DebugJDBCSQL','true')");
            interpreter.exec("save()");
            interpreter.exec("activate()");

        } catch (Exception e) {
            System.out.println("Exception "+e);
        }
    }
}

```

Using the WLST is a dynamic method and can be used to enable debugging while the server is running.

13.3.1.4 Changes to the config.xml File

Changes in debugging characteristics, through console, or WLST, or command line are persisted in the `config.xml` file. See [Example 13-1](#):

Example 13-1 Example Debugging Stanza for JDBC

```

.
.
.
<server>
<name>myserver</name>
<server-debug>
<debug-scope>
<name>weblogic.transaction</name>
<enabled>true</enabled>
</debug-scope>
<debug-jdbcsql>true</debug-jdbcsql>
</server-debug>
</server>
.
.
.

```

This sample `config.xml` fragment shows a transaction debug scope (set of debug attributes) and a single JDBC attribute.

13.3.2 JDBC Debugging Scopes

The following are registered debugging scopes for JDBC:

- DebugJDBCSQL (scope `weblogic.jdbc.sql`) - prints information about all JDBC methods invoked, including their arguments and return values, and thrown exceptions.
- DebugJDBCConn (scope `weblogic.jdbc.connection`) - trace all connection reserve and release operations in data sources as well as all application requests to get or close connections.
- DebugJDBCRMI (scope `weblogic.jdbc.rmi`) - similar to JDBC SQL but at the RMI level; turning on this one and JDBC SQL will get two sets of debug messages for each operation called from a client.
- DebugJDBCInternal (scope `weblogic.jdbc.internal`) - low level debugging in `weblogic/jdbc/common/internal` related to the data source, the connection environment, and the data source manager.
- DebugJDBCDriverLogging (scope `weblogic.jdbc.driverlogging`) - enables JDBC driver level logging (this replaces `ServerMBean JDBCLoggingEnabled` and `getJDBCLogFileName`). Note that to get driver level tracing for Oracle, you need to use `ojdbc14_g.jar` instead of `ojdbc14.jar`. Note that for this debug scope, it can be turned on once via the command line or configuration when the server is booted but cannot be turned on or off dynamically (due to the `DriverManager` interface).

Note: Oracle WebLogic JDBC Spy logs detailed information about JDBC calls issued by an application and then passes the calls to the wrapped Oracle Type 4 JDBC driver. You can use the information in the logs to help troubleshoot problems in your application. For more information about WebLogic JDBC Spy, see "Tracking JDBC Calls with WebLogic JDBC Spy" in *Type 4 JDBC Drivers for Oracle WebLogic Server*.

13.3.3 Request Dyeing

Another option for debugging is to trace the flow of an individual (typically "dyed") application request through the JDBC subsystem. For more information, see "Configuring the Dye Vector via the DyeInjection Monitor" in *Configuring and Using the Diagnostics Framework for Oracle WebLogic Server*.

Managing WebLogic JDBC Resources

Using the Administration Console, command line, JMX programs, or WebLogic Scripting Tool (WLST) scripts, you can manage the JDBC data sources in your domain. This section includes the following information:

- [Section 14.1, "Testing Data Sources and Database Connections"](#)
- [Section 14.2, "Managing the Statement Cache for a Data Source"](#)
- [Section 14.3, "Shrinking a Connection Pool"](#)
- [Section 14.4, "Resetting a Connection Pool"](#)
- [Section 14.5, "Suspending a Connection Pool"](#)
- [Section 14.6, "Resuming a Connection Pool"](#)
- [Section 14.7, "Shutting Down a Data Source"](#)
- [Section 14.8, "Starting a Data Source"](#)
- [Section 14.9, "Managing DBMS Network Failures"](#)

14.1 Testing Data Sources and Database Connections

```
JDBCDataSourceRuntimeMBean.testPool
```

To make sure that the database connections in a data source remain healthy, you should periodically test the connections. WebLogic Server includes two basic types of testing: automatic testing that you configure with attributes on the data source and manual testing that you can do to trouble-shoot a data source.

Allowing WebLogic Server to automatically maintain the integrity of pool connections should prevent most DBMS connection problems. For more information about configuring automatic connection testing, see [Section 10.2, "Connection Testing Options for a Data Source."](#)

To manually test a connection from a data source, you can use the Test Data Source feature on the JDBC Data Source: Monitoring: Testing page in the Administration Console (see "Test JDBC data sources") or the `testPool()` method in the `JDBCDataSourceRuntimeMBean`. To test a database connection from a data source, Test Reserved Connections must be enabled and Test Table Name must be defined in the data source configuration. Both are defined by default if you create the data source using the Administration Console.

When you test a data source, WebLogic Server reserves a connection, tests it using the query defined in Test Table Name, and then releases the connection.

14.2 Managing the Statement Cache for a Data Source

For each connection in a data source in your system, WebLogic Server creates a statement cache. When a prepared statement or callable statement is used on a connection, WebLogic Server caches the statement so that it can be reused. For more information about the statement cache, see [Section 10.1, "Increasing Performance with the Statement Cache."](#)

Each connection in the data source has its own statement cache, but configuration settings are made for all connections in the data source. You can clear the statement cache for *all* connections in a data source using the Administration Console or you can programmatically clear the statement cache for an *individual* connection.

14.2.1 Clearing the Statement Cache for a Data Source

```
JDBCDataSourceRuntimeMBean.clearStatementCache
```

You can manually clear the statement cache for all connections in a data source using the Administration Console (see "Clear the statement cache in a JDBC data source") or with the `clearStatementCache()` method on the `JDBCDataSourceRuntimeMBean`.

14.2.2 Clearing the Statement Cache for a Single Connection

```
weblogic.jdbc.extensions.WLConnection.clearStatementCache()
weblogic.jdbc.extensions.WLConnection.clearCallableStatement(java.lang.
String sql)
weblogic.jdbc.extensions.WLConnection.clearCallableStatement(java.lang.
String sql,int resType,int resConcurrency)
weblogic.jdbc.extensions.WLConnection.clearPreparedStatement(java.lang.
String sql)
weblogic.jdbc.extensions.WLConnection.clearPreparedStatement(java.lang.
String sql,int resType,int resConcurrency)
```

You can use methods in the `weblogic.jdbc.extensions.WLConnection` interface to clear the statement cache for a single connection or to clear an individual statement from the cache. These methods return `true` if the operation was successful and `false` if the operation fails because the statement was not found.

When prepared and callable statements are stored in the cache, they are stored (keyed) based on the exact SQL statement and result set parameters (type and concurrency options), if any. When clearing an individual prepared or callable statement, you must use the method that takes the proper result set parameters. For example, if you have callable statement in the cache with `resSetType` of `ResultSet.TYPE_SCROLL_INSENSITIVE` and a `resSetConcurrency` of `ResultSet.CONCUR_READ_ONLY`, you must use the method that takes the result set parameters:

```
clearCallableStatement(java.lang.String sql,int resSetType,int resSetConcurrency)
```

If you use the method that only takes the SQL string as a parameter, the method will not find the statement, nothing will be cleared from the cache, and the method will return `false`.

When you clear a statement that is currently in use by an application, WebLogic Server removes the statement from the cache, but does not close it. When you clear a statement that is not currently in use, WebLogic Server removes the statement from the cache and closes it.

For more details about these methods, see the Javadoc for `WLConnection`.

14.3 Shrinking a Connection Pool

```
JDBCDataSourceRuntimeMBean.shrink
```

A data source has a set of properties that define the initial and maximum number of connections in the pool (`initialCapacity` and `maxCapacity`), and the number of connections to add to the pool when all connections are in use (`capacityIncrement`). When the pool reaches its maximum capacity, the maximum number of connections are opened, and they remain opened unless you enable automatic shrinking on the data source or manually shrink the data source with the `shrink()` method.

You may want to drop some connections from the data source when a peak usage period has ended, freeing up WebLogic Server and DBMS resources. You can use the Shrink option on the JDBC Data Source: Control page in the Administration Console (see "Shrink the connection pool in a JDBC data source" in *Oracle WebLogic Server Administration Console Help*) or the `shrink()` method on the `JDBCDataSourceRuntimeMBean`. When you shrink a data source, WebLogic Server reduces the number of connections in the pool to the greater of either the initial capacity or the number of connections currently in use.

14.4 Resetting a Connection Pool

```
JDBCDataSourceRuntimeMBean.reset
```

To close and recreate all available database connections in a data source, you can use the Reset option on the JDBC Data Source: Control page in the Administration Console (see "Reset connections in a JDBC data source" in *Oracle WebLogic Server Administration Console Help*) or the `reset()` method on the `JDBCDataSourceRuntimeMBean`. This may be necessary after the DBMS has been restarted, for example. Often when one connection in a data source has failed, all of the connections in the pool are bad.

14.5 Suspending a Connection Pool

```
JDBCDataSourceRuntimeMBean.suspend  
JDBCDataSourceRuntimeMBean.forceSuspend
```

To suspend a data source, you can use the Suspend and Force Suspend options on the JDBC Data Source: Control page in the Administration Console (see "Suspend JDBC data sources" in *Oracle WebLogic Server Administration Console Help*) or the `suspend()` and `forceSuspend()` methods in the `JDBCDataSourceRuntimeMBean`.

When you suspend a data source (not forcibly suspend), the data source is marked as disabled and applications cannot use connections from the pool. Applications that already have a reserved connection from the data source when it is suspended will get an exception when trying to use the connection. WebLogic Server preserves all connections in the data source exactly as they were before the data source was suspended.

When you forcibly suspend a data source, all pool connections are destroyed and any subsequent attempt to use reserved connections fail. Any transactions on the connections that are closed are rolled back.

14.6 Resuming a Connection Pool

```
JDBCDataSourceRuntimeMBean.resume
```

To re-enable a data source that you suspended, you can use the Resume option on the JDBC Data Source: Control page in the Administration Console (see "Resume suspended JDBC data sources" in *Oracle WebLogic Server Administration Console Help*) or the `resume()` method on the `JDBCDataSourceRuntimeMBean`. When you resume a data source, WebLogic Server marks the data source as enabled and allows applications to use connections from the data source. If you suspended the data source (not forcibly suspended), all connections are preserved exactly as they were before the data source was suspended. Clients that had reserved a connection before the data source was suspended can continue exactly where they left off. If you forcibly suspended the data source, clients will have to reserve new connections to proceed.

Note: You cannot resume a data source that did not start correctly, for example, if the database server is unavailable.

14.7 Shutting Down a Data Source

`JDBCDataSourceRuntimeMBean.shutdown`
`JDBCDataSourceRuntimeMBean.forceShutdown`

To shut down a data source, you can use the Shutdown and Force Shutdown options on the JDBC Data Source: Control page in the Administration Console (see "Shut down JDBC data sources" in *Oracle WebLogic Server Administration Console Help*) or the `shutdown()` and `forceShutdown()` methods in the `JDBCDataSourceRuntimeMBean`.

When you shut down a data source (not forcibly shut down), WebLogic Server closes database connections in the data source and shuts down the data source. If any connections from the data source are currently in use, the operation will fail.

When you forcibly shut down a data source, WebLogic Server closes database connections in the data source and shuts down the data source. All current connection users are forcibly disconnected.

14.8 Starting a Data Source

`JDBCDataSourceRuntimeMBean.start`

After you shut down a data source, you can use the Start option on the JDBC Data Source: Control page in the Administration Console (see "Start JDBC data sources" in *Oracle WebLogic Server Administration Console Help*) or the `start()` method in the `JDBCDataSourceRuntimeMBean`. Invoking the Start operation re-initializes the data source, creates connections and transitions the data source to a health state of Running.

14.9 Managing DBMS Network Failures

`-Dweblogic.resourcepool.max_test_wait_secs=xx`

where `xx` is the amount of time, in seconds, WebLogic Server waits for connection test before considering the connection test failed. By default, a server instance is assigned a value of 10 seconds.

This command line flag manages failures, such as a DBMS network failure, which can cause connection tests and applications to hang for extended periods of time (for example, 10 minutes). If the assigned time period expires, the server instance purges

and disables the pool (closes all connections and blocks further reserve attempts) and re-enables the pool as soon as it is possible to reconnect.

A value of ten seconds provides a reasonable amount of time to allow for peak stress loads, when a DBMS may temporarily halt responses to clients, and then resume service on existing connections. However, if the wait time is too long or too short, add the flag to the `startWebLogic` script used for starting the server with a value that is more appropriate for your environment. Setting the value for the amount of time to zero (0) seconds, causes the server to wait indefinitely on a hanging connection test.

Configuring JDBC Application Modules for Deployment

When you package your enterprise application, you can include JDBC resources in the application by packaging JDBC modules in the archive and adding references to the JDBC modules in all applicable descriptor files. When you deploy the application, the JDBC resources are deployed, too. Depending on how you configure the JDBC modules, the JDBC data sources deployed with the application will either be restricted for use only by the containing application (*application-scoped modules*) or will be available to all applications and clients (*globally-scoped modules*).

The following sections describe the details of packaged JDBC modules:

- [Section A.1, "Packaging a JDBC Module with an Enterprise Application: Main Steps"](#)
- [Section A.2, "Creating Packaged JDBC Modules"](#)
- [Section A.3, "Referencing a JDBC Module in Java EE Descriptor Files"](#)
- [Section A.4, "Packaging an Enterprise Application with a JDBC Module"](#)
- [Section A.5, "Deploying an Enterprise Application with a JDBC Module"](#)
- [Section A.6, "Getting a Database Connection from a Packaged JDBC Module"](#)

A.1 Packaging a JDBC Module with an Enterprise Application: Main Steps

The main steps for creating, packaging, and deploying a JDBC module with an enterprise application are as follows:

1. Create the module. See [Section A.2, "Creating Packaged JDBC Modules."](#)
2. Add references to the module in all applicable descriptor files. See [Section A.3, "Referencing a JDBC Module in Java EE Descriptor Files."](#)
3. Package all application modules in an EAR. See [Section A.4, "Packaging an Enterprise Application with a JDBC Module."](#)
4. Deploy the application. See [Section A.5, "Deploying an Enterprise Application with a JDBC Module."](#)

A.2 Creating Packaged JDBC Modules

You can create JDBC application modules using any development tool that supports creating an XML descriptor file. You then deploy and manage JDBC modules using

JSR 88-based tools, such as the `weblogic.Deployer` utility, or the Administration Console.

Note: You can create a JDBC data source using the Administration Console, then copy the module as a template for use in your applications. You must change the name and `jndi-name` elements of the module before deploying it with your application to avoid a naming conflict in the namespace.

Each JDBC module represents a data source or a multi data source. Modules that represent a data source include all of the configuration parameters for the data source. Modules that represent a multi data source include configuration parameters for the multi data source, including a list of data source modules used by the multi data source.

A.2.1 Creating a JDBC Data Source Module Using the Administration Console

To create a data source module in the Administration Console that you can re-use as an application module, follow these steps.

1. Create a data source as described in [Section 3.3, "Creating a JDBC Data Source."](#) The data source module is created in the `config/jdbc` subdirectory of the domain directory.
2. Copy the `data-source-name.xml` file to a subdirectory within your application and rename the copy to include `-jdbc` as a suffix, such as `new-data-source-name-jdbc.xml`.
3. Open the file in an editor and change the following elements:
 - `name`—change the name to a name that is unique within the domain.
 - `jndi-name`—change the `jndi-name` to a name that you want the enterprise application to use to lookup the data source in the local application context.
 - `scope`—optionally, to limit access to the data source to only the containing application, add a `scope` element to the `jdbc-data-source-params` section of the module. For example, `<scope>Application</scope>`. See [Section A.2.7, "Application Scoping for a Packaged JDBC Module."](#)
4. Continue with adding references to the descriptor files in the enterprise application. See [Section A.3, "Referencing a JDBC Module in Java EE Descriptor Files."](#)

A.2.2 JDBC Packaged Module Requirements

A JDBC module must meet the following criteria:

- Conforms to the `jdbc-data-source.xsd` schema. The schema is available at <http://xmlns.oracle.com/weblogic/jdbc-data-source/1.0/jdbc-data-source.xsd>.
- Uses a file name that ends in `-jdbc.xml`.
- Includes a `name` element that is unique within the WebLogic domain.

Data source modules must also include the following JDBC driver parameters:

- `url`
- `driver-name`

- properties, including any properties required by the JDBC driver to create database connections, such as a user name and password.

Multi data source modules must also include the following data source parameters:

- `data-source-list`, which is a list of data source modules, separated by commas, that the multi data source uses to satisfy database connection requests from applications.

Note: All data sources listed in the `data-source-list` must have the same XA and transaction protocol settings.

All other configuration parameters are optional or have a default value that WebLogic Server uses if a value is not specified. However, to create a useful JDBC module, you will likely need to specify additional configuration options as required by your applications and your environment.

A.2.3 JDBC Application Module Limitations

Note the following limitations for JDBC application modules:

- The `LoggingLastResource` global-transactions-protocol is not permitted for use in JDBC application modules.
- When deploying an application in production with application-scoped JDBC resources, if the resource uses `EmulateTwoPhaseCommit` for the global-transactions-protocol, you cannot deploy multiple versions of the application at the same time.

A.2.4 Creating a JDBC Data Source Module

The main sections within a JDBC data source module are:

- `jdbc-driver-params`—includes entries for the JDBC driver used to create database connections, including `url`, `driver-name`, and individual driver property entries. See the `jdbc-data-source.xsd` schema for more valid entries. For an explanation of each element, see "JDBCDataSourceParamsBean" in the *Oracle WebLogic Server MBean Reference*.
- `jdbc-connection-pool-params`—includes entries for connection pool configuration, including connection testing options, statement cache options, and so forth. This element also inherits `connection-pool-params` from the `weblogic-javaee.xsd` schema, including `initial-capacity`, `max-capacity`, and other options common to pooled resources. For more information, see the following:
 - "JDBCConnectionPoolParamsBean" in the *Oracle WebLogic Server MBean Reference*
 - `jdbc-data-source.xsd` schema
- `jdbc-data-source-params`—includes entries for data source behavior options and transaction processing options, such as `jndi-name`, `row-prefetch-size`, and `global-transactions-protocol`. See the `jdbc-data-source.xsd` schema for more valid entries. For an explanation of each element, see "JDBCDataSourceParamsBean" in the *Oracle WebLogic Server MBean Reference*.
- `jdbc-xa-params`—includes entries for XA database connection handling options, such as `keep-xa-conn-till-tx-complete`, and

xa-transaction-timeout. For an explanation of each element, see "JDBCXAParamsBean" in the *Oracle WebLogic Server MBean Reference*.

Example A-1 shows an example of a JDBC module for a data source with some typical configuration options.

Example A-1 Sample JDBC Data Source Module

```
<jdbc-data-source xsi:schemaLocation="http://www.bea.com/ns/weblogic/90/domain.xsd"
xmlns="http://xmlns.oracle.com/weblogic/jdbc-data-source"
xmlns:sec="http://www.bea.com/ns/weblogic/90/security"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:wls="http://www.bea.com/ns/weblogic/90/security/wls">
  <name>examples-demoXA-2</name>
  <jdbc-driver-params>
    <url>jdbc:derby://localhost:1527/examples;create=true</url>
    <driver-name>org.apache.derby.jdbc.ClientXADataSource</driver-name>
    <properties>
      <property>
        <name>user</name>
        <value>examples</value>
      </property>
      <property>
        <name>DatabaseName</name>
        <value>examples</value>
      </property>
    </properties>
    <password-encrypted>{AES}MEK6bPum8M69KRP4FANx3TG/00iSWRYu2rZGUwnVo6U=</password-encrypted>
  </jdbc-driver-params>
  <jdbc-connection-pool-params>
    <max-capacity>100</max-capacity>
    <connection-reserve-timeout-seconds>25</connection-reserve-timeout-seconds>
    <test-table-name>SQL SELECT 1 FROM SYS.SYSTABLES</test-table-name>
  </jdbc-connection-pool-params>
  <jdbc-data-source-params>
    <global-transactions-protocol>TwoPhaseCommit</global-transactions-protocol>
  </jdbc-data-source-params>
</jdbc-data-source>
```

A.2.5 Creating a JDBC Multi Data Source Module

A JDBC multi data source module is much simpler than a data source module. Only one main section is required: `jdbc-data-source-params`. The `jdbc-data-source-params` element in a multi data source differs in that it contains options for multi data source behavior options instead of data source behavior options. Only the following parameters in the `jdbc-data-source-params` are valid for multi data sources:

- `jndi-name` (required)
- `data-source-list` (required)
- `scope`
- `algorithm-type`
- `connection-pool-failover-callback-handler`
- `failover-request-if-busy`

For an explanation of each element, see "JDBCDataSourceParamsBean" in the *Oracle WebLogic Server MBean Reference*.

[Example A–2](#) shows an example of a JDBC module for a data source with some typical configuration options.

Example A–2 Sample JDBC Multi Data Source Module

```
<jdbc-data-source xmlns="http://xmlns.oracle.com/weblogic/jdbc-data-source">
  <name>examples-demoXA-multi-data-source</name>
  <jdbc-data-source-params>
    <jndi-name>examples-demoXA -multi-data-source</jndi-name>
    <algorithm-type>Load-Balancing</algorithm-type>
    <data-source-list>examples-demoXA, examples-demoXA-2</data-source-list>
  </jdbc-data-source-params>
</jdbc-data-source>
```

A.2.6 Encrypting Database Passwords in a JDBC Module

Oracle recommends that you encrypt database passwords in a JDBC module to keep your data secure. To encrypt a database password, you process the password with the WebLogic Server `encrypt` utility, which returns an encrypted equivalent of the password that you include in the JDBC module as the `password-encrypted` element. For more details about using the WebLogic Server `encrypt` utility, see "encrypt" in the *WebLogic Scripting Tool Command Reference*.

A.2.6.1 Deploying JDBC Modules to New Domains

It is common practice for JDBC modules to be moved from one domain to another, such as moving an application from development to production. However, the encryption key generated by the WebLogic Server `encrypt` utility is not transferable to a new domain. When moving a JDBC module with an encrypted database password, you must do one of the following:

- Override the old encrypted password within a deployment plan that includes a password that was encrypted specifically for the new domain. See "Update a deployment plan" in *Oracle WebLogic Server Administration Console Help*.
- Re-encrypt the passwords for your new domain. See [Section A.2.6, "Encrypting Database Passwords in a JDBC Module."](#)

A.2.7 Application Scoping for a Packaged JDBC Module

By default, when you package a JDBC module with an application, the JDBC resource is globally scoped—that is, the resource is bound to the global JNDI namespace and is available to all applications and clients. To reserve the resource for use only by the enclosing application, you must include the `<scope>Application</scope>` parameter in the `jdbc-data-source-params` element in the JDBC module, which binds the resource to the local application namespace. For example:

```
<jdbc-data-source-params>
  <jndi-name>examples-demoXA-2</jndi-name>
  <scope>Application</scope>
</jdbc-data-source-params>
```

All data sources in a multi data source for an application-scoped JDBC module must also be application scoped.


A.3 Referencing a JDBC Module in Java EE Descriptor Files

When you package a JDBC module with an enterprise application, you must reference the module in all applicable descriptor files, including among others:

- `weblogic-application.xml`
- `ejb-jar.xml`
- `weblogic-ejb-jar.xml`
- `web.xml`
- `weblogic.xml`

Figure A-1 shows the relationship between entries in various descriptor files for an EJB application and how they refer to a JDBC module packaged with the application.

Figure A-1 Relationship Between JDBC Modules and Descriptors in an Enterprise Application



A.3.1 Packaged JDBC Module References in weblogic-application.xml

When including JDBC modules in an enterprise application, you must list each JDBC module as a `module` element of type `JDBC` in the `weblogic-application.xml` descriptor file packaged with the application. For example:

```

<module>
  <name>data-source-1</name>
  <type>JDBC</type>
  <path>datasources/data-source-1-jdbc.xml</path>
</module>

```

A.3.2 Packaged JDBC Module References in Other Descriptors

For other application modules in your enterprise application to use the JDBC modules packaged with your application, you must add the following entries in the descriptor files packaged with application modules:

- In the standard Java EE descriptor files packaged with your application modules, such as `ejb-jar.xml` for an EJB, you must add `resource-ref-name` references to specify the JNDI name of the data source as used in the application. For example:

```

<resource-ref>
  <res-ref-name>my-data-source</res-ref-name>
  <res-type>javax.sql.DataSource</res-type>
  <res-auth>Container</res-auth>
</resource-ref>

```

In this example, `my-data-source` is the data source name as used in the application module. Your application would look up the data source with the following code:

```

javax.sql.DataSource ds =
    (javax.sql.DataSource) ctx.lookup("java:comp/env/my-data-source");

```

- In the WebLogic-specific descriptor files, such as `weblogic-ejb-jar.xml` for an EJB, you must map each `resource-ref-name` reference to the `jndi-name` element of a data source. For example:

```

<resource-description>
  <res-ref-name>my-data-source</res-ref-name>
  <jndi-name>qa-database-1</jndi-name>
</resource-description>

```

In this example, the resource name (`<res-ref-name>my-data-source</res-ref-name>`) from the standard descriptor is mapped to the JNDI name (`<jndi-name>qa-database-1</jndi-name>`) of the data source in the JDBC module.

Figure A-1 shows the mapping of the of the data source name as used in the application module to the JNDI name of the JDBC data source in the JDBC module.

Note: For application-scoped data sources, if you do not add these entries to the descriptor files, your application will be unable to look up the data source to get a database connection.

A.4 Packaging an Enterprise Application with a JDBC Module

You package an application with a JDBC module as you would any other enterprise application. See "Packaging Applications Using `wlpackage`" in *Developing Applications for Oracle WebLogic Server*.

A.5 Deploying an Enterprise Application with a JDBC Module

You deploy an application with a JDBC module as you would any other enterprise application. See "Deploying Applications Using `wldeploy`" in *Developing Applications for Oracle WebLogic Server*.

Note: When deploying an application in production with application-scoped JDBC resources, if the resource uses `EmulateTwoPhaseCommit` for the `global-transactions-protocol`, you cannot deploy multiple versions of the application at the same time.

A.6 Getting a Database Connection from a Packaged JDBC Module

To get a connection from JDBC module packaged with an enterprise application, you look up the data source or multi data source defined in the JDBC module in the local environment (`java:comp/env`) or on the JNDI tree and then request a connection from the data source or multi data source. For example:

```
javax.sql.DataSource ds =  
    (javax.sql.DataSource) ctx.lookup("java:comp/env/my-data-source");  
java.sql.Connection conn = ds.getConnection();
```

When you are finished using the connection, make sure you close the connection to return it to the connection pool in the data source:

```
conn.close();
```

Using Multi Data Sources with Oracle RAC

Oracle continues to support multi data source configurations for legacy application environments using Oracle Real Application Clusters (RAC). The following sections provide information on how to configure and use multi data sources when using Oracle RAC with WebLogic Server:

- [Section B.1, "Overview of Oracle Real Application Clusters"](#)
- [Section B.2, "Software Requirements"](#)
- [Section B.3, "JDBC Driver Requirements"](#)
- [Section B.4, "Hardware Requirements"](#)
- [Section B.5, "Configuring Multi Data Sources with Oracle RAC"](#)
- [Section B.6, "Using Multi Data Sources with Global Transactions"](#)
- [Section B.7, "Using Multi Data Sources without Global Transactions"](#)
- [Section B.8, "Configuring Connections to Services on Oracle RAC Nodes"](#)
- [Section B.9, "XA Considerations and Limitations when using multi Data Sources with Oracle RAC"](#)
- [Section B.10, "JDBC Store Recovery with Oracle RAC"](#)

Both Oracle RAC and WebLogic Server are complex systems. To use them together requires specific configuration on both systems, as well as clustering software and a shared storage solution. This section describes the configuration required at a high level. For more details about configuring Oracle RAC, your clustering software, your operating system, and your storage solution, see the documentation from the respective vendors.


Note: Oracle recommends using GridLink data sources when developing new Oracle RAC applications and when legacy applications do not use multi data sources. See [Section 11, "Using WebLogic Server with Oracle RAC."](#)

B.1 Overview of Oracle Real Application Clusters

Oracle Real Application Clusters (Oracle RAC) is a software component you can add to a high-availability solution that enables users on multiple machines to access a single database with increased performance. Oracle RAC comprises two or more Oracle database instances running on two or more clustered machines and accessing a shared storage device via cluster technology. To support this architecture, the machines that host the database instances are linked by a high-speed interconnect to form the

cluster. The interconnect is a physical network used as a means of communication between the nodes of the cluster. Cluster functionality is provided by the operating system or compatible third party clustering software. [Figure B-1](#) shows an Oracle RAC configuration.

Figure B-1 Oracle Real Application Clusters Configuration



Oracle RAC offers features in the following areas:

- [Section B.1.1, "Oracle RAC Scalability with WebLogic Server Multi Data Sources."](#)
- [Section B.1.2, "Oracle RAC Availability with WebLogic Server Multi Data Sources."](#)
- [Section B.1.3, "Oracle RAC Load Balancing with WebLogic Server Multi Data Sources."](#)

B.1.1 Oracle RAC Scalability with WebLogic Server Multi Data Sources

An Oracle RAC installation appears like a single standard Oracle database and is maintained using the same tools and practices. All the nodes in the cluster execute transactions against the same database and Oracle RAC coordinates each node's access to the shared data to maintain consistency and ensure integrity. You can add nodes to the cluster easily and there is no need to partition data when you add them. This means that you can horizontally scale the database tier as usage and demand grows by adding Oracle RAC nodes, storage, or both.

As the number of nodes in an Oracle RAC increases, you scale the number of generic data sources by the number of nodes added to the Oracle RAC. This requires a complex configuration (requiring $n+1$ data sources where n is the number of generic data sources plus a multi data source) that requires administrative intervention when the Oracle RAC topology changes.

B.1.2 Oracle RAC Availability with WebLogic Server Multi Data Sources

A multi data source provides an ordered list of data sources to use to satisfy connection requests. Normally, every connection request to this kind of multi data source is served by the first data source in the list. If a database connection test fails and the connection cannot be replaced, or if the data source is suspended, a connection is sought sequentially from the next data source on the list. See [Section 5.3.1, "Failover"](#) and [Section B.5.3.1, "Multi Data Source-Managed Failover and Load Balancing."](#)

B.1.3 Oracle RAC Load Balancing with WebLogic Server Multi Data Sources

Multi data sources provide load balancing for XA and non-XA environments. The generic data sources that form a multi data source are accessed using a round-robin scheme. When switching connections, WebLogic Server selects a connection from the next generic data source in the order listed. For more information about using multi data sources with Oracle RAC, see [Section B, "Using Multi Data Sources with Oracle RAC."](#)

B.2 Software Requirements

To use WebLogic Server with Oracle RAC, you must install the following software on each Oracle RAC node:

- Operating system patches required to support Oracle RAC. See the release notes from Oracle for details.
- Oracle 10g or 11g database management system
- Clustering software for your operating system. See the Oracle documentation for supported clustering software and cluster configurations.
- Shared storage software, such as Veritas Cluster File System. Note that some clustering software includes a file storage solution, in which case additional shared storage software is not required.

Note: See the *Oracle Fusion Middleware Supported System Configurations* page at http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html for the latest WebLogic Server hardware platform and operating system support, and for the Oracle RAC versions supported by WebLogic Server versions and service packs. See the Oracle documentation for hardware and software requirements required for running the Oracle RAC software.

B.3 JDBC Driver Requirements

To use WebLogic Server with Oracle RAC, your WebLogic JDBC data sources must use the Oracle JDBC Thin driver 11g to create database connections.

B.4 Hardware Requirements

A typical WebLogic Server/Oracle RAC system includes a WebLogic Server cluster, an Oracle RAC cluster, and hardware for shared storage.

B.4.1 WebLogic Server Cluster

The WebLogic Server cluster can be configured in many ways and with various hardware options. See *Using Clusters for Oracle WebLogic Server* for more details about configuring a WebLogic Server cluster.

B.4.2 Oracle RAC Cluster

For the latest hardware requirements for Oracle RAC, see the Oracle RAC documentation. However, to use Oracle RAC with WebLogic Server, you must run Oracle RAC instances on robust, production-quality hardware. The Oracle RAC configuration must deliver database processing performance appropriate for reasonably-anticipated application load requirements. Unusual database response delays can lead to unexpected behavior during database failover scenarios.

B.4.3 Shared Storage

In an Oracle RAC configuration, all data files, control files, and parameter files are shared for use by all Oracle RAC instances. An HA storage solution that uses one of the following architectures is recommended:

- Direct Attached Storage (DAS), such as a dual ported disk array or a Storage Area Network (SAN)
- Network Attached Storage (NAS)

For a complete list of supported storage solutions, see your Oracle documentation.

B.5 Configuring Multi Data Sources with Oracle RAC

When using Multi data sources with Oracle RAC, you must configure your WebLogic Domain so that it can interact with Oracle RAC instances and so that it performs as expected. The following sections describe configuration options and requirements:

- [Section B.5.1, "Choosing a Multi Data Source Configuration for Use with Oracle RAC"](#)
- [Section B.5.2, "Configuring Multi Data Sources for use with Oracle RAC"](#)
- [Section B.5.3, "Configuration Considerations for Failover"](#)
- [Section B.5.4, "Configuring the Listener Process for Each Oracle RAC Instance"](#)
- [Section B.5.5, "Configuring Multi Data Sources When Remote Listeners are Enabled or Disabled"](#)
- [Section B.5.6, "Additional Configuration Considerations"](#)

B.5.1 Choosing a Multi Data Source Configuration for Use with Oracle RAC

WebLogic Server multi data sources support several configuration options for using Oracle RAC:

- To connect to multiple Oracle RAC 10g or Oracle RAC 11g instances when using global transactions (XA), see [Section B.6, "Using Multi Data Sources with Global Transactions."](#)
- To connect to multiple Oracle RAC 10g or Oracle RAC 11g instances when not using XA, see [Section B.7, "Using Multi Data Sources without Global Transactions."](#)


- You can also configure multi data sources to connect to specific services that are running on Oracle RAC nodes. Both XA and non-XA drivers are supported, see [Section B.8, "Configuring Connections to Services on Oracle RAC Nodes."](#)

B.5.2 Configuring Multi Data Sources for use with Oracle RAC

To connect WebLogic Server to multiple Oracle RAC nodes using multi data sources, first configure a JDBC data source for each Oracle RAC instance in your Oracle RAC cluster with the Oracle Thin driver. Then configure a multi data source, using either the algorithm for load balancing or the algorithm for failover, and add the data sources to it.

[Figure B-2](#) shows a typical multi data source configuration.

Figure B-2 Multi Data Source Configuration



You can use the Administration Console or any other means that you prefer to configure your domain, such as the WebLogic Scripting Tool (WLST) or a JMX program. For information about configuring a WebLogic JDBC multi data source see [Chapter 5, "Configuring JDBC Multi Data Sources."](#) For information on how to configure the data sources in a multi data source to connect to services running on

Oracle RAC nodes, see [Section B.8, "Configuring Connections to Services on Oracle RAC Nodes."](#)

To use a database connection in this configuration, your applications look up one multi data source on the JNDI tree and then request a connection. The multi data source determines which data source to use to satisfy the connection request based on the algorithm type specified in the configuration (that is, failover or load balancing).

B.5.2.1 Attributes of a Multi Data Source

The multi data source may have the following attributes, depending on the role of Oracle RAC in your system—load balancing or failover:

- `AlgorithmType="Load-Balancing"` or `AlgorithmType="Failover"`
With the Load-Balancing option, connection requests are distributed among available data sources; with the High-Availability option, connection requests are served by the first available pool in the list. When a data source becomes defunct, connection requests are served by the next data source in the list.
- `FailoverRequestIfBusy="true"`
With the Failover algorithm, this attribute enables failover when all connections in a data source are in use.
- `TestFrequencySeconds="120"`
This attribute controls the frequency at which WebLogic Server checks the health of data sources previously marked as unhealthy to see if connections can be recreated and if the data source can be re-enabled. For more details see [Chapter 5, "Configuring JDBC Multi Data Sources."](#)
For fast failover of Oracle RAC nodes, set this value to a smaller interval, for example, 10 (seconds).

B.5.3 Configuration Considerations for Failover

Consider the following information when configuring for failover.

B.5.3.1 Multi Data Source-Managed Failover and Load Balancing

Multi data sources offer failover and load balancing for global transactions. For a description of multi data source failover features, see [Section 5.5, "Multi Data Source Failover Enhancements."](#)

With this configuration, pictured in [Figure B-2](#), you get:

- Fast failover controlled by the multi data source
- Automatic failback by the WebLogic Server health monitor

The multi data source handles failover for database connections when an Oracle RAC node becomes unavailable. When WebLogic Server tests a connection and the connection fails, it attempts to recreate the connection. If that attempt fails, the server disables the data source and routes connection requests to other data sources (which correspond to other Oracle RAC nodes) in the multi data source. WebLogic Server periodically tries to recreate the database connections in the disabled data source. When WebLogic Server is successful in recreating the connections, it next re-enables the data source and begins routing connection requests to the data source again. Because of the connection request routing and automatic health checking features, there is minimal delay in satisfying connection requests after a failure.

B.5.3.2 Delays During Failover

Occasionally, when one Oracle RAC node fails over to another, there may be a delay before the data associated with a transaction branch in progress on the now failed node is available throughout the cluster. This prevents incomplete transactions from being properly completed, which could further result in data locking in the database. To protect against the potential consequences of such a delay, WebLogic Server provides two configuration attributes that enable XA call retry for Oracle RAC: `XARetryDurationSeconds` and `XARetryIntervalSeconds`.

`XARetryDurationSeconds` controls the period of time during which WebLogic Server will repeatedly retry XA operations such as recover, commit and rollback for pending transactions. `XARetryIntervalSeconds` controls the frequency of the retry attempts within the established time period.

To enable XA call retries, add a value for `XARetryDurationSeconds` to all JDBC data sources in your WebLogic domain that connect to an Oracle RAC instance. For example:

```
<jdbc-data-source xmlns="http://xmlns.oracle.com/weblogic/jdbc-data-source"
  xmlns:sec="http://xmlns.oracle.com/weblogic/security"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:wls="http://xmlns.oracle.com/weblogic"
  xsi:schemaLocation="http://xmlns.oracle.com/weblogic/domain/1.0/domain.xsd">
  <name>oracleRACXAPool</name>
  ...
  <jdbc-xa-params>
    ...
    <xa-retry-duration-seconds>300</xa-retry-duration-seconds>
  </jdbc-xa-params>
</jdbc-data-source>
```

Use the following formula to determine the value for `XARetryDurationSeconds`:

`XARetryDurationSeconds` = (longest transaction timeout for transactions that use connections from the data source) + (delay before XIDs are available on all Oracle RAC nodes, typically less than 5 minutes)

For example, if your application sets the longest transaction timeout as 180 seconds, you should set `XARetryDurationSeconds` to 180 seconds + 300 seconds, for a total of 480 seconds.

Note: It is generally better to set `XARetryDurationSeconds` higher than minimally necessary to make sure that all transactions are completed properly. Setting the value higher than minimally required should not affect application performance during normal operations. The additional processing only affects transactions that have been prepared but have failed to complete.

You can also optionally set a value for `XARetryIntervalSeconds`. This value determines the time between XA retry calls. By default, the value is 60 seconds. Decreasing the value will decrease the amount of time between XA retry attempts. The default value should suffice in most cases.

To enable `XARetryDurationSeconds` and `XARetryIntervalSeconds` from the Administration Console, use the following steps:

1. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
2. In the **Domain Structure** tree, expand **Services > JDBC**, then select **Data Sources**.

3. On the Summary of Data Sources page, click the data source name.
4. Select the **Configuration: Connection Pool** tab.
5. Scroll down and click **Advanced** to show the advanced connection pool options.
6. Update `XA Retry Duration` and `XA Retry Interval`.
7. Click **Save**.

Optionally, you can use WebLogic Scripting Tool (WLST) or a JMX program.

B.5.3.3 Failure Handling Walkthrough for Global Transactions

What happens to in-flight transactions to a database node if that node fails? When the primary Oracle RAC node fails? Does WebLogic Server support transparent failover? To answer these and other questions about how WebLogic Server handles failures, let's walk through the transaction processing steps and describe how a failure would be handled at each stage along the way.

The first stage at which a failure may occur is before the application calls for the transaction to be committed. If a database or Oracle RAC node fails at this stage, the application receives an exception and must get a new connection and make a new attempt at processing the transaction. WebLogic Server does not support transparent failover.

If a failure occurs after the application has called for the transaction to be committed, the handling of any in-flight transaction depends upon whether the `PREPARE` operation is complete. If the `PREPARE` operation is not complete, the transaction manager rolls back the transaction and sends the application an exception for the failed transaction. If the `PREPARE` operation is complete, the transaction manager attempts to drive the in-flight transaction to completion using another node.

If a failure occurs during the `COMMIT` operation, the transaction manager attempts to retry the `COMMIT` operation several times. Note that the connection is blocked during these attempts. If the `COMMIT` operation is not successful during the first set of retry attempts, the application receives an exception. The transaction manager then continues to retry the `COMMIT` operation periodically until it is successful; if the transaction cannot be completed successfully within the abandon time period, the transaction is driven to completion heuristically.


B.5.4 Configuring the Listener Process for Each Oracle RAC Instance

For Oracle RAC, the listener process establishes a communication pathway between a user process and an Oracle instance. When you use Oracle RAC with WebLogic Server, the user process is typically a JDBC data source.

When a multi data source is created, it attempts to create a pool of database connections for applications to borrow. If a pooled database connection becomes inoperative or if the data source is configured to grow in capacity, the data source attempts to create additional database connections up to the maximum specified in the configuration file. In all of these instances, the Oracle listener process handles the connection request on the Oracle RAC instance.

[Figure B-3](#) shows the Oracle listener process functionality.

Figure B-3 Oracle Listener Process Functionality



To enable this functionality, you have two options:

- **Use local listeners.** Configure the listener process for each Oracle RAC instance in the Oracle cluster. WLS requires that you configure a local listener on each Oracle RAC instance. Each database instance should be configured to register *only* with its local listener.

Oracle instances can be configured to register with the listener statically in the `listener.ora` file or registered dynamically using the instance initialization parameter `local_listener`, or both. Oracle recommends using dynamic registration.

A listener can start either a shared dispatcher process or a dedicated process. When using with WebLogic Server, Oracle recommends using dedicated processes.

- **Use remote listeners.** WLS requires that you specify both the `SERVICE_NAME` and the `INSTANCE_NAME` in the JDBC URL for the data sources in the multi data source. See [Section B.5.5, "Configuring Multi Data Sources When Remote Listeners are Enabled or Disabled."](#)

B.5.5 Configuring Multi Data Sources When Remote Listeners are Enabled or Disabled

If the server-side load balancing feature has been enabled for the Oracle RAC backend (using `remote_listeners`), the JDBC URL that you use in the data sources of a multi data source configuration should include the `INSTANCE_NAME`. For example, the URL can be specified in the following format:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=host-vip)(PORT=1521))
(CONNECT_DATA=(SERVICE_NAME=dbservice)(INSTANCE_NAME=inst1)))
```

If specifying the `INSTANCE_NAME` in the URL is not possible, remote listeners must be disabled. To disable remote listeners, delete any listed remote listeners in `spfile.ora` on each Oracle RAC node. For example:

```
*.remote_listener="
```

In this case, the recommended URL that you use in the data sources of a multi data source configuration is:

```
jdbc:oracle:thin:@host-vip:port/dbservice
```

or

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=host-vip)(PORT=1521))
(CONNECT_DATA=(SERVICE_NAME=dbservice)))
```

B.5.6 Additional Configuration Considerations

In some deployments of Oracle RAC, you may need to set parameters in addition to the out of the box configuration of a data source in an Oracle RAC configuration. The additional parameters are:

- Set `oracle.jdbc.ReadTimeout=300000` (300000 milliseconds) for each data source.

The actual value of the `ReadTimeout` parameter used may differ based on your application environment.

- If the network is not reliable, it is difficult for a client to detect the frequent disconnections when the server is abruptly disconnected. By default, a client running on Linux takes 7200 seconds (2 hours) to sense the abrupt disconnections. This value is equal to the value of the `tcp_keepalive_time` property. To configure the application to detect the disconnections faster, set the value of the `tcp_keepalive_time`, `tcp_keepalive_interval`, and `tcp_keepalive_probes` properties to a lower value at the operating system level.

Note: Setting a low value for the `tcp_keepalive_interval` property leads to frequent probe packets on the network, which can make the system slower. Set the value of this property based on system requirements of your application environment.

For example, set `tcp_keepalive_time=600` for a system running a WebLogic Server managed server.

- Specify the `ENABLE=BROKEN` parameter in the `DESCRIPTION` clause in the connection descriptor. For example:

```
jdbc:oracle:thin:@(DESCRIPTION=(enable=broken)(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=node1-vip.mycompany.com)(PORT
```

```
=1521)))(CONNECT_DATA=(SERVICE_NAME=orcl.us.oracle.com)(INSTANCE_
NAME=orcl1))
```

The following code snippet provides an example data source configuration:

```
<url>jdbc:oracle:thin:@(DESCRIPTION=(enable=broken) (ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=node1-vip.us.oracle.com) (PORT=1521))) (CONNECT_
DATA=(SERVICE_NAME=orcl.us.oracle.com) (INSTANCE_NAME=orcl1)))</url>
<driver-name>oracle.jdbc.xa.client.OracleXADataSource</driver-name>
<properties>
<property>
<name>oracle.jdbc.ReadTimeout</name>
<value>300000</value>
</property>
<property>
<name>user</name>
<value>jmsuser</value>
</property>
<property>
<name>oracle.net.CONNECT_TIMEOUT</name>
<value>10000</value>
</property>
</properties>
```

B.6 Using Multi Data Sources with Global Transactions

In this configuration, a multi data source "pins" a transaction to one and only one Oracle RAC instance. Individual transactions are load balanced with the initial connection request for the transaction. Failover is handled at the multi data source level when a Oracle RAC instance becomes unavailable. If there is a failure on a Oracle RAC instance before PREPARE, the operation is retried until the retry duration has expired. If there is a failure after PREPARE the transaction is failed over to another instance.

- [Section B.6.1, "Rules for Data Sources within a Multi Data Source Using Global Transactions"](#)
- [Section B.6.2, "Required Attributes of Data Sources within a Multi Data Source Using Global Transactions"](#)
- [Section B.6.3, "Sample Configuration Code"](#)

B.6.1 Rules for Data Sources within a Multi Data Source Using Global Transactions

The following rules apply to the XA data sources within a multi data source:

- All the data sources must be homogeneous. In other words, either all of them must use an XA driver or none of them can use an XA driver.
- If you choose to specify them, all XA-related attributes must be set to the same values for each data source. The attributes include the following:
 - XARetryDurationSeconds
 - SupportsLocalTransaction
 - KeepXAConnTillTxComplete
 - NeedTxCtxOnClose
 - XAEndOnlyOnce

- NewXAConnForCommit
- RollbackLocalTxUponConnClose
- RecoverOnlyOnce
- KeepLogicalConnOpenOnRelease

Note: If you are not using GridLink data sources, Oracle recommends the use of multi data sources for failover and load balancing across Oracle RAC instances for XA and non-XA environments. For more information on using multi data sources in non-XA environments, see [Section B.7, "Using Multi Data Sources without Global Transactions."](#)

B.6.2 Required Attributes of Data Sources within a Multi Data Source Using Global Transactions

Each data source within the multi data source should have the following attributes:

- Oracle JDBC Thin driver 11g. For example:

```
<url>jdbc:oracle:thin:@lcqsol24:1521:SNRAC1</url>  
<driver-name>oracle.jdbc.xa.client.OracleXADataSource</driver-name>
```
- KeepXAConnTillTxComplete="true"
 - Forces the data source to reserve a physical database connection and provide the same connection to an application throughout transaction processing until the distributed transaction is complete.
 - Required for proper transaction processing with Oracle RAC.
- XARetryDurationSeconds="300"
 - Enables the WebLogic Server transaction manager to retry XA recover, commit, and rollback calls for the specified amount of time.
- TestConnectionsOnReserve="true"
 - Enables testing of a database connection when an application reserves a connection from the data source. See [Section 5.4.1, "Test Connections on Reserve to Enable Fail-Over"](#) for more details about this attribute.
 - Required to enable failover to another Oracle RAC node.
- TestTableName="name_of_small_table" The name of the table used to test a physical database connection. For more details about this attribute, see [Section 10.2, "Connection Testing Options for a Data Source."](#)

B.6.3 Sample Configuration Code

Sample configuration code for a multi data source and two associated data sources is shown below.

```
<jdbc-data-source xmlns="http://xmlns.oracle.com/weblogic/jdbc-data-source"  
  xmlns:sec="http://xmlns.oracle.com/weblogic/security"  
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"  
  xmlns:wls="http://xmlns.oracle.com/weblogic"  
  xsi:schemaLocation="http://xmlns.oracle.com/weblogic/domain/1.0/domain.xsd">  
  <name>oracleRACXAPool</name>  
  <jdbc-driver-params>  
    <url>jdbc:oracle:thin:@lcqsol24:1521:SNRAC1</url>
```



```

<driver-name>oracle.jdbc.xa.client.OracleXADataSource</driver-name>
<properties>
  <property>
    <name>user</name>
    <value>wlsqa</value>
  </property>
</properties>
<password-encrypted>{3DES}aP/xScCS8uI=</password-encrypted>
</jdbc-driver-params>
<jdbc-connection-pool-params>
  <test-table-name>SQL SELECT 1 FROM DUAL</test-table-name>
  <profile-type>0</profile-type>
</jdbc-connection-pool-params>
<jdbc-data-source-params>
  <jndi-name>oracleRACXAJndiName</jndi-name>
  <global-transactions-protocol>TwoPhaseCommit
  </global-transactions-protocol>
</jdbc-data-source-params>
<jdbc-xa-params>
  <keep-xa-conn-till-tx-complete>true</keep-xa-conn-till-tx-complete>
  <xa-end-only-once>true</xa-end-only-once>
  <xa-set-transaction-timeout>true</xa-set-transaction-timeout>
  <xa-transaction-timeout>120</xa-transaction-timeout>
  <xa-retry-duration-seconds>300</xa-retry-duration-seconds>
</jdbc-xa-params>
</jdbc-data-source>
<jdbc-data-source xmlns="http://xmlns.oracle.com/weblogic/jdbc-data-source"
  xmlns:sec="http://xmlns.oracle.com/weblogic/security"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:wls="http://xmlns.oracle.com/weblogic"
  xsi:schemaLocation="http://xmlns.oracle.com/weblogic/domain/1.0/domain.xsd">
  <name>oracleRACXAPool2</name>
  <jdbc-driver-params>
    <url>jdbc:oracle:thin:@lcqsol25:1521:SNRAC2</url>
    <driver-name>oracle.jdbc.xa.client.OracleXADataSource</driver-name>
    <properties>
      <property>
        <name>user</name>
        <value>wlsqa</value>
      </property>
    </properties>
    <password-encrypted>{3DES}aP/xScCS8uI=</password-encrypted>
  </jdbc-driver-params>
  <jdbc-connection-pool-params>
    <test-table-name>SQL SELECT 1 FROM DUAL</test-table-name>
    <profile-type>0</profile-type>
  </jdbc-connection-pool-params>
  <jdbc-data-source-params>
    <jndi-name>oracleRACXAJndiName2</jndi-name>
    <global-transactions-protocol>TwoPhaseCommit
    </global-transactions-protocol>
  </jdbc-data-source-params>
  <jdbc-xa-params>
    <keep-xa-conn-till-tx-complete>true</keep-xa-conn-till-tx-complete>
    <xa-end-only-once>true</xa-end-only-once>
    <xa-set-transaction-timeout>true</xa-set-transaction-timeout>
    <xa-transaction-timeout>120</xa-transaction-timeout>
    <xa-retry-duration-seconds>300</xa-retry-duration-seconds>
  </jdbc-xa-params>
</jdbc-data-source>

```

```

<jdbc-data-source xmlns="http://xmlns.oracle.com/weblogic/jdbc-data-source"
  xmlns:sec="http://xmlns.oracle.com/weblogic/security"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:wls="http://xmlns.oracle.com/weblogic"
  xsi:schemaLocation="http://xmlns.oracle.com/weblogic/domain/1.0/domain.xsd">
  <name>oracleRACXAMDS</name>
  <jdbc-data-source-params>
    <jndi-name>oracleRACMDSJndiName</jndi-name>
    <algorithm-type>Load-Balancing</algorithm-type>
    <data-source-list>oracleRACXAPool,oracleRACXAPool2</data-source-list>
  </jdbc-data-source-params>
</jdbc-data-source>

```

B.7 Using Multi Data Sources without Global Transactions

The following sections describe a configuration that uses Oracle RAC with multi data sources in an application that does not require global transactions.

- [Section B.7.1, "Attributes of Data Sources within a Multi Data Source Not Using Global Transactions"](#)
-

B.7.1 Attributes of Data Sources within a Multi Data Source Not Using Global Transactions

Data sources must have the following attributes:

- Oracle JDBC Thin driver 11g. For example:


```

<url>jdbc:oracle:thin:@lcqsol24:1521:SNRAC1</url>
<driver-oracle.jdbc.OracleDriver/driver-name>

```
- TestConnectionsOnReserve="true"
 - Enables testing of a database connection when an application reserves a connection from the data source. [Section 5.4.1, "Test Connections on Reserve to Enable Fail-Over"](#) for more details about this attribute.
 - Required to enable failover and connection request routing within a multi data source (effectively, failover to another Oracle RAC node).
- TestTableName="name_of_small_table"
 - The name of the table used to test a physical database connection. For more details about this attribute, see [Section 10.2, "Connection Testing Options for a Data Source."](#)

B.7.2 Sample Configuration Code

Sample configuration code for a WebLogic JDBC multi data source and associated data sources is shown below.

```

<jdbc-data-source xmlns="http://xmlns.oracle.com/weblogic/jdbc-data-source"
  xmlns:sec="http://xmlns.oracle.com/weblogic/security"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:wls="http://xmlns.oracle.com/weblogic"
  xsi:schemaLocation="http://xmlns.oracle.com/weblogic/domain/1.0/domain.xsd">
  <name>jdbcPool</name>
  <jdbc-driver-params>

```

```

<url>jdbc:oracle:thin:@lcgsol24:1521:snrac1</url>
<driver-name>oracle.jdbc.OracleDriver</driver-name>
<properties>
  <property>
    <name>user</name>
    <value>wlsqa</value>
  </property>
</properties>
<password-encrypted>{3DES}aP/xScCS8uI=</password-encrypted>
</jdbc-driver-params>
<jdbc-connection-pool-params>
  <test-connections-on-reserve>true</test-connections-on-reserve>
  <test-table-name>SQL SELECT 1 FROM DUAL</test-table-name>
</jdbc-connection-pool-params>
<jdbc-data-source-params>
  <jndi-name>jdbcDataSource</jndi-name>
</jdbc-data-source-params>
</jdbc-data-source>
<jdbc-data-source xmlns="http://xmlns.oracle.com/weblogic/jdbc-data-source"
  xmlns:sec="http://xmlns.oracle.com/weblogic/security"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:wls="http://xmlns.oracle.com/weblogic"
  xsi:schemaLocation="http://xmlns.oracle.com/weblogic/domain/1.0/domain.xsd">
  <name>jdbcPool2</name>
  <jdbc-driver-params>
    <url>jdbc:oracle:thin:@lcgsol25:1521:SNRAC2</url>
    <driver-name>oracle.jdbc.OracleDriver</driver-name>
    <properties>
      <property>
        <name>user</name>
        <value>wlsqa</value>
      </property>
    </properties>
    <password-encrypted>{3DES}aP/xScCS8uI=</password-encrypted>
  </jdbc-driver-params>
  <jdbc-connection-pool-params>
    <test-connections-on-reserve>true</test-connections-on-reserve>
    <test-table-name>SQL SELECT 1 FROM DUAL</test-table-name>
  </jdbc-connection-pool-params>
  <jdbc-data-source-params>
    <jndi-name>jdbcDataSource2</jndi-name>
    <global-transactions-protocol>OnePhaseCommit
      </global-transactions-protocol>
  </jdbc-data-source-params>
</jdbc-data-source>
<jdbc-data-source xmlns="http://xmlns.oracle.com/weblogic/jdbc-data-source"
  xmlns:sec="http://xmlns.oracle.com/weblogic/security"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:wls="http://xmlns.oracle.com/weblogic"
  xsi:schemaLocation="http://xmlns.oracle.com/weblogic/domain/1.0/domain.xsd">
  <name>jdbcNonXAMultiPool</name>
  <jdbc-data-source-params>
    <jndi-name>jdbcDataSource</jndi-name>
    <algorithm-type>Failover</algorithm-type>
    <data-source-list>jdbcPool, jdbcPool2</data-source-list>
    <failover-request-if-busy>true</failover-request-if-busy>
  </jdbc-data-source-params>
</jdbc-data-source>

```

Note: Line breaks added for readability.

B.8 Configuring Connections to Services on Oracle RAC Nodes

If you rely on Oracle services in your Oracle RAC cluster for workload management, you must use multi data sources to connect to those services instead of you using a Service ID (SID). A WebLogic Server data source can be configured to connect only to a specific service on a specific Oracle RAC node, providing both workload management and load balancing.

In general, to connect to Oracle RAC services, you need to:

- Create a multi data source for each service to which you want to connect.
- Within each multi data source, create one data source for each Oracle RAC node in the cluster on which the service will be configured, whether or not the service will be actively running on each node.

[Section B.8.1, "Configuring a Data Source to Connect to a Service,"](#) describes special considerations for configuring these data sources. [Section B.8.2, "Service Connection Configurations,"](#) shows example configurations for either load balancing or workload management.

B.8.1 Configuring a Data Source to Connect to a Service

You configure a data source to connect to a service running on an Oracle RAC node in the same way as you configure any data source (using WLST, the Administration Console, or the Configuration Wizard), with the following exceptions:

- `initial-capacity="0"`

This prevents pool creation failure for inactive pools at WLS startup, and enables WLS to create the data source even if it can't connect to the service on the node. Without setting this option to 0, data source creation will fail and the server may fail to boot normally.

In the Administration Console, edit the data source after creating it, and set **Initial Capacity** to 0.

- Oracle JDBC Thin (or Thin XA) driver for 11g. For example:

For non-XA:

```
driver-name="oracle.jdbc.OracleDriver"
url="jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=RAC1)(PORT=1521))) (CONNECT_DATA=(SERVICE_
NAME=Service_1)(INSTANCE_NAME=DB_02)))"
```

If configuring via the Administration Console, select **Oracles's Driver (Thin) for RAC Service-Instance connections** from the **Database Driver** drop-down and specify the service in the **Service Name** field.

For XA:

```
driver-name="oracle.jdbc.xa.client.OracleXADataSource"
url="jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=RAC1)(PORT=1521))) (CONNECT_DATA=(SERVICE_
NAME=Service1)(INSTANCE_NAME=DBase1)))"
```

If configuring via the Administration Console, select **Oracle's Driver (Thin XA) for RAC Service-Instance connections** from the **Database Driver** drop-down and specify the service in the **Service Name** field.

Notes: The `SERVICE_NAME` must be the same for all data sources in a given multi data source.

Specify a different `HOST NAME` and/or port for each data source in a given multi data source.

- When specifying `max-capacity` (**Maximum Capacity** in the Administration Console) for the connection pool, you need to consider the connection capacity of each of the Oracle RAC nodes in your configuration, and the total number of connections from all data sources. See [Section B.8.3, "Connection Pool Capacity Planning,"](#) for more information.

Selecting the Appropriate Multi Data Source Algorithm

For service connection scenarios, Oracle recommends that you configure your multi data source with the **Load Balancing** algorithm. If the multi data source is configured with the **Load Balancing** algorithm, its connection pools are used in a round robin fashion. In this case, workload is load-balanced across all of the Oracle RAC nodes on which the associated service is currently active.

If the multi data source is configured with the **Failover** algorithm, the first data source is used to connect to the service on its associated Oracle RAC node, until a connection attempt fails for any reason (for example, the Oracle RAC node becomes unavailable or there are no more connections available in the data source). At that point, the second data source is used to connect to the service on its associated Oracle RAC node, and so on. In this case, the Oracle RAC node to which the first data source is connected will experience more use than the remaining nodes on which the service is running.

B.8.2 Service Connection Configurations

You can design your configuration to provide:

- [Section B.8.2.1, "Workload Management"](#)
- [Section B.8.2.2, "Load Balancing"](#)

B.8.2.1 Workload Management


In a workload management configuration, each multi data source has one data source configured for a given service on each Oracle RAC node, regardless of whether the service you are connecting to is active or inactive on a given Oracle RAC node. This lets you quickly start an inactive service on a node and create connections to that service should another node become unavailable due to unplanned downtime or scheduled maintenance. It also lets you quickly increase or decrease the available capacity for a given service based on workload demands.

When you start the service on a node, the associated data source detects that the service is now active, and the data source will then start making connections to that node as needed. When you stop a service on a given node, the associated data source can no longer make connections to that node, and will become inactive until the service is restarted on that node.

The WLS data source performs connection testing. This lets the data source adjust to changes in the topology of the Oracle RAC configuration. The data source performs

polling to see if its associated service is active or inactive. The connection test fails if the service is no longer available on the Oracle RAC node.

Figure B-4 Workload Management using Multi Data Sources



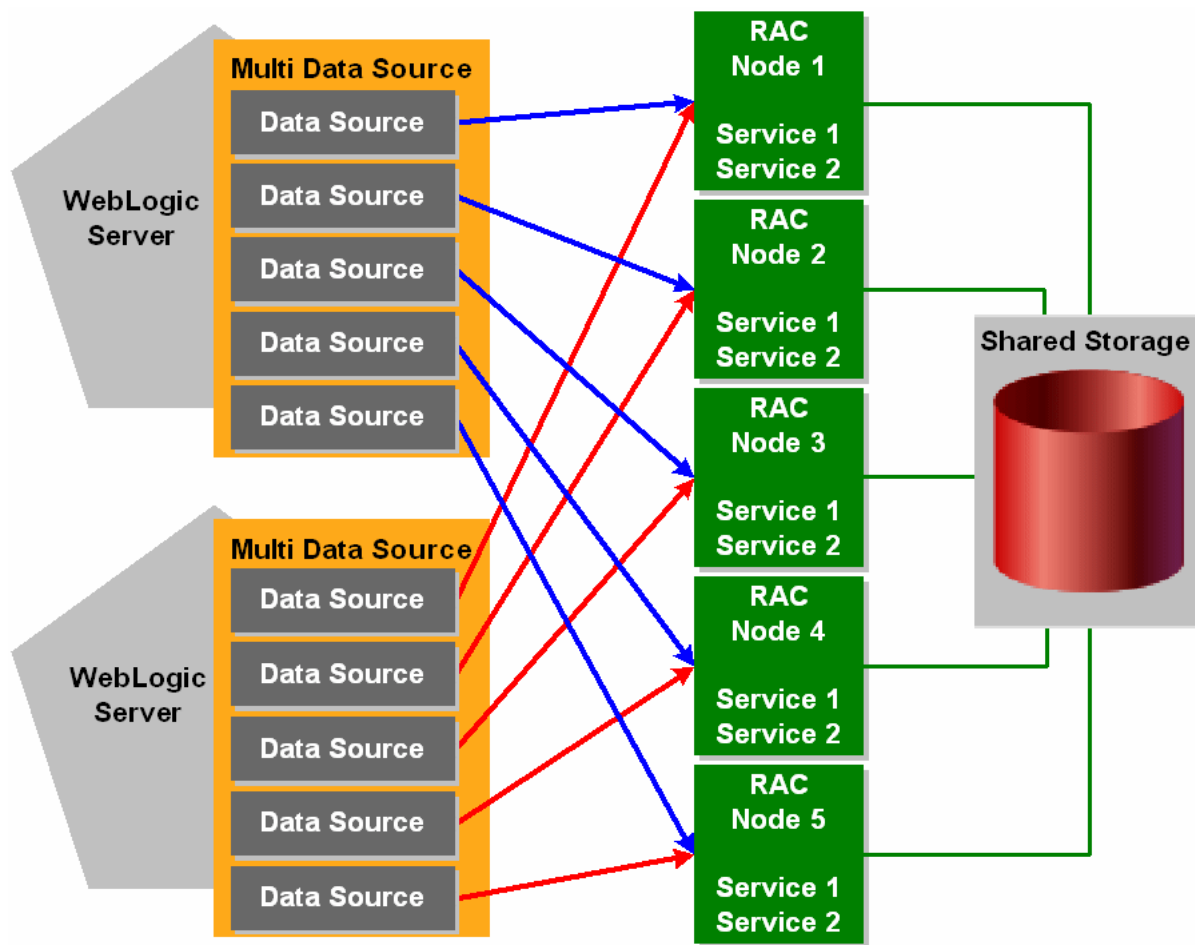
In this example, Service 1 is active on Oracle RAC Nodes 1, 2, and 3, while Service 2 is inactive on those nodes. Service 2 is active on Oracle RAC Nodes 4 and 5, while Service 1 is inactive on those nodes.

If Oracle RAC Node 1 becomes unavailable for any reason, you can start Service 1 on Oracle RAC Node 4. WebLogic Server will detect that the service is running on Node 4, and will begin making connections from the associated backup data source to Node 4 as needed.

B.8.2.2 Load Balancing

In a load balancing configuration, there are multiple services running concurrently on each Oracle RAC node. Each WLS multi data source has an active connection pool configured to connect to a given service on each of the nodes. In this scenario, you would configure each multi data source to use the Load Balancing algorithm.

Figure B-5 Load Balancing with Multi Data Sources



In this example, Service 1 and Service 2 are each actively running on all of the available Oracle RAC nodes. As a result, all of the connection pools in each multi data source will actively make connections in a round-robin fashion, balancing workload among the five nodes.

B.8.3 Connection Pool Capacity Planning

It is important to note the **Maximum Capacity** value you specify for a data source can cause the connection capacity to a given Oracle RAC node to be exceeded. You must consider the following factors when determining how to set this value for each of your data sources:

- The maximum number of simultaneous connections that a Oracle RAC node can support. This is based on the available memory on a given Oracle RAC node and the amount of memory consumed by each service connection (which can vary for each service). Memory consumption by each connection is a major limitation on the amount of work that can be generated from the WLS servers. Exceeding the amount of available memory by creating too many connections from your WLS data sources to a given Oracle RAC node can result in degraded performance on the Oracle RAC node, or could lead to failed connections.

Available memory for a node should be based on the PGA target attribute (per session memory).

- The *maximum* number of data sources that can potentially create connections to a given Oracle RAC node, and the number of WebLogic server instances to which each data source/multi data source is targeted. For example, if you have one data source that is targeted to three WLS servers, that data source counts as three data sources, as each server uses its own instance of the data source. This is the case whether the servers are part of a cluster or are independent server instances.
- The *maximum* number of services that may be actively running on a given Oracle RAC node, and the memory consumed on the node by each connection to each service.
- The expected relative workload for each service on a given node. For example, the expected workload of Service1 may be double that of the expected workload of Service2.

Regardless of whether or not a service is always active on a node, you should allocate resources for that service in the event you have to start it on the node.

- Always use the worst-case scenario when setting the **Maximum Capacity** value for your data sources. For example, assume that all available services will be actively running on the Oracle RAC node associated with each data source.

The following example explains how you could go about determining each data source's **Maximum Capacity** value. Keep in mind that this is a very simple example intended to illustrate the issue conceptually, and that real-world situations are much more complicated. In general, it is best to under-configure your data sources with a low **Maximum Capacity** value, monitor your Oracle RAC nodes for memory usage and performance, then adjust the **Maximum Capacity** values upward until you are approaching the maximum capacity of the associated Oracle RAC nodes.

Example

Suppose you have the following basic configuration:

- Five Oracle RAC nodes, each with 16 GB of memory.
- Two services actively running on each Oracle RAC node. Service1 uses 10MB per connection, Service2 uses 20MB per connection.
- Workload for each service is the same, that is, each service will generate an equivalent number of connections on a given Oracle RAC node.
- Two WebLogic Server clusters. Cluster1 has five servers. Cluster2 has four servers.
- For a given Oracle RAC node, one data source is targeted to Cluster1 and is configured to connect to Service1.
- For a given Oracle RAC node, one data source is targeted to Cluster 2 and is configured to connect to Service2.

Because Service2 uses twice as much memory per connection as Service1, you should allocate approximately 10GB of the node's memory for Service 2 and approximately 5GB for Service1.

Because Cluster1 has five WLS servers, there will be five data sources making connection requests to this Oracle RAC node. This gives you 1GB of memory available for connections from a given data source (5GB/5). Each connection requires 10MB of memory, so the **Maximum Capacity** value for each data source targeted to Cluster1 should be 100 or lower.

Because Cluster 2 has four WLS servers, there will be four data sources making connection requests to this Oracle RAC node. This gives you 2.5GB of memory available for connections from a given data source (10GB/4). Each connection requires

20MB, so the **Maximum Capacity** value for each data source targeted to Cluster2 should be 125 or lower.

If Service 2 generates more workload than Service1, you would have to adjust these values appropriately (increase the **Maximum Capacity** value for the data source connecting to Service2, decrease the value for the data source connecting to Service1). As long as:

$$\begin{aligned} &(\text{Max. connections to Service1} \times \text{memory used per connection}) + \\ &(\text{Max. connections to Service2} \times \text{memory used per connection}) < \\ &\text{Available memory} \end{aligned}$$

you can avoid the potential for performance degradation or connection failures.

Alternatively, in a simple configuration, such as is shown in [Figure B-6](#), the **Maximum Capacity** value you specify for each of your data sources can be loosely determined using the following formula:

$$\text{Maximum connection pool capacity} = \frac{\text{Maximum number of connections to Oracle RAC nodes}}{(\text{Number of WebLogic Server instances} \times \text{Number of data sources targeted to each instance} \times \text{Number of active Oracle RAC services configured} \times \text{Number of Oracle RAC Nodes})}$$

where:

Maximum number of connections to Oracle RAC nodes is determined by total memory available on all nodes divided by the memory consumed by each connection.

Number of WebLogic Server instances is the number of server instances to which the data sources are targeted. If the data sources is targeted to a WLS cluster, this is the number of servers in the cluster.

In the example in [Figure B-6](#):

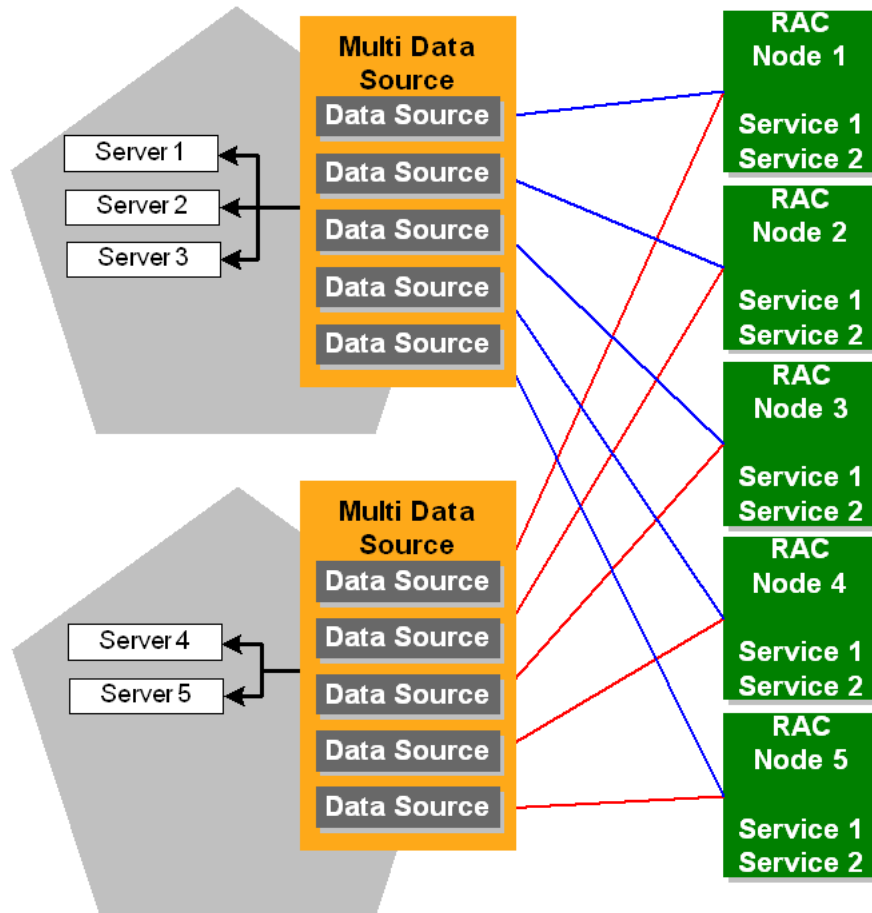
- assume that a maximum of 4000 total connections can be made to the group of Oracle RAC nodes, based on 8GB of available memory per Oracle RAC node, and 10MB of memory used per connection.
- there are a total of five server instances to which the data sources are targeted
- there are five data sources targeted to each server instance
- there are two services running on each Oracle RAC node, and
- there are five Oracle RAC nodes.

In this configuration, the **Maximum Capacity** value you would enter for each of your data sources would be:

$$\text{Maximum connection pool capacity} = \frac{4000}{(5 \text{ server instances} \times 5 \text{ data sources} \times 2 \text{ services} \times 5 \text{ Oracle RAC nodes})}$$

which would give you a **Maximum Capacity** value of 16 for each of your data sources.

Figure B-6 Example Multi Data Source Connection Configuration



Keep in mind that this formula is just a general guideline for configuring your data sources, as many configurations will be too complex for you to use such a simple calculation.

When calculating the **Maximum Capacity** value you should use, always consider the worst-case scenario that you will have in your overall configuration. It is best to under-configure this value for normal operation than to have it over-configured when a worst-case situation develops. You can always monitor your Oracle RAC nodes to determine if it is safe to increase the **Maximum Capacity** value for any of your data sources.

B.9 XA Considerations and Limitations when using multi Data Sources with Oracle RAC

When using XA (global transactions) with multi data sources for Oracle RAC, consider the following requirements and limitations.

- [Section B.9.1, "Oracle RAC XA Requirements when using multi Data Sources"](#)
- [Section B.9.2, "Known Limitations When Using Oracle RAC with multi Data Sources"](#)
- [Section B.9.3, "Known Issue Occurring After Database Server Crash"](#)

B.9.1 Oracle RAC XA Requirements when using multi Data Sources

Oracle RAC has the following requirements when using multi data sources with global transactions.

- [Section B.9.1.1, "Use Multi Data Sources"](#)
- [Section B.9.1.2, "A Global Transaction Must Be Initiated, Prepared, and Concluded in the Same Instance of the Oracle RAC Cluster"](#)
- [Section B.9.1.3, "Transaction IDs Must Be Unique Within the Oracle RAC Cluster"](#)

B.9.1.1 Use Multi Data Sources

Always use a multi data source when using XA transactions with multi data sources for Oracle RAC.

B.9.1.2 A Global Transaction Must Be Initiated, Prepared, and Concluded in the Same Instance of the Oracle RAC Cluster

Global transactions must be initiated, prepared, and concluded in the same instance of the Oracle RAC cluster. WebLogic Server data sources manage this for you when you set `KeepXAConnTillTxComplete="true"` in the JDBC data source configuration.

B.9.1.3 Transaction IDs Must Be Unique Within the Oracle RAC Cluster

When using global transactions, transaction IDs (XIDs) must be unique within the Oracle RAC cluster. However, neither the Oracle Thin driver nor an Oracle RAC instance can determine if an XID is unique within the Oracle RAC cluster. Transactions with the same XID can execute SQL code on different instances of the Oracle RAC cluster without any exception.

B.9.2 Known Limitations When Using Oracle RAC with multi Data Sources

The following sections describe known issues and limitations when using XA and multi data sources with Oracle RAC:

- [Section B.9.2.1, "Potential for Data Deadlocks in Some Failure Scenarios"](#)
- [Section B.9.2.2, "Potential for Transactions Completed Out of Sequence"](#)

Note: Some of these limitations are also described in Oracle's bug numbers 3428146 and 395790. Contact Oracle for more information about these issues.

B.9.2.1 Potential for Data Deadlocks in Some Failure Scenarios

There is a window of time in which transaction IDs are not available across the Oracle RAC cluster. Because of this known limitation, after some failure conditions, some incomplete transactions cannot be properly completed, which can result in deadlocks in the database. To prevent these failure conditions from arising, WebLogic Server provides two configuration attributes that enable XA call retry for Oracle RAC: `XARetryDurationSeconds` and `XARetryIntervalSeconds`. For more information about these configuration options, see [Section B.5.3.2, "Delays During Failover."](#)

B.9.2.2 Potential for Transactions Completed Out of Sequence

When using the Oracle DataBase Control, the order of transaction processing is not guaranteed. For example, if you implement a web service that uses DataBase Control do the following transaction sequence:

1. Create a table
2. Insert record 1
3. Insert record 2
4. Insert record 3
5. Select records

If the primary node goes down momentarily after the table is created, it is possible that transactions submitted to the database are performed out of sequence.

B.9.3 Known Issue Occurring After Database Server Crash

If, while a transaction is being processed, the database server instance crashes after the PREPARE operation is complete but before the results of that operation have been written to the transaction log, a COMMIT call from a client for that transaction may hang for several minutes and possibly until the TCP timeout period has expired. The window of time in which this might occur is small and the problem occurs rarely. There is no workaround for the issue at this time.

B.10 JDBC Store Recovery with Oracle RAC

If you are using a JDBC Store with Oracle RAC, there are features and limitations to consider that concern Oracle RAC node failover. See the following sections:

- [Configuring a JDBC Store for Use with Oracle RAC](#)
- [Automatic Retry](#)

For a list of the major services that use the JDBC store, see "Monitoring Store Connections" in *Configuring Server Environments for Oracle WebLogic Server*.

B.10.1 Configuring a JDBC Store for Use with Oracle RAC

The way that a JDBC Store works limits the options you have for configuring one for use with Oracle RAC. You cannot configure a JDBC store to use a JDBC data source that is configured to support global transactions. The JDBC store must use a JDBC data source that uses a non-XA JDBC driver. For more information about this configuration option, see [Section B.7, "Using Multi Data Sources without Global Transactions."](#)

A JDBC Store holds on to a connection until that connection fails, at which point it moves on to the next connection and repeats the process. Therefore you cannot implement load balancing with a JDBC Store, including using a load balancing multi data source. You should configure a multi data source for a JDBC store to use the Failover algorithm.

In short, for a JDBC store:

- Use a non-XA driver
- Configure the multi data source for Failover mode.

B.10.2 Automatic Retry

JMS has a limited connection retry mechanism which enables it to silently react to the failure of the Oracle RAC node that hosts its database connection. If the database has experienced either a minor network 'hiccup' or a Oracle RAC database has failed over to another node, the second connection attempt (the retry) will succeed to the next Oracle RAC node.

The time within which this retry is attempted and the number of retries attempted are limited to minimize the negative effects that an extended connection retry time could cause. If the database connection remains unavailable for a long period of time, the delay can impede the ability of JMS to properly continue its processing (for example, to maintain proper message ordering). Also, the transaction manager could declare the JMS resource of a transaction to be dead if there is not enough processing progress made within this time period, or out-of memory conditions could arise. There are system-level tuning guidelines that can help minimize the Oracle RAC failover time frame which is critical to the success of the automatic retry.

The tight loop on the automatic retry is particularly important when JMS processing occurs with transactions. If an I/O failure occurs in the JDBC Store, the store record is in an unknown state which will put the message itself in an unknown state. To prevent the message from being committed in this unknown state, JMS will mark the transaction associated with the message as a "failedTransaction." Any future attempts by the transaction manager to finishing committing the message will cause JMS to throw a `javax.transaction.xa.XAException` with an `errorCode` set to `XAException.XAER_RMERR`. This exception is an indication to the transaction manager that a transient error has occurred in the resource manager (JMS) and that the transaction manager should retry commit processing. The retry logic provides a second attempt to establish the connection before JMS communicates any failure to the upper layer which would translate into an RMERR. If the RMERR is generated, then the only way to recover the message and complete the transaction is to restart WebLogic Server.

The automatic retry logic is currently governed by an option on WebLogic Server as follows:

```
-Dweblogic.store.jdbc.IORetryDelaySeconds=X
```

Where X is the number of seconds that should elapse before the connection to the database is retried. The default is one second. This value is restricted to the range 0 to 15, and the retry will only be attempted once. If a failure occurs on the second attempt, an exception will be propagated up the call stack and a manual restart will be required to recover the messages associated with the failed transaction.

Note: In the event that an automatic retry attempt is not successful, you must restart WebLogic Server.

Using Connect-Time Failover with Oracle RAC


WebLogic Server provides Connect-Time Failover for legacy applications that use data sources configured to use connect-time failover and load balancing.

Note: New applications should consider using a GridLink Data Source, which provides the same capabilities. See [Section 4.1, "What is a GridLink Data Source."](#)

- [Section C.1, "Using Connect-Time Failover without Global Transactions"](#)
- [Section C.2, "Attributes of a Connect-Time Failover Configuration without Global Transactions"](#)
- [Section C.3, "Sample Configuration Code"](#)

C.1 Using Connect-Time Failover without Global Transactions

To connect WebLogic Server to multiple Oracle RAC nodes using data sources configured for connect-time failover and load balancing, configure a JDBC data source for each Oracle RAC instance in your Oracle RAC cluster with the Oracle Thin driver, as described in the sections that follow. [Figure C-1](#) shows an overview of the system.

Figure C-1 Data Source Configuration with Oracle Thin Driver Connect-Time Failover

You can use the Administration Console or any other means that you prefer to configure your domain, such as the Weblogic Scripting Tool (WLST) or a JMX program.

When connections are created in the data source, the Oracle Thin driver determines which Oracle RAC instance to use. When an application gets a connection, it looks up a data source on the JNDI tree and requests a connection from the data source. The data source delivers one of the available connections from the pool of connections in the data source.

The following sections describe a configuration which uses Oracle RAC's connect-time failover features to handle connection failures. With this configuration, in some failure cases, the failover time is as long as the TCP timeout, which can be several minutes, depending on your environment.

C.2 Attributes of a Connect-Time Failover Configuration without Global Transactions

To use this configuration, create JDBC data sources in your WebLogic domain with the following attributes.

- Oracle JDBC Thin driver 11g configured for connect-time failover. For example:


```
<url>jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(
  PROTOCOL=TCP)(HOST=lcqso124)(PORT=1521))(ADDRESS=(PROTOCOL=TCP)
  (HOST=lcqso125)(PORT=1521))(FAILOVER=on)(LOAD_BALANCE=off))
  (CONNECT_DATA=(SERVER=DEDICATED)(SERVICE_NAME=snrac))</url>
<driver-name>oracle.jdbc.OracleDriver</driver-name>
```
- ConnectionReserveTimeoutSeconds="120"
 - Enables application requests for a connection to wait 120 seconds for a connection to become available.
- TestConnectionsOnReserve="true"
 - Enables testing of a database connection when an application reserves a connection from the data source. See [Section 5.4.1, "Test Connections on Reserve to Enable Fail-Over"](#) for more details about this attribute.
 - Required to enable failover to another Oracle RAC node.
- TestTableName="name_of_small_table" The name of the table used to test a physical database connection. For more details about this attribute, see [Section 10.2, "Connection Testing Options for a Data Source."](#)

C.3 Sample Configuration Code

Sample configuration code is shown below.

Note: Line breaks added for readability.

```
<jdbc-data-source xmlns="http://xmlns.oracle.com/weblogic/jdbc-data-source"
  xmlns:sec="http://xmlns.oracle.com/weblogic/security"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:wls="http://xmlns.oracle.com/weblogic"
  xsi:schemaLocation="http://xmlns.oracle.com/weblogic/domain/1.0/domain.xsd">
  <name>oracleRACNonXAPool</name>
  <jdbc-driver-params>
    <url>jdbc:oracle:thin:@(DESCRIPTION=
      (ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)
        HOST=lcqso124)(PORT=1521))(ADDRESS=(PROTOCOL=TCP)
        (HOST=lcqso125)(PORT=1521))(FAILOVER=on)
        (LOAD_BALANCE=off))(CONNECT_DATA=(SERVER=DEDICATED)
        (SERVICE_NAME=snrac))</url>
    <driver-name>oracle.jdbc.OracleDriver</driver-name>
    <properties>
      <property>
        <name>user</name>
        <value>wlsqa</value>
      </property>
    </properties>
    <password-encrypted>{3DES}aP/xScCS8uI=</password-encrypted>
  </jdbc-driver-params>
```

```
<jdbc-connection-pool-params>
  <test-connections-on-reserve>true</test-connections-on-reserve>
  <test-table-name>SQL SELECT 1 FROM DUAL</test-table-name>
  <profile-type>4</profile-type>
</jdbc-connection-pool-params>
<jdbc-data-source-params>
  <jndi-name>oracleRACJndiName</jndi-name>
  <global-transactions-protocol>OnePhaseCommit
    </global-transactions-protocol>
</jdbc-data-source-params>
</jdbc-data-source>
```

Using Fast Connection Failover with Oracle RAC

WebLogic Server supports Fast Connection Failover, an Oracle feature which provides an application-independent method to implement Oracle RAC event notifications, such as detection and cleanup of invalid connections, load balancing of available connections, and work redistribution on active Oracle RAC instances.

For more information, see "Fast Connection Failover" at http://download-east.oracle.com/docs/cd/B19306_01/java.102/b14355/fstconfo.htm#CIHJBFFC in the *Oracle® Database JDBC Developer's Guide and Reference*.

D.1 JDBC Driver Configuration for use with Oracle Fast Connection Failover

To enable Fast Connection Failover on a data source, set the following connection pool properties:

- In Driver Class Name—set the class name to `oracle.jdbc.pool.OracleDataSource`.
- In Properties—set the ONS configuration string to remotely subscribe the Oracle RAC nodes to Oracle FAN/ONS events. For example:
`ONSConfiguration=nodes=hostname1:port1,hostname2:port2`

Note: Oracle's `OracleDataSource` class is not XA-capable, so the resulting data source does not implement a XA connection pool.

Because the Oracle driver uses ONS to implement Fast Connection Failover, you must add `ons.jar` to the WebLogic classpath. To do so, edit the `setDomainEnv` script in the domain's `bin` directory, and add `ons.jar` to the `POST_CLASSPATH` variable.

