

Oracle® Fusion Middleware

Administering Server Environments for Oracle WebLogic Server
12.1.3

12c (12.1.3)

E41942-07

August 2015

This document describes how to design, configure, and manage WebLogic Server 12.1.3 environments. It is a resource for system administrators and operators responsible for implementing a WebLogic Server installation.

Oracle Fusion Middleware Administering Server Environments for Oracle WebLogic Server 12.1.3, 12c (12.1.3)

E41942-07

Copyright © 2007, 2015, 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 is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware 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 that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about 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 unless otherwise set forth in an applicable agreement between you and Oracle. 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, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface	ix
Documentation Accessibility	ix
Conventions	ix
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 New and Changed Features in This Release	1-2
2 Using Work Managers to Optimize Scheduled Work	
2.1 Understanding How WebLogic Server Uses Thread Pools	2-1
2.2 Understanding Work Managers	2-2
2.2.1 Request Classes	2-3
2.2.2 Constraints	2-5
2.2.3 Stuck Thread Handling	2-6
2.2.4 Self-Tuning Thread Pool	2-6
2.3 Work Manager Scope	2-7
2.3.1 The Default Work Manager	2-8
2.3.1.1 Overriding the Default Work Manager	2-8
2.3.1.2 When to Use Work Managers	2-8
2.3.2 Global Work Managers	2-8
2.3.3 Application-scoped Work Managers	2-8
2.4 Using Work Managers, Request Classes, and Constraints	2-9
2.4.1 Dispatch Policy for EJB	2-9
2.4.2 Dispatch Policy for Web Applications	2-9
2.5 Deployment Descriptor Examples	2-9
2.6 Work Managers and Execute Queues	2-13
2.6.1 Enabling Execute Queues	2-13
2.6.2 Migrating from Execute Queues to Work Managers	2-14
2.7 Accessing Work Managers Using MBeans	2-14
2.8 Using CommonJ With WebLogic Server	2-14
2.8.1 Accessing CommonJ Work Managers	2-15
2.8.2 Mapping CommonJ to WebLogic Server Work Managers	2-15

3 Avoiding and Managing Overload

3.1	Configuring WebLogic Server to Avoid Overload Conditions	3-1
3.1.1	Limiting Requests in the Thread Pool	3-1
3.1.1.1	Work Managers and Thread Pool Throttling	3-2
3.1.2	Limiting HTTP Sessions	3-2
3.1.3	Exit on Out of Memory Exceptions	3-2
3.1.4	Stuck Thread Handling.....	3-3
3.2	WebLogic Server Self-Monitoring	3-3
3.2.1	Overloaded Health State.....	3-3
3.3	WebLogic Server Exit Codes	3-3

4 Configuring Network Resources

4.1	Overview of Network Configuration	4-1
4.2	Understanding Network Channels	4-2
4.2.1	What Is a Channel?	4-2
4.2.1.1	Rules for Configuring Channels.....	4-2
4.2.1.2	Custom Channels Can Inherit Default Channel Attributes	4-2
4.2.2	Why Use Network Channels?.....	4-3
4.2.2.1	Handling Channel Failures	4-3
4.2.2.2	Upgrading Quality of Service Levels for RMI.....	4-3
4.2.3	Standard WebLogic Server Channels	4-4
4.2.3.1	The Default Network Channel	4-4
4.2.3.2	Administration Port and Administrative Channel.....	4-4
4.2.3.2.1	Administration Port Capabilities	4-5
4.2.3.2.2	Administration Port Restrictions	4-5
4.2.3.2.3	Administration Port Requires SSL.....	4-5
4.2.3.2.4	Configure Administration Port	4-6
4.2.3.2.5	Booting Managed Servers to Use Administration Port	4-6
4.2.3.2.6	Booting Managed Servers to Use Administrative Channels.....	4-6
4.2.3.2.7	Custom Administrative Channels	4-6
4.2.4	Using Internal Channels	4-7
4.2.4.1	Channel Selection	4-7
4.2.4.2	Internal Channels Within a Cluster	4-7
4.3	Configuring a Channel.....	4-7
4.3.1	Guidelines for Configuring Channels	4-8
4.3.1.1	Channels and Server Instances	4-8
4.3.1.2	Dynamic Channel Configuration	4-8
4.3.1.3	Channels and Identity	4-8
4.3.1.4	Channels and Protocols	4-8
4.3.1.5	Reserved Names	4-8
4.3.1.6	Channels, Proxy Servers, and Firewalls	4-9
4.3.2	Configuring Network Channels For a Cluster	4-9
4.3.2.1	Create the Cluster	4-9
4.3.2.2	Create and Assign the Network Channel	4-9
4.3.2.3	Configuring a Replication Channel	4-10
4.3.2.4	Increase Packet Size When Using Many Channels	4-10
4.4	Assigning a Custom Channel to an EJB.....	4-10

4.5	Using IPv6 with IPv4.....	4-10
-----	---------------------------	------

5 Configuring Web Server Functionality

5.1	Overview of Configuring Web Server Components	5-1
5.2	Configuring the Server.....	5-1
5.2.1	Configuring the Listen Port	5-2
5.3	Web Applications.....	5-2
5.3.1	Web Applications and Clustering	5-3
5.4	Configuring Virtual Hosting.....	5-3
5.4.1	Virtual Hosting and the Default Web Application.....	5-3
5.4.2	Setting Up a Virtual Host	5-4
5.5	How WebLogic Server Resolves HTTP Requests	5-4
5.6	Setting Up HTTP Access Logs	5-5
5.6.1	Log Rotation	5-5
5.6.2	Common Log Format	5-5
5.6.3	Setting Up HTTP Access Logs by Using Extended Log Format.....	5-6
5.6.3.1	Creating the Fields Directive	5-7
5.6.3.2	Supported Field Identifiers	5-7
5.6.3.2.1	IP Address Related Fields	5-7
5.6.3.2.2	DNS Related Fields	5-8
5.6.3.2.3	Diagnostic Message Correlation Fields.....	5-8
5.6.3.3	Creating Custom Field Identifiers.....	5-9
5.6.3.3.1	Get Methods of the HttpAccountingInfo Object	5-10
5.7	Preventing POST Denial-of-Service Attacks.....	5-11
5.8	Setting Up WebLogic Server for HTTP Tunneling	5-12
5.8.1	Configuring the HTTP Tunneling Connection.....	5-12
5.8.2	Connecting to WebLogic Server from the Client	5-13
5.9	Using Native I/O for Serving Static Files (Windows Only).....	5-13

6 Using the WebLogic Persistent Store

6.1	Overview of the Persistent Store	6-1
6.1.1	Features of the Persistent Store.....	6-2
6.1.2	High-Performance Throughput and Transactional Support.....	6-2
6.1.3	Comparing File Stores and JDBC-accessible Stores.....	6-3
6.1.4	High Availability For Persistent Stores	6-3
6.1.4.1	Persistent Store Migration	6-4
6.1.4.2	High Availability Storage Solutions	6-4
6.2	Using the Default Persistent Store.....	6-5
6.2.1	Default Store Location.....	6-5
6.2.2	Example of a Default File Store	6-5
6.3	Using Custom File Stores and JDBC Stores	6-6
6.3.1	When to Use a Custom Persistent Store	6-6
6.3.2	Methods of Creating a Custom Persistent Store	6-6
6.3.3	Modifying Custom Persistent Store Parameters	6-7
6.4	Creating a Custom (User-Defined) File Store	6-7
6.4.1	Main Steps for Configuring a Custom File Store	6-7

6.4.2	Example of a Custom File Store.....	6-8
6.4.3	Guidelines for Configuring a Synchronous Write Policy	6-9
6.4.3.1	Direct-Write-With-Cache Policy	6-9
6.4.3.2	Direct-Write Policy	6-10
6.4.3.3	Cache-Flush Policy	6-10
6.4.3.4	Disabled Policy	6-11
6.5	Creating JDBC-accessible Stores	6-11
6.5.1	Using a JDBC TLog Store.....	6-11
6.5.1.1	Main Steps for Configuring a JDBC TLOG Store.....	6-11
6.5.1.1.1	Choosing a Data Source.....	6-12
6.5.1.2	Example of a JDBC TLOG Store	6-12
6.5.1.3	Configuration Guidelines.....	6-13
6.5.1.4	Additional Considerations	6-14
6.5.1.5	Server Migration when using a JDBC TLOG Store	6-14
6.5.1.6	Monitoring a JDBC TLOG Store.....	6-14
6.5.1.6.1	How to Monitor the JDBC TLOG Store Health State	6-15
6.5.1.6.2	How to Monitor Transaction Log Store Statistics.....	6-15
6.5.1.6.3	How to Monitor Transaction Log Store Connections	6-15
6.5.1.7	Security Considerations.....	6-15
6.5.2	Using a JDBC Store	6-15
6.5.2.1	Main Steps for Configuring a JDBC Store.....	6-15
6.5.2.2	Example of a JDBC Store	6-16
6.5.2.3	Supported JDBC Drivers	6-17
6.5.2.4	Creating a JDBC Store Table Using Default and Custom DDL Files.....	6-18
6.5.2.4.1	Creating a JDBC Store Table Using a Custom DDL File	6-18
6.5.2.4.2	Enabling Oracle BLOB Record Columns	6-19
6.5.2.5	Managing JDBC Store Tables.....	6-20
6.5.2.5.1	Using the utils.Schema Utility to Delete a JDBC Store Table.....	6-20
6.5.2.6	Guidelines for Configuring a JDBC Store	6-21
6.5.2.6.1	Using Prefixes with a JDBC Store	6-21
6.5.2.6.2	JDBC Store Table Rules	6-21
6.5.2.6.3	Prefix Name Format Guidelines.....	6-21
6.5.2.6.4	Recommended JDBC Data Source Settings for JDBC Stores.....	6-22
6.5.2.6.5	Automatic Reconnection to Failed Databases.....	6-22
6.5.2.6.6	Required Setting for Oracle DB2 Type 4 JDBC Drivers.....	6-22
6.5.2.6.7	Handling JMS Transactions with JDBC Stores.....	6-22
6.5.2.7	Enabling I/O Multithreading for JDBC Stores.....	6-23
6.5.2.7.1	Rebuilding the Store Table Index for an Oracle Database	6-24
6.5.2.7.2	Build a Reverse Index for an Oracle Database	6-24
6.5.2.7.3	Build a Non-Reverse Index for an Oracle Database.....	6-24
6.5.2.7.4	Reducing Contention in a Non-Oracle Database.....	6-25
6.6	Monitoring a Persistent Store.....	6-25
6.6.1	Monitoring Stores	6-25
6.6.2	Monitoring Store Connections.....	6-25
6.7	Administering a Persistent Store	6-26
6.7.1	Store Administration Using a Java Command Line	6-27
6.7.1.1	Accessing Store Administration Help	6-28

6.7.1.2	Dumping the Contents of a File Store	6-28
6.7.1.3	Compacting a File Store.....	6-28
6.7.2	Store Administration Using WLST	6-28
6.7.2.1	Accessing Store Administration Help	6-29
6.7.2.2	Dumping the Contents of a JDBC Store Using WLST.....	6-29
6.7.2.3	Compacting a File Store Using WLST	6-30
6.8	Security Considerations	6-30
6.9	Limitations and Considerations of the Persistent Store	6-30

7 Using the Plug-in for Oracle Virtual Assembly Builder

7.1	Overview	7-1
7.1.1	About Oracle Virtual Assembly Builder	7-2
7.1.2	About the WebLogic Server Introspection Plug-in.....	7-2
7.1.3	Setting Up the WebLogic Server Introspection Plug-in	7-2
7.2	Introspection Plug-in Parameters	7-3
7.3	Reference System Prerequisites	7-3
7.4	Plug-in Usage Requirements.....	7-3
7.5	Resulting Artifact Type	7-4
7.6	Wiring.....	7-5
7.7	Wiring Properties.....	7-5
7.8	WebLogic Server Appliance Properties.....	7-7
7.8.1	Assembly-Level System Properties.....	7-7
7.8.2	Properties Common to Administration Server and Managed Server Appliances.....	7-8
7.8.3	Administration Server Appliance Properties	7-9
7.9	Oracle Coherence*Web Introspection Extension	7-9
7.10	Supported Template Types	7-10

Preface

This preface describes the document accessibility features and conventions used in this guide—*Administering Server Environments for Oracle WebLogic Server 12.1.3*.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> 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.
monospace	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—*Administering Server Environments for Oracle WebLogic Server 12.1.3*.

This chapter includes the following sections:

- [Section 1.1, "Document Scope and Audience"](#)
- [Section 1.2, "Guide to This Document"](#)
- [Section 1.3, "Related Documentation"](#)
- [Section 1.4, "New and Changed Features in This Release"](#)

1.1 Document Scope and Audience

This document describes how you design, configure, and manage WebLogic Server environments. It is a resource for system administrators and operators responsible for implementing a WebLogic Server installation. This document is relevant to all phases of a software project, from development through test and production phases.

It is assumed that the reader is familiar with Java EE and Web technologies, object-oriented programming techniques, and the Java programming language.

1.2 Guide to This Document

The document is organized as follows:

- This chapter, [Chapter 1, "Introduction and Roadmap"](#) describes the scope of this guide and lists related documentation.
- [Chapter 2, "Using Work Managers to Optimize Scheduled Work"](#) describes the WebLogic Server execution model and the process of configuring application access to the execute queue.
- [Chapter 3, "Avoiding and Managing Overload"](#) describes detecting, avoiding, and recovering from overload conditions.
- [Chapter 4, "Configuring Network Resources"](#) describes optimizing your WebLogic Server domain for your network.
- [Chapter 5, "Configuring Web Server Functionality"](#) describes using WebLogic Server as a Web server.
- [Chapter 6, "Using the WebLogic Persistent Store"](#) describes configuring and monitoring the persistent store, a built-in, high-performance storage solution for WebLogic Server subsystems and services that require persistence.

- [Chapter 7, "Using the Plug-in for Oracle Virtual Assembly Builder"](#) describes the WebLogic Server introspection plug-in for Oracle Virtual Assembly Builder, which can be used to examine a single WebLogic domain and the Middleware home directory in which it resides.

1.3 Related Documentation

- *Understanding Domain Configuration for Oracle WebLogic Server*
- *Oracle WebLogic Server Administration Console Online Help*

1.4 New and Changed Features in This Release

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

Using Work Managers to Optimize Scheduled Work

This chapter describes how WebLogic Server 12.1.3 lets you configure how your application prioritizes the execution of its work using a Work Manager. Based on rules you define and by monitoring actual run-time performance, WebLogic Server can optimize the performance of your application and maintain service-level agreements. You define the rules and constraints for your application by defining a Work Manager and applying it either globally to a WebLogic Server domain or to a specific application component.

This chapter includes the following sections:

- [Section 2.1, "Understanding How WebLogic Server Uses Thread Pools"](#)
- [Section 2.2, "Understanding Work Managers"](#)
- [Section 2.3, "Work Manager Scope"](#)
- [Section 2.4, "Using Work Managers, Request Classes, and Constraints"](#)
- [Section 2.5, "Deployment Descriptor Examples"](#)
- [Section 2.6, "Work Managers and Execute Queues"](#)
- [Section 2.7, "Accessing Work Managers Using MBeans"](#)
- [Section 2.8, "Using CommonJ With WebLogic Server"](#)

2.1 Understanding How WebLogic Server Uses Thread Pools

In previous versions of WebLogic Server, processing was performed in multiple execute queues. Different classes of work were executed in different queues, based on priority and ordering requirements, and to avoid deadlocks. In addition to the default execute queue, `weblogic.kernel.default`, there were pre-configured queues dedicated to internal administrative traffic, such as `weblogic.admin.HTTP` and `weblogic.admin.RMI`.

You could control thread usage by altering the number of threads in the default queue, or configure custom execute queues to ensure that particular applications had access to a fixed number of execute threads, regardless of overall system load.

Now WebLogic Server uses a single thread pool, in which all types of work are executed. WebLogic Server prioritizes work based on rules you define, and run-time metrics, including the actual time it takes to execute a request and the rate at which requests are entering and leaving the pool.

The common thread pool changes its size automatically to maximize throughput. The queue monitors throughput over time and based on history, determines whether to adjust the thread count. For example, if historical throughput statistics indicate that a higher thread count increased throughput, WebLogic increases the thread count. Similarly, if statistics indicate that fewer threads did not reduce throughput, WebLogic decreases the thread count. This new strategy makes it easier for administrators to allocate processing resources and manage performance, avoiding the effort and complexity involved in configuring, monitoring, and tuning custom executes queues.

2.2 Understanding Work Managers

WebLogic Server prioritizes work and allocates threads based on an execution model that takes into account administrator-defined parameters and actual run-time performance and throughput.

Administrators can configure a set of scheduling guidelines and associate them with one or more applications, or with particular application components. For example, you can associate one set of scheduling guidelines for one application, and another set of guidelines for other applications. At run time, WebLogic Server uses these guidelines to assign pending work and enqueued requests to execution threads.

Note: Work requests from all Work Managers are executed by a single thread pool; separate thread pools are not created for each Work Manager.

To manage work in your applications, you define one or more of the following Work Manager components:

- Fair Share Request Class
- Response Time Request Class
- Min Threads Constraint
- Max Threads Constraint
- Capacity Constraint
- Context Request Class

For more information on these components, see [Section 2.2.1, "Request Classes"](#) or [Section 2.2.2, "Constraints"](#).

You can use any of these Work Manager components to control the performance of your application by referencing the name of the component in the application deployment descriptor. In addition, you may define a *Work Manager* that encapsulates all of the above components (except Context Request Class; see [Example 2-3](#)) and reference the name of the Work Manager in your application's deployment descriptor. You can define multiple Work Managers—the appropriate number depends on how many distinct demand profiles exist across the applications you host on WebLogic Server.

You can configure Work Managers at the domain level, application level, and module level in one of the following configuration files, or by using the WebLogic Server Administration Console:

- `config.xml`—Work Managers specified in `config.xml` can be assigned to any application, or application component, in the domain.

- `weblogic-application.xml`—Work Managers specified at the application level can be assigned to that application, or any component of that application.
- `weblogic-ejb-jar.xml` or `weblogic.xml`—Work Managers specified at the component level can be assigned to that component.
- `weblogic.xml`—Work Managers specified for a Web application.

[Example 2-1](#) is an example of a Work Manager definition.

Example 2-1 Work Manager Stanza

```
<work-manager>
<name>highpriority_workmanager</name>
  <fair-share-request-class>
    <name>high_priority</name>
    <fair-share>100</fair-share>
  </fair-share-request-class>
  <min-threads-constraint>
    <name>MinThreadsCountFive</name>
    <count>5</count>
  </min-threads-constraint>
</work-manager>
```

To assign the Work Manager in [Example 2-1](#) to control the dispatch policy of the entire Web application, add the code in [Example 2-2](#) to the Web application's `weblogic.xml` file:

Example 2-2 Referencing the Work Manager in a Web Application

```
<wl-dispatch-policy>highpriority-workmanager</wl-dispatch-policy>
```

To assign the Work Manager to control the dispatch policy of a particular servlet, add the following code to the Web application's `web.xml` file:

```
<servlet>
  ...
  <init-param>
    <param-name>wl-dispatch-policy</param-name>
    <param-value>highpriority_workmanager</param-value>
  </init-param>
  ...
</servlet>
```

The components you can define and use in a Work Manager are described in following sections:

- [Section 2.2.1, "Request Classes"](#)
- [Section 2.2.2, "Constraints"](#)
- [Section 2.2.3, "Stuck Thread Handling"](#)

2.2.1 Request Classes

A request class expresses a scheduling guideline that WebLogic Server uses to allocate threads to requests. Request classes help ensure that high priority work is scheduled before less important work, even if the high priority work is submitted after the lower priority work. WebLogic Server takes into account how long it takes for requests to each module to complete.

Request classes define a best effort. They do not guarantee that the configured ratio will be maintained consistently. The observed ratio may vary due to several factors during a period of sufficient demand, such as:

- The mixture of requests from different request classes in the queue at any particular time. For example, more requests than the configured ratio may be processed for a lower priority request class if there are not enough requests from a higher priority request class in the Work Manager queue.
- Because the ratio is specified in terms of thread-usage time, a larger number of shorter requests could be processed in the same amount of thread-usage time as a smaller number of time-consuming requests.

There are multiple types of request classes, each of which expresses a scheduling guideline in different terms. A Work Manager may specify only one request class.

- `fair-share-request-class`—Specifies the average thread-use time required to process requests. The default fair share value is 50.

For example, assume that WebLogic Server is running two modules. The Work Manager for `ModuleA` specifies a `fair-share-request-class` of 80 and the Work Manager for `ModuleB` specifies a `fair-share-request-class` of 20.

During a period of sufficient demand, with a steady stream of requests for each module such that the number requests exceed the number of threads, WebLogic Server will allocate 80% and 20% of the thread-usage time to `ModuleA` and `ModuleB`, respectively.

Note: The value of a fair share request class is specified as a relative value, not a percentage. Therefore, in the above example, if the request classes were defined as 400 and 100, they would still have the same relative values.

- `response-time-request-class`—Specifies a response time goal in milliseconds. Response time goals are not applied to individual requests. Instead, WebLogic Server computes a tolerable waiting time for requests with that class by subtracting the observed average thread use time from the response time goal, and schedules requests so that the average wait for requests with the class is proportional to its tolerable waiting time.

For example, given that `ModuleA` and `ModuleB` in the previous example, have response time goals of 2000 ms and 5000 ms, respectively, and the actual thread use time for an individual request is less than its response time goal. During a period of sufficient demand, with a steady stream of requests for each module such that the number of requests exceed the number of threads, and no "think time" delays between response and request, WebLogic Server will schedule requests for `ModuleA` and `ModuleB` to keep the average response time in the ratio 2:5. The actual average response times for `ModuleA` and `ModuleB` might be higher or lower than the response time goals, but will be a common fraction or multiple of the stated goal. For example, if the average response time for `ModuleA` requests is 1,000 ms., the average response time for `ModuleB` requests is 2,500 ms.

The previous sections described request classes based on fair share and response time by relating the scheduling to other work using the same request class. A mix of fair share and response time request classes is scheduled with a marked bias in favor of response time scheduling.

- `context-request-class`—Assigns request classes to requests based on context information, such as the current user or the current user's group.

For example, the `context-request-class` in [Example 2-3](#) assigns a request class to requests based on the value of the request's `subject` and `role` properties.

Example 2-3 Context Request Class

```
<work-manager>
  <name>context_workmanager</name>
  <context-request-class>
    <name>test_context</name>
    <context-case>
      <user-name>system</user-name>
      <request-class-name>high_fairshare</request-class-name>
    </context-case>
    <context-case>
      <group-name>everyone</group-name>
      <request-class-name>low_fairshare</request-class-name>
    </context-case>
  </context-request-class>
</work-manager>
```

The `high_fairshare` and `low_fairshare` request classes referenced by the `context_workmanager` in [Example 2-3](#) could be defined in the `config.xml` as follows:

```
<self-tuning>
  ...
  <fair-share-request-class>
    <name>high_fairshare</name>
    <target>myserver</target>
    <fair-share>75</fair-share>
  </fair-share-request-class>
  <fair-share-request-class>
    <name>low_fairshare</name>
    <target>myserver</target>
    <fair-share>25</fair-share>
  </fair-share-request-class>
  ...
</self-tuning>
```

Note: If a Web application's Work Manager references a context request class, the first user call will go through the default request class; subsequent calls in same session will go through the user-defined request class.

When using context request classes, set session timeout values to prevent sessions from expiring while requests wait in the Work Manager queue.

2.2.2 Constraints

A constraint defines minimum and maximum numbers of threads allocated to execute requests and the total number of requests that can be queued or executing before WebLogic Server begins rejecting requests.

You can define the following types of constraints:

- `max-threads-constraint`—Limits the number of concurrent threads executing requests from the constrained work set. The default is unlimited. For example, consider a constraint defined with maximum threads of 10 and shared by 3 entry points. The scheduling logic ensures that not more than 10 threads are executing requests from the three entry points combined.

You can define a `max-threads-constraint` in terms of the availability of the resource that requests depend upon, such as a connection pool.

A `max-threads-constraint` might, but does not necessarily, prevent a request class from taking its fair share of threads or meeting its response time goal. Once the constraint is reached the server does not schedule requests of this type until the number of concurrent executions falls below the limit. The server then schedules work based on the fair share or response time goal.

- `min-threads-constraint`—Guarantees the number of threads the server will allocate to affected requests to avoid deadlocks. The default is zero. A `min-threads-constraint` value of one is useful, for example, for a replication update request, which is called synchronously from a peer.

A `min-threads-constraint` might not necessarily increase a fair share. This type of constraint has an effect primarily when the server instance is close to a deadlock condition. In that case, the constraint will cause WebLogic Server to schedule a request even if requests in the service class have gotten more than its fair share recently.

- `capacity`—Causes the server to reject requests only when it has reached its capacity. The default is -1. Note that the capacity includes all requests, queued or executing, from the constrained work set. Work is rejected either when an individual capacity threshold is exceeded or if the global capacity is exceeded. This constraint is independent of the global queue threshold.

Note that the `capacity` constraint is not enforced if the request is made by a user belonging to the WebLogic Server Administrators group.

2.2.3 Stuck Thread Handling

In response to stuck threads, you can define a Stuck Thread Work Manager component that can shut down the Work Manager, move the application into admin mode, or mark the server instance as failed.

For example, the Work Manager defined in [Example 2-4](#) shuts down the Work Manager when two threads are stuck for longer than 30 seconds.

Example 2-4 Stuck-Thread Work Manager

```
<work-manager>
  <name>stuckthread_workmanager</name>
  <work-manager-shutdown-trigger>
    <max-stuck-thread-time>30</max-stuck-thread-time>
    <stuck-thread-count>2</stuck-thread-count>
  </work-manager-shutdown-trigger>
</work-manager>
```

2.2.4 Self-Tuning Thread Pool

WebLogic Server maintains three groups of threads for the self-tuning thread pool:

- **Running threads:** threads that are currently executing work requests submitted to Work Managers

- Idle threads: threads that are idly waiting for a work request

Idle threads include threads that have completed their previous work requests and are waiting for new requests, as well as threads that are created by the self-tuning thread pool based on usage statistics in order to anticipate future workload.

- Standby threads: threads that are not currently processing or waiting for work requests

Standby threads do not count toward the self-tuning thread pool thread count. When the self-tuning thread pool decides to decrease the thread count based on usage statistics, threads are moved from the group of idle threads into the group of standby threads. Conversely, when the self-tuning thread pool decides to increase the thread count, it first tries to find threads in the standby thread group to move to the idle thread group. The self-tuning thread pool only creates new threads if there are not enough threads in the standby group.

Threads are shut down when the number of standby threads reaches an internal maximum limit of 256. Ideally, a number of standby threads are ready if WebLogic Server needs to increase the self-tuning thread pool count occurs so that the WebLogic Server instance can avoid creating new threads at a time when workload is high. Standby threads can also be created and used to satisfy minimum threads constraints. For more information, see [Constraints](#).

By default, the self-tuning thread pool size limit is 400. This limit includes all running and idle threads, but does not include any standby threads. You can configure the limit using the `SelfTuningThreadPoolSizeMax` attribute in the `KernelMBean`. You may choose a higher size limit if your system can support additional workload even when the self-tuning thread pool has reached its upper thread count limit. Contrarily, you may choose to lower the limit if your system resources, such as CPU, become overloaded at a lower thread count. However, if lowering the `SelfTuningThreadPoolSizeMax` limit, note that if the value is set too low, the self-tuning thread pool may not be allowed to create enough threads to handle the system workload. This could result in a backlog of pending work requests on some Work Managers.

Note: Minimum threads constraints can affect the number of threads that are executing work requests for Work Managers, especially if the WebLogic Server instance is under heavy load.

The self-tuning thread pool does not consider the `SelfTuningThreadPoolSizeMax` attribute when creating a new standby thread to process incoming work requests for a Work Manager to satisfy its allocated minimum threads constraint. This is due to the importance of allocating threads for processing work requests for Work Managers with minimum threads constraints, which are designed to be used to avoid server-to-server deadlocks.

As a result, the maximum possible number of threads maintained by the self-tuning thread pool is the sum of the configured `SelfTuningThreadPoolSizeMax` attribute value and the sum of the values for all minimum threads constraints configured in the WebLogic Server instance, assuming a worst-case scenario where the configured number of threads are allocated to all configured minimum threads constraints.

2.3 Work Manager Scope

Essentially, there are three types of Work Managers, each one characterized by its scope and how it is defined and used.

- [Section 2.3.1, "The Default Work Manager"](#)
- [Section 2.3.2, "Global Work Managers"](#)
- [Section 2.3.3, "Application-scoped Work Managers"](#)

2.3.1 The Default Work Manager

To handle thread management and perform self-tuning, WebLogic Server implements a default Work Manager. This Work Manager is used by an application when no other Work Managers are specified in the application's deployment descriptors.

In many situations, the default Work Manager may be sufficient for most application requirements. WebLogic Server thread-handling algorithms assign each application its own fair share by default. Applications are given equal priority for threads and are prevented from monopolizing them.

2.3.1.1 Overriding the Default Work Manager

You can override the behavior of the default Work Manager by creating and configuring a global Work Manager called `default`. This allows you to control the default thread-handling behavior of WebLogic Server.

Note: When you override the default Work Manager, all instances are overridden.

2.3.1.2 When to Use Work Managers

Use the following are guidelines to help you determine when you might want to use Work Managers to customize thread management:

- The default fair share (50) is not sufficient.
This usually occurs in situations where one application needs to be given a higher priority over another.
- A response time goal is required.
- A minimum thread constraint needs to be specified to avoid server deadlock

2.3.2 Global Work Managers

You can create global Work Managers that are available to all applications and modules deployed on a server, in the WebLogic Server Administration Console and in `config.xml`.

An application uses a globally-defined Work Manager as a template. Each application creates its own instance which handles the work associated with that application and separates that work from other applications. This separation is used to handle traffic directed to two applications which are using the same dispatch policy. Handling each application's work separately, allows an application to be shut down without affecting the thread management of another application. Although each application implements its own Work Manager instance, the underlying components are shared.

2.3.3 Application-scoped Work Managers

In addition to globally-scoped Work Managers, you can also create Work Managers that are available only to a specific application or module. You can define

application-scoped Work Managers in the WebLogic Server Administration Console and in the following descriptors:

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

If you do not explicitly assign a Work Manager to an application, it uses the default Work Manager.

A method is assigned to a Work Manager, using the `<dispatch-policy>` element in the deployment descriptor. The `<dispatch-policy>` can also identify a custom execute queue, for backward compatibility. For an example, see [Example 2-2](#).

2.4 Using Work Managers, Request Classes, and Constraints

Work Managers, Request Classes, and Constraints require the following:

- A definition. You may define Work Managers, Request Classes, or Constraints globally in the domain's configuration using the WebLogic Server Administration Console, (see **Environments > Work Managers** in the WebLogic Server Administration Console) or you may define them in one of the deployment descriptors listed previously. In either case, you assign a name to each.
- A mapping. In your deployment descriptors you reference one of the Work Managers, Request Classes, or Constraints by its name.

2.4.1 Dispatch Policy for EJB

`weblogic-ejb-jar.xml`—The value of the existing `dispatch-policy` tag under `weblogic-enterprise-bean` can be a named `dispatch-policy`. For backwards compatibility, it can also name an `ExecuteQueue`. In addition, Oracle allows `dispatch-policy`, `max-threads`, and `min-threads`, to specify named (or unnamed with a numeric value for constraints) policy and constraints for a list of methods, analogously to the present `isolation-level` tag.

2.4.2 Dispatch Policy for Web Applications

`weblogic.xml`—Also supports mappings analogous to the `filter-mapping` of the `web.xml`, where named `dispatch-policy`, `max-threads`, or `min-threads` are mapped for `url-patterns` or `servlet names`.

2.5 Deployment Descriptor Examples

This section contains examples for defining Work Managers in various types of deployment descriptors.

For additional reference, see the schema for these deployment descriptors:

- `weblogic-ejb-jar.xml` schema:
<http://xmlns.oracle.com/weblogic/weblogic-ejb-jar/1.5/weblogic-ejb-jar.xsd>
- `weblogic-application.xml` schema:
<http://xmlns.oracle.com/weblogic/weblogic-application/1.6/weblogic-application.xsd>

- `weblogic.xml` schema: See "weblogic.xml Deployment Descriptor Elements" in *Developing Web Applications, Servlets, and JSPs for Oracle WebLogic Server*.

Example 2-5 `weblogic-ejb-jar.xml` With Work Manager Entries

```
<weblogic-ejb-jar xmlns="http://xmlns.oracle.com/weblogic/weblogic-ejb-jar"
  xmlns:j2ee="http://java.sun.com/xml/ns/j2ee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://xmlns.oracle.com/weblogic/weblogic-ejb-jar
  http://xmlns.oracle.com/weblogic/weblogic-ejb-jar/1.5/weblogic-ejb-jar.xsd">

<weblogic-enterprise-bean>
  <ejb-name>WorkEJB</ejb-name>
  <jndi-name>core_work_ejb_workbean_WorkEJB</jndi-name>
  <dispatch-policy>weblogic.kernel.System</dispatch-policy>
</weblogic-enterprise-bean>

<weblogic-enterprise-bean>
  <ejb-name>NonSystemWorkEJB</ejb-name>
  <jndi-name>core_work_ejb_workbean_NonSystemWorkEJB</jndi-name>
  <dispatch-policy>workbean_workmanager</dispatch-policy>
</weblogic-enterprise-bean>

<weblogic-enterprise-bean>
  <ejb-name>MinThreadsWorkEJB</ejb-name>
  <jndi-name>core_work_ejb_workbean_MinThreadsWorkEJB</jndi-name>
  <dispatch-policy>MinThreadsCountFive</dispatch-policy>
</weblogic-enterprise-bean>

<work-manager>
  <name>workbean_workmanager</name>
</work-manager>

<work-manager>
  <name>stuckthread_workmanager</name>
  <work-manager-shutdown-trigger>
    <max-stuck-thread-time>30</max-stuck-thread-time>
    <stuck-thread-count>2</stuck-thread-count>
  </work-manager-shutdown-trigger>
</work-manager>

<work-manager>
  <name>minthreads_workmanager</name>
  <min-threads-constraint>
    <name>MinThreadsCountFive</name>
    <count>5</count>
  </min-threads-constraint>
</work-manager>

<work-manager>
  <name>lowpriority_workmanager</name>
  <fair-share-request-class>
    <name>low_priority</name>
    <fair-share>10</fair-share>
  </fair-share-request-class>
</work-manager>

<work-manager>
  <name>highpriority_workmanager</name>
  <fair-share-request-class>
```

```

        <name>high_priority</name>
        <fair-share>100</fair-share>
    </fair-share-request-class>
</work-manager>

<work-manager>
<name>veryhighpriority_workmanager</name>
    <fair-share-request-class>
        <name>veryhigh_priority</name>
        <fair-share>1000</fair-share>
    </fair-share-request-class>
</work-manager>

```

The EJBs in [Example 2-6](#) are configured to get as many threads as there are instances of a resource they depend upon—a connection pool, and an application-scoped connection pool.

Example 2-6 *weblogic-ejb-jar.xml with Connection Pool Based Max Thread Constraint*

```

<weblogic-ejb-jar xmlns="http://xmlns.oracle.com/weblogic/weblogic-ejb-jar"
  xmlns:j2ee="http://java.sun.com/xml/ns/j2ee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://xmlns.oracle.com/weblogic/weblogic-ejb-jar
  http://xmlns.oracle.com/weblogic/weblogic-ejb-jar/1.5/weblogic-ejb-jar.xsd">

  <weblogic-enterprise-bean>
    <ejb-name>ResourceConstraintEJB</ejb-name>
    <jndi-name>core_work_ejb_resource_ResourceConstraintEJB</jndi-name>
    <dispatch-policy>test_resource</dispatch-policy>
  </weblogic-enterprise-bean>

  <weblogic-enterprise-bean>
    <ejb-name>AppScopedResourceConstraintEJB</ejb-name>
    <jndi-name>core_work_ejb_resource_AppScopedResourceConstraintEJB
    </jndi-name>
    <dispatch-policy>test_appscoped_resource</dispatch-policy>
  </weblogic-enterprise-bean>

  <work-manager>
    <name>test_resource</name>
    <max-threads-constraint>
      <name>pool_constraint</name>
      <pool-name>testPool</pool-name>
    </max-threads-constraint>
  </work-manager>

  <work-manager>
    <name>test_appscoped_resource</name>
    <max-threads-constraint>
      <name>appscoped_pool_constraint</name>
      <pool-name>AppScopedDataSource</pool-name>
    </max-threads-constraint>
  </work-manager>
</weblogic-ejb-jar>

```

Example 2-7 *weblogic-ejb-jar.xml with commonJ Work Managers*

For information using commonJ, see [Section 2.8, "Using CommonJ With WebLogic Server"](#) and the commonJ Javadocs.

Example 2–8 weblogic-application.xml

```

<weblogic-application
xmlns="http://xmlns.oracle.com/weblogic/weblogic-application"
  xmlns:j2ee="http://java.sun.com/xml/ns/j2ee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://xmlns.oracle.com/weblogic/weblogic-application
http://xmlns.oracle.com/weblogic/weblogic-application/1.6/weblogic-application.xsd
">

  <max-threads-constraint>
    <name>j2ee_maxthreads</name>
    <count>1</count>
  </max-threads-constraint>

  <min-threads-constraint>
    <name>j2ee_minthreads</name>
    count>1</count>
  </min-threads-constraint>

  <work-manager>
    <name>J2EEScopedWorkManager</name>
  </work-manager>
</weblogic-application>

```

The Web application in [Example 2–9](#) is deployed as part of the Enterprise application defined in [Example 2–8](#). This Web application's descriptor defines two Work Managers. Both Work Managers point to the same max threads constraint, `j2ee_maxthreads`, which is defined in the application's `weblogic-application.xml` file. Each Work Manager specifies a different response time request class.

Example 2–9 Web Application Descriptor

```

<weblogic xmlns="http://xmlns.oracle.com/weblogic"
  xmlns:j2ee="http://java.sun.com/xml/ns/j2ee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://xmlns.oracle.com/weblogic
http://xmlns.oracle.com/weblogic/1.0/weblogic.xsd">

  <work-manager>
    <name>fast_response_time</name>
    <response-time-request-class>
      <name>fast_response_time</name>
      <goal-ms>2000</goal-ms>
    </response-time-request-class>
    <max-threads-constraint-name>j2ee_maxthreads</max-threads-constraint-name>
  </work-manager>

  <work-manager>
    <name>slow_response_time</name>
    <max-threads-constraint-name>j2ee_maxthreads</max-threads-constraint-name>
    <response-time-request-class>
      <name>slow_response_time</name>
      <goal-ms>5000</goal-ms>
    </response-time-request-class>
  </work-manager>

</weblogic>

```


The descriptor in [Example 2–10](#) defines a Work Manager using the `context-request-class`.

Example 2–10 Web Application Descriptor

```
<?xml version="1.0" encoding="UTF-8"?>
<weblogic-web-app xmlns="http://xmlns.oracle.com/weblogic/weblogic-web-app"
xmlns:j2ee="http://java.sun.com/xml/ns/j2ee"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://xmlns.oracle.com/weblogic/weblogic-web-app
http://xmlns.oracle.com/weblogic/weblogic-web-app/1.7/weblogic-web-app.xsd">
  <work-manager>
    <name>foo-servlet-1</name>
    <request-class-name>test-fairshare2</request-class-name>
    <max-threads-constraint>
      <name>foo-mtc</name>
      <pool-name>oraclePool</pool-name>
    </max-threads-constraint>
  </work-manager>

  <work-manager>
    <name>foo-servlet</name>
    <context-request-class>
      <name>test-context</name>
      <context-case>
        <user-name>anonymous</user-name>
        <request-class-name>test-fairshare1</request-class-name>
      </context-case>

      <context-case>
        <group-name>everyone</group-name>
      </context-request-class>
    </work-manager>
</weblogic-web-app>
```

2.6 Work Managers and Execute Queues

This section describes how to enable backward compatibility with Execute Queues and how to migrate applications from using Execute Queues to Work Managers.

2.6.1 Enabling Execute Queues

WebLogic Server, Version 8.1, implemented Execute Queues to handle thread management in which you created thread-pools to determine how workload was handled. WebLogic Server still provides Execute Queues for backward compatibility, primarily to facilitate application migration. However, when developing new applications, you should use Work Managers to perform thread management more efficiently.

You can enable Execute Queues in the following ways:

- Using the command line option `-Dweblogic.Use81StyleExecuteQueues=true`
- Setting the `Use81StyleExecuteQueues` property via the Kernel MBean in `config.xml`.

Enabling Execute Queues disables all Work Manager configuration and thread self tuning. Execute Queues behave exactly as they did in WebLogic Server 8.1.

When enabled, Work Managers are converted to Execute Queues based on the following rules:

- If the Work Manager implements a minimum or maximum threads constraint, then an Execute Queue is created with the same name as the Work Manager. The thread count of the Execute Queue is based on the value defined in the constraint.
- If the Work Manager does not implement any constraints, the global default Execute Queue is used.

2.6.2 Migrating from Execute Queues to Work Managers

When an application is migrated from WebLogic Server 8.1, any Execute Queues defined in the server configuration before migration will still be present. WebLogic Server does not automatically convert the Execute Queues to Work Managers.

When an 8.1 application implementing Execute Queues is deployed on WebLogic Server 9.x, the Execute Queues are created to handle thread management for requests. However, only those requests whose dispatch-policy maps to an Execute Queue will take advantage of this feature.

2.7 Accessing Work Managers Using MBeans

Work Managers can be accessed using the `WorkManagerMBean` configuration MBean. For more information, see `WorkManagerMBean`.

`WorkManagerMBean` is accessed in the runtime tree or configuration tree depending on how the Work Manager is accessed by an application.

- If the Work Manager is defined at the module level, the `WorkManagerRuntimeMBean` is available through the corresponding `ComponentRuntimeMBean`.
- If a Work Manager is defined at the application level, then `WorkManagerRuntime` is available through `ApplicationRuntime`.
- If a Work Manager is defined globally in `config.xml`, each application creates its own instance of the Work Manager. Each application has its own corresponding `WorkManagerRuntime` available at the application level.

2.8 Using CommonJ With WebLogic Server

WebLogic Server Work Managers provide server-level configuration that allows administrators a way to set dispatch-policies to their servlets and EJBs.

WebLogic Server also provides a programmatic way of handling work from within an application. This is provided via the CommonJ API. WebLogic Server implements the `commonj.work` and `commonj.timers` packages of the CommonJ specification.

For specific information on the WebLogic Server implementation of CommonJ, see the CommonJ Javadocs.

The WebLogic Server implementation of CommonJ enables an application to break a single request task into multiple work items, and assign those work items to execute concurrently using multiple Work Managers configured in WebLogic Server. Applications that do not need to execute concurrent work items can also use configured Work Managers by referencing or creating Work Managers in their deployment descriptors or, for Java EE Connectors, using the JCA API.

The following are some differences between the WebLogic Server implementation and the CommonJ specification:

- The RemoteWorkItem interface is an optional interface provided by the CommonJ specification and is not supported in WebLogic Server. WebLogic Server implements its own cluster load balancing and failover policies. Workload management is based on these policies.
- WebLogic CommonJ timers behave differently than `java.util.Timer`. When the execution is greater than twice the period, the WebLogic CommonJ timer will skip some periods to avoid falling further behind. The `java.util.Timer` does not do this.
- In a WebLogic Server environment, the `WorkListener.WorkRejected` method is called when a thread becomes stuck.

2.8.1 Accessing CommonJ Work Managers

Unlike WebLogic Server Work Managers, which can only be accessed from an application via dispatch policies, you can access CommonJ Work Managers directly from an application. The following code example demonstrates how to lookup a CommonJ Work Manager using JNDI:

```
InitialContext ic = new InitialContext();
commonj.work.WorkManager wm =
(commonj.work.WorkManager)ic.lookup("java:comp/env/wm/myWM");
```

For more information on using CommonJ Work Managers, see the CommonJ Javadocs.

2.8.2 Mapping CommonJ to WebLogic Server Work Managers

You can map an externally defined CommonJ Work Manager to a WebLogic Server Work Manager. For example, if you have a CommonJ Work Manager defined in a descriptor, `ejb-jar.xml`, for example, as:

```
<resource-ref>
  <res-ref-name>minthreads_workmanager</res-ref-name>
  <res-type>commonj.work.WorkManager</res-type>
  <res-auth>Container</res-auth>
  <res-sharing-scope>Shareable</res-sharing-scope>
</resource-ref>
```

You can link this to a WebLogic Server Work Manager by ensuring that the name element is identical in the WebLogic Server descriptor such as `weblogic-ejb-jar.xml`:

```
<work-manager>
  <name>minthreads_workmanager</name>
  <min-threads-constraint>
    <count>5</count>
  </min-threads-constraint>
</work-manager>
```

This procedure is similar for a resource-ref defined in `web.xml`. The WebLogic Server Work Manager can be defined in either a module descriptor (`weblogic-ejb-jar.xml` or `weblogic.xml`, for example) or in the application descriptor (`weblogic-application.xml`).

Avoiding and Managing Overload

This chapter describes the features that WebLogic Server 12.1.3 has for detecting, avoiding, and recovering from overload conditions. WebLogic Server's overload protection features help prevent the negative consequences that can result from continuing to accept requests when the system capacity is reached. These consequences are degraded application performance and stability.

This chapter includes the following sections:

- [Section 3.1, "Configuring WebLogic Server to Avoid Overload Conditions"](#)
- [Section 3.2, "WebLogic Server Self-Monitoring"](#)
- [Section 3.3, "WebLogic Server Exit Codes"](#)

3.1 Configuring WebLogic Server to Avoid Overload Conditions

When system capacity is reached, if an application server continues to accept requests, application performance and stability can deteriorate. The following sections demonstrate how you can configure WebLogic Server to minimize the negative results of system overload.

3.1.1 Limiting Requests in the Thread Pool

In WebLogic Server, all requests, whether related to system administration or application activity—are processed by a single thread pool. An administrator can throttle the thread pool by defining a maximum queue length. Beyond the configured value, WebLogic Server will refuse requests, except for requests on administration channels.

Note: Administration channels allow access only to administrators. The limit you set on the execute length does not effect administration channel requests, to ensure that reaching the maximum thread pool length does not prevent administrator access to the system. To limit the number of administration requests allowed in the thread pool, you can configure an administration channel, and set the `MaxConnectedClients` attribute for the channel.

When the maximum number of enqueued requests is reached, WebLogic Server immediately starts rejecting:

- Web application requests.

- Non-transactional RMI requests with a low fair share, beginning with those with the lowest fair share.
If the overload condition continues to persist, higher priority requests will start getting rejected, with the exception of JMS and transaction-related requests, for which overload management is provided by the JMS and the transaction manager.

Throttle the thread pool by setting the `Shared Capacity For Work Managers` field in the WebLogic Server Administration Console (see **Environments > Servers > server_name > Configuration > Overload**). The default value of this field is 65536.

3.1.1.1 Work Managers and Thread Pool Throttling

An administrator can configure Work Managers to manage the thread pool at a more granular level, for sets of requests that have similar performance, availability, or reliability requirements. A Work Manager can specify the maximum requests of a particular request class that can be queued. The maximum requests defined in a Work Manager works with the global thread pool value. The limit that is reached first is honored.

See [Chapter 2, "Using Work Managers to Optimize Scheduled Work."](#)

3.1.2 Limiting HTTP Sessions

An administrator can limit the number of active HTTP sessions based on detection of a low memory condition. This is useful in avoiding out of memory exceptions.

WebLogic Server refuses requests that create new HTTP sessions after the configured threshold has been reached. In a WebLogic Server cluster, the proxy plug-in redirects a refused request to another Managed Server in the cluster. A non-clustered server instance can redirect requests to alternative server instance.

The Servlet container takes one of the following actions when maximum number of sessions is reached:

- If the server instance is in a cluster, the servlet container throws a `SessionCreationException`. Your application code should handle this run-time exception and send a relevant response.

To implement overload protection, you should handle this exception and send a 503 response explicitly. This response can then be handled by the proxy or load balancer.

You set a limit for the number of simultaneous HTTP sessions in the deployment descriptor for the Web application. For example, the following element sets a limit of 12 sessions:

```
<session-descriptor>  
  <max-in-memory-sessions>12</max-in-memory-sessions>  
</session-descriptor>
```

3.1.3 Exit on Out of Memory Exceptions

Administrators can configure WebLogic Server to exit upon an out of memory exception. This feature allows you to minimize the impact of the out of memory condition—automatic shutdown helps avoid application instability, and you can configure Node Manager or another high availability (HA) tool to automatically restart WebLogic Server, minimizing down-time.

You can configure this using the WebLogic Server Administration Console, or by editing the following elements in config.xml:

```
<overload-protection>
  <panic-action>system-exit</panic-action>
</overload-protection>
```

For more information, see the attributes of the OverloadProtectionMBean.

3.1.4 Stuck Thread Handling

WebLogic Server checks for stuck threads periodically. If all application threads are stuck, a server instance marks itself failed, if configured to do so, exits. You can configure Node Manager or a third-party high-availability solution to restart the server instance for automatic failure recovery.

You can configure these actions to occur when not all threads are stuck, but the number of stuck threads have exceeded a configured threshold:

- Shut down the Work Manager if it has stuck threads. A Work Manager that is shut down will refuse new work and reject existing work in the queue by sending a rejection message. In a cluster, clustered clients will fail over to another cluster member.
- Shut down the application if there are stuck threads in the application. The application is shutdown by bringing it into admin mode. All Work Managers belonging to the application are shut down, and behave as described above. Once the stuck thread condition is cleared, the application automatically returns to running mode.
- Mark the server instance as failed and shut it down if there are stuck threads in the server. In a cluster, clustered clients that are connected or attempting to connect will fail over to another cluster member.

For more information, see the attributes of the OverloadProtectionMBean.

3.2 WebLogic Server Self-Monitoring

The following sections describe WebLogic Server features that aid in determining and reporting overload conditions.

3.2.1 Overloaded Health State

WebLogic Server has a health state—`OVERLOADED`—which is returned by the `ServerRuntimeMBean.getHealthState()` when a server instance whose life cycle state is `RUNNING` becomes overloaded. This condition occurs as a result of low memory.

Upon entering the `OVERLOADED` state, server instances start rejecting requests from the Work Manager queue (if a Work Manager is configured), HTTP requests return a 503 Error (Service Unavailable), and RMI requests fail over to another server if clustered, otherwise, a remote exception is returned to the client.

The server instances health state returns to `OK` after the overload condition passes. An administrator can suspend or shut down an `OVERLOADED` server instance.

3.3 WebLogic Server Exit Codes

When WebLogic Server exits it returns an exit code. The exit codes can be used by shell scripts or HA agents to decide whether a server restart is necessary. See "WebLogic

Server Exit Codes and Restarting After Failure" in *Administering Server Startup and Shutdown for Oracle WebLogic Server*.

Configuring Network Resources

This chapter describes how to configure network resources, including network channels and domain-wide administration ports in WebLogic Server 12.1.3. These configurable network resources help you effectively use the network features of the machines that host your applications and manage quality of service.

This chapter includes the following sections:

- [Section 4.1, "Overview of Network Configuration"](#)
- [Section 4.2, "Understanding Network Channels"](#)
- [Section 4.3, "Configuring a Channel"](#)
- [Section 4.4, "Assigning a Custom Channel to an EJB"](#)
- [Section 4.5, "Using IPv6 with IPv4"](#)

4.1 Overview of Network Configuration

For many development environments, configuring WebLogic Server network resources is simply a matter of identifying a Managed Server listen address and listen port. However, in most production environments, administrators must balance finite network resources against the demands placed upon the network. The task of keeping applications available and responsive can be complicated by specific application requirements, security considerations, and maintenance tasks, both planned and unplanned.

WebLogic Server lets you control the network traffic associated with your applications in a variety of ways, and configure your environment to meet the varied requirements of your applications and end users. You can:

- Designate the Network Interface Cards (NICs) and ports used by Managed Servers for different types of network traffic.
- Support multiple protocols and security requirements.
- Specify connection and message time-out periods.
- Impose message size limits.

You specify these and other connection characteristics by defining a network channel—the primary configurable WebLogic Server resource for managing network connections. You configure a network channel in the WebLogic Server Administration Console (**Servers > Protocols > Channels**) or by using the `NetworkAccessPointMBean`.

4.2 Understanding Network Channels

The sections that follow describe network channels and the standard channels that WebLogic Server pre-configures, and discusses common applications for channels.

4.2.1 What Is a Channel?

A network channel is a configurable resource that defines the attributes of a network connection to WebLogic Server. For instance, a network channel can define:

- The protocol the connection supports.
- The listen address.
- The listen ports for secure and non-secure communication.
- Connection properties such as the login time-out value and maximum message sizes.
- Whether or not the connection supports tunneling.
- Whether the connection can be used to communicate with other WebLogic Server instances in the domain, or used only for communication with clients.

4.2.1.1 Rules for Configuring Channels

Follow these guidelines when configuring a channel.

- You can assign a particular channel to only one server instance.
- You can assign multiple channels to a server instance.
- Each channel assigned to a particular server instance must have a unique combination of listen address, listen port, and protocol.
- You can configure a custom identity keystore, and other channel-specific SSL attributes, that are separate from and that override the default keystore and SSL configuration settings for the Managed Server instance or the domain.
- If you assign non-SSL and SSL channels to the same server instance, make sure that they do not use the same port number.

4.2.1.2 Custom Channels Can Inherit Default Channel Attributes

If you do not assign a channel to a server instance, it uses the WebLogic Server default channel, which is automatically configured by WebLogic Server, based on the attributes in `ServerMBean` or `SSLMBean`; the operating system determines the network interface. The default channel is described in [Section 4.2.3.1, "The Default Network Channel"](#).

`ServerMBean` and `SSLMBean` represent a server instance and its SSL configuration. When you configure a server instance listen address, listen port, and SSL listen port, using the **Server > Configuration > General** page, those values are stored in the `ServerMBean` and `SSLMBean` for the server instance.

If you do not specify a particular connection attribute in a custom channel definition, the channel inherits the value specified for the attribute in `ServerMBean`. For example, if you create a channel, and do not define its listen address, the channel uses the listen address defined in `ServerMBean`. Similarly, if a Managed Server cannot bind to the listen address or listen port configured in a channel, the Managed Server uses the defaults from `ServerMBean` or `SSLMBean`.

4.2.2 Why Use Network Channels?

You use network channels to manage quality of service, meet varying connection requirements, and improve utilization of your systems and network resources. For example, network channels allow you to:

- **Segregate different types of network traffic**—You can configure whether or not a channel supports outgoing connections. By assigning two channels to a server instance—one that supports outgoing connections and one that does not—you can independently configure network traffic for client connections and server connections, and physically separate client and server network traffic on different listen addresses or listen ports.

You cannot create an outbound only network channel; there always has to be a corresponding inbound interface, port, and protocol associated with the channel. However, you can avoid directing your traffic to it or use a firewall to block it. Also remember that a custom channel is protocol specific, so you will need a network channel defined per protocol (HTTP, HTTPS, t3, t3s, and such). See, also `NetworkAccessPointMBean.OutboundEnabled`.

You can also segregate instance administration and application traffic by configuring a domain-wide administration port or administration channel. For more information, see [Section 4.2.3.2, "Administration Port and Administrative Channel"](#).

- **Support varied application or user requirements on the same Managed Server**—You can configure multiple channels on a Managed Server to support different protocols, or to tailor properties for secure versus non-secure traffic.

By configuring a network channel to use a custom identity keystore, you can assert an identity on that channel that is different from the identity configured for the Managed Server or domain.

- **Segregate internal application network traffic**—You can assign a specific channel to a an EJB.

If you use a network channel with a server instance on a multihomed machine, you must enter a valid listen address either in `ServerMBean` or in the channel. If the channel and `ServerMBean` listen address are blank or specify the localhost address (IP address 0.0.0.0 or 127.*.*.*), the server binds the network channel listen port and SSL listen ports to all available IP addresses on the multihomed machine. See [Section 4.2.3.1, "The Default Network Channel"](#) for information on setting the listen address in `ServerMBean`.

4.2.2.1 Handling Channel Failures

When initiating a connection to a remote server, and multiple channels with the same required destination, protocol and quality of service exist, WebLogic Server will try each in turn until it successfully establishes a connection or runs out of channels to try.

4.2.2.2 Upgrading Quality of Service Levels for RMI

For RMI lookups only, WebLogic Server may upgrade the service level of an outgoing connection. For example, if a T3 connection is required to perform an RMI lookup, but an existing channel supports only T3S, the lookup is performed using the T3S channel.

This upgrade behavior does not apply to server requests that use URLs, since URLs embed the protocol itself. For example, the server cannot send a URL request beginning with `http://` over a channel that supports only `https://`.

4.2.3 Standard WebLogic Server Channels

WebLogic Server provides pre-configured channels that you do not have to explicitly define.

- Default channel—Every Managed Server has a default channel.
- Administrative channel—If you configure a domain-wide administration port, WebLogic Server configures an administrative channel for each Managed Server in the domain.

4.2.3.1 The Default Network Channel

Every WebLogic Server domain has a default channel that is generated automatically by WebLogic Server. The default channel is based on the listen address and listen port defined in the `ServerMBean` and `SSLMBean`. It provides a single listen address, one port for HTTP (non-secure) communication (7001 by default), and one port for HTTPS (secure) communication (7002 by default). You can configure the listen address and listen port using the **Configuration > General** page in the WebLogic Server Administration Console; the values you assign are stored in attributes of the `ServerMBean` and `SSLMBean`.

The default configuration may meet your needs if:

- You are installing in a test environment that has simple network requirements.
- Your server uses a single NIC, and the default port numbers provide enough flexibility for segmenting network traffic in your domain.

Using the default configuration ensures that third-party administration tools remain compatible with the new installation, because network configuration attributes remain stored in `ServerMBean` and `SSLMBean`.

Even if you define and use custom network channels for your domain, the default channel settings remain stored in `ServerMBean` and `SSLMBean`, and are used if necessary to provide connections to a server instance.

Note: Unless specified, WebLogic Server uses the non-secure default channel for cluster communication to send session information among cluster members. If you disable the non-secure channel, there is no other channel available *by default* for the non-secure communication of cluster session information. To address this, you can:

- Enable the `secureReplicationEnabled` attribute of the `ClusterMBean` so that the cluster uses a secure channel for communication. See [Section 4.3.2.3, "Configuring a Replication Channel."](#)
 - Create a custom channel for non-secure communication. See [Section 4.2.1.2, "Custom Channels Can Inherit Default Channel Attributes."](#)
-

4.2.3.2 Administration Port and Administrative Channel

You can separate administration traffic from application traffic in your domain by defining an optional administration port. When configured, the administration port is used by each Managed Server in the domain exclusively for communication with the domain Administration Server. If an administration port is enabled, WebLogic Server automatically generates an administrative channel for your domain, based on the port

settings upon server instance startup. The administrative channel provides a listen address and listen port to handle administration traffic.

4.2.3.2.1 Administration Port Capabilities An administration port enables you to:

- Start a server in standby state. This allows you to administer a Managed Server, while its other network connections are unavailable to accept client connections. For more information on the standby state, see "STANDBY State" in *Administering Server Startup and Shutdown for Oracle WebLogic Server*.
- Separate administration traffic from application traffic in your domain. In production environments, separating traffic ensures that critical administration operations (starting and stopping servers, changing a server's configuration, and deploying applications) do not compete with high-volume application traffic on the same network connection.
- Administer a deadlocked server instance using WLST. If you do not configure an administration port, administrative commands such as `threadDump` and `shutdown` will not work on deadlocked server instances.

4.2.3.2.2 Administration Port Restrictions The administration port accepts only secure, SSL traffic, and all connections via the port require authentication. Enabling the administration port imposes the following restrictions on your domain:

- The Administration Server and all Managed Servers in your domain must be configured with support for the SSL protocol. Managed Servers that do not support SSL cannot connect with the Administration Server during startup—you will have to disable the administration port in order to configure them.
- Because all server instances in the domain must enable or disable the administration port at the same time, you configure the administration port at the domain level. You can change an individual Managed Server administration port number, but you cannot enable or disable the administration port for an individual Managed Server. The ability to change the port number is useful if you have multiple server instances with the same listen address.
- After you enable the administration port, you must establish an SSL connection to the Administration Server in order to start any Managed Server in the domain. This applies whether you start Managed Servers manually, at the command line, or using Node Manager. For instructions to establish the SSL connection, see [Section 4.2.3.2.3, "Administration Port Requires SSL"](#).
- After enabling the administration port, all WebLogic Server Administration Console traffic *must* connect via the administration port.
- If multiple server instances run on the same computer in a domain that uses a domain-wide administration port, you must either:
 - Host the server instances on a multihomed machine and assign each server instance a unique listen address, or
 - Override the domain-wide port on all but one of the servers instances on the machine. Override the port using the Local Administration Port Override option in the Advanced Attributes section of the **Server > Connections > SSL Ports** page in the WebLogic Server Administration Console.

4.2.3.2.3 Administration Port Requires SSL The administration port requires SSL, which is enabled by default when you install WebLogic Server. If SSL has been disabled for any server instance in your domain, including the Administration Server and all Managed

Servers, re-enable it using the **Server > Configuration > General** page in the WebLogic Server Administration Console.

Ensure that each server instance in the domain has a configured default listen port or default SSL listen port. The default ports are those you assign on the **Server > Configuration > General** page in the WebLogic Server Administration Console. A default port is required in the event that the server cannot bind to its configured administration port. If an additional default port is available, the server will continue to boot and you can change the administration port to an acceptable value.

By default WebLogic Server is configured to use demonstration certificate files. To configure production security components, follow the steps in "Configuring SSL" in *Administering Security for Oracle WebLogic Server*.

4.2.3.2.4 Configure Administration Port Enable the administration port as described in "Enabling the Domain-Wide Administration Port" in *Oracle WebLogic Server Administration Console Online Help*.

After configuring the administration port, you must restart the Administration Server and all Managed Servers to use the new administration port.

4.2.3.2.5 Booting Managed Servers to Use Administration Port If you reboot Managed Servers at the command line or using a start script, specify the administration port in the port portion of the URL. The URL must specify the `https://` prefix, rather than `http://`, as shown below.

```
-Dweblogic.management.server=https://host:admin_port
```

Note: If you use Node Manager for restarting the Managed Servers, it is not necessary to modify startup settings or arguments for the Managed Servers. Node Manager automatically obtains and uses the correct URL to start a Managed Server.

If the hostname in the URL is not identical to the hostname in the Administration Server's certificate, disable hostname verification in the command line or start script, as shown below:

```
-Dweblogic.security.SSL.ignoreHostnameVerification=true
```

4.2.3.2.6 Booting Managed Servers to Use Administrative Channels To allow a Managed Server to bind to an administrative channel during reboot, use the following command-line option:

```
-Dweblogic.admin.ListenAddress=<addr>
```

This allows the Managed Server to startup using an administrative channel. After the initial bootstrap connection, a standard administrative channel is used.

Note: This option is useful to ensure that the appropriate NIC semantics are used before `config.xml` is downloaded.

4.2.3.2.7 Custom Administrative Channels If the standard WebLogic Server administrative channel does not satisfy your requirements, you can configure a custom

channel for administrative traffic. For example, a custom administrative channel allows you to segregate administrative traffic on a separate NIC.

To configure a custom channel for administrative traffic, configure the channel as described in [Section 4.3, "Configuring a Channel"](#), and select "admin" as the channel protocol. Note the configuration and usage guidelines described in:

- [Section 4.2.3.2.3, "Administration Port Requires SSL"](#)
- [Section 4.2.3.2.5, "Booting Managed Servers to Use Administration Port"](#)

4.2.4 Using Internal Channels

Previous version of WebLogic Server allowed you to configure multiple channels for external traffic, but required you to use the default channel for internal traffic between server instances. WebLogic Server now allows you to create network channels to handle administration traffic or communications between clusters. This can be useful in the following situations:

- Internal administration traffic needs to occur over a secure connection, separate from other network traffic.
- Other types of network traffic, for example replication data, need to occur over a separate network connection.
- Certain types of network traffic need to be monitored.

4.2.4.1 Channel Selection

All internal traffic is handled via a network channel. If you have not created a custom network channel to handle administrative or clustered traffic, WebLogic Server automatically selects a default channel based on the protocol required for the connection. For more information on default channels, see [Section 4.2.3.1, "The Default Network Channel"](#).

4.2.4.2 Internal Channels Within a Cluster

Within a cluster, internal channels can be created to handle to the following types of server-to-server connections:

- Multicast traffic
- Replication traffic
- Administration traffic

For more information on configuring channels within a cluster, see [Section 4.3.2, "Configuring Network Channels For a Cluster"](#).

4.3 Configuring a Channel

You can configure a network channel using **Servers > Protocols > Channels** page in the WebLogic Server Administration Console or using the `NetworkAccessPointMBean`.

To configure a channel for clustered Managed Servers see, [Section 4.3.2, "Configuring Network Channels For a Cluster"](#).

For a summary of key facts about network channels, and guidelines related to their configuration, see [Section 4.3.1, "Guidelines for Configuring Channels"](#).

4.3.1 Guidelines for Configuring Channels

Follow these guidelines when configuring a channel.

4.3.1.1 Channels and Server Instances

- Each channel you configure for a particular server instance must have a unique combination of listen address, listen port, and protocol.
- A channel can be assigned to a single server instance.
- You can assign multiple channels to a server instance.
- If you assign non-SSL and SSL channels to the same server instance, make sure that they do not use the same combination of address and port number.

4.3.1.2 Dynamic Channel Configuration

- In WebLogic Server, you can configure a network channel without restarting the server. Additionally, you can start and stop dynamically configured channels while the server is running. However, when you shutdown a channel while the server is running, the server does not attempt to gracefully terminate any work in progress.

4.3.1.3 Channels and Identity

- By default, when you configure a network channel, the channel uses the SSL configuration that is set for the server instance. This means that the channel uses the same identity and trust that is established for the server. The server might use a custom identity that is specific to that server, or it might be a single domain-wide identity, depending on how the server instance and domain are configured.
- You can configure a network channel to use a custom identity keystore, and other SSL attributes, that are specific to that channel. This allows you to use an identity on that channel that is different from the one configured for the server. Using this capability, you can configure a server that can switch to a different identity when communicating with a particular client.

For more information, see "Configuring an Identity Keystore Specific to a Network Channel" in *Administering Security for Oracle WebLogic Server*.

4.3.1.4 Channels and Protocols

- Some protocols do not support particular features of channels. In particular the COM protocol does not support SSL or tunneling.
- You must define a separate channel for each protocol you wish the server instance to support, with the exception of HTTP.

HTTP is enabled by default when you create a channel, because RMI protocols typically require HTTP support for downloading stubs and classes. You can disable HTTP support on the Advanced Options portion of the **Servers > Protocols > Channels** page in the WebLogic Server Administration Console.

4.3.1.5 Reserved Names

- WebLogic Server uses the internal channel names `.WLDefaultChannel` and `.WLDefaultAdminChannel` and reserves the `.WL` prefix for channel names. Do not begin the name of a custom channel with the string `.WL`.

4.3.1.6 Channels, Proxy Servers, and Firewalls

If your configuration includes a firewall between a proxy Web server and a cluster (as described in "Firewall Between Proxy Layer and Cluster" in *Administering Clusters for Oracle WebLogic Server*), and the clustered servers are configured with two custom channels for segregating HTTPS and HTTP traffic, those channels must share the same listen address. Furthermore, if both HTTP and HTTPS traffic needs to be supported, there must be a custom channel for each—it is not possible to use the default configuration for one or the other.

If either of those channels has a `PublicAddress` defined, as is likely given the existence of the firewall, both channels must define `PublicAddress`, and they both must define the same `PublicAddress`.

4.3.2 Configuring Network Channels For a Cluster

To configure a channel for clustered Managed Servers, note the information in [Section 4.3.1, "Guidelines for Configuring Channels"](#), and follow the guidelines described in the following sections.

4.3.2.1 Create the Cluster

If you have not already configured a cluster you can:

- Use the Configuration Wizard to create a new, clustered domain, following the instructions in "Create a Clustered Domain" in *Administering Clusters for Oracle WebLogic Server*, or
- Use the WebLogic Server Administration Console to create a cluster in an existing domain, following the instructions "Create and configure clusters" in *Oracle WebLogic Server Administration Console Online Help*.

For information and guidelines about configuring a WebLogic Server cluster, see "Before You Start" in *Administering Clusters for Oracle WebLogic Server*.

4.3.2.2 Create and Assign the Network Channel

Use the instructions in "Configuring a Network Channel" in *Oracle WebLogic Server Administration Console Online Help* to create a new network channel for each Managed Server in the cluster. When creating the new channels:

- For each channel you want to use in the cluster, configure the channel identically, including its name, on each Managed Server in the cluster.
- Make sure that the listen port and SSL listen port you define for each Managed Server's channel are different than the Managed Server's default listen ports. If the custom channel specifies the same port as a Managed Server's default port, the custom channel and the Managed Server's default channel will each try to bind to the same port, and you will be unable to start the Managed Server.
- If a cluster address has been explicitly configured for the cluster, it will be appear in the Cluster Address field on the **Server > Protocols > Channels > Configuration** page.

If you are using dynamic cluster addressing, the Cluster Address field will be empty, and you do not need to supply a cluster address. For information about the cluster address, and how WebLogic Server can dynamically generate the cluster address, see "Cluster Address" in *Administering Clusters for Oracle WebLogic Server*.

Note: If you want to use dynamic cluster addressing, do not supply a cluster address on the **Server > Protocols > Channels > Configuration** page. If you supply a cluster address explicitly, that value will take precedence and WebLogic Server will not generate the cluster address dynamically.

4.3.2.3 Configuring a Replication Channel

A replication channel is a network channel that is designated to transfer replication information between clusters. If a replication channel is not explicitly defined, WebLogic Server uses a default network channel to communicate replication information.

When WebLogic Server uses a default network channel as the replication channel, it does not use SSL encryption by default. You must enable SSL encryption using the `secureReplicationEnabled` attribute of the `ClusterMBean`. You can also update this setting from the WebLogic Server Administration Console.

Enabling SSL encryption can have a direct impact on clustered application throughput as session replication is blocking by design. In other words, no response is sent to the client until replication is completed. You should consider this when deciding to enable SSL on the replication channel.

If a replication channel is explicitly defined, the channel's protocol is used to transmit replication traffic.

4.3.2.4 Increase Packet Size When Using Many Channels

Use of more than about twenty channels in a cluster can result in the formation of multicast header transmissions that exceed the default maximum packet size. The `MTUSize` attribute in the `Server` element of `config.xml` sets the maximum size for packets sent using the associated network card to 1500. Sending packets that exceed the value of `MTUSize` can result in a `java.lang.NegativeArraySizeException`. You can avoid exceptions that result from packet sizes in excess of `MTUSize` by increasing the value of `MTUSize` from its default value of 1500.

4.4 Assigning a Custom Channel to an EJB

You can assign a custom channel to an EJB. After you configure a custom channel, assign it to an EJB using the `network-access-point` element in `weblogic-ejb-jar.xml`. For more information, see "network-access-point" in *Developing Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

4.5 Using IPv6 with IPv4

WebLogic Server supports host machines configured to use either Internet Protocol (IP) versions 4 or 6 (IPv4 and IPv6).

If you have a domain that includes some machines that use IPv4 in network communications and others that use IPv6, and the Administration Server is hosted on a machine using IPv4, the status of the Managed Server instances hosted on the machines using IPv6 might be displayed as "unknown" in the WebLogic Server Administration Console.

To make the status of these Managed Server instances available in the WebLogic Server Administration Console, you must specify a listen address for them. If your server is running, you will have to restart it after specifying the listen address. For information on assigning the listen address for a Managed Server in an existing domain using the

WebLogic Server Administration Console, see "Configure listen addresses" in the *Oracle WebLogic Server Administration Console Online Help*.

You can also specify the listen address for your Managed Server when configuring it with the Configuration Wizard:

1. On the **Managed Servers** page, enter the physical IP address of each Managed Server in the **Listen Address** field.
2. Click **Next** to save your changes and continue configuring your domain.

Configuring Web Server Functionality

This chapter describes how to configure a Java EE Web application hosted on WebLogic Server 12.1.3 to function as a standard HTTP Web server hosting static content. Web applications also can host dynamic content such as JSPs and Servlets. See *Developing Web Applications, Servlets, and JSPs for Oracle WebLogic Server*.

This chapter includes the following sections:

- [Section 5.1, "Overview of Configuring Web Server Components"](#)
- [Section 5.2, "Configuring the Server"](#)
- [Section 5.3, "Web Applications"](#)
- [Section 5.4, "Configuring Virtual Hosting"](#)
- [Section 5.5, "How WebLogic Server Resolves HTTP Requests"](#)
- [Section 5.6, "Setting Up HTTP Access Logs"](#)
- [Section 5.7, "Preventing POST Denial-of-Service Attacks"](#)
- [Section 5.8, "Setting Up WebLogic Server for HTTP Tunneling"](#)
- [Section 5.9, "Using Native I/O for Serving Static Files \(Windows Only\)"](#)

5.1 Overview of Configuring Web Server Components

In addition to hosting dynamic Java-based distributed applications, WebLogic Server functions as a Web server that handles high-volume Web sites, serving static files such as HTML files and image files, as well as servlets and JavaServer Pages (JSP). WebLogic Server supports the HTTP 1.1 standard.

5.2 Configuring the Server

You can specify the port that each WebLogic Server listens on for HTTP requests. Although you can specify any valid port number, if you specify port 80, you can omit the port number from the HTTP request used to access resources over HTTP. For example, if you define port 80 as the listen port, you can use the form `http://hostname/myfile.html` instead of `http://hostname:portnumber/myfile.html`.

On UNIX systems, binding a process to a port lower than 1025 must be done from the account of a privileged user, usually root. Consequently, if you want WebLogic Server to listen on port 80, you must start WebLogic Server as a privileged user; yet it is undesirable from a security standpoint to allow long-running processes like WebLogic

Server to run with more privileges than necessary. WebLogic Server needs root privileges only until the port is bound.

WebLogic Server provides capabilities to switch its UNIX user ID (UID) and/or UNIX group ID (GID) after it binds to port 80. You can change the UID (or GID) either through the WebLogic Server Administration Console (see [Section 5.2.1, "Configuring the Listen Port"](#)) or by accessing `UnixMachineMBean` using WLST. Use `UnixMachineMBean.setPostBindUID()` to set the UID and enable the switch by setting `UnixMachineMBean.setPostBindUIDEnabled()` to `true`. Similarly, the GID can be changed using methods `UnixMachineMBean.setPostBindGID()` and `UnixMachineMBean.setPostBindGIDEnabled()`.

You can switch to the UNIX account "nobody," which is the least privileged user on most UNIX systems. If desired, you may create a UNIX user account expressly for running WebLogic Server. Make sure that files needed by WebLogic Server, such as log files and the WebLogic classes, are accessible by the non-privileged user. Once ownership of the WebLogic process has switched to the non-privileged user, WebLogic will have the same read, write, and execute permissions as the non-privileged user.

You define a separate listen port for non-SSL and secure (using SSL) requests. For additional information on configuring listen ports, see [Section 4.2, "Understanding Network Channels"](#).

5.2.1 Configuring the Listen Port

1. Use the WebLogic Server Administration Console to set the listen port to port 80. See "Configure Listen Ports".
2. If the machine hosting WebLogic Server is running Windows, skip to step 8.
3. Use the WebLogic Server Administration Console to create a new Unix Machine. See "Configure Machines".
4. Select the `Enable Post-Bind UID` field.
5. Enter the user name you want WebLogic Server to run as in the `Post-Bind UID` field.
6. Select the `Enable Post-Bind GID` fields.
7. Enter the group name you want WebLogic Server to run as in the `Post-Bind GID` field.
8. Click **Save**.
9. To activate these changes, in the Change Center of the WebLogic Server Administration Console, click **Activate Changes**.

5.3 Web Applications

HTTP and Web applications are deployed according to the Java EE 6 Servlet 3.0 and JSP 2.2 specifications, which describe *Web Applications* as a standard for grouping the components of a Web-based application. These components include JSP pages, HTTP servlets, and static resources such as HTML pages or image files. In addition, a Web application can access external resources such as EJBs and JSP tag libraries. Each server can host any number of Web applications. You typically use the name of the Web application as part of the URI you use to request resources from the Web application.

For more information, see *Developing Web Applications, Servlets, and JSPs for Oracle WebLogic Server*.

5.3.1 Web Applications and Clustering

Web applications can be deployed to a WebLogic Server cluster. When a user requests a resource from a Web application, the request is routed to one of the servers in the cluster that host the Web application. If an application uses a session object, then sessions must be replicated across the nodes of the cluster. Several methods of replicating sessions are provided.

For more information, see *Administering Clusters for Oracle WebLogic Server*.

5.4 Configuring Virtual Hosting

Virtual hosting allows you to define host names that servers or clusters respond to. When you use virtual hosting you use DNS to specify one or more host names that map to the IP address of a WebLogic Server instance or cluster, and you specify which Web applications are served by the virtual host. When used in a cluster, load balancing allows the most efficient use of your hardware, even if one of the DNS host names processes more requests than the others.

For example, you can specify that a Web application called `books` responds to requests for the virtual host name `www.books.com`, and that these requests are targeted to WebLogic Servers A,B, and C, while a Web application called `cars` responds to the virtual host name `www.autos.com` and these requests are targeted to WebLogic Servers D and E. You can configure a variety of combinations of virtual host, WebLogic Server instances, clusters, and Web applications, depending on your application and Web server requirements.

For each virtual host that you define you can also separately define HTTP parameters and HTTP access logs. The HTTP parameters and access logs set for a *virtual host* override those set for a *server*. You may specify any number of virtual hosts.

You activate virtual hosting by targeting the virtual host to a server or cluster of servers. Virtual hosting targeted to a cluster will be applied to all servers in the cluster.

5.4.1 Virtual Hosting and the Default Web Application

You can also designate a *default Web Application* for each virtual host. The default Web application for a virtual host responds to all requests that cannot be resolved to other Web applications deployed on the same server or cluster as the virtual host.

Unlike other Web applications, a default Web application does not use the Web application name (also called the *context path*) as part of the URI used to access resources in the default Web application.

For example, if you defined virtual host name `www.mystore.com` and targeted it to a server on which you deployed a Web application called `shopping`, you would access a JSP called `cart.jsp` from the `shopping` Web application with the following URI:

```
http://www.mystore.com/shopping/cart.jsp
```

If, however, you declared `shopping` as the default Web application for the virtual host `www.mystore.com`, you would access `cart.jsp` with the following URI:

```
http://www.mystore.com/cart.jsp
```

For more information, see [Section 5.5, "How WebLogic Server Resolves HTTP Requests"](#).

When using multiple Virtual Hosts with different default Web applications, you can not use single sign-on, as each Web application will overwrite the JSESSIONID cookies

set by the previous Web application. This will occur even if the CookieName, CookiePath, and CookieDomain are identical in each of the default Web applications.

5.4.2 Setting Up a Virtual Host

1. Use the Administration Console to define a virtual host. See "Virtual Host".
2. Add a line naming the virtual host to the `etc/hosts` file on your server to ensure that the virtual host name can be resolved.

5.5 How WebLogic Server Resolves HTTP Requests

When WebLogic Server receives an HTTP request, it resolves the request by parsing the various parts of the URL and using that information to determine which Web application and/or server should handle the request. [Table 5–1](#) demonstrates various combinations of requests for Web applications, virtual hosts, servlets, JSPs, and static files and the resulting response.

Note: If you package your Web application as part of an Enterprise application, you can provide an alternate name for a Web application that is used to resolve requests to the Web application. For more information, see *Developing Web Applications, Servlets, and JSPs for Oracle WebLogic Server*.

[Table 5–1](#) provides some sample URLs and the file that is served by WebLogic Server. The Index Directories Checked column refers to the Index Directories attribute that controls whether or not a directory listing is served if no file is specifically requested.

Table 5–1 Examples of How WebLogic Server Resolves URLs

URL	Index Directories Checked?	This file is served in response
<code>http://host:port/apples</code>	No	Welcome file* defined in the <code>apples</code> Web application.
<code>http://host:port/apples</code>	Yes	Directory listing of the top-level directory of the <code>apples</code> Web application.
<code>http://host:port/oranges/naval</code>	Does not matter	Servlet mapped with <code><url-pattern></code> of <code>/naval</code> in the <code>oranges</code> Web application. There are additional considerations for servlet mappings. For more information, see "Configuring Servlets" in <i>Developing Web Applications, Servlets, and JSPs for Oracle WebLogic Server</i> .
<code>http://host:port/naval</code>	Does not matter	Servlet mapped with <code><url-pattern></code> of <code>/naval</code> in the <code>oranges</code> Web application and <code>oranges</code> is defined as the default Web application. For more information, see "Configuring Servlets" in <i>Developing Web Applications, Servlets, and JSPs for Oracle WebLogic Server</i> .
<code>http://host:port/apples/pie.jsp</code>	Does not matter	<code>pie.jsp</code> , from the top-level directory of the <code>apples</code> Web application.
<code>http://host:port</code>	Yes	Directory listing of the top-level directory of the <i>default</i> Web application
<code>http://host:port</code>	No	Welcome file* from the <i>default</i> Web application.

Table 5–1 (Cont.) Examples of How WebLogic Server Resolves URLs

URL	Index Directories Checked?	This file is served in response
http://host:port/apples/myfile.html	Does not matter	myfile.html, from the top-level directory of the apples Web application.
http://host:port/myfile.html	Does not matter	myfile.html, from the top-level directory of the <i>default</i> Web application.
http://host:port/apples/images/red.gif	Does not matter	red.gif, from the images subdirectory of the top-level directory of the apples Web application.
http://host:port/myFile.html	Does not matter	Error 404
Where myfile.html does not exist in the apples Web application and a <i>default servlet</i> has not been defined.		
http://www.fruit.com/	No	Welcome file from the default Web application for a virtual host with a host name of www.fruit.com.
http://www.fruit.com/	Yes	Directory listing of the top-level directory of the default Web application for a virtual host with a host name of www.fruit.com.
http://www.fruit.com/oranges/myfile.html	Does not matter	myfile.html, from the oranges Web application that is targeted to a virtual host with host name www.fruit.com.

5.6 Setting Up HTTP Access Logs

WebLogic Server can keep a log of all HTTP transactions in a text file, in either *common log format* or *extended log format*. Common log format is the default. Extended log format allows you to customize the information that is recorded. You can set the attributes that define the behavior of HTTP access logs for each server instance or for each virtual host that you define.

To set up HTTP logging for a server or a virtual host, refer to the following topics in the *Oracle WebLogic Server Administration Console Online Help*:

- "Enabling and Configuring HTTP Access Logs"
- "Specifying HTTP Log File Settings for a Virtual Host"

5.6.1 Log Rotation

You can rotate the log file based on either the size of the file or after a specified amount of time has passed. When either criterion is met, the current access log file is closed and a new access log file is started. If you do not configure log rotation, the HTTP access log file grows indefinitely. You can configure the name of the access log file to include a time and date stamp that indicates when the file was rotated. If you do not configure a time stamp, each rotated file name includes a numeric portion that is incremented upon each rotation. Separate HTTP access logs are kept for each Virtual Host you have defined.

5.6.2 Common Log Format

The default format for logged HTTP information is the common log format (see <http://www.w3.org/Daemon/User/Config/Logging.html#common-logfile-format>). This standard format follows the pattern:

```
host RFC931 auth_user [day/month/year:hour:minute:second
    UTC_offset] "request" status bytes
```

where:**host**

Either the DNS name or the IP number of the remote client

RFC931

Any information returned by IDENTD for the remote client; WebLogic Server does not support user identification

auth_user

If the remote client user sent a userid for authentication, the user name; otherwise "-"

day/month/year:hour:minute:second UTC_offset

Day, calendar month, year and time of day (24-hour format) with the hours difference between local time and GMT, enclosed in square brackets

"request"

First line of the HTTP request submitted by the remote client enclosed in double quotes

status

HTTP status code returned by the server, if available; otherwise "-"

bytes

Number of bytes listed as the content-length in the HTTP header, not including the HTTP header, if known; otherwise "-"

5.6.3 Setting Up HTTP Access Logs by Using Extended Log Format

WebLogic Server also supports extended log file format, version 1.0, an emerging standard defined by the draft specification from the W3C at <http://www.w3.org/TR/WD-logfile.html>. The current definitive reference is on the W3C Technical Reports and Publications page at <http://www.w3.org/TR/>.

The extended log format allows you to specify the type and order of information recorded about each HTTP communication. To enable this format in the WebLogic Server Administration Console:

1. Navigate to the *server-name* > **Logging** > **HTTP** page.
2. Make sure that **HTTP access log file enabled** is checked.
3. Click **Advanced**.
4. In the field labeled **Format**, select **Extended**.

In the field labeled **Extended Logging Format Fields**, you can select one or more of the fields described in [Section 5.6.3.2, "Supported Field Identifiers"](#). If you want to add custom fields to an HTTP access log file, see [Section 5.6.3.3, "Creating Custom Field Identifiers"](#) for details.

You specify the information that should be recorded in the log file with directives, included in the actual log file itself. A directive begins on a new line and starts with a pound sign (#). If the log file does not exist, a new log file is created with default directives. However, if the log file already exists when the server starts, it must contain valid directives at the head of the file.

5.6.3.1 Creating the Fields Directive

The first line of your log file must contain a directive stating the version number of the log file format. You must also include a `Fields` directive near the beginning of the file:

```
#Version: 1.0
#Fields: xxxx xxxx xxxx ...
```

Where each `xxxx` describes the data fields to be recorded. Field types are specified as either simple identifiers, or may take a prefix-identifier format, as defined in the W3C specification. For example:

```
#Fields: date time cs-method cs-uri
```

This identifier instructs the server to record the date and time of the transaction, the request method that the client used, and the URI of the request for each HTTP access. Each field is separated by white space, and each record is written to a new line, appended to the log file.

Note: The `#Fields` directive must be followed by a new line in the log file, so that the first log message is not appended to the same line.

5.6.3.2 Supported Field Identifiers

The following identifiers are supported, and do not require a prefix.

date

Date at which transaction completed, field has type `<date>`, as defined in the W3C specification.

time

Time at which transaction completed, field has type `<time>`, as defined in the W3C specification.

time-taken

Time taken for transaction to complete in seconds, field has type `<fixed>`, as defined in the W3C specification.

bytes

Number of bytes transferred, field has type `<integer>`.

Note that the `cached` field defined in the W3C specification is not supported in WebLogic Server.

The following identifiers require prefixes, and cannot be used alone. The supported prefix combinations are explained individually.

5.6.3.2.1 IP Address Related Fields These fields give the IP address and port of either the requesting client, or the responding server. These fields have type `<address>`, as defined in the W3C specification. The supported fields are:

c-ip

The IP address of the client.

s-ip

The IP address of the server.

5.6.3.2.2 DNS Related Fields These fields give the domain names of the client or the server and have type <name>, as defined in the W3C specification. The supported fields are:

c-dns

The domain name of the requesting client.

s-dns

The domain name of the requested server.

sc-status

Status code of the response, for example (404) indicating a "File not found" status. This field has type <integer>, as defined in the W3C specification.

sc-comment

The comment returned with status code, for instance "File not found". This field has type <text>.

cs-method

The request method, for example GET or POST. This field has type <name>, as defined in the W3C specification.

cs-uri

The full requested URI. This field has type <uri>, as defined in the W3C specification.

cs-uri-stem

Only the stem portion of URI (omitting query). This field has type <uri>, as defined in the W3C specification.

cs-uri-query

Only the query portion of the URI. This field has type <uri>, as defined in the W3C specification.

5.6.3.2.3 Diagnostic Message Correlation Fields These fields give message correlation information for diagnostic messages, helping you to determine relationships between messages across components. These fields are logged if the diagnostic context is present and populated for the executed request. The diagnostic context may be present if it is propagated into the server with the incoming request, or it may be created for the request by WebLogic Server if the diagnostic context is enabled. The supported fields are:

ctx-ecid

The Execution Context ID (ECID). The ECID is a globally unique identifier associated with the execution of a particular request.

ctx-rid

The Relationship ID (RID). The RID distinguishes the work done in one thread on one process, from work done by any other threads on this and other processes on behalf of the same request.

If the diagnostic context does not exist, or the values of the ECID and RID are not available in the diagnostic context, a hyphen (-) is logged as their values. For more information about the ECID and RID, see "Correlating Messages Across Log Files and Components" in *Administering Oracle Fusion Middleware*.

5.6.3.3 Creating Custom Field Identifiers

You can also create user-defined fields for inclusion in an HTTP access log file that uses the extended log format (ELF). To create a custom field, you identify the field in the ELF log file using the `Fields` directive and then you create a matching Java class that generates the desired output. You can create a separate Java class for each field, or the Java class can output multiple fields. For a sample of the Java source for such a class, see [Section 5-1, "Java Class for Creating a Custom ELF Field"](#).

To create a custom field:

1. Include the field name in the `Fields` directive, using the form:

```
x-myCustomField.
```

Where `myCustomField` is a fully-qualified class name.

For more information about the `Fields` directive, see [Section 5.6.3.1, "Creating the Fields Directive"](#).

2. Create a Java class with the same fully-qualified class name as the custom field you defined with the `Fields` directive (for example `myCustomField`). This class defines the information you want logged in your custom field. The Java class must implement the following interface:

```
weblogic.servlet.logging.CustomELFLogger
```

In your Java class, you must implement the `logField()` method, which takes a `HttpAccountingInfo` object and `FormatStringBuffer` object as its arguments:

- Use the `HttpAccountingInfo` object to access HTTP request and response data that you can output in your custom field. Getter methods are provided to access this information. For a complete listing of these get methods, see [Section 5.6.3.3.1, "Get Methods of the HttpAccountingInfo Object"](#).
 - Use the `FormatStringBuffer` class to create the contents of your custom field. Methods are provided to create suitable output.
3. Compile the Java class and add the class to the `CLASSPATH` statement used to start WebLogic Server. You will probably need to modify the `CLASSPATH` statements in the scripts that you use to start WebLogic Server.

Note: Do not place this class inside of a Web application or Enterprise application in exploded or jar format.

4. Configure WebLogic Server to use the extended log format. For more information, see [Section 5.6.3, "Setting Up HTTP Access Logs by Using Extended Log Format"](#).

Notes: When writing the Java class that defines your custom field, do not execute any code that is likely to slow down the system (For instance, accessing a DBMS or executing significant I/O or networking calls.) Remember, an HTTP access log file entry is created for *every* HTTP request.

Note: If you want to output more than one field, delimit the fields with a tab character. For more information on delimiting fields and other ELF formatting issues, see "Extended Log Format" at <http://www.w3.org/TR/WD-logfile-960221.html>.

5.6.3.3.1 Get Methods of the `HttpAccountingInfo` Object The following methods return various data regarding the HTTP request. These methods are similar to various methods of `javax.servlet.ServletRequest`, `javax.servlet.http.Http.ServletRequest`, and `javax.servlet.http.HttpServletResponse`.

The Javadoc for these interfaces is available at the following locations:

- <http://docs.oracle.com/javaee/6/api/javax/servlet/ServletRequest.html>
- <http://docs.oracle.com/javaee/6/api/javax/servlet/ServletResponse.html>
- <http://docs.oracle.com/javaee/6/api/javax/servlet/http/HttpServletRequest.html>

For details about these methods, see the corresponding methods in the Java interfaces listed in the following table, or refer to the specific information contained in this table.

Table 5–2 Getter Methods of `HttpAccountingInfo`

HttpAccountingInfo Methods	Method Information
<code>Object getAttribute(String name);</code>	<code>javax.servlet.ServletRequest</code>
<code>Enumeration getAttributeNames();</code>	<code>javax.servlet.ServletRequest</code>
<code>String getCharacterEncoding();</code>	<code>javax.servlet.ServletRequest</code>
<code>int getResponseContentLength();</code>	<code>javax.servlet.ServletResponse.setContentLength()</code> This method gets the content length of the response, as set with the <code>setContentLength()</code> method.
<code>String getContentType();</code>	<code>javax.servlet.ServletRequest</code>
<code>Locale getLocale();</code>	<code>javax.servlet.ServletRequest</code>
<code>Enumeration getLocales();</code>	<code>javax.servlet.ServletRequest</code>
<code>String getParameter(String name);</code>	<code>javax.servlet.ServletRequest</code>
<code>Enumeration getParameterNames();</code>	<code>javax.servlet.ServletRequest</code>
<code>String[] getParameterValues(String name);</code>	<code>javax.servlet.ServletRequest</code>
<code>String getProtocol();</code>	<code>javax.servlet.ServletRequest</code>
<code>String getRemoteAddr();</code>	<code>javax.servlet.ServletRequest</code>
<code>String getRemoteHost();</code>	<code>javax.servlet.ServletRequest</code>
<code>String getScheme();</code>	<code>javax.servlet.ServletRequest</code>
<code>String getServerName();</code>	<code>javax.servlet.ServletRequest</code>
<code>int getServerPort();</code>	<code>javax.servlet.ServletRequest</code>
<code>boolean isSecure();</code>	<code>javax.servlet.ServletRequest</code>
<code>String getAuthType();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>String getContextPath();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>Cookie[] getCookies();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>long getDateHeader(String name);</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>String getHeader(String name);</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>Enumeration getHeaderNames();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>Enumeration getHeaders(String name);</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>int getIntHeader(String name);</code>	<code>javax.servlet.http.HttpServletRequest</code>

Table 5–2 (Cont.) Getter Methods of HttpAccountingInfo

HttpAccountingInfo Methods	Method Information
<code>String getMethod();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>String getPathInfo();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>String getPathTranslated();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>String getQueryString();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>String getRemoteUser();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>String getRequestURI();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>String getRequestedSessionId();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>String getServletPath();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>Principal getUserPrincipal();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>boolean isRequestedSessionIdFromCookie();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>boolean isRequestedSessionIdFromURL();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>boolean isRequestedSessionIdFromUrl();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>boolean isRequestedSessionIdValid();</code>	<code>javax.servlet.http.HttpServletRequest</code>
<code>byte[] getURIAsBytes();</code>	Returns the URI of the HTTP request as byte array, for example: If GET /index.html HTTP/1.0 is the first line of an HTTP Request, /index.html is returned as an array of bytes.
<code>long getInvokeTime();</code>	Returns the starting time of <code>currentTimeMillis()</code> . To get the length of time taken by the servlet to send the response to the client, use the following code: <pre>long milsec = System.currentTimeMillis() - metrics.getInvokeTime(); Float sec = new Float(milsec / 1000.0);</pre>
<code>int getResponseStatusCode();</code>	<code>javax.servlet.http.HttpServletResponse</code>
<code>String getResponseHeader(String name);</code>	<code>javax.servlet.http.HttpServletResponse</code>

Example 5–1 Java Class for Creating a Custom ELF Field

```
import weblogic.servlet.logging.CustomELFLogger;
import weblogic.servlet.logging.FormatStringBuffer;
import weblogic.servlet.logging.HttpAccountingInfo;
/* This example outputs the User-Agent field into a
   custom field called MyCustomField
*/
public class MyCustomField implements CustomELFLogger{
public void logField(HttpAccountingInfo metrics,
   FormatStringBuffer buff) {
   buff.appendValueOrDash(metrics.getHeader("User-Agent"));
}
}
```

5.7 Preventing POST Denial-of-Service Attacks

A Denial-of-Service attack is a malicious attempt to overload a server with phony requests. One common type of attack is to send huge amounts of data in an HTTP POST method. You can set three attributes in WebLogic Server that help prevent this type of

attack. These attributes are set in the Console, under *Servers* or *Virtual Hosts*. If you define these attributes for a virtual host, the values set for the virtual host override those set under *Servers*.

PostTimeoutSecs

Amount of time that WebLogic Server waits between receiving chunks of data in an HTTP POST.

The default value for `PostTimeoutSecs` is 30.

MaxPostTimeSecs

Maximum time that WebLogic Server spends receiving post data. If this limit is triggered, a `PostTimeoutException` is thrown and the following message is sent to the server log:

```
Post time exceeded MaxPostTimeSecs.
```

The default value for `MaxPostTimeSecs` is 30.

MaxPostSize

Maximum number of bytes of data received in a POST from a single request. If this limit is triggered, a `MaxPostSizeExceeded` exception is thrown and the following message is sent to the server log:

```
POST size exceeded the parameter MaxPostSize.
```

An HTTP error code 413 (Request Entity Too Large) is sent back to the client.

If the client is in listening mode, it gets these messages. If the client is not in listening mode, the connection is broken.

The default value for `MaxPostSize` is -1.

5.8 Setting Up WebLogic Server for HTTP Tunneling

HTTP tunneling provides a way to simulate a stateful socket connection between WebLogic Server and a Java client when your only option is to use the HTTP protocol. It is generally used to *tunnel* through an HTTP port in a security firewall. HTTP is a stateless protocol, but WebLogic Server provides tunneling functionality to make the connection appear to be a regular `T3Connection`. However, you can expect some performance loss in comparison to a normal socket connection.

5.8.1 Configuring the HTTP Tunneling Connection

Under the HTTP protocol, a client may only make a request, and then accept a reply from a server. The server may not voluntarily communicate with the client, and the protocol is stateless, meaning that a continuous two-way connection is not possible.

WebLogic HTTP tunneling simulates a `T3Connection` via the HTTP protocol, overcoming these limitations. There are attributes that you can configure in the Administration Console to tune a tunneled connection for performance. It is advised that you leave them at their default settings unless you experience connection problems. These properties are used by the server to determine whether the client connection is still valid, or whether the client is still alive.

Enable Tunneling

Enables or disables HTTP tunneling. HTTP tunneling is disabled by default.

Note that the server must also support both the HTTP and T3 protocols in order to use HTTP tunneling.

Tunneling Client Ping

When an HTTP tunnel connection is set up, the client automatically sends a request to the server, so that the server may volunteer a response to the client. The client may also include instructions in a request, but this behavior happens regardless of whether the client application needs to communicate with the server. If the server does not respond (as part of the application code) to the client request within the number of seconds set in this attribute, it does so anyway. The client accepts the response and automatically sends another request immediately.

Default is 45 seconds; valid range is 20 to 900 seconds.

Tunneling Client Timeout

If the number of seconds set in this attribute have elapsed since the client last sent a request to the server (in response to a reply), then the server regards the client as dead, and terminates the HTTP tunnel connection. The server checks the elapsed time at the interval specified by this attribute, when it would otherwise respond to the client's request.

Default is 40 seconds; valid range is 10 to 900 seconds.

5.8.2 Connecting to WebLogic Server from the Client

When your client requests a connection with WebLogic Server, all you need to do in order to use HTTP tunneling is specify the HTTP protocol in the URL. For example:

```
Hashtable env = new Hashtable();
env.put(Context.PROVIDER_URL, "http://wlhost:80");
Context ctx = new InitialContext(env);
```

On the client side, a special tag is appended to the http protocol, so that WebLogic Server knows this is a tunneling connection, instead of a regular HTTP request. Your application code does not need to do any extra work to make this happen.

The client must specify the port in the URL, even if the port is 80. You can set up your WebLogic Server instance to listen for HTTP requests on any port, although the most common choice is port 80 since requests to port 80 are customarily allowed through a firewall.

You specify the listen port for WebLogic Server in the Administration Console under the "Servers" node, under the "Network" tab.

5.9 Using Native I/O for Serving Static Files (Windows Only)

When running WebLogic Server on Windows NT/2000/XP you can specify that WebLogic Server use the native operating system call `TransmitFile` instead of using Java methods to serve static files such as HTML files, text files, and image files. Using native I/O can provide performance improvements when serving larger static files.

To use native I/O, add two parameters to the `web.xml` deployment descriptor of a Web application containing the files to be served using native I/O. The first parameter, `weblogic.http.nativeIOEnabled` should be set to `TRUE` to enable native I/O file serving. The second parameter, `weblogic.http.minimumNativeFileSize` sets the minimum file size for using native I/O. If the file being served is larger than this value, native I/O is used. If you do not specify this parameter, a value of 4K is used.

Generally, native I/O provides greater performance gains when serving larger files; however, as the load on the machine running WebLogic Server increases, these gains diminish. You may need to experiment to find the correct value for `weblogic.http.minimumNativeFileSize`.

The following example shows the complete entries that should be added to the web.xml deployment descriptor. These entries must be placed in the web.xml file after the <distributable> element and before <servlet> element.

```
<context-param>
  <param-name>weblogic.http.nativeIOEnabled</param-name>
  <param-value>TRUE</param-value>
</context-param>
<context-param>
  <param-name>weblogic.http.minimumNativeFileSize</param-name>
  <param-value>500</param-value>
</context-param>
```

weblogic.http.nativeIOEnabled can also be set as a context parameter in the FileServlet.

Using the WebLogic Persistent Store

This chapter explains how to configure and monitor the WebLogic Server 12.1.3 persistent store, which provides a built-in, high-performance storage solution for WebLogic Server subsystems and services that require persistence.

This chapter includes the following sections:

- [Section 6.1, "Overview of the Persistent Store"](#)
- [Section 6.2, "Using the Default Persistent Store"](#)
- [Section 6.3, "Using Custom File Stores and JDBC Stores"](#)
- [Section 6.4, "Creating a Custom \(User-Defined\) File Store"](#)
- [Section 6.5, "Creating JDBC-accessible Stores"](#)
- [Section 6.6, "Monitoring a Persistent Store"](#)
- [Section 6.8, "Security Considerations"](#)
- [Section 6.9, "Limitations and Considerations of the Persistent Store"](#)

6.1 Overview of the Persistent Store

The persistent store provides a built-in, high-performance storage solution for WebLogic Server subsystems and services that require persistence. For example, it can store persistent JMS messages or temporarily store messages sent using the Store-and-Forward feature. The persistent store supports persistence to a file-based store or to a JDBC-accessible store in a database.

[Table 6–1](#) defines many of the WebLogic services and subsystems that can create connections to the persistent store. Each subsystem that uses the persistent store specifies a unique connection ID that identifies that subsystem.

Table 6–1 Persistent Store Users

Subsystem/Service	What It Stores	More Information
Diagnostic Service	Log records, data events, and harvested metrics.	"Understanding WLDF Configuration" in <i>Configuring and Using the Diagnostics Framework for Oracle WebLogic Server</i>
JMS Messages	Persistent messages and durable subscribers.	"Understanding the Messaging Models" in <i>Developing JMS Applications for Oracle WebLogic Server</i>
JMS Paging Store	One per JMS server. Paged persistent and non-persistent messages.	"Main Steps for Configuring Basic JMS System Resources" in <i>Administering JMS Resources for Oracle WebLogic Server</i> .

Table 6–1 (Cont.) Persistent Store Users

Subsystem/Service	What It Stores	More Information
JTA Transaction Log (TLOG)	Information about committed transactions coordinated by the server that may not have been completed. TLOGs can be stored in the default persistent store or a JDBC TLOG store.	<ul style="list-style-type: none"> ■ "Managing Transactions" in <i>Developing JTA Applications for Oracle WebLogic Server</i>. ■ Section 6.5.1, "Using a JDBC TLog Store"
Path Service	The mapping of a group of messages to a messaging resource.	"Using the WebLogic Path Service" in <i>Administering JMS Resources for Oracle WebLogic Server</i>
Store-and-Forward (SAF) Service Agents	Messages for a sending SAF agent for retransmission to a receiving SAF agent	"Understanding the Store-and-Forward Service" in <i>Administering the Store-and-Forward Service for Oracle WebLogic Server</i>
Web Services	Request and response SOAP messages from an invocation of a reliable WebLogic Web Service.	"Using Web Services Reliable Messaging" in <i>Developing JAX-RPC Web Services for Oracle WebLogic Server</i>
EJB Timer Services	EJB Timer objects.	"Understanding Enterprise JavaBeans" in <i>Developing Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server</i>

For more information about the store connection IDs, see [Section 6.6.2, "Monitoring Store Connections"](#).

6.1.1 Features of the Persistent Store

The key features of the persistent store include:

- Default file store for each server instance that requires no configuration.
- The Default and custom stores are shareable by multiple subsystems, as long as they are all targeted to the same server instance, cluster, or migratable target.
- When configured, a JDBC TLOG store which contains information about committed transactions coordinated by the server that may not have been completed. You can select to persist TLOG information either in the default store or the JDBC TLOG store, depending on your application needs. See [Section 6.5.1, "Using a JDBC TLog Store."](#)
- High-performance throughput and transactional support.
- Modifiable parameters that let you create customized file stores and JDBC stores.
- Monitoring capabilities for persistent store statistics and open store connections.
- In a clustered environment, the JDBC TLOG store and customized stores can be migrated from an unhealthy server to a backup server, either on the whole-server level or on the service level.

6.1.2 High-Performance Throughput and Transactional Support

Throughput is the main performance goal of the persistent store. Multiple subsystems can share the same default or custom store, as long as they are all targeted to the same server instance, cluster, or migratable target.

Note: The JDBC TLOG store is only used to persist information about committed transactions coordinated by the server that may not have been completed. It can not be shared by other subsystems.

This is a performance advantage because the persistent store is treated as a single resource by the transaction manager for a particular transaction, even if the transaction involves multiple services that use the same store. For example, if the TLOG, JMS and EJB timers share a file store, and a JMS message and an EJB timer are created in a single transaction, the transaction will be one-phase and incur a single resource write, rather than two-phase, which incurs four resource writes (two on each resource), plus a transaction entry write (on the transaction log).

Both a file store and a JDBC store can survive a process crash or hardware power failure without losing any committed updates. Uncommitted updates may be retained or lost, but in no case will a transaction be left partially complete after a crash.

6.1.3 Comparing File Stores and JDBC-accessible Stores

The following are some similarities and differences between file stores and JDBC-accessible stores:

- The default persistent store can only be a file store. Therefore, a JDBC store cannot be used as a default persistent store.
- Both have the same transaction semantics and guarantees. As with JDBC store writes, file store writes are guaranteed to be persisted to disk and are not simply left in an intermediate (that is, unsafe) cache.
- Both have the same application interface (no difference in application code).
- All things being equal, file stores generally offer better throughput than a JDBC store.

Note: If a database is running on high-end hardware with very fast disks, and WebLogic Server is running on slower hardware or with slower disks, then you may get better performance from the JDBC store.

- File stores are generally easier to configure and administer, and do not require that WebLogic subsystems depend on any external component.
- File stores generate no network traffic; whereas, JDBC stores generate network traffic if the database is on a different machine from WebLogic Server.
- JDBC stores may make it easier to handle failure recovery since the JDBC interface can access the database from any machine on the same network. With the file store, the disk must be shared or migrated.

6.1.4 High Availability For Persistent Stores

For high availability, a persistent file-based store (default, or custom) can be migrated along with its parent server as part of the "whole server-level" migration feature, which provides both automatic and manual migration at the server level, rather than on the service level. For more information, see "Whole Server Migration" in *Administering Clusters for Oracle WebLogic Server*. However, file-based stores must be

configured on a shared disk that is available to the migratable target servers in the cluster.

6.1.4.1 Persistent Store Migration

File-based stores *and* JDBC-accessible stores can also be migrated as part of a "service-level" migration for JMS-related services, such as JMS servers, SAF agents, and the path service, which rely on stores to maintain data. Service-level migration is controlled by a *migratable target*, which serves as a grouping of JMS-related services, and which is hosted on only one physical server in a cluster. Such hosted services can be automatically migrated from the current unhealthy hosting server to a healthy active server with the help of the Health Monitoring subsystem. JMS services hosted by a migratable target can also be manually migrated, either in response to a server failure or as part of regularly scheduled server maintenance. When the migratable target is migrated, all pinned services hosted by that target are also migrated. For more information on service-level migration, see "Service Migration" in *Administering Clusters for Oracle WebLogic Server*.

Migratable JMS-related services cannot use the default file store, so you must configure a custom file store or JDBC store and target it to the same migratable target as the JMS server, SAF agent, or path service associated with the store.

Note: As a best practice, a path service should use its own custom store and migratable target.

Migratable file stores must also either be configured on a shared disk that is available to the migratable targets in the cluster or you can use pre/post-migration scripts to migrate a file store's data to a backup server. See "Custom Store Availability for JMS Services" in *Administering Clusters for Oracle WebLogic Server*.

6.1.4.2 High Availability Storage Solutions

If you have applications that need access to persistent stores that reside on remote machines after the migration of a JMS server or JTA transaction log, then you should implement one of the following highly-available storage solutions:

- File-based stores (default or custom)—Implement a hardware solution, such as a dual-ported SCSI disk or Storage Area Network (SAN) to make a file store available from shareable disks or remote machines.

Note: If a file store is disconnected and re-connected again, its host server instance must be rebooted to successfully continue sending/receiving persistent JMS messages. For example, if for some reason the file system containing a file store is unmounted and then remounted, attempts to send persistent JMS messages will generate JMS exceptions until the host server is rebooted.

- JDBC-accessible stores—Configure a JDBC store or JDBC TLOG store and use JDBC to access this store, which can be on yet another server. Applications can then take advantage of any high-availability or failover solutions offered by your database vendor. In addition, JDBC stores support GridLink data sources and multi data sources, which provide failover between nodes of a highly available database system, such as Oracle Real Application Clusters (Oracle RAC). For more information, see:

- "Configuring JDBC Multi Data Sources" in *Administering JDBC Data Sources for Oracle WebLogic Server*
- "Using GridLink Data Sources" in *Administering JDBC Data Sources for Oracle WebLogic Server*
- Any persistent store—Use high-availability clustering software which provides an integrated, out-of-the-box solution for WebLogic Server-based applications.

6.2 Using the Default Persistent Store

Each server instance, including the Administration Server, has a default persistent store that requires no configuration. The default store is a file-based store that maintains its data in a group of files in a server instance `data\store\default` directory. A directory for the default store is automatically created if one does not already exist. This default store is available to subsystems that do not require explicit selection of a particular store and function best by using the system's default storage mechanism. For example, a JMS Server with no persistent store configured will use the default store for its Managed Server and will support persistent messaging.

The default store can be configured by directly manipulating `DefaultFileStoreMBean` parameters. If this MBean is not defined in the domain configuration file, then the configuration subsystem ensures that the `DefaultFileStoreMBean` always exists with the default values.

Also, the WebLogic Server Administration Console enables you to change the default store parameters, such as its default directory location and Synchronous Write Policy, as described in "Modify the Default Store Settings" in the *Oracle WebLogic Server Administration Console Online Help*.

6.2.1 Default Store Location

The default store maintains its data in a `data\store\default` directory inside the `servername` subdirectory of a domain's root directory

For example, if no directory name is specified for the default file store, it defaults to:

```
ORACLE_HOME\user_projects\domains\domain-name\servers\server-name\data\store\default
```

where `domainname` is the root directory of your domain, typically `c:\oracle\user_projects\domains\domainname`, which is parallel to the directory in which WebLogic Server program files are stored, typically `c:\oracle\wlserver_12.1`.

You can, however, specify another location for the default store by directly manipulating the `DefaultFileStoreMBean` parameters or by using the WebLogic Server Administration Console, as described in "Modify the Default Store Settings" in the *Oracle WebLogic Server Administration Console Online Help*.

6.2.2 Example of a Default File Store

Here's an example of how a default file store may look in a domain's configuration file, with the default directory location and Synchronous Write Policy settings overridden:

```
<server
  <name>myserver</name>
  <default-file-store>
  <directory>C:/store</directory>
  <synchronous-write-policy>Disabled</synchronous-write-policy>
</default-file-store>
```

</server>

6.3 Using Custom File Stores and JDBC Stores

In addition to using the default file store, you can also configure a file store or JDBC store to suit your specific needs. A custom file store, like the default file store, maintains its data in a group of files in a directory. However, you may want to create a custom file store so that the file store's data is persisted to a particular storage device or when you want a JMS service that accesses a file store to be able to migrate with the store to another server member in a cluster. When configuring a file store directory, the directory must be accessible to the server instance on which the file store is located.

A JDBC store can be configured when a relational database is used for storage. A JDBC store enables you to store persistent messages in a standard JDBC-accessible database, which is accessed through a designated JDBC data source. The data is stored in the JDBC store's database table, which has a logical name of `WLStore`. It is up to the database administrator to configure the database for high availability and performance. JDBC stores also support migratable targets for automatic or manual JMS service migration.

For more information about configuring a persistent store, see [Section 6.3.1, "When to Use a Custom Persistent Store"](#).

6.3.1 When to Use a Custom Persistent Store

WebLogic Server provides configuration options for creating a custom file store or JDBC-accessible store. For example, you may want to:

- Place a file store's files on a particular device.
- Use a JDBC store rather than a file store for a particular server instance. If you want to persist transaction logs, use a JDBC TLOG store. See [Section 6.5.1, "Using a JDBC TLog Store."](#)
- Allow all physical stores in a cluster to share the same logical name.
- Logically separate different services to use different files or tables. (This may simplify administration and maintenance at the expense of reduced performance.)
- Migratable JMS-related services cannot use the default persistent store, so you must configure a custom store and target it to the same migratable target as the migratable JMS service. For more information, see "Service Migration" in *Administering Clusters for Oracle WebLogic Server*.

6.3.2 Methods of Creating a Custom Persistent Store

A user-defined persistent store can be configured in the following ways:

- Use the WebLogic Server Administration Console. To configure a custom file store or JDBC store, see "Configure Persistent Stores" in the *Oracle WebLogic Server Administration Console Online Help*. To configure a JDBC TLOG store, see "Configure the Transaction Log Store" in the *Oracle WebLogic Server Administration Console Online Help*.
- Directly edit the configuration file (`config.xml`) of the server instance that is hosting a persistent store.
- Use the WebLogic Java Management Extensions (JMX) to create persistent stores. JMX is the Java EE solution for monitoring and managing resources on a network.

For more information see, *Developing Custom Management Utilities Using JMX for Oracle WebLogic Server*.

- Use the WebLogic Scripting Tool (WLST) to create persistent stores. WLST is a command-line scripting interface that you use to interact with and configure WebLogic Server instances and domains. For more information, see *Understanding the WebLogic Scripting Tool*.
- Use the WebLogic Configuration Wizard to change the options of the default persistent store. For detailed information on how to use the Configuration Wizard to configure a persistent store, see "Creating a WebLogic Domain" in *Creating WebLogic Domains Using the Configuration Wizard*.

6.3.3 Modifying Custom Persistent Store Parameters

Modifying certain custom store configuration options, such as a JDBC store's prefix or a file store's directory, do not necessarily require a server restart if you do the following:

1. Set the targets of any dependent services to null (such as a JMS server that uses the custom store), and then setting the custom store target to null. (Setting a service's target to null implicitly shuts down the service.)
2. Reverse the process by setting the custom store target back to its original value and then setting the dependent resource targets back to their original values.

In cases where the custom store and JMS servers share a migratable target, you can administratively restart the migratable target.

6.4 Creating a Custom (User-Defined) File Store

The following sections provide an example of a custom file store and configuration guidelines for choosing a synchronous write policy.

To create a custom file store, you can directly modify the default FileStoreMBean parameters. For instructions on using the WebLogic Server Administration Console to create a custom file store, see "Create File Stores" in the *Oracle WebLogic Server Administration Console Online Help*.

6.4.1 Main Steps for Configuring a Custom File Store

The main steps for creating a custom file store are as follows:

1. Create a directory where the file store's data will be persisted.
2. Create a custom file store and specify the directory location that you created.
3. Associate the custom file store with the subsystem(s) or migratable target that will be accessing it, such as:
 - For JMS servers, select the custom file store on the General Configuration page.
 - For Store-and-Forward agents, select the custom file store on the General Configuration page.
 - For a Path Service, select the custom file store on the General Configuration page.

6.4.2 Example of a Custom File Store

Here's an example of how a custom file store may look in a domain's configuration file with its files kept in a `/disk1/jmslog` directory.

```
<file-store>
  <name>SampleFileStore</name>
  <target>myserver</target>
  <directory>/disk1/jmslog</directory>
</file-store>
```

Table 6–2 briefly describes the file store configuration parameters that can be modified.

Table 6–2 Custom File Store Configuration Options

Option	Required	What It Does
Name	Yes	The name of the file store, which must be unique across all stores in the domain.
Targets	Yes	The server instance, cluster, or migratable target where a file store is targeted. Multiple subsystems can share the same file store, as long as they are all targeted to the same server instance or migratable target. Note: <ul style="list-style-type: none"> When using a cluster to host a JMS Server, you must target the file store to the same cluster used by the JMS Server. See When using migratable targets for JMS services, you must target the file store to the same migratable target used by the JMS service. See "Service Migration" in <i>Administering Clusters for Oracle WebLogic Server</i>.
Directory	Yes	The path name to the directory on the file system where the file store is kept. Note: When targeting a file store to a migratable target, the store directory must be accessible from all candidate server members in the migratable target. For highest availability, use either a SAN (Storage Area Network) or a dual-ported SCSI disk. See "Service Migration" in <i>Administering Clusters for Oracle WebLogic Server</i> . Modifying an existing file store's directory does not necessarily require a server restart, as described in Section 6.3.3, "Modifying Custom Persistent Store Parameters" .
CacheDirectory	No	This setting only applies for the <code>Direct-Write-With-Cache</code> file store synchronous write policy. See Section 6.4.3, "Guidelines for Configuring a Synchronous Write Policy."
Logical Name	No	Optionally used with subsystems, like EJBs, when deploying a module to an entire cluster, but also require a different physical store on each server instance in the cluster. In such a configuration, each physical store would have its own name, but all the persistent stores would share the same logical name.
Synchronous Write Policy	No	Defines the IO behavior of a file store including immediate durability of individual write operations. Values are: <code>Direct-Write</code> (default), <code>Direct-Write-With-Cache</code> , <code>Cache-Flush</code> , and <code>Disabled</code> . For more information, see Section 6.4.3, "Guidelines for Configuring a Synchronous Write Policy" .

For instructions on configuring a custom file store using the WebLogic Server Administration Console, see "Create File Stores" in the *Oracle WebLogic Server Administration Console Online Help*.

6.4.3 Guidelines for Configuring a Synchronous Write Policy

There are several Synchronous Write Policies available for file stores. The Synchronous Write Policy determines the behavior of the write operation of the file store. You should select a policy that best suits your environment and meets your needs for runtime performance and data integrity after a possible crash. See "Tuning the WebLogic Persistent Store" in *Tuning Performance of Oracle WebLogic Server* for more details about tuning and performance specifics of Synchronous Write Policy and other file store options.

Note: To view a running custom or default file store's synchronous write policy and driver, check the WL-280008 and WL-280009 messages in the server log.

6.4.3.1 Direct-Write-With-Cache Policy

For most scenarios, Oracle recommends using the `Direct-Write-With-Cache` policy. When this policy is selected, WebLogic Server writes synchronously to a primary set of files in the location defined by the `Directory` attribute of the file store configuration using a native I/O `wlfileio` driver. WebLogic Server also asynchronously writes to a corresponding cache file in the location defined by the `CacheDirectory` attribute of the file store configuration, which is done implicitly by using OS memory caching the cache file blocks as output buffers for the primary data file. The cache files are used for performance optimizations at runtime and boot time and for recovery. This combination of direct writing with a native file driver and the use of corresponding cache files typically provides the best overall performance with the most safe disk writes.

This option uses approximately twice as much disk space as other policies and stores files in two locations. You may need to consider disk space allocations in these locations and you may need to secure both of these locations.

When configuring file locations with the `Direct-Write-With-Cache` policy, the location of the `CacheDirectory` attribute should be a local directory, even when configuring for high availability (Whole Server Migration or Automatic Service Migration). The cache files are used for performance optimizations only. The true persistent storage for messages is defined by the `Directory` attribute of the file store configuration. Only that directory needs to be available to the migrated WebLogic Server instance or JMS service after migration. The same applies to disaster recovery scenarios: only the files defined in the `Directory` location need to be replicated to the backup site.

Notes: If the file store native `wlfileio` driver cannot be loaded, the store automatically runs in an alternate specialized `Direct-Write` policy mode. To view a running custom or default file store's configured and actual synchronous write policy and driver, examine the server log for `WL-280008` and `WL-280009` messages.

Certain older versions of Microsoft Windows may incorrectly report storage device synchronous write completion if the Windows default `Write Cache Enabled` setting is used. This violates the transactional semantics of transactional products (not specific to Oracle), including file stores configured with a `Direct-Write` (default) or `Direct-Write-With-Cache` policy, as a system crash or power failure can lead to a loss or a duplication of records/messages. One of the visible symptoms is that this problem may manifest itself in high persistent message/transaction throughput exceeding the physical capabilities of your storage device. You can address the problem by applying a Microsoft supplied patch, disabling the Windows `Write Cache Enabled` setting, or by using a power-protected storage device. See <http://support.microsoft.com/kb/281672> and <http://support.microsoft.com/kb/332023>.

6.4.3.2 Direct-Write Policy

When the `Direct-Write` policy is selected, WebLogic Server writes synchronously to a primary set of files in the location defined by the `Directory` attribute of the file store configuration using a native I/O `wlfileio` driver. This policy typically performs slower than the `Direct-Write-With-Cache` policy, but it uses less disk space and may have fewer environmental considerations to manage. The `Direct-Write` policy is typically faster than the `Cache-Flush` policy.

Note: Certain older versions of Microsoft Windows may incorrectly report storage device synchronous write completion if the Windows default `Write Cache Enabled` setting is used. This violates the transactional semantics of transactional products (not specific to Oracle), including file stores configured with a `Direct-Write` (default) or `Direct-Write-With-Cache` policy, as a system crash or power failure can lead to a loss or a duplication of records/messages. One of the visible symptoms is that this problem may manifest itself in high persistent message/transaction throughput exceeding the physical capabilities of your storage device. You can address the problem by applying a Microsoft supplied patch, disabling the Windows `Write Cache Enabled` setting, or by using a power-protected storage device. See <http://support.microsoft.com/kb/281672> and <http://support.microsoft.com/kb/332023>.

6.4.3.3 Cache-Flush Policy

When the `Cache-Flush` policy is selected, WebLogic Server enables the default file write behavior of the operating system and storage device, which typically includes caching and scheduling file writes, but forces a flush of the cache to disk before completing a transaction. Transactions cannot complete until all of their writes have been flushed down to disk. This policy is reliable and scales well as the number of simultaneous users increases. It is transactionally safe, but tends to provide lower

runtime performance than the direct-write policies in typical use cases, except in those cases with large numbers of simultaneous producers or consumers.

6.4.3.4 Disabled Policy

When the `Disabled` policy is selected, WebLogic Server relies on the default file write behavior of the operating system and storage device. In most cases, file writes are cached in memory and are scheduled for writing instead of being directly written to disk. The `Disabled` policy generally improves file store performance, often quite dramatically, but at the expense of possibly losing sent messages or generating duplicate received messages (even if messages are transactional) in the event of an operating system crash or a hardware failure. This is because transactions are complete as soon as their writes are cached in memory, instead of waiting for the writes to successfully reach the disk. Simply shutting down an operating system or killing a WebLogic Server process does not generate these failures, as an OS flushes all outstanding writes under these circumstances during a normal shutdown. Instead, these failures can be emulated by abruptly shutting the power off to a busy server.

6.5 Creating JDBC-accessible Stores

The following sections provide information on how to configure and use JDBC-accessible stores:

- JDBC TLog Stores: to persist transaction logs (TLOGs) in a database. See [Section 6.5.1, "Using a JDBC TLog Store."](#)
- JDBC Stores: to persist WebLogic Server instance services and subsystem information, excluding TLOGs, in a database. See [Section 6.5.2, "Using a JDBC Store."](#)

6.5.1 Using a JDBC TLog Store

You can configure a JDBC TLOG store to persist transaction logs to a database, which provides the following benefits:

- Leverages replication and HA characteristics of the underlying database.
- Simplifies disaster recovery by allowing the easy synchronization of the state of the database and TLOGs.
- Improved Transaction Recovery service migration as the transaction logs do not need to be migrated (copied) to a new location.

6.5.1.1 Main Steps for Configuring a JDBC TLOG Store

The main steps for creating a JDBC TLOG store are as follows:

1. Create a JDBC data source, GridLink data source, or multi data source to interface with the JDBC store. See [Section 6.5.1.1.1, "Choosing a Data Source."](#)
2. Create a JDBC TLOG store and associate it with the JDBC data source, GridLink data source, or multi data source created in Step 1. See "Configure the Transaction Log Store" in the *Oracle WebLogic Server Administration Console Online Help*.
3. Optional. It is highly recommended that you configure the Prefix option to a unique value for each configured JDBC TLOG store.
4. For high availability, make your data source available to backup servers. See [Chapter 6.5.1.5, "Server Migration when using a JDBC TLOG Store."](#)

6.5.1.1.1 Choosing a Data Source You can choose one of the following data source types, depending on your WebLogic Server license and application needs:

- Generic Data Sources—See "Creating a JDBC Data Source" and "Using Connect-Time Failover with Oracle RAC" in *Administering JDBC Data Sources for Oracle WebLogic Server*.
- GridLink Data Sources—See "Using GridLink Data Sources" in *Administering JDBC Data Sources for Oracle WebLogic Server*.
- Multi data sources—Backed by a fully replicated, zero-latency database, such as Oracle RAC. See "Configuring JDBC Multi Data Sources" and "Using Multi Data Sources with Oracle RAC" in *Administering JDBC Data Sources for Oracle WebLogic Server*.

6.5.1.2 Example of a JDBC TLOG Store

Here's an example of how a JDBC TLOG store may look in the configuration file, using the JDBC data source `MyDataSource`, and with a logical name specified:

```
<server>
  <transaction-log-jdbc-store>
    <data-source>MyDataSource</data-source>
    <prefix-name>TLOG_MS1</prefix-name>
    <create-table-ddl-file>myDDL/myCreateTable.sql</create-table-ddl-file>
    <max-retry-seconds-before-tlog-fail>120</max-retry-seconds-before-tlog-fail>
  </transaction-log-jdbc-store>
</server>
```

Table 6–3 describes the JDBC TLOG store configuration parameters that can be modified.

Table 6–3 JDBC TLOG Store Configuration Options

Option	Required	What It Does
Data Source	Yes	<p>The JDBC data source or multi data source used by this JDBC store to access the store's database table (<code>WLStore</code>). This data source or multi data source must be targeted to the same server instance as the JDBC store.</p> <p>Note: The JDBC store must use a JDBC data source that uses a non-XA JDBC driver and has Supports Global Transactions disabled. This limitation does not remove the XA capabilities of layered subsystems that use JDBC stores. For example, WebLogic JMS is fully XA-capable regardless of whether it uses a file store or any JDBC store.</p>
Prefix Name	No	<p>The prefix for the JDBC store's table is generally entered in the following format: <code>[[[catalog.]schema.]prefix]</code></p> <p>When using multiple JDBC stores, it is required to set this option to a unique value for each configured JDBC store. When no prefix is specified, the JDBC store table name is simply <code>WLStore</code> and the database implicitly determines the schema according the current user of the JDBC connection. Also, two JDBC stores cannot share the same database table. For more information, see Section 6.5.2.6.1, "Using Prefixes with a JDBC Store".</p> <p>Modifying an existing JDBC store's prefix does not necessarily require a server restart, as described in Section 6.3.3, "Modifying Custom Persistent Store Parameters".</p>

Table 6–3 (Cont.) JDBC TLOG Store Configuration Options

Option	Required	What It Does
Create Table from DDL File	No	Optionally used with supported DDL (data definition language) files to create the JDBC store's database table (WLStore). This option is ignored when the JDBC store's database table already exists. For more information, see Section 6.5.2.4, "Creating a JDBC Store Table Using Default and Custom DDL Files" .
Deletes Per Batch Maximum	Default is 20.	The maximum number of table rows that are deleted per database call.
Inserts Per Batch Maximum	Default is 20.	The maximum number of table rows that are inserted per database call.
Deletes Per Statement Maximum	Default is 20	The maximum number of table rows that are deleted per database call.
MaxRetrySecondsBeforeTLogFail	Default is 300.	The maximum amount of time, in seconds, WebLogic Server tries to recover from a JDBC TLog store failure. If store remains unusable after this period, WebLogic Server set the health state to HEALTH_FAILED. A value of 0 indicates WebLogic Server does not conduct a retry and and immediately sets the health state as HEALTH_FAILED.
MaxRetrySecondsBeforeTXRollback	Default is 60.	The maximum amount of time, in seconds, WebLogic Server waits before trying to recover from a JDBC TLog store failure while processing a transaction. If store remains unusable after this amount of time, WebLogic Server rolls back the affected transaction. A value of 0 indicates WebLogic Server does not conduct a retry and rolls back the transaction immediately. The practical maximum value is a value less than the current value of MaxRetrySecondsBeforeTLogFail.
RetryIntervalSeconds	Default is 5.	The amount of time, in seconds, WebLogic Server waits before attempting to verify the health of the TLOG store after a store failure has occurred.

For instructions on configuring a JDBC TLOG store using the WebLogic Server Administration Console, see "Configure the Transaction Log Store" in the *Oracle WebLogic Server Administration Console Online Help*.]

6.5.1.3 Configuration Guidelines

The following section provides guidelines for configuring JDBC TLOG stores.

- Only globally-scoped (not application-scoped) data sources can be targeted to a JDBC TLOG store.
- Only one JDBC TLOG store can be configured per WebLogic Server. Conversely, multiple WebLogic Servers can not share a JDBC TLOG store.
- You must configure a JDBC TLOG store. The default is to persist TLOG information to the server's default persistent store.
- You cannot use a data source that is configured to use an XA JDBC driver or is configured to support global transactions. Use a non-XA data source.
- For general rules on JDBC-accessible stores, see [Section 6.5.2.6, "Guidelines for Configuring a JDBC Store."](#)

6.5.1.4 Additional Considerations

The following section provides additional information on JDBC TLOG store behavior and limitations:

- The database used to store the TLOG information must be available at server startup. If the database is not available, the WebLogic Server instance will fail to boot.
- Only the JTA sub-system can use the JDBC TLOG store to persist information about committed transactions coordinated by the server that may not have been completed. No other systems can access the JDBC TLOG store.
- Using a JDBC TLOG store does not change LLR behavior. A JDBC TLOG store can be used with or without LLR. When used in tandem with LLR transactions, the transaction committing information is stored in a LLR table but the checkpoint records and heuristic logs are stored in the JDBC TLOG store.
- If the TLOG store is changed from one store type to another or from one location to another, the change takes effect only after reboot and all pending transactions in the old store are not be copied to the new store. You must ensure there are no pending transactions before changing the TLOG store type or location.
- If the JDBC TLOG store becomes unavailable, the JTA health state transitions to FAILED and any global transactions will fail. In turn, the server life-cycle changes to FAILED. The JTA Transaction Recovery System then attempts to recover from transient runtime errors if possible and resolves any in-doubt transactions. See [Section 6.5.1.5, "Server Migration when using a JDBC TLOG Store."](#)
- If the database used to store TLOG is corrupted and can not be restored, than all pending transaction information is lost.
- If database tables or rows used by the JDBC TLOG store are locked for some reason in the database, the database administrator must resolve these locks manually. Otherwise, the JTA subsystem is blocked and will be suspended until the lock(s) are released, or encounters an exception due to lock. The JTA subsystem will remain unable to operate correctly until the lock(s) are released or the value of `MaxRetrySecondsBeforeTLOGFail` is exceeded.

Note: Different databases have different features for locked local transactions. Some databases may have trouble resolving database locks in a timely manner. You may need to contact your database administrator for more information on basic row locking issues that may occur in your application environment.

6.5.1.5 Server Migration when using a JDBC TLOG Store

WebLogic Server supports both manual and automatic migration of the Transaction Recovery Service when using a JDBC TLOG store. The data source used by the JDBC TLOG store must be targeted on both the primary server instance and a backup server instance. Oracle recommends targeting the data source to all the server instances of a cluster. For more information, see "Transaction Recovery After a Server Fails" in *Developing JTA Applications for Oracle WebLogic Server*.

6.5.1.6 Monitoring a JDBC TLOG Store

You can monitor statistics for Transaction Log Store statistics and for each open store connection. For general information on how to monitor persistent stores, see [Section 6.6, "Monitoring a Persistent Store."](#)

6.5.1.6.1 How to Monitor the JDBC TLOG Store Health State When you configure WebLogic Server to use a JDBC TLOG store, the store is registered with the Health system as a non-critical subsystem using a name with the following pattern:

```
PersistentStore.TLOG_servername
```

where *servername* is the name of the WebLogic Server instance hosting the primary TLOG store.

You can monitor the JDBC TLOG store health state in the WebLogic Server Administration Console, see "Monitor server health" in *Oracle WebLogic Server Administration Console Online Help*.

6.5.1.6.2 How to Monitor Transaction Log Store Statistics You can monitor Transaction Log Store statistics in the WebLogic Server Administration Console, see "View transaction log statistics for a server" in *Oracle WebLogic Server Administration Console Online Help*.

6.5.1.6.3 How to Monitor Transaction Log Store Connections You can monitor Transaction Log Store connection statistics in the WebLogic Server Administration Console, see "View statistics for TLOG store connections" in *Oracle WebLogic Server Administration Console Online Help*.

6.5.1.7 Security Considerations

Properly secure your application environment, including the JDBC TLOG store table. Failure to do so may allow a process to:

- Delete information, maliciously or unintentionally. Such a deletion can cause transaction information to be lost and cause affected global transactions to be completed heuristically.
- Modify information, maliciously or unintentionally. Such modification can cause unexpected behavior.
- Read confidential transaction information, such as when transaction starts and what resources are involved.
- Access the database instance or database machine.
- Access the network between JTA and the database, potentially intercepting, viewing, or modifying data.

See [Chapter 6.8, "Security Considerations."](#)

6.5.2 Using a JDBC Store

The following sections provide an example of a JDBC store, and information about creating a database table for a JDBC store, either using an existing DDL, a custom DDL, and using Oracle blob record columns in a DDL file.

To create a JDBC store, you can directly modify the default JDBCStoreMBean parameters. For instructions on using the WebLogic Server Administration Console to create a JDBC store, see "Create JDBC Stores" in the *Oracle WebLogic Server Administration Console Online Help*.

For configuration guidelines on using prefixes with JDBC stores and recommended JDBC data source settings, see [Section 6.5.2.6, "Guidelines for Configuring a JDBC Store"](#).

6.5.2.1 Main Steps for Configuring a JDBC Store

The main steps for creating a JDBC store are as follows:

1. Create a JDBC data source or multi data source to interface with the JDBC store.
2. Create a JDBC store and associate it with the JDBC data source or multi data source.
3. It is highly recommended that you configure the Prefix option to a unique value for each configured JDBC store table.
4. Associate the JDBC store with the subsystem(s) that will be using it, such as:
 - For JMS servers, select the JDBC store on the General Configuration page.
 - For Store-and-Forward agents, select the JDBC store on the General Configuration page.
 - For a Path Service, select the custom file store on the General Configuration page.

6.5.2.2 Example of a JDBC Store

Here's an example of how a JDBC store may look in the configuration file, using the JDBC data source `MyDataSource`, and with a logical name specified:

```
<jdbc-store>
  <name>SampleJDBCStore</name>
  <target>yourserver</target>
  <data-source>MyDataSource</data-source>
  <logical-name>Baz</logical-name>
</jdbc-store>
```

Table 6–4 describes the JDBC store configuration parameters that can be modified.

Table 6–4 JDBC Store Configuration Options

Option	Required	What It Does
Name	Yes	The name of the JDBC store, which must be unique across all stores in the domain.
Targets	Yes	The server instance, cluster, or migratable target where a JDBC store is targeted. Multiple subsystems can share the same JDBC store, as long as they are all targeted to the same server instance or migratable target. Note: <ul style="list-style-type: none"> ■ When using a cluster to host a JMS Server, you must target the JDBC store to the same cluster used by the JMS Server. See "Configuring Dynamic Clustered JMS" in <i>Administering JMS Resources for Oracle WebLogic Server</i>. ■ When using migratable targets for JMS services, you must target the JDBC store to the same migratable target used by the JMS service. See "Service Migration" in <i>Administering Clusters for Oracle WebLogic Server</i>.

Table 6–4 (Cont.) JDBC Store Configuration Options

Option	Required	What It Does
Data Source	Yes	<p>The JDBC data source or multi data source used by this JDBC store to access the store's database table (<code>WLStore</code>). This data source or multi data source must be targeted to the same target as the JDBC store.</p> <p>Note: The JDBC store must use a JDBC data source that uses a non-XA JDBC driver and has Supports Global Transactions disabled. This limitation does not remove the XA capabilities of layered subsystems that use JDBC stores. For example, WebLogic JMS is fully XA-capable regardless of whether it uses a file store or any JDBC store.</p>
Prefix Name	No	<p>The prefix for the JDBC store's table is generally entered in the following format: <code>[[catalog.]schema.]prefix</code></p> <p>When using multiple JDBC stores, it is required to set this option to a unique value for each configured JDBC store. When no prefix is specified, the JDBC store table name is simply <code>WLStore</code> and the database implicitly determines the schema according the current user of the JDBC connection. Also, two JDBC stores cannot share the same database table. For more information, see Section 6.5.2.6.1, "Using Prefixes with a JDBC Store".</p> <p>Modifying an existing JDBC store's prefix does not necessarily require a server restart, as described in Section 6.3.3, "Modifying Custom Persistent Store Parameters".</p>
Logical Name	No	<p>Optionally used with WebLogic Server subsystems, like EJBs, when deploying a module to an entire cluster, but also require a different physical store on each server instance in the cluster. In such a configuration, each physical store would have its own name, but all the persistent stores would share the same logical name.</p>
Create Table from DDL File	No	<p>Optionally used with supported DDL (data definition language) files to create the JDBC store's database table (<code>WLStore</code>). This option is ignored when the JDBC store's database table already exists. For more information, see Section 6.5.2.4, "Creating a JDBC Store Table Using Default and Custom DDL Files".</p>

For instructions on configuring a JDBC store using the WebLogic Server Administration Console, see "Create JDBC Stores" in the *Oracle WebLogic Server Administration Console Online Help*.]

6.5.2.3 Supported JDBC Drivers

When using a JDBC store, the backing database can be any database that is accessible through a JDBC driver. WebLogic Server detects some drivers for supported databases.

For each of these databases, there are corresponding DDL (data definition language) files within the `ORACLE_HOME\wlserver\modules\com.bea.core.store.jdbc_1.0.0.0.jar` file, in the `weblogic/store/io/jdbc/ddl` directory, where `ORACLE_HOME` is the top-level installation directory of your WebLogic Server installation.

Table 6–5 Supported JDBC Drivers and Corresponding DDL Files

Database	DDL Files
IBM DB2	db2.ddl db2v6.ddl

Table 6–5 (Cont.) Supported JDBC Drivers and Corresponding DDL Files

Database	DDL Files
Informix	informix.ddl
Microsoft SQL (MSSQL) Server	mssql.ddl
MySQL	mysql.ddl
Oracle	oracle.ddl oracle_blob.ddl oracle_blob_securefile.ddl
Sybase	sysbase.ddl

The DDL files are actually text files containing the SQL commands (terminated by semicolons) that create the JDBC store's database table (`WLStore`). Therefore, if you are using a database that is not included in this list, you can copy and edit any one of the existing DDL files to suit your specific database, as described in [Section 6.5.2.4.1, "Creating a JDBC Store Table Using a Custom DDL File"](#).

6.5.2.4 Creating a JDBC Store Table Using Default and Custom DDL Files

The JDBC Store Configuration page provides an optional Create Table from DDL File option, through which you can access a pre-configured DDL file that is used to create the JDBC store's backing table (`WLStore`). This option is ignored when the JDBC store's backing table already exists. It is mainly used to specify a custom DDL file created for an unsupported database, or when upgrading JMS JDBC store tables from a prior release to a current JDBC Store table.

If a DDL file name is *not* specified in the Create Table from DDL File field, and the JDBC store detects that its backing table does not already exist, the JDBC store automatically creates the table by executing a pre-configured DDL file that is specific to the database vendor (see [Section 6.5.2.3, "Supported JDBC Drivers"](#)).

If a DDL file name is specified in the Create Table from DDL File field, and the JDBC store detects that its backing table does not already exist, the JDBC store searches for the specified DDL file in the file path first, and then, if not found, searches for the DDL file in the `CLASSPATH`. Once found, the SQL within the DDL file is executed to create the JDBC store's backing table. If the configured file is not found and the table doesn't already exist, the JDBC store will fail to boot.

6.5.2.4.1 Creating a JDBC Store Table Using a Custom DDL File To use a different database from those listed in [Section 6.5.2.3, "Supported JDBC Drivers"](#), you can copy and edit any one of the existing DDL template files to suit your specific database.

1. Use the JAR utility supplied with the JDK to extract the DDL files to the `/weblogic/store/io/jdbc/ddl` directory using the following command:

```
jar xf com.bea.core.store.jdbc_1.0.0.0.jar /weblogic/store/io/jdbc/ddl
```

Note: If you omit the `weblogic/store/io/jdbc/ddl` parameter, the entire jar file is extracted.

2. Edit the DDL file for your database. An SQL command can span several lines and is terminated with a semicolon (;). Lines beginning with pound signs (#) are comments.

3. Save your changes and rename the new DDL appropriately (for example, `mydatabase.ddl`)
4. Create a JDBC store, as explained in "Create JDBC Stores" in the *Oracle WebLogic Server Administration Console Online Help*.
5. Use the Create Table from DDL File option on the General Configuration page to specify your custom DDL file (for example, `/mydatabase.ddl`).

Note: On Windows systems, for full path names always include the drive letter.

6.5.2.4.2 Enabling Oracle BLOB Record Columns For Oracle databases, you can use the `oracle_blob.ddl` or `oracle_blob_securefile.ddl` file to create a JDBC store table with a BLOB record column type rather than the default `LONG RAW` record column type. The `oracle_blob.ddl` is used to create Oracle basic file BLOBs and the `oracle_blob_securefile.ddl` file is used to create Oracle secure file BLOBs. Both files types are pre-configured and supplied in the WebLogic CLASSPATH, as described in [Section 6.5.2.3, "Supported JDBC Drivers"](#).

Oracle Database 11g Release 2 includes a zero copy I/O performance enhancement for Secure Files and a logical cache for BLOBs. Use of these enhancements can improve throughput with a JDBC store when message sizes are large and when network connections to the database are slow. The Oracle `LONG RAW` datatype is typically better performing than BLOBs when using a fast connection to the database.

Note: If you need to preserve data already in a Oracle `LONG RAW` column, but still want to switch the column to BLOB, *do not* use this method. Instead, consult the Oracle documentation for the `SQL ALTER TABLE` command.

To use the Oracle BLOB DDL with a JDBC store:

1. Shut down the server instance that uses the JDBC store.
2. Delete the current JDBC table, as explained in [Section 6.5.2.5, "Managing JDBC Store Tables"](#).
3. Reboot the server instance.
4. Create a new JDBC store, as explained in "Create JDBC Stores" in the *Oracle WebLogic Server Administration Console Online Help*.
5. In the Create Table from DDL File field on the General Configuration page, enter the location of:
 - the `oracle_blob.ddl` file as: `/oracle_blob.ddl`
 - the `oracle_blob_securefile.ddl` file as: `/oracle_blob_securefile.ddl`
6. Click **Finish** to create the JDBC store's backing table.

When using Oracle BLOBs, you may be able to improve performance by tuning the `ThreeStepThreshold` value.

When the JDBC store schema contains an Oracle BLOB column (basic file or secure file), the JDBC store populates the BLOB data using one of the following implementations based on the size of the BLOB data:

- The JDBC store inserts a row with BLOB data directly into the store table in a single step. Because only a single step is involved, JDBC batch insert is also adopted and performs best when the data size is small. This implementation is used when the BLOB data to be inserted is less than or equal to the value of the `ThreeStepThreshold`.
- The JDBC store inserts a row with BLOB data into the store table in three steps using the Oracle LOB API. This implementation provides better performance when the data size is large. This implementation is used when the BLOB data to be inserted is greater than the value of the `ThreeStepThreshold`.

The default value of `ThreeStepThreshold` is 200K.

6.5.2.5 Managing JDBC Store Tables

The JDBC `utils.Schema` utility allows you to regenerate a new JDBC store database table (`WLStore`) by deleting the existing version. Running this utility is usually not necessary, since WebLogic Server automatically creates this table for you. However, if your existing JDBC store database table somehow becomes corrupted, you can delete it using the `utils.Schema` utility.

The `utils.Schema` utility is a Java program that takes command-line arguments to specify the following:

- JDBC driver
- Database connection information
- Name of a file containing the SQL Data Definition Language (DDL) commands that create the database table

6.5.2.5.1 Using the `utils.Schema` Utility to Delete a JDBC Store Table Enter the `utils.Schema` command, as follows:

```
$ java utils.Schema url JDBC_driver [options] DDL_file
```

Note: To execute `utils.Schema`, your `CLASSPATH` must contain the `weblogic.jar` file.

Table 6–6 lists the `utils.Schema` command-line arguments.

Table 6–6 Command-line arguments for `utils.Schema`

Argument	Description
<code>url</code>	Database connection URL. This value must be a colon-separated URL as defined by the JDBC specification.
<code>JDBC_driver</code>	Full package name of the JDBC Driver class.
<code>options</code>	Optional command options. If required by the database, you can specify: <ul style="list-style-type: none"> ■ The user name and password as follows: <code>-u <username> -p <password></code> ■ The Domain Name Server (DNS) name of the JDBC database server as follows: <code>-s <dbserver></code> You can also specify the <code>-verbose</code> option, which causes <code>utils.Schema</code> to echo SQL commands as they are executed.

Table 6–6 (Cont.) Command-line arguments for *utils.Schema*

Argument	Description
DDL_file	The full pathname of the DDL text file containing the SQL commands that you want to execute. For more information, see Section 6.5.2.3, "Supported JDBC Drivers" .

For example, the following command deletes a JDBC table named `MYWLStore` in an Oracle server named `DEMO`, with the user name `user1` and password `foobar`:

```
$ echo "drop MYWLStore;" > drop.ddl

$ java utils.Schema
jdbc:weblogic:oracle:DEMO \
weblogic.jdbc.oci.Driver -u user1 -p foobar -verbose \
drop.ddl
```

6.5.2.6 Guidelines for Configuring a JDBC Store

The following sections provide guidelines for using JDBC store prefixes, recommended WebLogic JDBC data source settings for JDBC stores, and handling JMS transactions with JDBC stores.

6.5.2.6.1 Using Prefixes with a JDBC Store The JDBC store database contains a database table, named `WLStore`, that is generated automatically and is used internally by WebLogic Server. The JDBC store provides an optional Prefix Name parameter, which can be used to provide more precise access to the database table.

It is always a best practice to configure a prefix for the JDBC `WLStore` table name, especially when:

- The database requires fully-qualified names. (You should verify this with your database administrator.)
- There is more than one JDBC store instance sharing a database, since no two JDBC stores can share the same table.
- There are many tables in the database. Setting the prefix reduces the number of tables the JDBC store must search through to find its table during boot.

6.5.2.6.2 JDBC Store Table Rules To avoid potential data loss, follow these rules:

- Each JDBC store table name must be unique.
- If multiple JDBC stores share a table, the behavior is undefined and data loss is likely.
- There is no procedure for combining two database tables into a single table.

6.5.2.6.3 Prefix Name Format Guidelines For most databases, the Prefix Name option for the JDBC store's backing database table should be set in the following format for each configured JDBC store, which will result in a valid table name when prepended to the JDBC store table name:

```
[[[catalog.]schema.]prefix]
```

Note that each period in the `[[[catalog.]schema.]prefix]` format is significant. Generally, `catalog` identifies the set of system tables being referenced by the DBMS, and `schema` generally corresponds to ID of the table owner (`username`). When no prefix

is specified, the JDBC store table name is simply `WLStore` and the database implicitly determines the schema according to the current user of the JDBC connection.

For example, in a production database, the database administrator could maintain a unique table for the Sales department, as follows:

```
[[[Production.]JMSAdmin.]Sales]
```

The resulting table will be created in the Production catalog, under the JMSAdmin schema, and will be named `SalesWLStore`.

For some DBMS vendors, such as Oracle, there is no catalog to set or choose, so the format simplifies to `[schema.]prefix`. For more information, refer to your DBMS documentation for instructions on fully-qualified table names, but note that the syntax specified by the DBMS may differ from the format required for this option.

Caution: If the Prefix Name setting is changed, but the `WLStore` database table already exists in the database, take care to preserve existing table data. In this case, the existing database table must be renamed by a database administrator to match the new configured table name.

6.5.2.6.4 Recommended JDBC Data Source Settings for JDBC Stores The following settings are recommended when you use a JDBC data source or multi data source for JDBC stores.

6.5.2.6.5 Automatic Reconnection to Failed Databases WebLogic Server provides robust JDBC data sources that can automatically reconnect to failed databases after they come back online, without requiring you to restart WebLogic Server. To take advantage of this capability, and make your use of JDBC stores more robust, configure the following options on the JDBC data source associated with the JDBC store:

```
TestConnectionsOnReserve="true"  
TestTableName="SYSTABLES"  
ConnectionCreationRetryFrequencySeconds="600"
```

For more information about JDBC default Test Table Names, see "Connection Testing Options for a Data Source" in the *Administering JDBC Data Sources for Oracle WebLogic Server*. For more information about setting the number of database reconnection attempts, see the "Enabling Connection Creation Retries" section in *Administering JDBC Data Sources for Oracle WebLogic Server*.

6.5.2.6.6 Required Setting for Oracle DB2 Type 4 JDBC Drivers For data sources used as a JDBC store that use the Oracle Type 4 JDBC driver for DB2, the `BatchPerformanceWorkaround` property must be set to "true" due to internal JMS batching requirements.

6.5.2.6.7 Handling JMS Transactions with JDBC Stores The JDBC store must use a JDBC data source that uses a non-XA JDBC driver and has **Supports Global Transactions** disabled. This limitation does not remove the XA capabilities of layered subsystems that use JDBC stores. For example, WebLogic JMS is fully XA-capable regardless of whether it uses a file store or any JDBC store.

Because the JDBC store implements the `XAResource` interface, it acts as its own resource manager and handles the transactions above the JDBC driver level. That is, the store itself implements the `XAResource` and handles the transactions without depending on the database (even when the messages are stored in the database).

This means that whenever you are using a JDBC store and a database (even if it is the same database where the JMS messages are stored), then it is two-phase commit transaction.

For more information about using JMS transactions with a JDBC store, see "Using Transactions with WebLogic JMS" in *Developing JMS Applications for Oracle WebLogic Server*.

From a performance perspective, you may also boost your performance as follows:

- Ensure that the JDBC data source used for the database work exists on the same server instance as the JMS destination—the transaction will still be two-phase, but it will be handled with less network overhead.
- Use file stores rather than JDBC stores.
- Configure multiple services to share the same store if they will commonly be invoked within the same transaction.
- If an application directly performs database operations in addition to invoking store services (such as JMS) within the same transaction, consider using a JDBC data source with Logging Last Resource (LLR) enabled for the database operations.

With the LLR optimization, the transaction will follow the two-phase commit protocol, but the database operations will be handled in a single local transaction, which may improve overall transaction performance. For more information on using the LLR optimization, see "Understanding the Logging Last Resource Transaction Option" in *Administering JDBC Data Sources for Oracle WebLogic Server*.

6.5.2.7 Enabling I/O Multithreading for JDBC Stores

Under heavy JDBC store I/O loads, you can improve performance by configuring a JDBC store to use multiple JDBC connections to concurrently process I/O operations.

Note: Enabling I/O multithreading under light loads may actually reduce performance. Oracle recommends that you tune your applications appropriately.

To enable I/O multithreading, set the `Worker Count` attribute to an integer value greater than 1. The default value of this configuration property is 1 and disables this option. The `Worker Count` attribute specifies the number of worker threads the JDBC store uses to process store I/O. Each worker thread acquires one JDBC connection from the configured data source pool when the store is opened. In many cases, benefits of multithreading tends to decrease after 4 concurrent threads. When using a slow connection to the database, multithreading may not improve performance.

Note: If you set the `Worker Count` to a value where there are not enough connections available in the connection pool, the JDBC store will fail to open.

You can tune the workload for each worker thread by changing the value of the `Worker Preferred Batch Size` attribute. Increasing the value of this attribute incrementally increases the workload assigned to each worker thread. The workload consists of store I/O requests, which are grouped and pushed to each JDBC worker thread for processing. If the size of individual I/O requests is commonly very large (for example,

requests to store 1 MB JMS messages), then tune the value of `Worker Preferred Batch Size` to a smaller value for better performance.

6.5.2.7.1 Rebuilding the Store Table Index for an Oracle Database When I/O multithreading is enabled, multiple JDBC connections are used to concurrently process store I/O operations which can result in database resource contention. To reduce contention on Oracle databases, Oracle recommends rebuilding the primary key index into a reverse key index when I/O multithreading is used. If you use and then disable I/O multithreading, Oracle recommends rebuilding the primary key index as a non-reverse index. For more information on reverse key indexes, see "Indexes and Index-Organized Tables" in *Oracle Database Concepts*.

Use the following basic steps to rebuild the Store table index for Oracle database:

1. Login to the Oracle database under the Store schema name. The Store schema name may or may not be the same as the data source user name.
2. Use the PL/SQL script found in [Section 6.5.2.7.2, "Build a Reverse Index for an Oracle Database"](#) or [Section 6.5.2.7.3, "Build a Non-Reverse Index for an Oracle Database"](#) to rebuild the Store table index as needed. Replace `<Store Table Name>` in each script with the Store table name as described in [Section 6.5.2.6.1, "Using Prefixes with a JDBC Store."](#) For more information on PL/SQL, see "Execution of PL/SQL Subprograms" in *Oracle Database Concepts*.

Note: Oracle recommends reverse indexes for I/O multithreading and non-reverse indexes for single threaded I/O.

6.5.2.7.2 Build a Reverse Index for an Oracle Database To rebuild the Store table index as a reverse index for an Oracle database, run the following PL/SQL block as the store database user:

```
declare
  idx          user_ind_columns.index_name%TYPE;
  alter_stmt   VARCHAR2(200);
begin
  select index_name into idx from user_ind_columns where table_name =
  <Store Table Name> and column_name = 'ID';
  alter_stmt := 'alter index ' || idx || ' rebuild reverse';
  dbms_output.put_line(alter_stmt);
  execute immediate alter_stmt;
end;
/
```

6.5.2.7.3 Build a Non-Reverse Index for an Oracle Database To rebuild a reverse Store table index as a non-reverse index for Oracle database, run the following PL/SQL block as the store database user:

```
declare
  idx          user_ind_columns.index_name%TYPE;
  alter_stmt   VARCHAR2(200);
begin
  select index_name into idx from user_ind_columns where table_name =
  <Store Table Name> and column_name = 'ID';
  alter_stmt := 'alter index ' || idx || ' rebuild noreverse';
  dbms_output.put_line(alter_stmt);
  execute immediate alter_stmt;
end;
```

6.5.2.7.4 Reducing Contention in a Non-Oracle Database For non-Oracle databases, refer to the database provider's documentation on how to reduce the contention.

6.6 Monitoring a Persistent Store

You can monitor statistics for each existing persistent store and for each open store connection.

6.6.1 Monitoring Stores

Each persistent store is represented at run time by an instance of the `PersistentStoreRuntimeMBean`, which provides the following options.

Table 6–7 Persistent Store Run-time Options

Option	What It Does
CreateCount	Number of create requests issued to this persistent store.
ReadCount	Number of read requests issued to this persistent store.
UpdateCount	Number of update requests issued by this persistent store.
DeleteCount	Number of delete requests issued by this persistent store.
ObjectCount	Number of objects contained in the persistent store.
Connections	Number of active connections in the persistent store.
PhysicalWriteCount	Number of times the persistent store flushes its data to durable storage.

6.6.2 Monitoring Store Connections

For each open persistent store connection, the persistent store also registers a `PersistentStoreConnectionRuntimeMBean`, which provides the following options.

Table 6–8 Persistent Store Connection Runtime Options

Option	What It Does
CreateCount	Number of create requests issued to this connection.
ReadCount	Number of read requests issued to this connection.
UpdateCount	Number of update requests issued by this connection.
DeleteCount	Number of delete requests issued by this connection.
ObjectCount	Number of objects contained in the connection.

[Table 6–9](#) defines most of the run-time prefix names of the WebLogic services and subsystems that can create a connection to the persistent store.

Table 6–9 Persistent Store Run-Time Prefix Names

Subsystem/Service	Run-Time Prefix Name
Deployment	<code>weblogic.deploy.internal</code> where <i>internal</i> is the name of the deployment connection

Table 6–9 (Cont.) Persistent Store Run-Time Prefix Names

Subsystem/Service	Run-Time Prefix Name
Diagnostic Service	<code>weblogic.diagnostics.internal</code> where <i>internal</i> is the logical name of the diagnostic archive connection
EJB Timer Services	<code>weblogic.ejb.timer.internal</code> where <i>internal</i> uniquely identifies EJB deployments in a server instance
JMS Service	JMS server: <code>weblogic.messaging.jmsServer.internal</code> where <i>internal</i> is the name of the JMS server connection JMS durable subscriber: <code>weblogic.messaging.jmsServer.durablesubs.internal</code> where <i>internal</i> is the name of the durable subscriber connection
JTA Transaction Log (TLOG)	<code>weblogic.transaction.internal</code> where <i>internal</i> is the name of the TLOG connection
Path Service	<code>weblogic.messaging.PathService.internal</code> where <i>internal</i> is the name of the path service connection
SAF Service	SAF agent <code>weblogic.messaging.SAFAgent@server1.internal</code> where <i>internal</i> is the name of the SAF agent's connection SAF durable subscriber: <code>weblogic.messaging.SAFAgent@server1.durablesubs.internal</code> where <i>internal</i> is the name of the durable subscriber connection
Web Services	<code>weblogic.wsee.server.store.internal</code> where <i>internal</i> is the name of the Web Service's connection

6.7 Administering a Persistent Store

The WebLogic Store administration utility enables administrators to troubleshoot a WebLogic persistent store. The store utility operates only on a store that is not currently opened by a running server instance. This utility can be run from a Java command line or from WebLogic Scripting Tool (WLST), as described in [Section 6.7.1, "Store Administration Using a Java Command Line"](#) and [Section 6.7.2, "Store Administration Using WLST"](#).

The most common uses-cases for store administration are for compacting a file store to reduce its size and for dumping the contents of a file store or JDBC store to an XML file for troubleshooting purposes. Examples of these use cases are provided later in this section.

[Table 6–10](#) defines the available store administration commands for Java and WLST.

Table 6–10 Persistent Store Administration Options

Java Command	WLST Method	What It Does
<code>help</code>	<code>helpstore</code>	Displays available commands, usage, and examples.

Table 6–10 (Cont.) Persistent Store Administration Options

Java Command	WLST Method	What It Does
compact	compactstore	Compacts and defragments the space occupied by a file store. This command only works offline and does not work for JDBC stores. Note: Compacting a file store is usually not necessary if you know that file store will likely grow to the current size again. File stores automatically re-use space freed by deleted records and expand only when there is insufficient internal space for new records. Also, file stores do not normally become fragmented as most persistent records are short-lived.
openfile	openfilestore	Opens an existing file store for further operations. If a file store does not exist, a new one is created in an open state using the <code>-create</code> parameter.
openjdbc	openjdbcstore	Opens an existing JDBC store for further operations. If a JDBC store does not exist, a new one is created in an open state
dump	dumpstore	Dumps store or connection contents in a human-readable format to user-specified XML file. The XML file format is the same format used by the diagnostic image of the persistent store.
list	liststore	Lists store names, open stores, or connections in a store.
n/a	getstoreconns	Returns a list of connections in the specified store (for script access)
n/a	getopenstores	Returns a list of opened stores (for script access).
opts	n/a	Lists invocation options for the store administration tool.
verbose	n/a	Controls display of additional information, such as stack traces.
close	closestore	Closes a previously opened store.
quit	exit	Ends the store administration session.

A persistent store can be backed by the file system (file store) or by a JDBC-capable database (JDBC store). Except for the `openfile/openfilestore()` and `openjdbc/openjdbcstore()` options, there is no difference in the options to operate on these two different types of stores.

Most commands and methods work in terms of store names, while others also work in terms of connection names. Store connections are logical groups of records within persistent stores. For example, the JMS and JTA subsystems persist their respective records in different connections in the same file store.

6.7.1 Store Administration Using a Java Command Line

To open the persistent store administration utility from a Java command line, type the following:

```
> java weblogic.store.Admin
> storeadmin->
```

6.7.1.1 Accessing Store Administration Help

Type `help` for detailed descriptions on available store administration commands, as well as examples of typical command usage. For example, the following comprehensive help is provided for the `list` command, which lists store names, open stores, or connections in a store.

```
storeadmin->help list
  Command:
    list
  Description:
    lists store names, open stores, or connections in a store
  Usage:
    list [-store storename|-dir dir]
  Examples:
    list #lists all opened stores by storename
    list -store store1 #lists all connections in store1
    list -dir dir1 #lists all storenames found in dir1
```

6.7.1.2 Dumping the Contents of a File Store

Here's an example of using a series of store administration commands to ultimately export the contents of a file store named `myfilestore` into a human-readable XML file format in a temporary directory. This does not include store connection names or the actual record contents, which require the optional `-conn` and `-deep` parameters.

```
> storeadmin-> list -dir .
> storeadmin-> openfile -store myfilestore -dir .
> storeadmin-> dump -store myfilestore -out d:\tmp\filestore1-out
> storeadmin-> close -store myfilestore
```

The `list` command shows all the store names in the current directory. The `openfile` and `openjdbc` commands must be used to open and/or create a file or JDBC store first before calling certain administration functions, like `dump` and `list` (only when listing open stores). After administering an open store, you must close it using the `close` command.

6.7.1.3 Compacting a File Store

Here is an example of using the `compact` command to compact the space occupied by a file store in the `mystores` directory.

```
> storeadmin->compact -dir c:\mystores -tempdir c:\tmp
```

Since the `compact` command can only be used on an unopened file store, none of the stores that have files in the source `-dir` directory should be open. Also, the temporary `-tempdir` directory should have at least enough extra space as the source directory and should also not be under the source directory. When `compact` successfully completes, the newly compacted store files will be in the `mystores` directory. In addition, a new, uniquely-named directory will be created under `tmp` containing the original uncompact store files.

6.7.2 Store Administration Using WLST

The WLST interface has a couple of additional methods (compared to the Java command line) such as `getopenstores` and `getstoreconns`, that return relevant Java objects and can be used for scripting in WLST.

Note: In this release, `ThreeStepThreshold`, `Worker Count`, and `Worker Preferred Batch Size` are not supported when using the WebLogic Scripting Tool (WLST) offline.

6.7.2.1 Accessing Store Administration Help

To access the persistent store administration utility from WLST, type the following command:

```
> java weblogic.WLST
```

Type `helpstore()` for detailed descriptions on available store administration commands, as well as examples of typical command usage. For example, the following help is provided for the `list` command, which lists store names, open stores, or connections in a store.

```
> wls:/offline> helpstore(liststore)
lists storenames, opened stores, or connections (for interactive access)
Parameters store and dir cannot both be specified concurrently.
```

```
Usage: liststore(store='null',dir='null')
```

```
@param store [optional] a previously opened JDBC or File store's name.
```

```
    If store is specified, all connections in the store are listed.
```

```
@param dir [optional] directory for which to list available store names
```

```
    If dir is specified, all store names in the directory are listed.
```

```
If neither store nor dir are specified, all open store names are listed.
```

```
@return 1 on success, 0 on failure
```

Note that the parameters with an equal sign "=" are optional. For example, the `compactstore` method can be invoked as either `compactstore(dir='storename', tempdir='/tmp')` or `compactstore(store='storename')`, where `tempdir` takes the default value. Default values for optional parameters are listed in the command-specific help.

6.7.2.2 Dumping the Contents of a JDBC Store Using WLST

Here is an example of using the `dumpstore` method (`store`, `outfile`, `conn='null'`, `deep='false'`) to export the contents of a JDBC store named `myJDBCStore` in a human-readable XML file format out to a file named `mystoredump-out.xml`. This does not include store connection names or the actual record contents, which require the optional `conn` and `deep` parameters.

```
> wls:/offline>
openjdbcstore('myJDBCStore', 'oracle.jdbc.OracleDriver',
'jdbc:oracle:thin:@test2k31:1521:test120a', './wlstoreadmin-dump.props',
'jmstest', 'jmstest', '', 'jdbcstoreprefix')
dumpstore('myJDBCStore', 'mystoredump-out')
closestore('myJDBCStore')
```

The `openjdbcstore` and `openfilestore` methods must be used to open and/or create a store first before calling certain administration functions, like `dumpstore` and `liststore` (only when listing open stores). After administering an open store, you must close it using the `closestore` method.

6.7.2.3 Compacting a File Store Using WLST

Here is an example of a WLST script that uses the `compactstore` method (`dir, tempdir= 'null'`) to compact the space occupied by a file store files in the `mystores` directory.

```
> wls:/offline> compactstore('c:\mystores', 'c:\tmpmystore.dir')
```

Since the `compactstore()` method can only be used on unopened file stores, none of the stores that have files in the source `'dir'` directory should be open. Also, the temporary `'tempdir'` directory should have at least enough extra space as the source directory and should also not be under the source directory. When compact successfully completes, the newly compacted store files will be in the `mystores` directory. In addition, a new, uniquely-named directory will be created under `tmpmystore` containing the original uncompact store files.

6.8 Security Considerations

In order to properly secure file store data, you must set appropriate directory permissions on all your file store directories. If you require data encryption, you must use appropriate third-party encryption software.

6.9 Limitations and Considerations of the Persistent Store

The following limitations apply to the persistent store:

- A persistent file store should not be opened simultaneously by two server instances; otherwise, there is no guarantee that the data in the file will not be corrupted. If possible, the persistent store will attempt to return an error in this case, but it will not be possible to detect this condition in every case. It is the responsibility of the administrator to ensure that the persistent store is being used in an environment in which multiple servers will not try to access the same store at the same time. (Two file stores are considered the "same store" if they have the same name and the same directory.)
- Two JDBC stores must not share the same database table, because this will result in data corruption.
- A persistent store may not survive arbitrary corruption. If the disk file is overwritten with arbitrary data, then the results are undefined. The store may return inconsistent data in this case, or even fail to recover at all.
- A file store may return exceptions when its disk is full. However, it will resume normal operation by no longer throwing an exception when disk space has been made available. Also, the data in the persistent store must remain intact as described in the previous points.
- When using MySQL as the backing database for a JDBC store, Oracle recommends using the InnoDB engine because it provides safe writes. If the MyISAM engine is used, data may be lost in some cases.

Using the Plug-in for Oracle Virtual Assembly Builder

This chapter describes the WebLogic Server introspection plug-in for Oracle Virtual Assembly Builder, which examines a single domain in WebLogic Server 12.1.3. The domain specified and the Oracle home are captured.

Note: This chapter describes the WebLogic Server plug-in that is included with the WebLogic Server installation program. Older versions of this plug-in are included with the Oracle Virtual Assembly Builder installation program. For information about using an older plug-in, see *Release Notes for Oracle Virtual Assembly Builder*.

This chapter includes the following sections:

- [Overview](#)
- [Introspection Plug-in Parameters](#)
- [Reference System Prerequisites](#)
- [Plug-in Usage Requirements](#)
- [Resulting Artifact Type](#)
- [Wiring](#)
- [Wiring Properties](#)
- [WebLogic Server Appliance Properties](#)
- [Oracle Coherence*Web Introspection Extension](#)
- [Supported Template Types](#)

7.1 Overview

The following topics provide a brief overview of Oracle Virtual Assembly Builder and the WebLogic Server introspection plug-in:

- [About Oracle Virtual Assembly Builder](#)
- [About the WebLogic Server Introspection Plug-in](#)

7.1.1 About Oracle Virtual Assembly Builder

Oracle Virtual Assembly Builder is a tool for virtualizing installed Oracle components, modifying those components, and then deploying them into your own environment. Using Oracle Virtual Assembly Builder, you capture the configuration of existing software components in artifacts called software appliances. Appliances can then be grouped, and their relationships defined into artifacts called software assemblies which provide a blueprint describing a complete multi-tier application topology.

Oracle Virtual Assembly Builder allows the logical connections between appliances within an assembly to be reconfigured by a process known as assembly editing. When a desired assembly configuration has been achieved, you use Oracle Virtual Assembly Builder to prepare the assembly for deployment and then deploy it into your environment.

For more information, see *Using Oracle Virtual Assembly Builder*.

7.1.2 About the WebLogic Server Introspection Plug-in

Introspection is an operation performed on a software component or a group of related components to create an appliance or assembly. During introspection, Oracle Virtual Assembly Builder creates an XML description of the component and captures a component-specific set of configuration files. This information forms a snapshot of the component's configuration at the time of introspection.

The introspection architecture is plug-in based. A plug-in is available for each supported component type, including a plug-in for WebLogic Server. Typically the result of introspecting a component is an appliance. However, when you use Oracle Virtual Assembly Builder to introspect a WebLogic Server domain, the introspection plug-in generates an assembly. The generated assembly contains:

- An appliance representing the Administration Server.
- One appliance for each standalone Managed Server.
- One appliance for each cluster. (This appliance encompasses all Managed Server instances contained in that cluster.)

The WebLogic Server introspection plug-in fully supports WebLogic domains that are configured with dynamic clusters and dynamic servers.

7.1.3 Setting Up the WebLogic Server Introspection Plug-in

The following steps summarize the process of making the WebLogic Server introspection plug-in available for us with Oracle Virtual Assembly Builder:

1. Install the Oracle Fusion Middleware Infrastructure software (`wls_jrf_generic.jar`), as described in *Planning an Installation of Oracle Fusion Middleware*.
For installation type, select **Fusion Middleware Infrastructure**.
2. Install Oracle Virtual Assembly Builder, as described in *Installing Oracle Virtual Assembly Builder*.
3. Install the WebLogic Server plug-in for Oracle Virtual Assembly Builder using the following command:

```
$AB_INSTANCE/bin/abctl installPlugins -productRoot $ORACLE_HOME/wlserver
```

In the preceding command, `$AB_INSTANCE` is the environment variable set to the Studio instance home directory.

7.2 Introspection Plug-in Parameters

Table 7–1 lists the introspection parameters for the WebLogic Server plug-in for Oracle Virtual Assembly Builder. The specific set of parameters that are available for use depend on whether you are using the `abctl` command or Oracle Virtual Assembly Builder Studio. The column labeled Command-Line Support indicates whether the corresponding parameter may be specified using the `abctl` command. The column labeled Required indicates whether the parameter must be specified for the plug-in.

Table 7–1 WebLogic Server Plug-in Introspection Parameters

Parameter	Description	Command-Line Support	Required
<code>domainRoot</code>	The fully qualified path to the domain you want to introspect. This should be the directory that contains the <code>config</code> directory.	Yes	Yes
<code>wlsHome</code>	The fully qualified path to the WebLogic home directory. For example, <code>/u01/Oracle/Middleware/Oracle_Home/wlserver</code> .	Yes	Yes
<code>adminUser</code>	The administrative user for the WebLogic domain.	Yes	Yes
<code>adminPassword</code>	The password for the administrative user corresponding to the <code>adminUser</code> parameter.	No	Yes
<code>nodeManagerUser</code>	The Node Manager user name. If you specify this parameter during introspection time, the specified user name is used when the domain is rehydrated. If you specify this option in the <code>abctl</code> command, you are automatically prompted for the <code>nodeManagerPassword</code> , which is used during dehydration. If you do not specify this parameter, the value of the <code>adminUser</code> parameter is used by default as the Node Manager user.	Yes	No
<code>nodeManagerPassword</code>	The password for the user specified with the <code>nodeManagerUser</code> parameter.	No	Yes ¹
<code>nodeManagerDir</code>	The non-default directory that contains the Node Manager configuration. If this parameter is not specified, the following value is used by default: <code>\$WL_HOME/common/nodemanager</code>	Yes	No
<code>identityKeystorePassPhrase</code>	The password for the identity keystore that is used if the domain is configured to have a custom identity keystore.	No	No
<code>privateKeyPassPhrase</code>	The password for the private key contained in the custom identity keystore.	No	No

¹ Required only if the Node Manager user name is specified.

7.3 Reference System Prerequisites

The Administration Server for the domain must be running and introspection must target the Administration Server host machine.

7.4 Plug-in Usage Requirements

Table 7–2 lists and describes the requirements for using the introspection plug-in. Some requirements listed in this table pertain to the WebLogic Server plug-in, while

others pertain overall to Oracle Virtual Assembly Builder. For more information, see *Installing Oracle Virtual Assembly Builder*.

Table 7–2 Requirements for Using the Introspection Plug-in

Requirement	Description
Editable domain	You must ensure that any WebLogic domain being introspected is configured to be editable. This allows edits to be performed successfully during deployment. For more information, see <i>Understanding Domain Configuration for Oracle WebLogic Server</i> .
Oracle Internet Directory	Oracle recommends the use of Oracle Internet Directory as the LDAP identity store on the reference system, not a file-based LDAP server. File-based LDAP cannot work properly in a deployed system due to synchronization issues.
Remote user	The remote user specified for remote introspection of WebLogic Server must be able to access files created by the user who owns the WebLogic Server process. When possible Oracle recommends that the remote user specified be the same as the user who owns the WebLogic Server process.
Applications accessing Web services	For each application that accesses a Web service hosted on the WebLogic Server reference system, you must: <ol style="list-style-type: none"> 1. Update the application to access the Web service WSDL on the new Oracle VM host. 2. Redeploy the application to the Oracle VM WebLogic Server environment, using any of the administration tools listed in "Summary of System Administration Tools and APIs" in <i>Understanding Oracle WebLogic Server</i>.
Avoid creating templates on individual servers	You must not create a template on an individual server in WebLogic Server. Such templates cannot be deployed because they lack certain deployment artifacts (the domain template jar in content, and data at the assembly level).
Administration Server URL	If you want to perform manual starts from the context of the Guest-OS, you must manually modify the <code>StartManagedServer</code> script to provide the correct Administration Server URL (Administration Server host name). This is required to provide the default Administration Server URL the correct value (the machine name of the Administration Server is not known at the time of template creation).

7.5 Resulting Artifact Type

An atomic assembly that contains an appliance for the Administration Server and appliances for any clusters found and any standalone (that is, non-clustered) Managed Servers found. One appliance is created for a cluster regardless of the number of Managed Servers in that cluster. The WebLogic Server plug-in assumes that every Managed Server in a cluster is configured identically. The names and the number of Managed Servers in the cluster are saved as 'scale out' information in the appliance metadata.

Note: An atomic assembly cannot be edited to add or remove appliances. To wire other appliances to an atomic WebLogic Server assembly a non-atomic assembly must be created and the WebLogic Server assembly must be added to the non-atomic assembly.

7.6 Wiring

Inputs are created on the WebLogic Server assembly for all the channels the servers in the domain are listening on. Typically Oracle HTTP Server outputs would be connected to the WebLogic Server inputs.

Outputs are created on the WebLogic Server assembly for the following types of configuration that are found:

- JDBC
- LDAP
- JMS messaging bridges
- Foreign JMS

These outputs must all be connected to either an external resource or to an appliance before deployment. The description on the output and the protocol supported by the output provide hints about the type of appliance to connect the output to.

7.7 Wiring Properties

All input endpoints have the following wiring properties:

- Two editable properties: port and description.
- One non-editable property: a list of protocols. The protocols indicate the type of outputs that can be connected to the input.

All output endpoints have the following wiring properties:

- One editable property: description.
- The following non-editable properties:
 - protocol — Indicates the type of input that can be connected to the output.
 - singleton — Indicates the type of appliance that the output can be connected to. If singleton is true, the output can be connected only to an input on an appliance that has a scalability absolute max value of 1.

[Table 7–3](#) describes common WebLogic Server appliance input user properties:

Table 7–3 Common WebLogic Server Appliance Input User Properties

Name	Type	Req'd	Default	Description
keepLocalHost	Boolean	false	none	If this input was originally bound to localhost explicitly, this property will exist and be set to true. Connections should not be made to this input if this property exists and its value is not overridden to false.
readymetric-naming-password	String	false	none	The password to use for the connection made to the server when doing the ready metric check.

Table 7-3 (Cont.) Common WebLogic Server Appliance Input User Properties

Name	Type	Req'd	Default	Description
readymetric-naming-protocol	String	false	none	Optional protocol you can specify for naming connections used for the ready metric check (for example, iiop).
readymetric-naming-user	String	false	none	The user to use for the connection made to the server when doing the ready metric check (for example, weblogic).
readymetric-server-protocol	String	false	none	The protocol to use for the connection made to the server when doing the ready metric check (for example, iiop).

[Table 7-4](#) describes common WebLogic Server appliance input properties. These are not user-modifiable properties.

Table 7-4 Common WebLogic Server Appliance Input Properties

Name	Type	Req'd	Default	Description
originalBindAddresses	String	false	none	The original address of the system that was introspected.
originalDefaultHostname	String	false	none	The original host name of the system that was introspected. (For example, example.com).

[Table 7-5](#) describes the Administration Server appliance input system properties:

Table 7-5 Administration Server Appliance Input Properties

Name	Type	Req'd	Default	Description
admin-password	String	true	none	The admin user's password. This is a user-modified property, specified in the deployment plan, that allows a deployer to introspect a well-known domain and deploy it multiple times, each deployed instance having a custom password.
admin-username	String	true	none	The admin user name for connecting to the Administration Server (for example, weblogic).

[Table 7-6](#) through [Table 7-9](#) describe Administration Server appliance output user properties for JDBC, foreign JMS, JMS messaging bridge, and LDAP.

[Table 7-6](#) describes Administration Server appliance output user and system properties for JDBC. The password and username properties are user properties, and original-url is a system property.

Table 7-6 Administration Server Appliance Output Properties: JDBC

Name	Type	Req'd	Default	Description
password	String	false	<empty>	The password for the user needed for the data source connection.
username	String	false	none	The user needed for the data source connection. The value will be the original user for the data source connection.
original-url	String	false	none	The original JDBC URL from the introspected WebLogic domain. For example, jdbc:oracle:thin:@adc2100927.example.com:1521:orc1.

[Table 7-7](#) describes Administration Server appliance output user properties for foreign JMS:

Table 7-7 Administration Server Appliance Output Properties: Foreign JMS

Name	Type	Req'd	Default	Description
original-connection-url	String	false	none	The original URL for the foreign JMS server.

[Table 7-8](#) describes Administration Server appliance output system properties for JMS messaging bridge:

Table 7-8 Administration Server Appliance Output Properties: JMS Messaging Bridge

Name	Type	Req'd	Default	Description
original-url	String	false	none	The original URL for the JMS messaging bridge server.
original-username	String	false	none	The original username for the JMS messaging bridge server.
original-password	String	false	none	The original password for the JMS messaging bridge server, encrypted.

[Table 7-9](#) describes Administration Server appliance output system properties for LDAP:

Table 7-9 Administration Server Appliance Output Properties: LDAP

Name	Type	Req'd	Default	Description
original-name	String	false	none	The original name for the LDAP security provider.
original-host	String	false	none	The original host for the LDAP security provider.
original-port	String	false	none	The original port for the LDAP security provider.
original-user	String	false	none	The original user for the LDAP security provider.

7.8 WebLogic Server Appliance Properties

This section discusses the following properties for assemblies with a WebLogic Server appliance. Those properties include assembly-level properties, properties on the inputs and outputs of each application, and properties of the appliances themselves. This section contains the following subsections:

- [Assembly-Level System Properties](#)
- [Properties Common to Administration Server and Managed Server Appliances](#)
- [Administration Server Appliance Properties](#)

7.8.1 Assembly-Level System Properties

[Table 7-10](#) describes assembly-level system properties:

Table 7–10 Assembly-level System Properties

Name	Type	Req'd	Default	Description
admin-password	String	true	none	The admin user password for the domain.
admin-username	String	true	none	The admin user for the domain (for example, weblogic).
admsvr-jmx-input	String	true	none	Indicates what input on the Administration Server appliance should be used when making JMX connections (for example, Default).
admsvr-jmx-protocol	String	true	none	The protocol to use when making a JMX connection to the Administration Server (for example, iiop).
domain-name	String	false	none	The domain name of the domain that was introspected (for example, test_domain).
usesOracleHomes	boolean	true	none	Indicates that this is not a core WebLogic Server installation and as such has an Oracle home associated with it. This will be true for SOA and WebCenter domains. Allowable values are true and false.

7.8.2 Properties Common to Administration Server and Managed Server Appliances

The following information describes properties common to Administration Server and Managed Server appliances.

[Table 7–11](#) describes common WebLogic Server appliance user properties:

Table 7–11 Common WebLogic Server Appliance User Properties

Name	Type	Req'd	Default	Description
NodeManagerPort	Integer	true	5556	The port the Node Manager should listen on (for example, 5556). This will only be present if the Node Manager was found to be configured on the reference system.
readymetric-attribute-compare-type	String	false	EQUALS	The comparison to make between the readymetric-attribute's value and the value specified for the property readymetric-attribute-value. Valid values are EQUALS, LESSER_THAN, GREATER_THAN, LESSER_THAN_OR_EQUAL, and GREATER_THAN_OR_EQUAL.
readymetric-attribute-name	String	false	State	The MBean attribute to check.
readymetric-attribute-type	String	false	STRING	The type of the MBean attribute. Valid values (but specific to the attribute being examined) are STRING, INTEGER, SHORT, LONG, DOUBLE, FLOAT, and BOOLEAN.
readymetric-attribute-value	String	false	RUNNING	The value the property readymetric-attribute-name must have for the check to be considered successful.
readymetric-instance-name-0	String	false	com.bea:Name=AdminServer,Type=ServerRuntime	The instance name to use for the JMX ready metric check.
readymetric-max-wait-period	String	false	600	The maximum time in seconds to wait for a successful ready metric check.

Table 7–11 (Cont.) Common WebLogic Server Appliance User Properties

Name	Type	Req'd	Default	Description
readymetric-naming-input	String	false	none	The input to use for the ready metric check (for example, "Default").
readymetric-polling-period	String	false	none	The time between connection attempts, in seconds, for the ready metric check.
readymetric-server-input	String	false	none	The input to use for the ready metric check (for example, "Default").
readymetric-trust-store-0	String	false	none	The location of the trust store to use if the ready metric check is using an SSL enabled port.
readymetric-type	String	false	JMX	The type of ready metric to use for the appliance.
readymetric-verify	String	false	true	If this property is set to true the ready metric check will be performed. Otherwise it will be skipped.
useTemplate	String	false	OEL	Specifies the template type to use by default when creating a template for the appliance.

7.8.3 Administration Server Appliance Properties

Table 7–12 describes Administration Server appliance system properties:

Table 7–12 Administration Server Appliance System Properties

Name	Type	Req'd	Default	Description
admin-input-name	String	false	none	The input to use for connecting to the Administration Server admin-input-protocol (for example, "Default").
admin-input-protocol	String	false	none	The protocol to use for connecting to the Administration Server (for example, "http").

Table 7–13 describes Administration Server appliance user properties:

Table 7–13 Administration Server Appliance User Properties

Name	Type	Req'd	Default	Description
<cluster name>-cluster-address	String	false	<empty>	The cluster address for the cluster named by the first part of the property name.
<cluster name>-frontend-host	String	false	<empty>	The front-end host for the cluster named by the first part of the property name.
<cluster name>-frontend-http-port	String	false	<empty>	The non-secure front-end port for the cluster named by the first part of the property name.
<cluster name>-frontend-https-port	String	false	<empty>	The secure front-end port for the cluster named by the first part of the property name.

7.9 Oracle Coherence*Web Introspection Extension

The Oracle Coherence*Web introspection extension extends the functionality of the WebLogic Server Introspector. It examines the configuration of Coherence cache

clusters and servers configured as part of a WebLogic domain. For more information, see the Oracle Coherence documentation.

For more information, see

7.10 Supported Template Types

The supported template type is Oracle Enterprise Linux (OEL).