

Oracle® Fusion Middleware

Performance and Tuning Guide

11g Release 1 (11.1.1.7.0)

E10108-14

November 2014

Describes how to monitor and optimize performance, configure components for optimal performance in the Oracle Fusion Middleware environment.

Copyright © 2013, 2014 Oracle and/or its affiliates. All rights reserved.

Primary Authors: Lisa Jamen, Sreetama Ghosh, Hannah Cheng

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

The information contained in this document is for informational sharing purposes only and should be considered in your capacity as a customer advisory board member or pursuant to your beta trial agreement only. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle.

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle Software License and Service Agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced, or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

Contents

Preface	xvii
Audience	xvii
Documentation Accessibility	xvii
Conventions	xvii

Part I Introduction

1 Introduction and Roadmap

1.1 Document Scope and Audience	1-1
1.2 Guide to this Document	1-1
1.3 Related Documentation	1-4

2 Top Performance Areas

2.1 About Identifying Top Performance Areas	2-1
2.2 Securing Sufficient Hardware Resources	2-2
2.3 Tuning the Operating System	2-3
2.4 Tuning Java Virtual Machines (JVMs)	2-3
2.4.1 Configuring Garbage Collection	2-4
2.4.2 Logging Low Memory Conditions	2-7
2.4.3 Monitoring and Profiling the JVM	2-8
2.5 Tuning the WebLogic Server	2-8
2.6 Tuning Database Parameters	2-8
2.6.1 Tuning Database Parameters	2-8
2.6.2 Tuning Redo Logs Location and Sizing	2-10
2.6.3 Tuning Automatic Segment-Space Management (ASSM)	2-10
2.7 Reusing Database Connections	2-11
2.8 Enabling Data Source Statement Caching	2-11
2.9 Controlling Concurrency	2-12
2.9.1 Setting Server Connection Limits	2-12
2.9.2 Configuring Connection Pools	2-14
2.9.3 Tuning the WebLogic Server Thread Pool	2-14
2.9.4 Tuning Oracle WebCenter Concurrency	2-16
2.9.5 Tuning BPEL Concurrency	2-16
2.10 Setting Logging Levels	2-16

3 Performance Planning

3.1	About Oracle Fusion Middleware Performance Planning	3-1
3.2	Performance Planning Methodology	3-1
3.2.1	Define Your Performance Objectives.....	3-1
3.2.2	Design Applications for Performance and Scalability	3-4
3.2.3	Monitor and Measure Your Performance Metrics	3-4

4 Monitoring Oracle Fusion Middleware

4.1	About Oracle Fusion Middleware Management Tools.....	4-1
4.1.1	Measuring Your Performance Metrics.....	4-2
4.2	Oracle Enterprise Manager 11g Fusion Middleware Control	4-2
4.2.1	Viewing Performance Metrics Using Fusion Middleware Control.....	4-3
4.3	Oracle WebLogic Server Administration Console	4-4
4.4	WebLogic Diagnostics Framework (WLDF).....	4-5
4.5	WebLogic Scripting Tool (WLST).....	4-6
4.5.1	Using Custom WLST Commands	4-6
4.6	DMS Spy Servlet.....	4-7
4.6.1	Viewing Performance Metrics Using the Spy Servlet	4-7
4.6.2	Using the DMS Spy Servlet	4-7
4.7	Oracle Process Manager and Notification Server	4-9
4.8	Oracle Enterprise Manager Cloud Control	4-9
4.9	Native Operating System Performance Commands.....	4-10
4.10	Network Performance Monitoring Tools	4-10

Part II Core Components

5 Understanding the Oracle Dynamic Monitoring Service

5.1	About Dynamic Monitoring Service (DMS)	5-1
5.1.1	Understanding Common DMS Terms and Concepts	5-1
5.2	Understanding DMS Availability.....	5-6
5.3	Understanding DMS Architecture	5-6
5.4	Viewing DMS Metrics	5-7
5.4.1	Viewing Metrics Using the Spy Servlet.....	5-7
5.4.2	Viewing Metrics with WLDF (WebLogic Diagnostic Framework).....	5-8
5.4.3	Viewing metrics with WLST (Oracle WebLogic Server).....	5-8
5.4.4	Viewing metrics with JConsole.....	5-8
5.4.5	Viewing metrics with Oracle Enterprise Manager	5-8
5.4.6	Viewing metrics using WSADMIN (IBM WebSphere)	5-9
5.5	Accessing DMS Metrics with WLDF.....	5-9
5.6	DMS Execution Context	5-9
5.6.1	DMS Execution Requests and Sub-Tasks	5-10
5.6.2	DMS Execution Context Usage.....	5-11
5.6.3	DMS Execution Context Communication	5-11
5.7	DMS Tracing and Events	5-11
5.7.1	Configuring the DMS Event System.....	5-12
5.7.2	Configuring Destinations	5-15

5.7.3	Understanding DMS Event Output	5-24
5.7.4	Understanding DMS Event Actions.....	5-27
5.8	DMS Best Practices	5-28

6 Oracle HTTP Server Performance Tuning

6.1	About Oracle HTTP Server.....	6-1
6.2	Monitoring Oracle HTTP Server Performance	6-1
6.3	Basic Tuning Considerations	6-2
6.3.1	Tuning Oracle HTTP Server Directives.....	6-2
6.3.2	Reducing Httpd Process Availability with Persistent Connections.....	6-7
6.3.3	Logging Options for Oracle HTTP Server.....	6-8
6.4	Advanced Tuning Considerations	6-9
6.4.1	Tuning Oracle HTTP Server Security	6-9
6.4.2	Tuning Oracle HTTP Server.....	6-12

7 Oracle Metadata Service (MDS) Performance Tuning

7.1	About Oracle Metadata Services (MDS).....	7-1
7.2	Monitoring Oracle Metadata Service Performance	7-1
7.3	Basic Tuning Considerations.....	7-2
7.3.1	Tuning the Database Repository	7-2
7.3.2	Tuning Cache Configuration	7-3
7.3.3	Purging Document Version History	7-4
7.3.4	Using Database Polling Interval for Change Detection	7-5
7.4	Advanced Tuning Considerations	7-6
7.4.1	Analyzing Performance Impact from Customization	7-6

Part III Oracle Fusion Middleware Server Components

8 Oracle Application Development Framework Performance Tuning

8.1	About Oracle ADF	8-1
8.2	Basic Tuning Considerations.....	8-1
8.2.1	Oracle ADF Faces Configuration and Profiling	8-2
8.2.2	Tuning ADF Faces	8-2
8.2.3	Tuning ADF Faces Component Attributes	8-9
8.2.4	Tuning Table and Tree Components	8-12
8.2.5	Tuning Data Visualization Tool (DVT) Components.....	8-15
8.3	Advanced Tuning Considerations	8-16
8.3.1	Tuning ADF Server Performance.....	8-16
8.3.2	Tuning Groovy Usage.....	8-28

9 Oracle TopLink (EclipseLink) JPA Performance Tuning

9.1	About Oracle TopLink and EclipseLink.....	9-1
9.2	Monitoring TopLink Performance	9-2
9.3	Basic Tuning Considerations.....	9-2
9.3.1	Using Efficient SQL Statements and Queries	9-3

9.3.2	Tuning Cache Configuration	9-10
9.3.3	Integrating Oracle Toplink with Coherence	9-16
9.4	Advanced Tuning Considerations	9-17
9.4.1	Configuring Mappings	9-17
9.4.2	Configuring Data Partitioning	9-17

10 Oracle Web Cache Performance Tuning

10.1	About Oracle Web Cache.....	10-1
10.2	Performance Considerations	10-1
10.2.1	Optimizing Hardware Resources.....	10-1
10.2.2	Optimizing Platform Connections	10-4
10.3	Basic Tuning Considerations.....	10-4
10.3.1	Optimizing Network Connections	10-4
10.3.2	Increasing Cache Hit Rates.....	10-7
10.3.3	Optimizing Response Time	10-9
10.4	Advanced Tuning Considerations	10-10
10.4.1	Optimizing Performance with Oracle ADF	10-10

Part IV SOA Suite Components

11 General Tuning for SOA Suite Components

11.1	About SOA Suite Configuration Properties.....	11-1
11.2	SOA Infrastructure Configurations.....	11-1
11.2.1	Audit Level	11-2
11.2.2	Instance Tracking Audit Trail Threshold	11-2
11.2.3	Logging Level.....	11-2
11.3	Modifying SOA Configuration Parameters	11-2
11.4	JVM Tuning Parameters.....	11-3
11.5	Database Settings	11-3
11.5.1	Configuring Data Sources for SOA	11-3
11.5.2	Managing Tables and Indexes	11-3
11.5.3	Weblogic Server Performance Tuning.....	11-4

12 Oracle Business Rules Performance Tuning

12.1	About Oracle Business Rules	12-1
12.2	Basic Tuning Considerations.....	12-1
12.2.1	Use Java Beans.....	12-1
12.2.2	Assert Child Facts instead of Multiple Dereferences	12-2
12.2.3	Avoid Side Affects in Rule Conditions.....	12-2
12.2.4	Avoid Expensive Operations in Rule Conditions.....	12-2
12.2.5	Consider Pattern Ordering	12-2
12.2.6	Consider the Ordering of Tests in Rule Conditions	12-2
12.2.7	Enable assertXPath Support.....	12-3

13 Oracle BPEL Process Manager Performance Tuning

13.1	About BPEL Process Manager	13-1
------	----------------------------------	------

13.2	Basic Tuning Considerations.....	13-1
13.2.1	Tuning Audit Levels.....	13-1
13.2.2	Tuning Database Persistence for BPEL.....	13-3
13.2.3	Tuning Invoke Messages	13-4
13.2.4	Tuning Processed Requests List	13-4
13.2.5	Tuning XML Document Persistence	13-5
13.2.6	Validating XML.....	13-5
13.2.7	Tuning Wait Time.....	13-5
13.2.8	Tuning Instance Key Block Size.....	13-5
13.2.9	Tuning Automatic Recovery Attempts	13-6
13.3	Advanced Tuning Considerations	13-6
13.3.1	Tuning BPEL Properties Set Inside a Composite	13-6
13.3.2	Identifying Tables Impacted By Instance Data Growth.....	13-6
14	Oracle Business Activity Monitoring Performance Tuning	
14.1	About Oracle Business Activity Monitoring.....	14-1
14.2	Basic Tuning Considerations.....	14-1
14.2.1	BAM Server Tuning.....	14-1
14.2.2	BAM Dashboard Tuning	14-2
14.2.3	BAM Database Tuning.....	14-3
14.2.4	Internet Browser Tuning.....	14-3
14.2.5	Enterprise Message Source Tuning.....	14-3
15	Oracle Mediator Performance Tuning	
15.1	About Oracle Mediator	15-1
15.2	Basic Tuning Considerations.....	15-1
15.2.1	Tuning metricsLevel.....	15-2
15.2.2	Using Domain-Value Maps.....	15-2
15.2.3	Deploying Deferred Routing Rules.....	15-2
15.2.4	Tuning Error and Retry Parameters.....	15-3
15.2.5	Setting the Audit Level	15-3
15.2.6	Using Resequencer for Messages	15-3
15.3	Tuning Event Delivery Network (EDN).....	15-4
16	Oracle Business Process Management Performance Tuning	
16.1	About Oracle Business Process Management.....	16-1
16.2	Basic Tuning Considerations.....	16-1
16.2.1	Audit Level	16-2
16.2.2	LargeDocumentThreshold	16-2
16.2.3	Dispatcher System Threads.....	16-2
16.2.4	Dispatcher Engine Threads	16-3
16.2.5	Dispatcher Invoke Threads	16-3
16.3	Tuning Oracle Workspace and Worklist Applications	16-3
16.4	Tuning Process Analytics.....	16-4
16.4.1	Process Measurement.....	16-5
16.4.2	Tuning Process Cubes	16-5

17 Oracle Human Workflow Performance Tuning

17.1	About Oracle Human Workflow	17-1
17.2	Basic Tuning Considerations.....	17-1
17.2.1	Minimize Client Response Time.....	17-2
17.2.2	Choose the Right Workflow Service Client	17-2
17.2.3	Narrow Qualifying Tasks Using Precise Filters	17-2
17.2.4	Retrieve Subset of Qualifying Tasks (Paging)	17-3
17.2.5	Fetch Only the Information That Is Needed for a Qualifying Task	17-3
17.2.6	Reduce the Number of Return Query Columns	17-4
17.2.7	Use the Aggregate API for Charting Task Statistics	17-4
17.2.8	Use the Count API Methods for Counting the Number of Tasks	17-5
17.2.9	Create Indexes On Demand for Flexfields	17-5
17.2.10	Use the doesTaskExist Method.....	17-5
17.3	Improving Server Performance.....	17-5
17.3.1	Archive Completed Instances Periodically	17-6
17.3.2	Select the Appropriate Workflow Callback Functionality.....	17-6
17.3.3	Minimize Performance Impacts from Notification.....	17-6
17.3.4	Deploy Clustered Nodes	17-6
17.4	Completing Workflows Faster	17-7
17.4.1	Use Workflow Reports to Monitor Progress	17-7
17.4.2	Specify Escalation Rules	17-7
17.4.3	Specify User and Group Rules for Automated Assignment	17-7
17.4.4	Use Task Views to Prioritize Work	17-8
17.5	Tuning Identity Provider.....	17-8
17.6	Tuning the Database.....	17-8

18 Oracle Adapters Performance Tuning

18.1	About Oracle Adapters	18-1
18.2	Oracle JCA Adapters for Files/FTP	18-1
18.2.1	Inbound Throttling Best Practices	18-2
18.2.2	Outbound Throttling Best Practices.....	18-2
18.2.3	Outbound Performance Best Practices	18-3
18.3	Oracle JCA Adapter for Database Tuning.....	18-4
18.3.1	JCA Adapter Basic Tuning Considerations	18-4
18.3.2	Existence Checking.....	18-6
18.3.3	Throttling	18-7
18.4	Oracle Socket Adapter Tuning.....	18-7
18.5	Oracle SOA JMS Adapter Tuning.....	18-8
18.5.1	adapter.jms.receive.threads Property	18-8
18.6	Oracle AQ Adapter Tuning.....	18-8
18.6.1	adapter.aq.dequeue.threads Property	18-8
18.7	Oracle MQ Adapter Tuning	18-9

19 User Messaging Service Performance Tuning

19.1	About Oracle User Messaging Services.....	19-1
19.2	Basic Tuning Considerations.....	19-1

19.2.1	SMPP Driver Performance Tuning	19-1
19.2.2	Email Driver Polling Frequency	19-2
19.3	Database Tuning for Optimal Throughput.....	19-2

20 Oracle B2B Performance Tuning

20.1	About Oracle B2B.....	20-1
20.2	Basic Tuning Considerations.....	20-1
20.2.1	Tuning Data Storage Configurations for B2B.....	20-1
20.2.2	Tuning MDS Cache Size	20-2
20.2.3	Tuning Number of Threads	20-2
20.2.4	Tuning the JMS Multiple Out Queues Setting.....	20-2

21 Oracle Service Bus Performance Tuning

21.1	About Oracle Service Bus	21-1
21.2	Monitoring Oracle Service Bus	21-1
21.3	Basic Tuning Considerations.....	21-2
21.3.1	Tuning JVM Memory	21-2
21.3.2	Tuning WebLogic Server for OSB	21-2
21.3.3	Tuning OSB Operational Settings	21-3
21.4	Advanced Tuning Considerations	21-4
21.4.1	Transport Tuning (Oracle WebLogic Server and Oracle Service Bus).....	21-4
21.4.2	Design Time Considerations for Proxy Applications.....	21-5
21.4.3	Design Considerations for XQuery Tuning	21-6

22 Oracle Business Intelligence Performance Tuning

22.1	About Oracle Business Intelligence.....	22-1
22.2	Oracle BI Server Query Performance Tuning.....	22-1
22.3	Oracle BI Server Query Cache Performance Tuning	22-2
22.4	Oracle BI Web Client Performance Tuning.....	22-2

Part V Identity Management Suite Components

23 Oracle Internet Directory Performance Tuning

23.1	About Oracle Internet Directory.....	23-1
23.2	Monitoring Oracle Internet Directory Performance	23-2
23.2.1	Monitoring Performance on UNIX and Windows Systems	23-2
23.2.2	Obtaining Recommendations by Using the Tuning and Sizing Wizard	23-3
23.2.3	Updating Database Statistics by Using oidstats.sql.....	23-4
23.2.4	Setting Performance-Related Replication Configuration Attributes	23-5
23.2.5	Managing System Configuration Attributes	23-5
23.2.6	Setting Garbage Collection Configuration Attributes.....	23-6
23.3	Basic Tuning Considerations.....	23-6
23.3.1	Database Parameters	23-7
23.3.2	LDAP Server Attributes.....	23-7
23.3.3	Database Statistics.....	23-9

23.3.4	Low-Priority Tuning Considerations.....	23-9
23.4	Advanced Tuning Considerations	23-10
23.4.1	Replication or Oracle Directory Integration Platform.....	23-10
23.4.2	Replication Server Configuration.....	23-11
23.4.3	Garbage Collection Configuration	23-12
23.4.4	Oracle Internet Directory with Cluster Configuration.....	23-12
23.4.5	Password Policies and Verifier Profiles.....	23-13
23.4.6	Server Entry Cache	23-13
23.4.7	Result Set Cache.....	23-15
23.4.8	Tuning Security Event Tracking.....	23-16
23.4.9	Optimizing Searches.....	23-16
23.5	Specific Use Cases That Require Additional Tuning.....	23-20
23.5.1	Bulk Load Operations	23-20
23.5.2	Bulk Delete Operations	23-20
23.5.3	High LDAP Write Operations Load	23-21
23.5.4	Sparc T4 Hardware Tuning.....	23-21

24 Oracle Unified Directory Performance Tuning

24.1	About Oracle Unified Directory	24-1
24.2	Performance Considerations	24-1
24.3	Monitoring Unified Directory Performance	24-2
24.3.1	Examining Log Files.....	24-2
24.3.2	Monitoring the Server With LDAP	24-3
24.3.3	Monitoring the Server With SNMP.....	24-3
24.4	Basic Tuning Considerations.....	24-3
24.4.1	Tuning Java Virtual Machine Settings.....	24-3
24.4.2	Tuning the Server Configuration	24-5
24.5	Advanced Tuning Recommendations	24-7

25 Oracle Virtual Directory Performance Tuning

25.1	About Oracle Virtual Directory	25-1
25.2	Basic Tuning Considerations.....	25-1
25.2.1	Tuning the Ping Interval.....	25-2
25.2.2	Tuning Worker Threads	25-3
25.2.3	Tuning Work Queue Capacity	25-3
25.2.4	Tuning the LDAP Connection Pool	25-3
25.2.5	Tuning Heap Size.....	25-4
25.3	Advanced Tuning Considerations	25-4
25.3.1	Tuning Database Adapters.....	25-4
25.3.2	Tuning Join Adapters.....	25-5
25.3.3	Tuning Filters	25-5
25.3.4	Tuning Load Balancer Local Store Adapter	25-6
25.3.5	Tuning the Cache Plug-In.....	25-6
25.3.6	Tuning LDAP Listener	25-7
25.3.7	Tuning the Server for OVD	25-8

26 Oracle Identity Federation Performance Tuning

26.1	About Oracle Identity Federation.....	26-1
26.2	Basic Tuning Considerations.....	26-1
26.2.1	Tuning Database Parameters for Identity Federation.....	26-1
26.2.2	Tuning the Oracle HTTP Server	26-3
26.3	Advanced Tuning Considerations	26-3
26.3.1	Tuning the LDAP Servers.....	26-3
26.3.2	Tuning SAML Protocol	26-5

27 Oracle Fusion Middleware Security Performance Tuning

27.1	About Security Services	27-1
27.2	Basic Tuning Considerations.....	27-2
27.3	Tuning Oracle Platform Security Services.....	27-2
27.3.1	JVM Tuning Parameters	27-3
27.3.2	JDK Tuning Parameters	27-3
27.3.3	LDAP Tuning Parameters	27-3
27.3.4	Authentication Tuning Parameters.....	27-3
27.3.5	Authorization Tuning Properties	27-3
27.3.6	OPSS PDP Service Tuning Parameters.....	27-6
27.4	Oracle Web Services Security Tuning.....	27-8
27.4.1	Choosing the Right Policy	27-9
27.4.2	Policy Manager.....	27-9
27.4.3	Configuring the Log Assertion to Record SOAP Messages	27-9
27.4.4	Configuring Connection Pooling	27-9
27.4.5	Monitoring the Performance of Web Services.....	27-10

Part VI Oracle WebCenter Components

28 Oracle WebCenter Portal Performance Tuning

28.1	About Oracle WebCenter Portal.....	28-1
28.2	Basic Tuning Considerations.....	28-1
28.2.1	Setting System Limit.....	28-2
28.2.2	Setting JDBC Data Source	28-2
28.2.3	Setting JRockit Virtual Machine (JVM) Arguments.....	28-3
28.2.4	Using Content Compression to Reduce Downloads	28-4
28.3	Tuning WebCenter Portal Application Configuration.....	28-4
28.3.1	Setting Session Timeout for a Spaces Application.....	28-5
28.3.2	Setting HTTP Session Timeout for a Framework Application.....	28-5
28.3.3	Setting JSP Page Timeout.....	28-5
28.3.4	Setting ADF Client State Token	28-6
28.3.5	Setting ADF View State Compression	28-6
28.3.6	Setting MDS Cache Size and Purge Rate.....	28-6
28.3.7	Configuring Concurrency Management	28-7
28.4	Tuning Back-End Component Configuration.....	28-11
28.4.1	Tuning Performance of the Announcements Service.....	28-11
28.4.2	Tuning Performance of the Discussions Service	28-12

28.4.3	Tuning Performance of the Instant Messaging and Presence (IMP) Service	28-12
28.4.4	Tuning Performance of the Mail Service	28-13
28.4.5	Tuning Performance of the Personal Events Service	28-13
28.4.6	Tuning Performance of the RSS News Feed Service	28-14
28.4.7	Tuning Performance of the Search Service	28-14
28.4.8	Tuning Policy Store Parameters	28-15
28.5	Tuning Identity Store Configuration	28-16
28.5.1	Tuning the Identity Store when Using SSL	28-16
28.5.2	Tuning Performance when Using OVD	28-16
28.5.3	Tuning Performance when Using Active Directory	28-17
28.6	Tuning Portlet Configuration	28-17
28.6.1	Tuning Performance of the Portlet Service	28-18
28.6.2	Configuring Portlet Cache Size	28-19
28.6.3	Enabling Java Object Cache for WSRP Producers	28-19
28.6.4	Suppressing Optimistic Rendering for WSRP Portlets	28-19
28.6.5	Tuning Performance of Oracle PDK-Java Producers	28-20
28.6.6	Setting Portlet Container Runtime Options	28-20
28.6.7	Excluding Request Attributes for Portlets	28-20
28.6.8	Setting WSRP Attribute for Portlet-served Resources	28-21
28.6.9	Setting WSRP Attribute for Resources Not Served by the Portlet	28-21
28.6.10	Configuring Portlet Timeout	28-22
28.6.11	Tuning Performance of OmniPortlet	28-22

Part VII Capacity Planning, Scalability, and Availability

29 Capacity Planning

29.1	About Capacity Planning for Oracle Fusion Middleware	29-1
29.1.1	Capacity Planning Factors to Consider	29-2
29.2	Determining Performance Goals and Objectives	29-2
29.3	Measuring Your Performance Metrics	29-3
29.4	Identifying Bottlenecks in Your System	29-3
29.4.1	Using Clustered Configurations	29-3
29.4.2	Using Connection Pooling	29-4
29.4.3	Setting the Max HeapSize on JVM	29-4
29.4.4	Increasing Memory or CPU	29-4
29.4.5	Segregation of Network Traffic	29-4
29.4.6	Segregation of Processes and Hardware Interrupt Handlers	29-4
29.5	Implementing a Capacity Management Plan	29-5
29.5.1	Hardware Configuration Requirements	29-5
29.5.2	JVM Requirements	29-6
29.5.3	Managed Servers	29-6
29.5.4	Database Configuration	29-6

30 Using Clusters and High Availability Features

30.1	About Clusters and High Availability Features	30-1
30.2	Using Clusters with Oracle Fusion Middleware	30-2

30.3 Using High Availability Features with Oracle Fusion Middleware 30-3

Preface

This guide describes how to monitor and optimize performance, review the key components that impact performance, use multiple components for optimal performance, and design applications for performance in the Oracle Fusion Middleware environment.

This preface contains these topics:

- [Audience](#)
- [Documentation Accessibility](#)
- [Conventions](#)

Audience

Oracle Fusion Middleware Performance and Tuning Guide is aimed at a target audience of Application developers, Oracle Fusion Middleware administrators, database administrators, and Web masters.

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.

Convention	Meaning
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Part I

Introduction

This part describes basic performance concepts, how to measure performance, and designing applications for performance and scalability. It contains the following chapters:

- [Chapter 1, "Introduction and Roadmap"](#)
- [Chapter 2, "Top Performance Areas"](#)
- [Chapter 3, "Performance Planning"](#)
- [Chapter 4, "Monitoring Oracle Fusion Middleware"](#)

Introduction and Roadmap

This section describes the contents and organization of this guide.

- [Section 1.1, "Document Scope and Audience"](#)
- [Section 1.2, "Guide to this Document"](#)
- [Section 1.3, "Related Documentation"](#)

1.1 Document Scope and Audience

Oracle Fusion Middleware Performance and Tuning Guide is for a target audience of Application developers, Oracle Fusion Middleware administrators, database administrators, and Web masters. This Guide assumes knowledge of Fusion Middleware Administration and hardware performance tuning fundamentals, WebLogic Server, XML, and the Java programming language.

1.2 Guide to this Document

- This chapter, [Chapter 1, "Introduction and Roadmap,"](#) introduces the objectives and organization of this guide.
- [Chapter 2, "Top Performance Areas,"](#) describes top tuning areas for Oracle Fusion Middleware and serves as a 'quick start' for tuning applications.
- [Chapter 3, "Performance Planning,"](#) describes the performance planning methodology and tuning concepts for Oracle Fusion Middleware.
- [Chapter 4, "Monitoring Oracle Fusion Middleware,"](#) describes how to monitor Oracle Fusion Middleware and its components to obtain performance data that can assist you in tuning the system and debugging applications with performance problems.
- [Chapter 6, "Oracle HTTP Server Performance Tuning,"](#) discusses the techniques for optimizing Oracle HTTP Server performance, the Web server component for Oracle Fusion Middleware. It provides a listener for Oracle WebLogic Server and the framework for hosting static pages, dynamic pages, and applications over the Web.
- [Chapter 5, "Understanding the Oracle Dynamic Monitoring Service"](#) provides an overview and features available in the Oracle Dynamic Monitoring Service (DMS).
- [Chapter 7, "Oracle Metadata Service \(MDS\) Performance Tuning,"](#) provides tuning tips for Oracle Metadata Service (MDS). MDS is used by components such as Oracle WebCenter Framework and Oracle Application Development Framework to manage metadata.

- [Chapter 8, "Oracle Application Development Framework Performance Tuning,"](#) provides basic guidelines on how to maximize the performance and scalability of the ADF stack in applications. Oracle ADF is an end-to-end application framework that builds on Java Platform, Enterprise Edition (Java EE) standards and open-source technologies to simplify and accelerate implementing service-oriented applications. This chapter covers design time, configuration time, and deployment time performance considerations.
- [Chapter 9, "Oracle TopLink \(EclipseLink\) JPA Performance Tuning,"](#) provides some of the available performance options for Java Persistence API (JPA) entity architecture. Oracle TopLink includes EclipseLink as the JPA implementation.
- [Chapter 10, "Oracle Web Cache Performance Tuning,"](#) provides methods and guidelines for improving the performance of Oracle Application Server Web Cache (Oracle Web Cache). Oracle Web Cache is a content-aware server accelerator or reverse proxy that improves the performance, scalability, and availability of Web sites that run on Oracle Fusion Middleware.
- [Chapter 11, "General Tuning for SOA Suite Components,"](#) describes the common SOA infrastructure tuning parameters for configuring Oracle Service-Oriented Architecture (SOA) Suite components to improve performance. Oracle SOA Suite provides a complete set of service infrastructure components for designing, deploying, and managing SOA composite applications. Oracle SOA Suite enables services to be created, managed, and orchestrated into SOA composite applications. Composites enable you to easily assemble multiple technology components into one SOA composite application.
- [Chapter 12, "Oracle Business Rules Performance Tuning"](#) describes the technology that enables automation of business rules; it also discusses the extraction of business rules from procedural logic such as Java code or BPEL processes.
- [Chapter 13, "Oracle BPEL Process Manager Performance Tuning,"](#) provides several BPEL property settings that can be configured to optimize performance at the process, domain, and application server levels. This chapter describes these property settings and provides recommendations on how to use them.
- [Chapter 15, "Oracle Mediator Performance Tuning,"](#) describes how to tune Oracle Mediator, a service engine within the Oracle SOA Service Infrastructure, for optimal performance. Oracle Mediator provides the framework to mediate between various providers and consumers of services and events. The Mediator service engine runs with the SOA Service Infrastructure Java EE application.
- [Chapter 16, "Oracle Business Process Management Performance Tuning"](#) describes how to tune Oracle Service Bus (OSB) which provides connectivity, routing, mediation, management and also some process orchestration capabilities.
- [Chapter 17, "Oracle Human Workflow Performance Tuning,"](#) describes how to tune Oracle Human Workflow for optimal performance. Oracle Human Workflow is a service engine running in Oracle SOA Service Infrastructure that allows the execution of interactive human driven processes. A human workflow provides the human interaction support such as approve, reject, and reassign actions within a process or outside of any process. The Human Workflow service consists of a number of services that handle various aspects of human interaction with a business process.
- [Chapter 18, "Oracle Adapters Performance Tuning,"](#) describes how to tune Oracle Adapters for optimal performance. Oracle technology adapters integrate Oracle Application Server and Oracle Fusion Middleware components such as Oracle BPEL Process Manager (Oracle BPEL PM) or Oracle Mediator components to file

systems, FTP servers, database queues (advanced queues, or AQ), Java Message Services (JMS), database tables, and message queues (MQ Series).

- [Chapter 14, "Oracle Business Activity Monitoring Performance Tuning,"](#) describes how to tune the Oracle Business Activity Monitoring dashboard application for optimal performance. Oracle Business Activity Monitoring (BAM) provides the tools for monitoring business services and processes in the enterprise.
- [Chapter 19, "User Messaging Service Performance Tuning,"](#) describes tips for tuning the User Messaging Service. Oracle User Messaging Service (Oracle UMS) enables two way communications between users and deployed applications. It has support for a variety of channels, such as E-mail, IM, SMS, and text-to-voice messages. Oracle UMS is integrated with Oracle Fusion Middleware components, such as Oracle BPEL PM, Oracle Human Workflow, Oracle BAM and Oracle WebCenter.
- [Chapter 20, "Oracle B2B Performance Tuning"](#) provides tuning tips for Oracle B2B. Oracle B2B is an e-commerce gateway that enables the secure and reliable exchange of business documents between an enterprise and its trading partners. Oracle B2B supports business-to-business document standards, security, transports, messaging services, and trading partner management. With Oracle B2B used as a binding component within an Oracle SOA Suite composite application, end-to-end business processes can be implemented.
- [Chapter 21, "Oracle Service Bus Performance Tuning"](#) provides basic and advanced tuning tips and design considerations for Oracle Service Bus.
- [Chapter 22, "Oracle Business Intelligence Performance Tuning"](#) provides basic and advanced tuning tips for Oracle Business Intelligence.
- [Chapter 23, "Oracle Internet Directory Performance Tuning,"](#) provides guidelines on Oracle Internet Directory tuning and configuration requirements. Oracle Internet Directory is an LDAP Version 3-enabled service that enables fast retrieval and centralized management of information about dispersed users, network configuration, and other resources.
- [Chapter 25, "Oracle Virtual Directory Performance Tuning,"](#) provides tuning tips for Oracle Virtual Directory. Oracle Virtual Directory is an LDAP Version 3-enabled service that provides an abstracted view of one or more enterprise data sources. Oracle Virtual Directory consolidates multiple data sources into a single directory view, enabling you to integrate LDAP-aware applications with diverse directory server data stores.
- [Chapter 26, "Oracle Identity Federation Performance Tuning,"](#) provides tuning tips for Oracle Identity Federation, a standalone, self-contained federation server that enables single sign-on (SSO) and authentication in a multiple-domain identity network.
- [Chapter 27, "Oracle Fusion Middleware Security Performance Tuning,"](#) describes Oracle Platform Security for Java. Oracle Platform Security for Java is the Oracle Fusion Middleware security implementation for Java features such as Java Authentication and Authorization Service (JAAS) and Java EE security. This chapter describes how you can configure it for optimal performance.
- [Chapter 28, "Oracle WebCenter Portal Performance Tuning,"](#) provides suggested tuning tips for Oracle WebCenter including: Environment Configuration, Application Configuration and Back-End Services and Server Configuration.
- [Chapter 29, "Capacity Planning,"](#) discusses the process of determining what type of hardware and software configuration is required to meet application needs.

- [Chapter 30, "Using Clusters and High Availability Features,"](#) discusses the architecture, interaction, and dependencies of Oracle Fusion Middleware components, and explains how they can be deployed in a high availability architecture to maximize performance.

1.3 Related Documentation

For more information, see the following documents in the Oracle Fusion Middleware 11g Release 1 (11.1.1.7.0) documentation set:

- *Oracle Fusion Middleware Administrator's Guide*
- *Oracle Fusion Middleware 2 Day Administration Guide*
- *Oracle Fusion Middleware Concepts*
- *Oracle Fusion Middleware Application Security Guide*
- *Oracle Fusion Middleware High Availability Guide*
- *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*
- *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*
- *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*
- *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
- *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache*
- *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*
- *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*
- *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*
- *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation*

Top Performance Areas

This chapter describes the top tuning areas for Oracle Fusion Middleware. It covers critical Oracle Fusion Middleware performance areas and provides a quick start for tuning Java EE applications in the following sections:

- [Section 2.1, "About Identifying Top Performance Areas"](#)
- [Section 2.2, "Securing Sufficient Hardware Resources"](#)
- [Section 2.3, "Tuning the Operating System"](#)
- [Section 2.4, "Tuning Java Virtual Machines \(JVMs\)"](#)
- [Section 2.5, "Tuning the WebLogic Server"](#)
- [Section 2.6, "Tuning Database Parameters"](#)
- [Section 2.7, "Reusing Database Connections"](#)
- [Section 2.8, "Enabling Data Source Statement Caching"](#)
- [Section 2.9, "Controlling Concurrency"](#)
- [Section 2.10, "Setting Logging Levels"](#)

2.1 About Identifying Top Performance Areas

One of the most challenging aspects of performance tuning is knowing where to begin. This chapter serves as a 'quick start' guide to performance tuning your Oracle Fusion Middleware applications.

[Table 2-1](#) provides a list of common performance considerations for Oracle Fusion Middleware. While the list is a useful tool in starting your performance tuning, it is not meant to be comprehensive list of areas to tune. You must monitor and track specific performance issues within your application to understand where tuning can improve performance. See [Chapter 4, "Monitoring Oracle Fusion Middleware"](#) for more information.

Table 2–1 Top Performance Areas for Oracle Fusion Middleware

Performance Area	Description and Reference
Hardware Resources	<p>Ensure that your hardware resources meet or exceed the application's resource requirements to maximize performance.</p> <p>See Section 2.2, "Securing Sufficient Hardware Resources" for information on how to determine if your hardware resources are sufficient.</p>
Operating System	<p>Each operating system has native tools and utilities that can be useful for monitoring purposes.</p> <p>See Section 2.3, "Tuning the Operating System"</p>
Java Virtual Machines (JVMs)	<p>This section discusses best practices and provides practical tips to tune the JVM and improve the performance of a Java EE application. It also discusses heap size and JVM garbage collection options.</p> <p>See Section 2.4, "Tuning Java Virtual Machines (JVMs)".</p>
Database	<p>For applications that access a database, ensure that your database is properly configured to support your application's requirements.</p> <p>See Section 2.6, "Tuning Database Parameters" for more information on garbage collection.</p>
WebLogic Server	<p>If your Oracle Fusion Middleware applications are using the WebLogic Server, see Section 2.5, "Tuning the WebLogic Server".</p>
Database Connections	<p>Pooling the connections so they are reused is an important tuning consideration.</p> <p>See Section 2.7, "Reusing Database Connections"</p>
Data Source Statement Caching	<p>For applications that use a database, you can lower the performance impact of repeated statement parsing and creation by configuring statement caching properly.</p> <p>See Section 2.8, "Enabling Data Source Statement Caching"</p>
Oracle HTTP Server	<p>Tune the Oracle HTTP Server directives to set the level of concurrency by specifying the number of HTTP connections.</p> <p>See Section 2.9, "Controlling Concurrency".</p>
Concurrency	<p>This section discusses ways to control concurrency with Oracle Fusion Middleware components.</p> <p>See Section 2.9, "Controlling Concurrency"</p>
Logging Levels	<p>Logging levels are thresholds that a system administrator sets to control how much information is logged. Performance can be impacted by the amount of information that applications log therefore it is important to set the logging levels appropriately.</p> <p>See Section 2.10, "Setting Logging Levels".</p>

2.2 Securing Sufficient Hardware Resources

A key component of managing the performance of Oracle Fusion Middleware applications is to ensure that there are sufficient CPU, memory, and network resources to support the user and application requirements for your installation.

No matter how well you tune your applications, if you do not have the appropriate hardware resources, your applications cannot reach optimal performance levels. Oracle Fusion Middleware has minimum hardware requirements for its applications and database tier. For details on Oracle Fusion Middleware supported configurations, see "System Requirements and Prerequisites" in the *Oracle Fusion Middleware Installation Planning Guide* for your platform.

Sufficient hardware resources should meet or exceed the acceptable response times and throughputs for applications without becoming saturated. To verify that you have sufficient hardware resources, you should monitor resource utilization over an extended period to determine if (or when) you have occasional peaks of usage or whether a resource is consistently saturated. For more information on monitoring, see [Chapter 4, "Monitoring Oracle Fusion Middleware"](#).

Tip: Your target CPU usage should not reach 100% utilization. You should determine a target CPU utilization based on your application needs, including CPU cycles for peak usage.

If your CPU utilization is optimized at 100% during normal load hours, you have no capacity to handle a peak load. In applications that are latency sensitive and maintaining a fast response time is important, high CPU usage (approaching 100% utilization) can increase response times while throughput stays constant or even decreases. For such applications, a 70% - 80% CPU utilization is recommended. A good target for non-latency sensitive applications is about 90%.

If any of the hardware resources are saturated (consistently at or near 100% utilization), one or more of the following conditions may exist:

- The hardware resources are insufficient to run the application.
- The system is not properly configured.
- The application or database must be tuned.

For a consistently saturated resource, the solutions are to reduce load or increase resources. For peak traffic periods when the increased response time is not acceptable, consider increasing resources or determine if there is traffic that can be rescheduled to reduce the peak load, such as scheduling batch or background operations during slower periods.

Oracle Fusion Middleware provides a variety of mechanisms to help you control resource concurrency; this can limit the impact of bursts of traffic. However, for a consistently saturated system, these mechanisms should be viewed as temporary solutions. For more information see [Section 2.9, "Controlling Concurrency"](#).

2.3 Tuning the Operating System

Each operating system has native tools and utilities that can be useful for monitoring and tuning purposes. Native operating system commands enable you to monitor CPU utilization, paging activity, swapping, and other system activity information.

For details on operating system commands, and guidelines for performance tuning of the network or operating system, refer to the documentation provided by the operating system vendor.

2.4 Tuning Java Virtual Machines (JVMs)

How you tune your Java virtual machine (JVM) greatly affects the performance of Oracle Fusion Middleware and your applications. This section discusses the tuning options that have the greatest impact on performance.

To maximize performance from your JVM, be sure that you use only production JVMs on which your applications have been certified and that your operating system patches are up-to-date.

The Supported Configurations pages at http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html are frequently updated and contain the latest certification information on various platforms.

Note: For additional information about tuning the JVM, see the following:

- Java Performance Documentation:
(<http://java.sun.com/docs/performance/>)
 - The Java Tuning White Paper
(<http://java.sun.com/performance/reference/whitepapers/tuning.html>)
 - Garbage Collection Tuning
(<http://www.oracle.com/technetwork/java/javase/tech/index-jsp-136373.html>)
 - "Tuning Java Virtual Machines (JVMs)" in *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*
-
-

This section covers the following performance tuning areas for your JVM:

- [Configuring Garbage Collection](#)
- [Logging Low Memory Conditions](#)
- [Monitoring and Profiling the JVM](#)

2.4.1 Configuring Garbage Collection

Garbage collection is the JVM process of freeing up unused Java objects in the Java heap. JVM garbage collection can be a resource-intensive operation and may effect application performance. In some cases, inefficient garbage collection can severely degrade application performance. Therefore, it is important to understand how applications create and destroy objects.

This section cover the following Garbage Collection tuning options:

- [Specifying Heap Size Values](#)
- [Selecting a Garbage Collection Scheme](#)
- [Disabling Explicit Garbage Collection](#)

An acceptable rate for garbage collection is application-specific and should be adjusted after analyzing the actual time and frequency of garbage collections. If you set a large heap size, full garbage collection is slower, but it occurs less frequently. If you set your heap size in accordance with your memory needs, full garbage collection is faster, but occurs more frequently.

To tune the JVM garbage collection options you must analyze garbage collection data and check for the frequency and type of garbage collections, the size of the memory pools, and the time spent on garbage collection.

Before you configure JVM garbage collection, analyze the following data points:

1. How often is garbage collection taking place? Compare the time stamps around the garbage collection.
2. How long is a full garbage collection taking?

3. What is the heap size after each full garbage collection? If the heap is always 85 percent free, for example, you might set the heap size smaller.
4. Do the young generation heap sizes (Sun) or Nursery size (Jrockit) need tuning?

You can manually log garbage collection and memory pool sizes using verbose garbage collection logging:

- Sun JVM command line options:

```
-verbose:gc
-XX:+PrintGCDetails
-XX:+PrintGCTimeStamps
```

Look for "Full GC" to identify major collections.

- Additional Sun Tools:

- JStat
- JConsole
- Visualgc

For more information on Sun's options, see

<http://java.sun.com/javase/technologies/hotspot/gc/index.jsp>

- Jrockit JVM command line options:

```
-XXverbose:gc
```

NOTE: Oracle provides other command-line options to improve the performance of your JRockit VM. For detailed information, see "JRockit JDK Command Line Options by Name" at http://download.oracle.com/docs/cd/E13150_01/jrockit_jvm/jrockit/webdocs/index.html

- Additional JRockit Tools:

- JRockit Runtime Analyzer (jra recording)
- JRockit Management Console (jrmc)
- JRockit Memory Leak Detector

2.4.1.1 Specifying Heap Size Values

The goal of tuning your heap size is to minimize the time that your JVM spends doing garbage collection while maximizing the number of clients that the Fusion Middleware stack can handle at a given time.

Specifically the Java heap is where the objects of a Java program live. It is a repository for live objects, dead objects, and free memory. When an object can no longer be reached from any pointer in the running program, it is considered "garbage" and ready for collection. A best practice is to tune the time spent doing garbage collection to within 5% of execution time.

The JVM heap size determines how often and how long the virtual machine spends collecting garbage. An acceptable rate for garbage collection is application-specific and should be adjusted after analyzing the actual time and frequency of garbage collections. If you set a large heap size, full garbage collection is slower, but it occurs less frequently. If you set your heap size in accordance with your memory needs, full garbage collection is faster, but occurs more frequently.

In production environments, set the minimum heap size and the maximum heap size to the same value to prevent wasting virtual machine resources used to constantly grow and shrink the heap. Ensure that the sum of the maximum heap size of all the JVMs running on your system does not exceed the amount of available physical RAM. If this value is exceeded, the Operating System starts paging and performance degrades significantly. The virtual machine always uses more memory than the heap size. The memory required for internal virtual machine functionality, native libraries outside of the virtual machine, and permanent generation memory (memory required to store classes and methods) is allocated in addition to the heap size settings.

For example, you can use the following JVM options to tune the heap:

- If you run out of heap memory (not due to a memory leak), increase `-Xmx`.
- If you run out of native memory, you may need to decrease `-Xmx`.
- For Oracle JRockit, modify `-Xns:<nursery size>` to tune the size of the nursery.
- For Sun JVM, modify `-Xmn` to tune the size of the heap for the young generation.

If you receive `java.lang.OutOfMemoryError: PermGen space` errors, you may also need to increase the permanent generation space.

See Also: For more information on how to specify heap size values for Oracle WebLogic Server, see "Specifying Heap Size Values" in *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*

For more information on tuning the young generation see the "Young Generation" section of the *Java SE 6 HotSpot Virtual Machine Garbage Collection Tuning* at

http://java.sun.com/javase/technologies/hotspot/gc/gc_tuning_6.html#generation_sizing.young_gen

For more information on Oracle JRockit heap configurations, see "Setting the Heap and Nursery Size" in *Diagnostics Guide* at

http://download.oracle.com/docs/cd/E13188_01/jrockit/geninfo/diagnos/memman.html

For the Sun java virtual machine see the "Insufficient Memory" section of *Monitoring and Managing Java SE 6 Platform Applications* at

http://java.sun.com/developer/technicalArticles/J2SE/monitoring/index.html#Insufficient_Memory.

"Out of Memory" Frequently Asked Questions section at

http://java.sun.com/docs/hotspot/HotSpotFAQ.html#gc_oom

2.4.1.2 Selecting a Garbage Collection Scheme

Depending on which JVM you are using, you can choose from several garbage collection schemes to manage your system memory. Some garbage collection schemes are more appropriate for a given type of application. Once you have an understanding of the workload of the application and the different garbage collection algorithms utilized by the JVM, you can optimize the configuration of the garbage collection.

Refer to the following links for garbage collection options for your JVM:

- For an overview of the garbage collection schemes available with Sun's HotSpot VM, see "Java SE 6 HotSpot Virtual Machine Garbage Collection Tuning" at http://java.sun.com/javase/technologies/hotspot/gc/gc_tuning_6.html.

- For a comprehensive explanation of the collection schemes available, see "Memory Management in the Java HotSpot™ Virtual Machine" at http://java.sun.com/j2se/reference/whitepapers/memorymanagement_whitepaper.pdf.
- For a discussion of the garbage collection schemes available with the JRockit JDK, see "Using the JRockit Memory Management System" at http://download.oracle.com/docs/cd/E13150_01/jrockit_jvm/jrockit/webdocs/index.html.

2.4.1.3 Disabling Explicit Garbage Collection

The following parameters are used to help diagnose whether explicit garbage collections are occurring. They can also be used to disable the explicit garbage collections if necessary until the code is fixed:

- For Sun virtual machines use `-XX:+DisableExplicitGC`
For more information on using the explicit garbage collections, see "Java SE 6 HotSpot Virtual Machine Garbage Collection Tuning " at http://java.sun.com/javase/technologies/hotspot/gc/gc_tuning_6.html.
- For Oracle JRockit virtual machines use `-XXnoSystemGC`
For more information on tuning the Oracle JRockit, see at http://download.oracle.com/docs/cd/E13188_01/jrockit/geninfo/diagnos/bestpractices.html

These parameters disable explicit garbage collection. Applications should avoid the use of `system.gc()` calls. If you suspect an application may be explicitly triggering garbage collection, set this parameter and observe the differences in your garbage collection behavior. If you detect that performance is affected by explicit collections, check the code to determine where explicit garbage collections are used and why, and the impact of disabling the calls. Application developers sometimes use `system.gc()` calls to trigger finalizers. This is not a recommended practice and can yield indeterminate behavior.

2.4.2 Logging Low Memory Conditions

WebLogic Server enables you to automatically log low memory conditions observed by the server. WebLogic Server detects low memory by sampling the available free memory a set number of times during a time interval. At the end of each interval, an average of the free memory is recorded and compared to the average obtained at the next interval. If the average drops by a user-configured amount after any sample interval, the server logs a low memory warning message in the log file and sets the server health state to "warning."

See Also: For more information on using WebLogic Server to detect low memory conditions refer to the following:

"Log low memory conditions" in *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*.

"Automatically Logging Low Memory Conditions" in *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*

2.4.3 Monitoring and Profiling the JVM

Monitoring the performance of your JVM is crucial to achieving optimal performance. Depending on your platform, the following tools can be used to monitor and profile your JVM:

- **Oracle JRockit® Mission Control**

Oracle JRockit Mission Control is a suite of tools designed to monitor, manage, profile, and eliminate memory leaks in your Java application without the performance impacts normally associated with these types of tools.

For more information on the Oracle JRockit Mission Control see:

http://download.oracle.com/docs/cd/E13188_01/jrockit/tools/index.html

- **Sun JVM**

The Java™ Platform comes with the following monitoring facilities built-in:

- Java Virtual Machine Monitoring and Management API
- JConsole
- Hprof Tools
- Logging Monitoring and Management Interface
- Java Management Extensions (JMX)

For more information on the Java platform monitoring tools, see:

<http://java.sun.com/developer/technicalArticles/J2SE/monitoring/>

2.5 Tuning the WebLogic Server

If your Oracle Fusion Middleware applications are using the WebLogic Server, see "Tuning Java Virtual Machines (JVMs)" in *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*.

2.6 Tuning Database Parameters

To achieve optimal performance for applications that use the Oracle database, the database tables you access must be designed with performance in mind. Monitoring and tuning the database ensures that you get the best performance from your applications.

This section covers the following:

- [Tuning Database Parameters](#)
- [Tuning Redo Logs Location and Sizing](#)
- [Tuning Automatic Segment-Space Management \(ASSM\)](#)

Note: Always review the tuning guidelines in your database-specific vendor documentation. For more information on tuning the Oracle database, see the *Oracle Database Performance Tuning Guide*.

2.6.1 Tuning Database Parameters

The following tables provide common **init.ora** parameters and their descriptions. Consider following these guidelines to set the database parameters. Ultimately,

however, the DBA should monitor the database health and tune parameters based on the need. See the following tables for more information:

- [Table 2–2, " Important init.ora Oracle 10g Database Tuning Parameters"](#)
- [Table 2–3, " Important inti.ora Oracle 11g Database Tuning Parameters"](#)

Note: Consider applying Patch Set Release (PSR) 11.1.0.7 and upgrade the database prior to attempting the following modifications.

2.6.1.1 Initialization Parameters for Oracle 10g

The following table describes several performance-related database initialization parameters for Oracle 10g database. The tuning considerations listed below are applicable to most scenarios. Always set your database parameters based on your own use case scenarios.

Table 2–2 Important init.ora Oracle 10g Database Tuning Parameters

Database Parameter	Description
_b_tree_bitmap_plans	Consider setting this parameter to FALSE to prevent optimizer from attempting bitmap operations as there are no bitmap indexes in Fusion Middleware.
DB_BLOCK_SIZE	DB_BLOCK_SIZE specifies (in bytes) the size of Oracle database blocks. The default block size of 8K is optimal for most systems. Set this parameter at the time of database creation.
NLS_SORT	Consider setting NLS_SORT to BINARY, otherwise sort will do full table scan and performance can be impacted.
OPEN_CURSORS	Consider using a value of 500 open cursors (handles to private SQL areas) a session can have at once.
SESSION_CACHED_CURSORS	Consider using a value of 500 session cursors to cache.
SESSION_MAX_OPEN_FILES	SESSION_MAX_OPEN_FILES specifies the maximum number of BFILEs that can be opened in any session. Once this number is reached, subsequent attempts to open more files in the session by using DBMS_LOB.FILEOPEN() or OCILobFileOpen() may fail. The maximum value for this parameter depends on the equivalent parameter defined for the underlying operating system.
JOB_QUEUE_PROCESSES	JOB_QUEUE_PROCESSES specifies the maximum number of processes that can be created for the execution of jobs. It specifies the number of job queue processes per instance. NOTE: The value for the JOB_QUEUE_PROCESSES parameter should NEVER exceed the number of CPU cores available on the database server.
LOG_BUFFER	LOG_BUFFER specifies the amount of memory (in bytes) that Oracle uses when buffering redo entries to a redo log file. Redo log entries contain a record of the changes that have been made to the database block buffers. The LGWR process writes redo log entries from the log buffer to a redo log file.
UNDO_MANAGEMENT	UNDO_MANAGEMENT specifies which undo space management mode the system should use. When set to AUTO, the instance starts in automatic undo management mode. In manual undo management mode, undo space is allocated externally as rollback segments.
PL_SQL_CODE_TYPE	Consider setting PL_SQL_CODE_TYPE to NATIVE
PROCESSES	Consider using a value of 5000 operating system processes to be connected to Oracle concurrently.
PGA_AGGREGATE_TARGET	Consider setting PGA_AGGREGATE_TARGET to 1G of PGA memory available to all server processes attached to the instance.

Table 2–2 (Cont.) Important init.ora Oracle 10g Database Tuning Parameters

Database Parameter	Description
SGA_MAX_SIZE	Consider setting the SGA_MAX_SIZE to 2G initially and then monitor the production database on daily basis and adjust SGA and PGA accordingly.
SGA_TARGET	Consider setting the SGA_TARGET to 2G initially and then monitor the production database on daily basis and adjust SGA and PGA accordingly.
TRACE_ENABLED	TRACE_ENABLED controls tracing of the execution history, or code path, of Oracle. Oracle Support Services uses this information for debugging. Although the performance impacts incurred from processing is not excessive, you may improve performance by setting TRACE_ENABLED to FALSE.

2.6.1.2 Initialization Parameters for Oracle 11g

The following table provides information on some important performance-related database initialization parameters for Oracle 11g database.

Table 2–3 Important inti.ora Oracle 11g Database Tuning Parameters

Database Parameter	Description
AUDIT_TRAIL	If there is NO policy to audit db activity, consider setting this parameter to NONE. Enabling auditing can impact performance.
MEMORY_MAX_TARGET	MEMORY_MAX_TARGET specifies the maximum value to which a DBA can set the MEMORY_TARGET initialization parameter.
MEMORY_TARGET	Consider setting the MEMORY_TARGET to NONE. Set SGA and PGA separately as setting MEMORY_TARGET does not allocate sufficient memory to SGA and PGA as needed.
PGA_AGGREGATE_TARGET	Consider using a value of 1G for PGA initially and then monitor the production database on daily basis and adjust SGA and PGA accordingly. If the database server has more memory, consider setting PGA_AGGREGATE_TARGET to a value higher than 1G based on usage needs.
SGA_MAX_SIZE	Consider setting MEMORY_TARGET instead of setting SGA and the PGA separately.
SGA_TARGET	Consider using a value of 2G for SGA is 2G to start with and initially and then monitor the production database on daily basis and adjust SGA and PGA accordingly. If the database server has more memory, consider setting SGA_TARGET to a value higher than 2G based on usage needs.

2.6.2 Tuning Redo Logs Location and Sizing

Tuning the redo log options can provide performance improvement for applications running in an Oracle Fusion Middleware environment, and in some cases, you can significantly improve I/O throughput by moving the redo logs to a separate disk.

Consider having at least 3 redo log groups with 2G of size each. Redo log files should be placed on a disk separate from data files to improve I/O performance.

For more information see the *Oracle Database Performance Tuning Guide*.

2.6.3 Tuning Automatic Segment-Space Management (ASSM)

For permanent tablespaces, consider using automatic segment-space management. Such tablespaces, often referred to as bitmap tablespaces, are locally managed tablespaces with bitmap segment space management.

For backward compatibility, the default local tablespace segment-space management mode is MANUAL.

While configuring tablespaces, consider setting the extent allocation type to SYSTEM. If the allocation type is set to UNIFORM, it might impact performance.

For more information, see "Free Space Management" in *Oracle Database Concepts*, and "Specifying Segment Space Management in Locally Managed Tablespaces" in *Oracle Database Administrator's Guide*.

2.7 Reusing Database Connections

Creating a database connection is a relatively resource intensive process in any environment. Typically, a connection pool starts with a small number of connections. As client demand for more connections grow, there may not be enough in the pool to satisfy the requests. WebLogic Server creates additional connections and adds them to the pool until the maximum pool size is reached.

One way to avoid connection creation delays is to initialize all connections at server startup, rather than on-demand as clients need them. This may be appropriate if your load is predictable and even. Set the initial number of connections equal to the maximum number of connections in the Connection Pool tab of your data source configuration. Determine the optimal value for the Maximum Capacity as part of your pre-production performance testing.

If your load is uneven, and has a much higher number of connections at peak load than at typical load, consider setting the initial number of connections equal to your typical load. In addition, consider setting the maximum number of connections based on your supported peak load. With these configurations, WebLogic server can free up some connections when they are not used for a period of time.

For more information, see "Tuning Data Source Connection Pool Options" in *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server*.

2.8 Enabling Data Source Statement Caching

When you use a prepared statement or callable statement in an application or EJB, there may be a performance impact associated with the processing of the communication between the application server and the database server and on the database server. To minimize the processing impact, enable the data source to cache prepared and callable statements used in your applications. When an application or EJB calls any of the statements stored in the cache, the server reuses the statement stored in the cache. Reusing prepared and callable statements reduces CPU usage on the database server, improving performance for the current statement and leaving CPU cycles for other tasks.

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

- **Statement Cache Type**—The algorithm that determines which statements to store in the statement cache.
- **Statement Cache Size**—The number of statements to store in the cache for each connection. The default value is 10. You should analyze your database's statement parse metrics to size the statement cache sufficiently for the number of statements you have in your application.

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

For more information on using statement caching, see "Increasing Performance with the Statement Cache" in the *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server*.

2.9 Controlling Concurrency

Limiting concurrency, at multiple layers of the system to match specific usage needs, can greatly improve performance. This section discusses a few of the areas within Oracle Fusion Middleware where concurrency can be controlled.

When system capacity is reached, and a web server or application server continues to accept requests, application performance and stability can deteriorate. There are several places within Oracle Fusion Middleware where you can throttle the requests to avoid overloading the mid-tier or database tier systems and tune for best performance.

- [Setting Server Connection Limits](#)
- [Configuring Connection Pools](#)
- [Tuning the WebLogic Server Thread Pool](#)
- [Tuning Oracle WebCenter Concurrency](#)
- [Tuning BPEL Concurrency](#)

2.9.1 Setting Server Connection Limits

Oracle HTTP Server uses directives in `httpd.conf`. This configuration file specifies the maximum number of HTTP requests that can be processed simultaneously, logging details, and certain limits and time outs.

For more information on modifying the `httpd.conf` file, see "Configuring Oracle HTTP Server" in *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*.

You can use the `MaxClients` and `ThreadsPerChild` directives to limit incoming requests to WebLogic instances from the Oracle HTTP Server based on your expected client load and system resources. The following sections describe some Oracle HTTP Server tuning parameters related to connection limits that you typically need to tune based on your expected client load. See [Chapter 6, "Oracle HTTP Server Performance Tuning"](#) for more information and a more complete list of tunable parameters.

2.9.1.1 MaxClients/ThreadsPerChild

Note: The `MaxClients` parameter is applicable only to UNIX platforms and on Microsoft Windows (`mpm_winnt`), the same is achieved through the `ThreadsPerChild` and `ThreadLimit` parameters.

The `MaxClients` property specifies a limit on the total number of server threads running, that is, a limit on the number of clients who can simultaneously connect. If the number of client connections reaches this limit, then subsequent requests are queued in the TCP/IP system up to the limit specified (in the `ListenBackLog` directive).

You can configure the `MaxClients` directive in the `httpd.conf` file up to a maximum of 8K (the default value is 150). If your system is not resource-saturated and you have a user population of more than 150 concurrent HTTP connections, you can improve your performance by increasing `MaxClients` to increase server concurrency. Increase `MaxClients` until your system becomes fully utilized (85% is a good threshold).

When system resources are saturated, increasing `MaxClients` does not improve performance. In this case, the `MaxClients` value could be reduced as a throttle on the number of concurrent requests on the server.

If the server handles persistent connections, then it may require sufficient concurrent `httpd` server processes to handle both active and idle connections. When you specify `MaxClients` to act as a throttle for system concurrency, you need to consider that persistent idle `httpd` connections also consume `httpd` processes. Specifically, the number of connections includes the currently active persistent and non-persistent connections and the idle persistent connections. When there are no `httpd` server threads available, connection requests are queued in the TCP/IP system until a thread becomes available, and eventually clients terminate connections.

You can define a number of server processes and the threads per process (`ThreadsPerChild`) to handle the incoming connections to Oracle HTTP Server. The `ThreadsPerChild` property specifies the upper limit on the number of threads that can be created under a server (child) process.

Note: `ThreadsPerChild`, `StartServers`, and `ServerLimit` properties are inter-related with the `MaxClients` setting. All of these properties must be set appropriately to achieve the number of connections as specified by `MaxClients`. See [Table 6-1, "Oracle HTTP Server Configuration Properties"](#) for a description of all the HTTP configuration properties.

2.9.1.2 KeepAlive

A persistent, `KeepAlive`, HTTP connection consumes an `httpd` child process, or thread, for the duration of the connection, even if no requests are currently being processed for the connection.

If you have sufficient capacity, `KeepAlive` should be enabled; using persistent connections improves performance and prevents wasting CPU resources re-establishing HTTP connections. Normally, you should not need to change `KeepAlive` parameters.

Note: The default maximum requests for a persistent connection is 100, as specified with the `MaxKeepAliveRequests` directive in `httpd.conf`. By default, the server waits for 15 seconds between requests from a client before closing a connection, as specified with the `KeepAliveTimeout` directive in `httpd.conf`.

2.9.1.3 Tuning HTTP Server Modules

The Oracle HTTP Server (OHS) uses the `mod_wl_ohs` module to route requests to the underlying Weblogic server or the Weblogic Server cluster. The configuration details for `mod_wl_ohs` are available in the `mod_wl_ohs.conf` file in the `config` directory.

For more information on the tuning parameters for `mod_wl_ohs` see, "Understanding Oracle HTTP Server Modules" in *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*.

2.9.2 Configuring Connection Pools

Connection pooling is configured and maintained per Java runtime. Connections are not shared across different runtimes. To use connection pooling, no configuration is required. Configuration is necessary only if you want to customize how pooling is done, such as to control the size of the pools and which types of connections are pooled.

You configure connection pooling by using a number of system properties at program startup time. Note that these are system properties, not environment properties and that they affect all connection pooling requests.

For applications that use a database, performance can improve when the connection pool associated with a data source limits the number of connections. You can use the `MaxCapacity` attribute to limit the database requests from Oracle Application Server so that incoming requests do not saturate the database, or to limit the database requests so that the database access does not overload the Oracle Application Server-tier resource.

The connection pool `MaxCapacity` attribute specifies the maximum number of connections that a connection pool allows. By default, the value of `MaxCapacity` is set to 15. For best performance, you should specify a value for `MaxCapacity` that matches the number appropriate to your database performance characteristics.

Limiting the total number of open database connections to a number your database can handle is an important tuning consideration. You should check to make sure that your database is configured to allow at least as large a number of open connections as the total of the values specified for all the data sources `MaxCapacity` option, as specified in all the applications that access the database.

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

"Tuning Data Source Connection Pool Options" in *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server*.

2.9.3 Tuning the WebLogic Server Thread Pool

By default WebLogic Server uses a single thread pool, in which all types of work are executed. WebLogic Server uses Work Managers to prioritize work based on rules you can 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. There is a default work manager that manages the common thread pool.

The common thread pool changes its size automatically to maximize throughput. WebLogic Server 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.

Since the WebLogic Server thread pool by default is sized automatically, in most situations you do not need to tune this. However, for special requirements, an administrator can configure custom Work Managers to manage the thread pool at a more granular level for sets of requests that have similar performance, availability, or reliability requirements. With custom work managers, you can define priorities and

guidelines for how to assign pending work (including specifying a min threads or max threads constraint, or a constraint on the total number of requests that can be queued or executing before WebLogic Server begins rejecting requests).

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

- The default fair share 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.
- You use MDBs in your application.

To ensure MDBs use a well-defined share of server thread resources, and to tune MDB concurrency, most MDBs should be modified to reference a custom work manager that has a max-threads-constraint. In general, a custom work manager is useful when you have multiple MDB deployments, or if you determine that a particular MDB needs more threads.

See Also: For more information on how to use custom Work Managers to customize thread management, and when to use custom work managers, see the following:

- "Tune Pool Sizes" in *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*
- "Thread Management" in *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*
- "MDB Thread Management" in *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*
- "Using Work Managers to Optimize Scheduled Work" in *Oracle Fusion Middleware Configuring Server Environments for Oracle WebLogic Server*
- "Avoiding and Managing Overload" in *Oracle Fusion Middleware Configuring Server Environments for Oracle WebLogic Server*

You can use Oracle WebLogic Administration Console to view general information about the status of the thread pool (such as active thread count, total thread count, and queue length.) You can also use the Console to view each application's scoped work manager metrics from the Workload tab on the Monitoring page. The metrics provided include the number of pending requests and number of completed requests.

For more information, see "Servers: Monitoring: Threads" and "Deployments: Monitoring: Workload" in the *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*.

The work manager and thread pool metrics can also be viewed from the Oracle Fusion Middleware Control. For more information, see [Section 4.2.1, "Viewing Performance Metrics Using Fusion Middleware Control"](#).

2.9.4 Tuning Oracle WebCenter Concurrency

Oracle WebCenter has its own controls for managing concurrency. See "Advanced Systems Administration for Oracle WebCenter Portal" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.

2.9.5 Tuning BPEL Concurrency

The Oracle BPEL Process Manager has its own thread controls and specialized tuning.

2.10 Setting Logging Levels

The amount of information that applications log depends on how the environment is configured and how the application code is instrumented. To maximize performance it is recommended that the logging level is not set higher than the default `INFO` level logging. If the logging setting does not match the default level, reset the logging level to the default for best performance.

Once the application and server logging levels are set appropriately, ensure that the debugging properties or other application level debugging flags are also set to appropriate levels or disabled. To avoid performance impacts, do not set log levels to levels that produce more diagnostic messages, including the `FINE` or `TRACE` levels.

Each component may have specific recommendations for logging levels. See the component chapters in this book for more information.

Performance Planning

This chapter discusses performance and tuning concepts for Oracle Fusion Middleware. This chapter contains the following sections:

- [Section 3.1, "About Oracle Fusion Middleware Performance Planning"](#)
- [Section 3.2, "Performance Planning Methodology"](#)

3.1 About Oracle Fusion Middleware Performance Planning

To maximize Oracle Fusion Middleware performance, you must monitor, analyze, and tune all the components that are used by your applications. This guide describes the tools that you can use to monitor performance and the techniques for optimizing the performance of Oracle Fusion Middleware components.

Performance tuning usually involves a series of trade-offs. After you have determined what is causing the bottlenecks, you may have to modify performance in some other areas to achieve the expected results. However, if you have a clearly defined plan for achieving your performance objectives, the decision on what to trade for higher performance is easier because you have identified the most important areas.

3.2 Performance Planning Methodology

The Fusion Middleware components are built for performance and scalability. To maximize the performance capabilities of your applications, you must build performance and scalability into your design. The performance plan should address the current performance requirements, the existing issues (such as bottlenecks or insufficient hardware resources) and any anticipated variances in load, users or processes. The performance plan should also address how the components scale during peak usage without impacting performance.

The following sections of this chapter discuss the steps you should take to help create a plan to tune your application environment and optimize performance:

- Step 1: [Define Your Performance Objectives](#)
- Step 2: [Design Applications for Performance and Scalability](#)
- Step 3: [Monitor and Measure Your Performance Metrics](#)

3.2.1 Define Your Performance Objectives

Before you can begin performance tuning your applications, you must first identify the performance objectives you hope to achieve. To determine your performance

objectives, you must understand the applications deployed and the environmental constraints placed on the system.

To understand what your performance objectives are, you must complete the following steps:

- [Define Operational Requirements](#)
- [Identify Performance Goals](#)
- [Understand User Expectations](#)
- [Conduct Performance Evaluations](#)

Performance objectives are limited by constraints, such as:

- The configuration of hardware and software such as CPU type, disk size, disk speed, and sufficient memory.

There is no single formula for determining your hardware requirements. The process of determining what type of hardware and software configuration is required to meet application needs adequately is called *capacity planning*.

Capacity planning requires assessment of your system performance goals and an understanding of your application. Capacity planning for server hardware should focus on maximum performance requirements. For more information on capacity planning, see [Chapter 29, "Capacity Planning"](#).

- The configuration of high availability architecture to address peak usage and response times. For more information on implementing high availability features in Oracle Fusion Middleware applications, see [Chapter 30, "Using Clusters and High Availability Features"](#).
- The ability to interoperate between domains, use legacy systems, support legacy data.
- Development, implementation, and maintenance costs.

Understanding these constraints - and their impacts - ensure that you set realistic performance objectives for your application environment, such as response times, throughput, and load on specific hardware.

3.2.1.1 Define Operational Requirements

Before you begin to deploy and tune your application on Oracle Fusion Middleware, it is important to clearly define the operational environment. The operational environment is determined by high-level constraints and requirements such as:

- Application Architecture
- Security Requirements
- Hardware Resources

3.2.1.2 Identify Performance Goals

Whether you are designing a new system or maintaining an existing system, you should set specific performance goals so that you know how and what to optimize. To determine your performance objectives, you must understand the application deployed and the environmental constraints placed on the system.

Gather information about the levels of activity that components of the application are expected to meet, such as:

- Anticipated number of users

- Number and size of requests
- Amount of data and its consistency
- Target CPU utilization


3.2.1.3 Understand User Expectations

Application developers, database administrators, and system administrators must be careful to set appropriate performance expectations for users. When the system carries out a particularly complicated operation, response time may be slower than when it is performing a simple operation. Users should be made aware of which operations might take longer.

For example, you might want to ensure that 90% of the users experience response times no greater than 5 seconds and the maximum response time for all users is 20 seconds. Usually, it's not that simple. Your application may include a variety of operations with differing characteristics and acceptable response times. You need to set measurable goals for each of these.

You also need to determine how variances in the load can affect the response time. For example, users might access the system heavily between 9:00am and 10:00am and then again between 1:00pm and 2:00pm, as illustrated by the graph in [Figure 3-1](#). If your peak load occurs on a regular basis, for example, daily or weekly, the conventional wisdom is to configure and tune systems to meet your peak load requirements. The lucky users who access the application in off-time can experience better response times than your peak-time users. If your peak load is infrequent, you may be willing to tolerate higher response times at peak loads for the cost savings of smaller hardware configurations.

Figure 3-1 Adjusting Capacity and Functional Demand



3.2.1.4 Conduct Performance Evaluations

With clearly defined performance goals and performance expectations, you can readily determine when performance tuning has been successful. Success depends on the functional objectives you have established with the user community, your ability to measure whether the criteria are being met, and your ability to take corrective action to overcome any exceptions.

Ongoing performance monitoring enables you to maintain a well-tuned system. Keeping a history of the application's performance over time enables you to make useful comparisons. With data about actual resource consumption for a range of loads, you can conduct objective scalability studies and from these predict the resource requirements for anticipated load volumes. For more information on evaluating performance, see [Chapter 4, "Monitoring Oracle Fusion Middleware"](#).

3.2.2 Design Applications for Performance and Scalability

The key to good performance is good design. The design phase of the application development cycle should be an on-going process. Cycling through the planning, monitoring and tuning phases of the application development cycle is critical to achieving optimal performance across Fusion Middleware deployments. Using an iterative design methodology enables you to accommodate changes in your work loads without impacting your performance objectives.

See the following Oracle Fusion Middleware developer's documentation for more information on recommended design techniques:

- *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*
- *Oracle Fusion Middleware Developer's Guide for Oracle WebCenter Portal*
- *Oracle Fusion Middleware Developer's Guide for Oracle TopLink*
- *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*
- *Oracle Fusion Middleware Application Developer's Guide for Oracle Identity Management*

3.2.3 Monitor and Measure Your Performance Metrics

Oracle Fusion Middleware provides a variety of technologies and tools that can be used to monitor Server and Application performance. Monitoring enables you to evaluate Server activity, watch trends, diagnose system bottlenecks, debug applications with performance problems and gather data that can assist you in tuning the system. For more information, see [Chapter 4, "Monitoring Oracle Fusion Middleware"](#).

Performance tuning is specific to the applications and resources that you have deployed on your system. Some common tuning areas are included in [Chapter 2, "Top Performance Areas"](#).

See Also: *Oracle Database Performance Tuning Guide*

Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server

Oracle Fusion Middleware Administrator's Guide

Monitoring Oracle Fusion Middleware

Oracle Fusion Middleware provides a variety of technologies and tools that can be used to monitor Server and Application performance. Monitoring is an important step in performance tuning and enables you to evaluate server activity, watch trends, diagnose system bottlenecks, debug applications with performance problems and gather data that can assist you in tuning the system.

This chapter contains the following sections:

- [Section 4.1, "About Oracle Fusion Middleware Management Tools"](#)
- [Section 4.2, "Oracle Enterprise Manager 11g Fusion Middleware Control"](#)
- [Section 4.3, "Oracle WebLogic Server Administration Console"](#)
- [Section 4.4, "WebLogic Diagnostics Framework \(WLDF\)"](#)
- [Section 4.5, "WebLogic Scripting Tool \(WLST\)"](#)
- [Section 4.6, "DMS Spy Servlet"](#)
- [Section 4.7, "Oracle Process Manager and Notification Server"](#)
- [Section 4.8, "Oracle Enterprise Manager Cloud Control"](#)
- [Section 4.9, "Native Operating System Performance Commands"](#)
- [Section 4.10, "Network Performance Monitoring Tools"](#)

Note: Additional monitoring information is included for most products in the product-specific chapters of this guide.

4.1 About Oracle Fusion Middleware Management Tools

After you install and configure Oracle Fusion Middleware, you can use the graphical user interfaces or command-line tools to manage your environment.

You can use the following tools to manage your Oracle Fusion Middleware installations:

- Oracle Enterprise Manager Fusion Middleware Control. See [Section 4.2](#).
- Oracle WebLogic Server Administration Console. See [Section 4.3](#).
- Oracle WebLogic Diagnostics Framework (WLDF). See [Section 4.4](#).
- Oracle WebLogic Scripting Tool (WLST). See [Section 4.5](#).
- DMS Spy Servlet. See [Section 4.6](#).
- Oracle Process Manager and Notification Server. See [Section 4.7](#).

- Oracle Enterprise Manager Cloud Control. See [Section 4.8](#).
- Operating System Performance Commands. See [Section 4.9](#).
- Network Performance Monitoring Tools. See [Section 4.10](#).

Use these tools, rather than directly editing configuration files, to perform all administrative tasks unless a specific procedure requires you to edit a file. Editing a file may cause the settings to be inconsistent and generate problems.

Both Fusion Middleware Control and Oracle WebLogic Server Administration Console are graphical user interfaces that you can use to monitor and administer your Oracle Fusion Middleware environment. You can perform some tasks with either tool, but, for other tasks, you can only use one of the tools.

For more information on using WebLogic Server Administration Console for monitoring your domain, see the *Oracle Fusion Middleware Administrator's Guide*.

4.1.1 Measuring Your Performance Metrics

Metrics are the criteria you use to measure your scenarios against your performance objectives. You can use performance metrics to help locate bottlenecks, identify resource availability issues, or help tune your components to improve throughput and response times. After you have determined your performance criteria, take measurements of the metrics used to quantify your performance objectives.

For example, you might use response time, throughput, and resource utilization as your metrics. The performance objective for each metric is the value that is acceptable. You match the actual value of the metrics to your objectives to verify that you are meeting, exceeding, or failing to meet your performance objectives.

When you manage or monitor an Oracle Fusion Middleware component or application with Fusion Middleware Control, you may see performance metrics that provide insight into the current performance of the component or application. In many cases, these metrics are shown in interactive charts; other times they are presented in tabular format. The best way to use and correlate the performance metrics is from the Performance Summary page for the component or application you are monitoring.

The next sections of this chapter provide an overview of the Oracle Fusion Middleware technologies and tools that can be used to monitor Server and Application performance. If you are new to Oracle Fusion Middleware or if you need additional information about monitoring your environment using the Performance Summary pages, see "Viewing the Performance of Oracle Fusion Middleware" in the *Oracle Fusion Middleware Administrator's Guide*. In addition, the Fusion Middleware Control online help provides definitions and other information about specific performance metrics that are available on its management and monitoring pages. See [Section 4.2.1, "Viewing Performance Metrics Using Fusion Middleware Control"](#).

4.2 Oracle Enterprise Manager 11g Fusion Middleware Control

Fusion Middleware Control is a Web browser-based, graphical user interface that you can use to monitor and administer a farm. Fusion Middleware Control organizes a wide variety of performance data and administrative functions into distinct, Web-based home pages for the farm, domain, servers, components, and applications. The Fusion Middleware Control home pages make it easy to locate the most important monitoring data and the most commonly used administrative functions—all from your Web browser.

In addition, Fusion Middleware Control provides a set of MBean browsers that allow you to browse the MBeans for a WebLogic Server or for a selected application and perform specific monitoring and configuration tasks from the MBean browser.

See Also: "Getting Started Using Oracle Enterprise Manager Fusion Middleware Control" in *Oracle Fusion Middleware Administrator's Guide*.

Use Fusion Middleware Control to:

- Monitor and administer a single Fusion Middleware Farm
- Monitor all elements of the farm - including deployed applications and Fusion Middleware components such as:
 - WebLogic Domain
 - Cluster and Managed Servers
 - SOA components
 - Web Center
 - Web Cache
 - Oracle HTTP Server
 - Oracle Identity Management
- Monitor the state and performance of each of these targets by providing out-of-the-box performance metrics
- Monitor CPU usage, heap usage, Work Manager, JMS servers, and JDBC and JTA usage for Oracle WebLogic Server
- Monitor JVM performance in terms of heap versus non-heap usage, garbage collection, and threads performance
- Monitor applications and Web services deployed to WebLogic Server
- Monitor a wide range of application metrics for servlets, JSPs, and EJBs are available, as well as Web services metrics for faults, invocations, and violations. Such metrics are accessible from a target's home page.
- Access customizable performance summary pages to help administrators monitor performance and diagnose problems. These charts can be modified to display content that is relevant to your domain. A target or component might be added to the chart so that you can compare the performance information for two targets in one chart.

See Also: For more information about monitoring your environment using the Performance Summary pages, see "Viewing the Performance of Oracle Fusion Middleware" in *Oracle Fusion Middleware Administrator's Guide*.

4.2.1 Viewing Performance Metrics Using Fusion Middleware Control

When you manage or monitor an Oracle Fusion Middleware component or application with Fusion Middleware Control, you often see performance metrics that provide insight into the current performance of the component or application. In many cases, these metrics are shown in interactive charts; other times they are presented in tabular

format. The best way to use and correlate the performance metrics is from the Performance Summary page for the component or application you are monitoring.

Use the Fusion Middleware Control online help to obtain a definition of a specific performance metric. There are two ways to access this information:

- Browse or search for the metric in the Fusion Middleware Control online help.
- Navigate to the Performance Summary page for your Oracle Fusion Middleware component or application and do the following:
 1. Click **Show Metric Palette**.
 2. Browse the list of metrics available for the component or application to locate a specific metric.
 3. Right-click the name of the metric and select **Help** from the context menu.

If you encounter a problem, such as an application that is running slowly or is hanging, you can view more detailed performance information, including performance metrics for a particular target, to find out more information about the problem.

Oracle Fusion Middleware automatically and continuously measures run-time performance. The performance metrics are automatically enabled; you do not need to set options or perform any extra configuration to collect them.

4.3 Oracle WebLogic Server Administration Console

Oracle WebLogic Server Administration Console is a Web browser-based, graphical user interface that you use to manage an Oracle WebLogic Server domain. It is accessible from any supported Web browser with network access to the Administration Server.

See Also: For general information on using the WebLogic Server console, see "Getting Started Using Oracle WebLogic Server Administration Console" in *Oracle Fusion Middleware Administrator's Guide*.

Use the WebLogic Server Administration Console to:

- Configure, start, and stop WebLogic Server instances
- Configure and Monitor WebLogic Server clusters
- Configure and Monitor WebLogic Server services, such as database connectivity (JDBC) and messaging (JMS)
- Configure security parameters, including creating and managing users, groups, and roles
- Configure and deploy Java EE applications
- Monitor server and application performance
- View server and domain log files
- View application deployment descriptors
- Edit selected run-time application deployment descriptor elements

Oracle WebLogic Server contains a Java Management Extensions (JMX) server implementation and provides its own set of Management Beans (MBeans). Oracle

management tools described in this chapter use the MBeans provided by WebLogic Server to allow you to configure, monitor, and manage WebLogic Server resources.

Additional WebLogic Server Console Resources:

For details on the content contained in each summary table, see "Monitor Servers" in WebLogic Administration Console Online Help.

For detailed information on using the WebLogic Server to monitor your domain, see the *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*.

The Oracle Technology Network at <http://www.oracle.com/technology/index.html> provides product downloads, articles, sample code, product documentation, tutorials, white papers, news groups, and other key content for WebLogic Server.


4.4 WebLogic Diagnostics Framework (WLDF)

The WebLogic Diagnostic Framework (WLDF) is a monitoring and diagnostic framework that can collect diagnostic data that servers and applications generate. The WLDF can be configured to collect the data and store it in various sources, including log records, data events, and harvested metrics.

WLDF includes several components for collecting and analyzing data:

- **Data Creators**— data publishers and data providers that are distributed across WLDF components.
- **Diagnostic Image Capture**—Creates a diagnostic snapshot from the server that can be used for post-failure analysis.
- **Archive**—Captures and persists data events, log records, and metrics from server instances and applications.
- **Instrumentation**—Adds diagnostic code to WebLogic Server instances and the applications running on them to execute diagnostic actions at specified locations in the code. The Instrumentation component provides the means for associating a diagnostic context with requests so they can be tracked as they flow through the system.
- **Harvester**—Captures metrics from run-time MBeans, including WebLogic Server MBeans and custom MBeans, which can be archived and later accessed for viewing historical data.
- **Watches and Notifications**—Provides the means for monitoring server and application states and sending notifications based on criteria set in the watches. (A watch rule can monitor log data, event data from the Instrumentation component, or metric data from a data provider that is harvested by the Harvester. The Watch Manager is capable of managing watches that are composed of several watch rules.)
- **Logging services**—Manage logs for monitoring server, subsystem, and application events.

The relationship among these components is shown in [Figure 4-1](#).

Figure 4–1 Major WLDF Components

All of the framework components operate at the server level and are only aware of server scope. All the components exist entirely within the server process and participate in the standard server lifecycle. All artifacts of the framework are configured and stored on a per server basis.

Note: For more information on the WebLogic Diagnostics Framework and how it can be leveraged for monitoring Oracle Fusion Middleware components, see *Oracle Fusion Middleware Configuring and Using the Diagnostics Framework for Oracle WebLogic Server*.

4.5 WebLogic Scripting Tool (WLST)

The Oracle WebLogic Scripting Tool (WLST) is a command-line scripting environment that you can use to create, manage, and monitor Oracle WebLogic Server domains. It is based on the Java scripting interpreter, Jython. In addition to supporting standard Jython features such as local variables, conditional variables, and flow control statements, WLST provides a set of scripting functions (commands) that are specific to WebLogic Server. You can extend the WebLogic scripting language to suit your needs by following the Jython language syntax.

You can use any of the following techniques to invoke WLST commands:

- Interactively, on the command line
- In script mode, supplied in a file
- Embedded in Java code

See Also:

- *Oracle Fusion Middleware WebLogic Scripting Tool Command Reference*
 - "Using Custom WLST Commands" in *Oracle Fusion Middleware Administrator's Guide*
-

4.5.1 Using Custom WLST Commands

Many components, such as Oracle SOA Suite, Oracle Platform Security Services (OPSS), Oracle Fusion Middleware Audit Framework, and MDS, and services such as SSL and logging, supply custom WLST commands.

To use these custom WLST commands, you must invoke WLST from the Oracle home in which the component has been installed. See "Using Custom WLST Commands" in the *Oracle Fusion Middleware Administrator's Guide* for more information.

4.5.1.1 Using WLST Commands for System Components

In addition to the commands provided by WLST for Oracle WebLogic Server, WLST provides a subset of commands to monitor and manage system components. These commands are:

- `startproc(componentName [, componentType] [, componentSet)`: Starts the specified component.
- `stopproc(componentName [, componentType] [, componentSet)`: Stops the specified component.
- `status(componentName [, componentType] [, componentSet)`: Obtains the status of the specified component.
- `proclist()`: Obtains the list of components.
- `dumpMetrics([servers,] [format])`: Displays available metrics in the internal format, PDML, or in XML.
- `displayMetricTables([metricTable_1], [metricTable_2], [...], [servers] [variables])`: Displays the content of the DMS metric tables.
- `displayMetricTableNames([servers])`: Displays the names of the available DMS metric tables. The returned value is a string array containing metric table names.

Note: The `dmstool` command has been replaced with the following commands: `dumpMetrics`, `displayMetricTables`, `displayMetricTableNames`.

4.6 DMS Spy Servlet

The DMS Spy servlet provides access to DMS metric data from a web browser. Data that is created and updated by DMS-enabled applications and components is accessible through the DMS Spy Servlet.

4.6.1 Viewing Performance Metrics Using the Spy Servlet

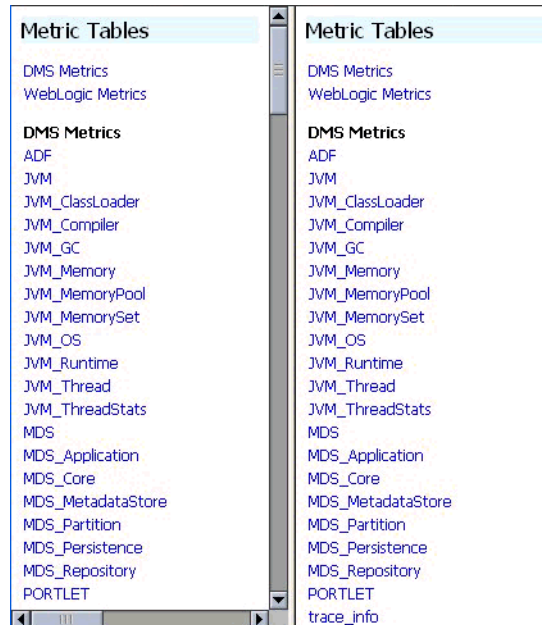
The DMS Spy Servlet is part of the DMS web application. The DMS web application's web archive file is `dms.war`, and can be found in the same directory as `dms.jar`:
`<ORACLE_HOME>/modules/oracle.dms_11.1.1/dms.war.`

The DMS web application is deployed by default as part of a JRF-enabled server instance. The URL is: `http://host:port/dms/Spy`.

Only users who have Administrator role access can view this URL as access is controlled by standard Java EE elements in `web.xml`.

4.6.2 Using the DMS Spy Servlet

Figure 4–2 shows the initial page of the Spy servlet: both sides show the same list of metric tables.

Figure 4–2 Spy Servlet Page - Metrics Tables

Note that the Spy servlet can display metric tables for WebLogic Server and also for non-Java EE components that are deployed.

For metric tables to appear in the Spy servlet, the component that creates and updates that table must be installed and running. Metric tables for components that are not running are not displayed. Metric tables with ":" in their name (for example, `weblogic_j2eeserver:app_overview`) are aggregated metric tables generated by metric rules.

To view the contents of a metric table, click the table name. For example, [Figure 4–3](#) shows the `MDS_Partition` table.

Figure 4–3 MDS Partition Table

MDS_Partition								
Name	Host	Process	readDocument	writeDocument	MDS_Application	MDS_Repository	ServerName	
oracle		WLS_Spaces:8888	active, threads avg, msec completed, ops maxActive, threads maxTime, msec minTime, msec time, msec	0 active, threads 0.106 avg, msec 254 completed, ops 1 maxActive, threads 5 maxTime, msec 0 minTime, msec 27 time, msec	0 0 0 0 0 0	webcenter(11.1.1.2.0)		WLS_Spaces
owsm		WLS_Spaces:8888	active, threads avg, msec completed, ops maxActive, threads maxTime, msec minTime, msec time, msec	0 active, threads 0 avg, msec 0 completed, ops 0 maxActive, threads 0 maxTime, msec 0 minTime, msec 0 time, msec	0 10.66 100 1 69 4 1066	wsm-pm	oracle	WLS_Spaces
webcenter		WLS_	active,	0 active,	0	webcenter(11.	mids-	WLS_

To get a description of the fields in a metric table, click the Metric Definitions link below the table.

4.7 Oracle Process Manager and Notification Server

Oracle Process Manager and Notification Server (OPMN) monitors the status of Oracle Fusion Middleware components. You can also start and stop system components, monitor system components, and perform many other tasks related to process management. For example, you can use OPMN to start and stop OPMN-managed processes, such as Oracle HTTP Server and Oracle Web Cache. For more information on OPMN commands, see "[Section 6.2, "Monitoring Oracle HTTP Server Performance"](#)".

Note: For more information on using OPMN, refer to *Oracle Fusion Middleware Oracle Process Manager and Notification Server Administrator's Guide*.

4.8 Oracle Enterprise Manager Cloud Control

Oracle Enterprise Manager is Oracle's integrated enterprise information technology (IT) management product line, which provides the industry's only complete, integrated, and business-driven enterprise cloud management solution. Oracle Enterprise Manager creates business value for IT by leveraging the built-in management capabilities of the Oracle stack for traditional and cloud environments, enabling customers to achieve unprecedented efficiency gains while dramatically increasing service levels.

The key capabilities of Enterprise Manager include:

- A complete cloud lifecycle management solution enabling you to quickly set up, manage, and support enterprise clouds and traditional Oracle IT environments from applications to disk.
- Maximum return on IT management investment through the best solutions for intelligent management of the Oracle stack and engineered systems with real-time integration of Oracle's knowledge base with each customer environment.
- Best service levels for traditional and cloud applications through business-driven application management

For more information about the Oracle Enterprise Manager Cloud Control, refer to *Oracle Enterprise Manager Cloud Control Introduction*.

4.9 Native Operating System Performance Commands

Each operating system has native tools and utilities that can be useful for monitoring purposes. Native operating system commands enable you to gather and monitor for example CPU utilization, paging activity, swapping, and other system activity information.

For details on operating system commands, refer to the documentation provided by the operating system vendor.

4.10 Network Performance Monitoring Tools

Your operating system's network monitoring tools can be used to monitor utilization, verify that the network is not becoming a bottleneck, or detect packet loss or other network performance issues. For details on network performance monitoring, refer to your operating system documentation.

Part II

Core Components

This part describes configuring core components to improve performance. It contains the following chapters:

- [Chapter 5, "Understanding the Oracle Dynamic Monitoring Service"](#)
- [Chapter 6, "Oracle HTTP Server Performance Tuning"](#)
- [Chapter 7, "Oracle Metadata Service \(MDS\) Performance Tuning"](#)

Note: For information on performance tuning the Oracle WebLogic Server, see *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*.

Understanding the Oracle Dynamic Monitoring Service

This chapter provides an overview and features available in the Oracle Dynamic Monitoring Service (DMS).

- [Section 5.1, "About Dynamic Monitoring Service \(DMS\)"](#)
- [Section 5.2, "Understanding DMS Availability"](#)
- [Section 5.3, "Understanding DMS Architecture"](#)
- [Section 5.4, "Viewing DMS Metrics"](#)
- [Section 5.5, "Accessing DMS Metrics with WLDF"](#)
- [Section 5.6, "DMS Execution Context"](#)
- [Section 5.7, "DMS Tracing and Events"](#)
- [Section 5.8, "DMS Best Practices"](#)

5.1 About Dynamic Monitoring Service (DMS)

The Oracle Dynamic Monitoring Service (DMS) enables Oracle Fusion Middleware components to provide administration tools, such as Oracle Enterprise Manager, with data regarding the component's performance, state and on-going behavior. Fusion Middleware Components push data to DMS and in turn DMS publishes that data through a range of different components. Specifically, DMS is used by Oracle WebCache, Oracle HTTP Server (OHS), Oracle Application Development Framework (ADF), WebLogic Diagnostic Framework (WLDF), and JDBC. DMS measures and reports metrics, trace events and system performance and provides a context correlation service for these components.

5.1.1 Understanding Common DMS Terms and Concepts

This section defines common DMS terms and concepts related to the following:

- [DMS Tracing and Events](#)
- [DMS Nouns](#)
- [DMS Sensors](#)

5.1.1.1 DMS Tracing and Events

[Table 5–1](#) provides a list of DMS tracing and event terminology.

Table 5–1 DMS Tracing and Event Terminology

DMS Term	Definition
Condition	<p>A condition is the logic behind a condition filter. It determines which events may pass through a filter, based on the rules defined in the condition. Every condition filter has zero or one root condition, but conditions may include AND or OR arguments together to create compound conditions. The single root condition can describe a relatively complex rule.</p> <p>Two types of condition exist:</p> <ul style="list-style-type: none"> ▪ Noun Type Condition - operates on the name of the noun type associated with a sensor or noun event. ▪ Context Condition - operates on the values currently set within the current Execution Context. <p>For more information on using conditions, see Section 5.7, "DMS Tracing and Events".</p>
Destination	<p>A destination implements a mechanism for reacting to events that are passed to it. For example, a destination could log events to a file, another could send transformed copies of event to the JRockit Flight Recorder, yet another might render information gleaned from incoming events as data in an MBean.</p>
Event Route	<p>An event route connects a filter to a destination. Event routes may be enabled or disabled. For event tracing to be activated for a specific filter, one or more event routes must exist for that filter and must be enabled.</p>
Filter	<p>An event tracing filter selectively passes a subset of all possible DMS runtime events. Filters can be configured with rules that determine which events are passed and which are blocked.</p> <p>For example it is possible to define filters to:</p> <ul style="list-style-type: none"> ▪ Only pass sensor updates that are made when the execution context has a key-value pair of "role"-"admin" ▪ Only pass sensor updates from nouns of type "JDBC_Statement" <p>For more information on using filters, see Section 5.7, "DMS Tracing and Events".</p>
Listener	<p>A DMS listener is also known as the destination. See Section 5.7.2, "Configuring Destinations" for more information.</p>

5.1.1.2 DMS Nouns

DMS **nouns** organize performance data. Sensors, with their associated metrics, are organized in an hierarchy according to nouns. Nouns enable you to organize DMS metrics in a manner comparable to a directory structure in a file system. For example, nouns can represent classes, methods, objects, queues, connections, applications, databases, or other objects that you want to measure.

A **noun type** is a name that reflects the set of metrics being collected.

5.1.1.2.1 General DMS Naming A **noun name** is a simple string, not including a delimiter. For example, `BasicBinomial` is a noun name. A noun full name consists of the noun name with the namespace and localpart. The noun name is preceded by the full name of its parent, and a delimiter.

`/dmsDemo/BasicBinomial/"{http://mynamespace/}JAXWSHelloService"` is a noun full name.

A **sensor name** is a simple string, not including the "." or the derivation. For example, `computeSeries`, `loops`, and `lastComputed` are sensor names.

A **sensor full name** consists of the sensor name, preceded by the name of its associated noun, and a delimiter. Examples: `/dmsDemo/BasicBinomial/computeSeries`, `/dmsDemo/BasicBinomial/loops`, `/dmsDemo/BasicBinomial/lastComputed`.

A **DMS metric name** consists of a sensor name plus the "." character plus the metric. For example, `computeSeries.time`, `loops.count`, and `lastComputed.value` are valid DMS metric names.

Note: The suffixes `.time`, `.count`, and `.value` are immutable. Sensor and noun names, however, can be modified as needed.

5.1.1.2.2 General DMS Naming Conventions and Character Sets DMS names should be as compact as possible. When you define noun and sensor names, avoid special characters such as white space, slashes, periods, parenthesis, commas, and control characters.

Table 5–2 shows DMS replacement for special characters in names.

Table 5–2 Replacement for Special Characters in DMS Names

Character	DMS Replacement Character
Space character	Underscore character: <code>_</code>
Period character: <code>.</code>	Underscore character: <code>_</code>
Control character	Underscore character: <code>_</code>
Less than character: <code><</code>	Open parenthesis: <code>(</code>
Greater than character: <code>></code>	Close parenthesis: <code>)</code>
Ampersand: <code>&</code>	Caret: <code>^</code>
Double quote: <code>"</code>	Backquote: <code>'</code>
Single quote: <code>'</code>	Backquote: <code>'</code>

Note: Oracle Fusion Middleware includes several built-in metrics. The Oracle Fusion Middleware built-in metrics do not always follow the DMS naming conventions.

5.1.1.2.3 Noun and Noun Type Naming Conventions The following conventions are used when naming noun and noun types:

- A noun name should be unique.
- A noun name should identify a specific entity of interest.
- Noun types should have names that clearly reflect the set of metrics being collected. For example, `Servlet` is the type for a noun under which the metrics that are specific to a given servlet fall.
- Noun type names should start with a capital letter to distinguish them from other DMS names. All nouns of a given type should contain the same set of sensors.
- The noun naming scheme uses a '/' as the root of the hierarchy, with each noun acting as a container under the root, or under its parent noun.

5.1.1.3 DMS Sensors

DMS **sensors** measure performance data and enable DMS to define and collect a set of metrics. Certain metrics are always included with a sensor and other metrics are optionally included with a sensor.

DMS has three different kinds of sensors:

- [Section 5.1.1.3.1, "DMS PhaseEvent Sensors"](#)
- [Section 5.1.1.3.2, "DMS Event Sensors"](#)
- [Section 5.1.1.3.3, "DMS State Sensors"](#)

5.1.1.3.1 DMS PhaseEvent Sensors A DMS **PhaseEvent sensor** measures the time spent in a specific section of code that has a beginning and an end. Use a PhaseEvent sensor to track time in a method or in a block of code.

DMS can calculate optional metrics associated with a PhaseEvent, including the average, maximum, and minimum time that is spent in the PhaseEvent sensor.

[Table 5–3](#) lists the metrics available with PhaseEvent sensors.

Table 5–3 DMS PhaseEvent Sensor Metrics

Metric	Description
<code>sensor_name.time</code>	Specifies the total time spent in the phase <code>sensor_name</code> . Default metric: <code>time</code> is a default PhaseEvent sensor metric.
<code>sensor_name.completed</code>	Specifies the number of times the phase <code>sensor_name</code> has completed since the process was started. Optional metric
<code>sensor_name.minTime</code>	Specifies the minimum time spent in the phase <code>sensor_name</code> , for all the times the <code>sensor_name</code> phase completed. Optional metric
<code>sensor_name.maxTime</code>	Specifies the maximum time spent in the phase <code>sensor_name</code> , for all the times the <code>sensor_name</code> phase completed. Optional metric
<code>sensor_name.avg</code>	Specifies the average time spent in the phase <code>sensor_name</code> , computed as the (total time)/(number of times the phase completed). Optional metric
<code>sensor_name.active</code>	Specifies the number of threads in the phase <code>sensor_name</code> , at the time the DMS statistics are gathered (the value may change over time). Optional metric
<code>sensor_name.maxActive</code>	Specifies the maximum number of concurrent threads in the phase <code>sensor_name</code> , since the process started. Optional metric

5.1.1.3.2 DMS Event Sensors A DMS **event sensor** counts system events. Use a DMS event sensor to track system events that have a short duration, or where the duration of the event is not of interest but the occurrence of the event is of interest.

[Table 5–4](#) describes the metric that is associated with an event sensor.

Table 5–4 DMS Event Sensor Metrics

Metric	Description
<code>sensor_name.count</code>	Specifies the number of times the event has occurred since the process started, where <code>sensor_name</code> is the name of the Event sensor as specified in the DMS instrumentation API. Default: <code>count</code> is the default metric for an event sensor. No other metrics are available for an event sensor.

5.1.1.3.3 DMS State Sensors A DMS **state sensor** tracks the value of Java primitives or the content of a Java object. Supported types include integer, double, long, and object. Use a state sensor when you want to track system status information or when you need a metric that is not associated with an event. For example, use state sensors to track queue lengths, pool sizes, buffer sizes, or host names. You assign a precomputed value to a state sensor.

Table 5–5 describes the state sensor metrics. State sensors support a default metric value, as well as optional metrics. The optional `minValue` and `maxValue` metrics only apply for state sensors if the state sensor represents a numeric Java primitive (of type integer, double, or long).

Table 5–5 DMS State Sensor Metrics

Metric	Description
<code>sensor_name.value</code>	Specifies the metric value for <code>sensor_name</code> , using the type assigned when <code>sensor_name</code> is created. Default: <code>value</code> is the default State metric.
<code>sensor_name.count</code>	Specifies the number of times <code>sensor_name</code> is updated. Optional metric
<code>sensor_name.minValue</code>	Specifies the minimum value for <code>sensor_name</code> since startup. Optional metric
<code>sensor_name.maxValue</code>	Specifies the maximum value this <code>sensor_name</code> since startup. Optional metric

5.1.1.3.4 Sensor Naming Conventions The following list describes DMS sensor naming conventions:

- Sensor names should be descriptive, but not redundant. Sensor names should not contain any part of the noun name hierarchy, or type, as this is redundant.
- Sensor names should avoid containing the value for the individual metrics.
- Where multiple words are required to describe a sensor, the first word should start with a lowercase letter, and the following words should start with uppercase letters. Example: `computeSeries`
- In general, avoid using a "/" character in a sensor name. However, there are cases where it makes sense to use a name that contains "/". If "/" is used in a noun or sensor name, then when you use the sensor in a string with DMS methods, you need to use an alternative delimiter, such as ";" or "_", which does not appear anywhere in the path; this enables the "/" to be properly understood as part of the noun or sensor name rather than as a delimiter.

For example, a child noun can have a name such as:

```
examples/jsp/num/numguess.jsp
```

and you can look this up using the string:

```
,default,WEBs,defaultWebApp,JSPs,example/jsp/num/numguess.jsp,service
```

where the delimiter is the "," character.

- Event sensor and PhaseEvent sensor names should have the form *verbnoun*. Examples: `activateInstance` and `runMethod`. When a PhaseEvent monitors a function, method, or code block, it should be named to reflect the task performed as clearly as possible.
- The name of a state sensor should be a noun, possibly preceded by an adjective, which describes the semantics of the value which is tracked with this state sensor. Examples: `lastComputed`, `totalMemory`, `port`, `availableThreads`, `activeInstances`
- To avoid confusion, do not name sensors with strings such as ".time", ".value", or ".avg", which are names of sensor metrics, as shown in [Table 5-3](#), [Table 5-4](#), and [Table 5-5](#).

5.2 Understanding DMS Availability

DMS functionality is available on all certified Java EE servers. This includes both the runtime features and supporting commands. Also, several features of DMS will operate in JSE applications and standalone C applications.

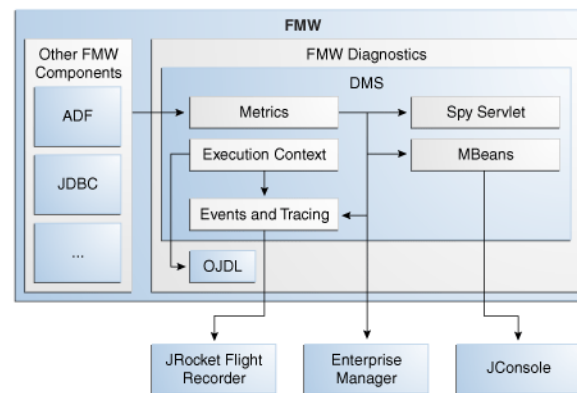
For more information, see the Oracle Fusion Middleware Certification Matrix at http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html.

5.3 Understanding DMS Architecture

DMS consists of the following features:

- **DMS Metrics** - The DMS metrics feature provides Java and C APIs that are also used by other Oracle Fusion Middleware components for instrumenting code with performance measurements and other useful state metrics. In addition, the metrics feature provides an aggregation language for computing derived metrics and tools for accessing the metrics.
- **Execution Context** - Execution Context supports the maintenance and propagation of a specific context structure throughout the Oracle stack. By making the context structure available consistently across all Oracle code, the potential for cross component and cross product correlation of diagnostic data increases. For more information see [Section 5.6, "DMS Execution Context"](#).
- **Events and Tracing** - Event Tracing enables you to configure live tracing with no restarts. DMS metrics updated during the course of using Oracle Fusion Middleware products may be traced using the DMS Event Tracing feature. The system has been designed to facilitate not only tracing, but also to support other functionality that may be driven from DMS activity.

[Figure 5-1](#) shows the components of DMS and how they interact with other Oracle Fusion Middleware components. Arrows show the direction in which information flows from one component to the next.

Figure 5–1 DMS Interactions with Oracle Fusion Middleware Components

5.4 Viewing DMS Metrics

Oracle Fusion Middleware components are instrumented with DMS metrics in order to collect information that developers, system administrators, and support analysts can use to analyze system performance or monitor system status. The Fusion Middleware Control online help provides information on each of the specific metrics. See [Section 4.2.1, "Viewing Performance Metrics Using Fusion Middleware Control"](#) for information on accessing metric information.

The Oracle Fusion Middleware metrics come from various sources and locations. They include MBean attributes and DMS metrics. They also come from non-Java EE servers, such as Oracle HTTP servers and Oracle WebCache.

The following sections describe how to use various tools to view the DMS metrics:

- [Viewing Metrics Using the Spy Servlet](#)
- [Viewing Metrics with WLDF \(WebLogic Diagnostic Framework\)](#)
- [Viewing metrics with WLST \(Oracle WebLogic Server\)](#)
- [Viewing metrics with JConsole](#)
- [Viewing metrics with Oracle Enterprise Manager](#)
- [Viewing metrics using WSADMIN \(IBM WebSphere\)](#)

5.4.1 Viewing Metrics Using the Spy Servlet

The Spy Servlet is part of the DMS Application that is deployed by default on JRF-extended installations. The Spy Servlet is launched from `http://<host>:<port>/dms/Spy`. The default port for WebLogic is 7001.

The DMS Application's web archive file is `dms.war`, and can be found in the same directory as `dms.jar`: `oracle_common/modules/oracle.dms_11.1.1/dms.war`.

For more information see [Section 4.6, "DMS Spy Servlet"](#).

Note: The Spy Servlet is secured using standard Java EE declarative security in the web-application's `web.xml` file, and will only grant access to the Spy Servlet to members of the Administrator's group.

5.4.2 Viewing Metrics with WLDF (WebLogic Diagnostic Framework)

You can use WebLogic Diagnostic Framework (WLDF) to harvest DMS metrics from DMS metric MBeans. You can also use WLDF to monitor changes to the attribute value of an MBean. For more information see "Configuring the Harvester for Metric Collection" in *Oracle Fusion Middleware Configuring and Using the Diagnostics Framework for Oracle WebLogic Server*.

5.4.3 Viewing metrics with WLST (Oracle WebLogic Server)

DMS provides three commands to view metrics in WLST:

Use this command...	To do this...
displayMetricTableNames	List the names of the available metric tables.
displayMetricTables	Show the content of the DMS metric tables.
dumpMetrics	Display metrics in the internal format. Valid formats for the dumpMetrics command include raw, xml and pdml.

As well as displaying textual output, these commands also return a structured object or single value that you can use in a script to process.

For more information on using these commands, see [Section 4.5.1.1, "Using WLST Commands for System Components"](#).

5.4.4 Viewing metrics with JConsole

To provide a standards-based way to access metrics, DMS exposes them through MBeans. An MBean will be created and registered for each typed noun with the runtime MBean Server. The DMS sensors contained by the noun are exposed as the attributes of the MBean. Exposing the DMS metrics as MBeans allows administrators to use tools such as JConsole (the Java monitoring and management console), and other Java Management Extension (JMX) clients, to access the DMS metrics.

MBeans also allow for integration with other Oracle diagnostics software such as WLDF (WebLogic Diagnostics Framework), which is described in [Section 5.5](#). The noun name and noun type are exposed as the name and type properties of the metric MBean object name. The MBean domain name is "oracle.dms". The object name also reflects the DMS noun hierarchy.

Note: You can use JConsole to view DMS generated MBeans on a Java EE server either locally or remotely. DMS generates an MBean for each Java DMS noun that has a valid noun type. It does not generate MBeans for the non-Java EE component's metrics and the DMS nouns that have no noun types. Each DMS metric contained under the noun is mapped to an attribute in the metric MBean.

5.4.5 Viewing metrics with Oracle Enterprise Manager

Oracle Fusion Middleware automatically and continuously measures data regarding the component's performance, state and on-going behavior. The metrics are automatically enabled; there is no need to set options or perform any extra configuration to collect them. For more information see [Section 4.2.1, "Viewing Performance Metrics Using Fusion Middleware Control"](#).

5.4.6 Viewing metrics using WSADMIN (IBM WebSphere)

The following commands can be used with IBM WebSphere to display the following:

Use this command...	To do this...
OracleDMS.displayMetricTableNames()	List the names of the available metric tables.
OracleDMS.displayMetricTables()	Show the content of the DMS metric tables.
OracleDMS.dumpMetrics()	Display metrics in the internal format. Valid formats include raw, xml and pdml.

For more information on using IBM WebSphere, see "Managing Oracle Fusion Middleware on IBM WebSphere" in the *Oracle Fusion Middleware Third-Party Application Server Guide*.

5.5 Accessing DMS Metrics with WLDF

The WebLogic Diagnostics Framework (WLDF) provides a diagnostic feature that allows MBean attributes to be harvested and monitored for specific conditions. This provides a proactive way of monitoring activity in your environment and creating E-mail and JMX notifications when a condition is triggered.

The following steps describe how to configure WLDF to send an E-mail notification using the WebLogic Administration Console:

1. Select an existing or create a new Diagnostics Module from the Diagnostics screen.
2. Click on the **Watches and Notifications** tab.
3. Click **New**.
4. Enter a Watch Name and click **Next**.
5. Enter the text as the Watch Rule and click **Next**.

```
{ServerRuntime//[NOUNTYPE]oracle.dms:name=/starWars/alliance,type=NounType//forceBalance_value} = 'BAD'}
```

6. Select **Use a manual reset alarm** and click **Next**. The manual reset option means that once an E-mail is triggered, you must reset the watch using the WebLogic Administration Console.
7. Select the E-mail notification type and click **Finish**.

It is also possible to configure WLDF to collect the MBean data for offline storage and analysis. This is achieved by configuring a WLDF Diagnostic Module to collect specific MBean attributes, and can be done so using the WebLogic Administration Console.

For more information on using WLDF to harvest and monitor MBean data see *Oracle Fusion Middleware Configuring and Using the Diagnostics Framework for Oracle WebLogic Server*.

5.6 DMS Execution Context

The DMS execution context is the mechanism by which requests (such as HTTP or RMI requests) can be uniquely identified and thus tracked as they flow through the system. It also provides a means by which context information can be communicated between cooperating Fusion Middleware components involved in fulfilling requests.

5.6.1 DMS Execution Requests and Sub-Tasks

The DMS execution context has been developed with the understanding that a single request (or task) may form the root of a tree of sub-tasks that are coordinated to complete the request or root task. Consider the following examples of requests and their associated sub-tasks:

1. An HTTP request sent directly to Oracle WebLogic Server from a browser:
 - Root task only on Oracle WebLogic Server
2. An HTTP request sent through Oracle HTTP Server (acting as a reverse proxy) to Oracle WebLogic Server:
 - Root task on Oracle HTTP Server
 - Single sub-task on Oracle WebLogic Server
3. An HTTP request sent from Oracle HTTP Server (acting as a reverse proxy) to Oracle WebLogic Server that then requires invocation of two remote web services from Oracle WebLogic Server in order to fulfill the request:
 - Root task on Oracle HTTP Server
 - Single sub-task on Oracle WebLogic Server
 - Two sub-sub-tasks, one on each web service

A DMS execution context is composed of the following:

- A unique identifier, the ECID

The Execution Context ID (ECID) is unique for each new root task and is shared across the tree of tasks associated with the root task.
- A relationship identifier, the RID

The Relationship ID (RID) is an ordered set of numbers that describes the location of each task in the tree of tasks. The leading number is usually a zero. A leading number of 1 indicates that it has not been possible to track the location of the sub-task within the overall sub-task tree.
- A set of name-value pairs by which globally relevant data can be shared among Oracle Fusion Middleware components.

The following three scenarios illustrate how ECID and RID are used when an HTTP request is sent from Oracle HTTP Server (acting as a reverse proxy) to an Oracle WebLogic Server and the server requires invocation of two remote web services from Oracle WebLogic Server.

1. Root task on Oracle HTTP Server:
 - New ECID = B5C094FA...BE4AE8
 - Root RID = 0
2. Single sub-task on Oracle WebLogic Server:
 - Same ECID = B5C094FA...BE4AE8
 - Sub-task RID = 0:1
3. Two Sub-tasks, one on each web service:
 - First web service invoked
 - Same ECID = B5C094FA...BE4AE8
 - Sub-task RID = 0:1:1

- Second web service invoked
Same ECID = B5C094FA...BE4AE8
Sub-task RID = 0:1:2

5.6.2 DMS Execution Context Usage

The most immediate benefits of the DMS execution context are realized when attempting to correlate log messages between servers. The Oracle standard format for logging involves a field dedicated to the ECID. Once the ECID is known, when its read from an ERROR level log message for example, it is possible to locate all other log messages associated with that task by querying the log files for messages containing that ECID.

The following example shows a very specific case of using the command:


```
displayLogs (ecid="B5C094FA...BE4AE8");
```

In this example, any log files with messages that contain the ECID B5C094FA...BE4AE8 will be displayed.

5.6.3 DMS Execution Context Communication

Figure 5–2 shows the components that cooperate in order to communicate the DMS execution context between each other. Arrows pointing to a component indicate the protocols that are inspected for incoming context information. Outgoing arrows show protocols to which context information is added. It is possible for a single component to send requests to itself, passing context information in that request.

Figure 5–2 DMS Execution Context Communication Protocols



5.7 DMS Tracing and Events

Starting with Oracle Fusion Middleware 11g Release 1 (11.1.1.3.0), DMS can selectively trace the following:

- DMS sensor lifecycle events (create, update, delete of state sensors, event sensors and phase sensors)
- Context events (start, stop)
- HTTP events (start, stop)

The configuration that controls which of these types of events are traced, and how those events are processed, is recorded in the `dms_config.xml` file. The DMS trace configuration is split into three parts:

1. Filter Configuration
Defines the rules that select the events that are of interest
2. Destination Configuration
Defines how the events are used
3. eventRoute Configuration
Defines which filters are wired to which destinations

A filter can be associated with one or more destinations thus granting the administrator the ability to define a filter rule once and have the resulting subset of all possible events processed on one or more different destinations.

The configuration can be modified using the DMS configuration MBean or WLST commands at runtime; this makes the DMS tracing feature invaluable for diagnosing issues within a specific time period or collecting specific data at a specific time for a specific set of criteria.

For more information, see "Configuring Selective Tracing Using WLST" in *Oracle Fusion Middleware Administrator's Guide*.

The following types of filter rules are supported:

- Event Type Conditions
Used to identify if an event was triggered from the START or STOP of a PHASE_SENSOR
- Context Type Conditions
Used to identify if the event was generated from a unit of work whose context contains a value (for example, "USER")
- Noun Type Conditions
Used to identify if the event was triggered from a sensor whose noun is of a specific type (for example, JDBC_CONNECTION)
- Logical AND and OR combinations of the above conditions

5.7.1 Configuring the DMS Event System

Configuration is recorded in each server's `dms_config.xml` file. MBean updates can be made at runtime using command line interface (CLI) commands and through the Event Configuration Mbean. Configuration updates are applied to the running system in a thread safe, but non-atomic, manner.

The object name of the DMS Event configuration MBean is:

```
oracle.dms.event.config:name=DMSEventConfigMBean,type=JMSEventConfig
```

To review the current state of your system's DMS event configuration, use the following command:

```
listDMSEventConfiguration([server=<server>])
```

The resulting output will look similar to this:

```
Event routes:
    FILTER      : auto662515911
    DESTINATION : destination1
    ENABLED     : true
    FILTER      : filter0
    DESTINATION : q
    ENABLED     : true
Filters with no event route:
    Fred

Destinations with no event route:
    des4
```

5.7.1.1 Adding and Editing Filters

Filters define the rules that select which events are considered for tracing.

The following example shows how to add a filter that selects all events related to JDBC operations:

```
addDMSEventFilter(id='myJDBCFilter', props={'condition': 'NOUNTYPE sw JDBC_'})
```

Or:

```
addDMSEventFilter(id='myJDBCFilter', props={'condition': 'NOUNTYPE startsWith
JDBC_'})
```

This filter assumes that all DMS sensor updates associated with JDBC operations are performed on nouns of types whose names begin "JDBC_".

If the rule must be modified, the filter may be updated as shown in the following example:

```
updateDMSEventFilter(id="myJDBCFilter", props={'condition': 'NOUNTYPE startsWith
JDBC_ OR NOUNTYPE startsWith MDS_'});
```

As of Oracle Fusion Middleware 11.1.1.6.0, the following shortened convenience operators have been added. Operators can be specified using either the shortened or longer name.

Note that operators with an underscore have been deprecated in favor of the ODL format, which is to use mixed case. For example, `not_equals` becomes `notEquals` or `ne`. The old format will still work, but is discouraged.

Noun Type Operators

equals, eq	notEquals, ne
contains	in
startsWith, sw	

Context Operators

equals, eq	notequals, ne
isnull	isnotnull

Context Operators

startswith, sw	contains
lt	gt

Example:

```
addDMSEventFilter(id='mdsbruce', name='MyFilter', props={'condition':
'NOUNTYPE eq MDS_Connections AND CONTEXT user ne bruce'})
```

```
addDMSEventFilter(id='mdsbruce', name='MyFilter', props={'condition':
'NOUNTYPE equals MDS_Connections AND CONTEXT user notequals bruce'})
```

For more information about the syntax used to describe a filter's rule (the condition property), refer to the WebLogic Scripting Tool Command Reference or the command help.

5.7.1.2 Adding and Editing Destinations

Destinations encapsulate logic for responding to events. For example, a basic destination will log the event, a different destination may transform an event and pass it to another system for further processing.

The following example shows how to add a destination that will log events:

```
addDMSEventDestination(id="myLoggerDestination",
class="oracle.dms.trace2.runtime.LoggerDestination",
props={"loggerName": "myLogger"});
```

Note that merely adding the destination is not sufficient for events to be logged; to log the events, you must associate a filter with a destination using an eventRoute, and the eventRoute must be enabled (default).

The types of destination available, and their configuration options, are described in [Section 5.7.2](#). The following example shows how to edit an existing destination:

```
updateDMSEventDestination(id="myLoggerDestination",
props={"loggerName": "myTraceLogger"});
```

5.7.1.3 Adding and Editing Event Routes

The following example shows how to join the filter and destination created above:

```
addDMSEventRoute(filterid='myJDBCFilter', destinationid='myLoggerDestination')
```

Note that you can invoke `addDMSEventRoute` without an explicit `filterId`. In these scenarios, all events are passed to the destination without filtering.

To remove a filter or destination, you must first remove the event routes associated with the filter or destination (even if the event route is disabled). For example, if you wanted to remove `myJDBCFilter`, you would first need to remove the eventRoute created in the previous example, and then remove the filter as shown in the following example:

```
removeDMSEventRoute(filterid='myJDBCFilter', destinationid='myLoggerDestination')
removeDMSEventFilter(id='myJDBCFilter')
```

5.7.1.4 Compound Operations

It is possible to create a filter and an eventRoute based on that filter using a single command (rather than using two separate commands as shown in [Section 5.7.1.3](#)).

Note, however, that the destination to be used by the event route must already be defined:

```
enableDMSEventTrace (destinationid='myLoggerDestination', condition='NOUNTYPE
starts_with JDBC_')
```

In the example above, `enableDMSEventTrace` automatically creates a filter with the specified condition, and also creates and enables an event route using the new filter and the nominated destination. The output is shown in the following example:

```
Filter "auto605449842" using Destination "myLoggerDestination" added, and
event-route enabled for server "AdminServer"
```

5.7.2 Configuring Destinations

DMS offers the following types of destinations:

- [LoggerDestination](#)
- [MBean Creator Destination](#)
- [HTTP Request Tracker Destination](#)
- [JRockit Flight Recorder Destination](#)

5.7.2.1 LoggerDestination

Description	The <code>LoggerDestination</code> writes each event to the associated logger.
Implementing Class	<code>oracle.dms.trace2.runtime.LoggerDestination</code>
Properties	
<code>loggerName</code>	The name of the ODL logger to which events will be written.

Instances of logger destinations write events to the named logger at a log level of `FINER`.

The `loggerName` property specifies the name of a logger, but the logger does not necessarily have to be described in `logging.xml`, though it can be. If the logger name refers to a logger that is explicitly named in `logging.xml`, then the logger is referred to as a static logger (see [Section 5.7.2.1.1](#)). If the logger name refers to a logger that is not explicitly named in `logging.xml`, then the logger is referred to as a dynamic logger (see [Section 5.7.2.1.2](#)).

Use in the default configuration: the default configuration defines a logger destination, with an identification of `LoggerDestination`. This particular instance does not form part of any `eventRoute` and therefore is not active. It is provided for convenience, and uses a dynamic logger.

5.7.2.1.1 Static Loggers and Handlers Loggers are the objects to which log records are presented. Log handlers are the objects through which log records are written to log files.

For complete control over the log file to which DMS trace data is written, define the logger named in the logger destination in `logging.xml`. Doing this allows you to explicitly define the name of the log file, the maximum size, format, file rotation and policies.

Oracle recommends using commands (like the example below) to update the configuration.

```
setLogLevel(logger="myTraceLogger", level="FINER", addLogger=1);

configureLogHandler(name="my-trace-handler", addToLogger=["myTraceLogger"],
path="/tmp/myTraceLogFiles/trace", maxFileSize="10m", maxLogSize="50m",
handlerType="oracle.core.ojdl.logging.ODLHandlerFactory", addHandler=1,
useParentHandlers=0);

configureLogHandler(name="my-trace-handler",
propertyName="useSourceClassandMethod", propertyValue="false", addProperty=1);
```

For more information on logging configuration, see "Managing Log Files and Diagnostic Data" in the *Oracle Fusion Middleware Administrator's Guide*.]

The use of the optional property `useSourceClassandMethod` set to `FALSE` prevents the 'SRC_CLASS' and 'SRC_METHOD' from appearing in every message and will marginally improve performance by reducing file output times.

For static loggers, consider setting the `useParentHandlers` parameter to `FALSE`, otherwise duplicate event messages will be logged to `[server]-diagnostics.log`, and shown in a log query.

See [Section 5.7.3, "Understanding DMS Event Output"](#) for more information about interpreting logger output.

5.7.2.1.2 Dynamic Loggers and Handlers If the named logger has no associated handler defined in `logging.xml`, then the logger destination will dynamically create a handler object that will write to a file in the server's default log output directory. (Instances of logger destinations write events to the named logger at a log level of `FINER`.) The file name will be the logger's name followed by `-event.log`. For instance, in the example in [Section 5.7.2.1.1](#), DMS events would be written to `myTraceLogger-event.log`.

5.7.2.1.3 Default Locations of the logging.xml File The `logging.xml` file can typically be found in one of the following platform locations:

Platform	Server	Location
Oracle WebLogic Server	AdminServer	ORACLE_HOME/Middleware/user_projects/domains/base_domain/config/fmwconfig/servers/AdminServer/logging.xml
WAS ND	OracleAdminServer	ORACLE_HOME/Middleware/was_profiles/DefaultTopology/DefaultServer/config/cells/DefaultCell/nodes/<nodename>/servers/OracleAdminServer/fmwconfig/logging.xml

5.7.2.1.4 Using a CLI Command to Query the Trace Log File If the logger destination's logger and handler are defined in `logging.xml` then you can take advantage of the `displayLogs()` command to conveniently access logged trace data without having to manually locate or search for it.

Examples:

- To display all the log messages for the `myTraceLogger`:

```
displayLogs(query='MODULE equals myTraceLogger')
```


- To display only the log messages for myTraceLogger which have an ECID of '0000HpmSpLWEkJQ6ub3FEH194kwB000004':

```
displayLogs(query='MODULE equals myTraceLogger and ECID equals
0000HpmSpLWEkJQ6ub3FEH194kwB000004')
```

- To display only the log messages for myTraceLogger which have an ECID of '0000HpmSpLWEkJQ6ub3FEH194kwB000004' in the last 10 minutes:

```
displayLogs(query='MODULE equals myTraceLogger and ECID equals
0000HpmSpLWEkJQ6ub3FEH194kwB000004', last=10)
```

- To display all the log messages from a dynamic logger the log's file name must be included:

```
displayLogs(disconnected=1, log=DOMAIN_
ROOT+"/servers/AdminServer/logs/myTraceLogger-event.log")
```

5.7.2.2 MBean Creator Destination

Description	The MBean creator destination make nouns accessible as MBeans, exposing their metrics as attributes, for access via WLDF, JConsole, etc.
Implementing Class	oracle.dms.jmx.MetricMBeanFactory

Use in the default configuration: An instance of the MBean Creator destination is configured and active by default, and will create MBeans for all nouns created in the server.

By associating an instance of this destination type with a filter based on a noun-type rule, it is possible to expose (as MBeans) only those noun types that are of interest to the administrator.

Although it is possible to modify the configuration associated with an MBean creator destination at runtime, it must be understood that the reinitialization process for this type of destination may impact performance. Frequent runtime reconfiguration is therefore discouraged.

Note that WebLogic Diagnostic Framework (WLDF) can be used to harvest DMS metrics exposed by the MBean creator destination. For more information about WLDF, see *Oracle Fusion Middleware Configuring and Using the Diagnostics Framework for Oracle WebLogic Server*.

5.7.2.2.1 Metric MBean Object Name The noun name and noun type are exposed as the name and type properties of the metric MBean object name. The MBean domain name is "oracle.dms". The object name also reflects the DMS noun hierarchy.

For example if the noun's full path name is:

```
/oracle/dfw/ofm/base_domain/AdminServer
```

and the noun type is DFW_Incident, the object name of the MBean representing the noun is

```
oracle.dms:Location=AdminServer,name=/oracle/dfw/ofm/base_
domain/AdminServer,type=DFW_Incident.
```

5.7.2.3 HTTP Request Tracker Destination

Description	The HTTP Request Tracker destinations maintains a list of active HTTP requests, and makes the requests accessible to other Diagnostic Framework (DFW) components.
Implementing Class	oracle.dms.event.HTTPRequestTrackerDestination
Properties	
excludeHeaderNames	Comma separated list of header names to exclude from tracking

Use in the default configuration: An instance of the HTTP request tracker destination is enable by default. In the case of a DFW incident being generated the active HTTP request list will be dumped automatically, allowing an administrator to correlate the failure with a specific request.

For each HTTP request the following information will be dumped:

- URI (such as /webcenter/home)
- Start time of the request
- ECID
- Query string
- HTTP Headers

When the HTTP request tracker is not enabled the HTTP Request Dump will output the following:

HTTP Requests are not being tracked. To enable HTTP request tracking enable the DMS oracle.dms.event.HTTPRequestTrackerDestination in dms_config.xml

5.7.2.3.1 Executing the HTTP Request Tracker Dump The information being maintained by the HTTP request tracker can be accessed manually. In order to execute the dump that reports the HTTP request information the WLST `executeDump` command can be used, when connected to a server, as follows:

```
> executeDump(name="http.requests")
Active Requests:

StartTime: 2009-12-14 02:24:41.870
ECID: 0000IMChyqEC8xT6uBF9EH1B9X9^000009,0
URI: /myApp/Welcome.jsp
QueryString:
Headers:
  Host: myHost.myDomain.com:7001
  Connection: keep-alive
  User-Agent: Mozilla/5.0 (Windows; U; Windows NT 5.1; en-US) AppleWebKit/532.5
(KHTML, like Gecko) Chrome/4.0.249.30 Safari/532.5
  Accept:
application/xml,application/xhtml+xml,text/html;q=0.9,text/plain;q=0.8,image/png,*
/*;q=0.5
  Accept-Encoding: gzip,deflate
  Cookie: ORA_MOS_LOCALE=en%7CGB; s_nr...
  Accept-Language: en-GB,en-US;q=0.8,en;q=0.6
  Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.3
```

5.7.2.4 JRockit Flight Recorder Destination

The JRockit Flight Recorder (JFR) records information regarding the runtime status and behavior of the JRockit JVM. JFR also exposes an API through which third party events can be reported. JFR is available in JRockit R28 and beyond.

By themselves DMS traces and JFR traces only show part of the picture of the actions being performed in the server. DMS integration with JFR enhances the diagnostic information available to administrators and developers as follows:

1. Application level events and JVM level events can be reported as a single sequence therefore avoiding the need to combine such events from separate log files based only on timestamp (which may not tick over fast enough to accurately order events created at or around the same time).
2. Recent DMS activity can be dumped, retroactively, from the JVM at will.
3. Recent DMS and JVM events can be dumped to disk in the event of a fatal error that causes the JVM to exit gracefully.
4. The DMS ECID can be used to correlate activity relating to the same request, or unit of work, across the span of a JFR recording.
5. The DMS ECID can be used to collect diagnostic information from all systems involved with an event, or series of events, recorded by JFR.

5.7.2.4.1 Dynamically Derived JFR Event Types – Names, Values and Descriptions A DMS noun type will be associated with a JFR InstantEvent event type:

- The name of the JFR event type for a noun type will be the noun type's name with the suffix " state".
- The path of the JFR event type for a noun type will be "dms/" followed by the producer-name, followed by the event type name.
- Event sensors will not contribute any values to the noun type's JFR event type.
- The values of the JFR event for a noun type are described in [Table 5–6](#):

Table 5–6 Values of the JFR Event for a Noun Type

Value Name	Description	Relational	Notes
ECID	The Execution Context ID (ECID) associated with the action.	Yes	
RID	The RID associated with the action.	Yes	

Table 5–6 (Cont.) Values of the JFR Event for a Noun Type

Value Name	Description	Relational	Notes
<noun type> name	The full path of the noun.		This field will be populated with the full path of the noun. The field's name assumes that the noun_type meaningfully categorizes all objects being measured by the nouns of that type.
<state-sensor-name>	The value of the state sensor.	No	Each state sensor belonging to the noun will contribute one of these values to the instant event. There may be more than one value in each noun.
event name	The name of the event sensor that was updated, left null otherwise.	No	The event name field is required for being able to count the number of times a DMS event sensor has been updated in a recording (event sensors do not contribute values to an event type).

A DMS phase sensor will be associated with a JFR DurationEvent event type:

- The name of the JFR event type for a phase sensor belonging to a noun of a particular noun type will be the noun type's name following by the phase sensor's name.
- The path of the JFR event for a noun type will be "dms/" followed by the producer-name, followed by the event type name.
- The values of the duration event will be as above (except for the sensorName value). For example the "stop" of a phase event will result in a JFR duration event being reported to JFR that contains the state information of the phase event's parent noun.

Several DMS objects allow integrators to add descriptions. Descriptions from DMS objects will be used as follows:

- Noun type description will be used in creation of the JFR event type
- State and event sensor descriptions will not be applied – there is nowhere to apply them.
- Phase sensor descriptions will be applied to their JFR event type.

5.7.2.4.2 Examples of Dynamically Derived Producers and Events [Table 5–7](#) provides examples for the rules described in [Section 5.7.2.4.1](#):

Table 5-7 Examples of Dynamically Derived Producers and Events

DMS	JRockit Flight Recorder (JFR)
<p>Noun type: JDBC_Connection</p> <p>Noun path: /JDBC/Driver/CONNECTION_7</p> <p>Sensors: CreateStatement (P) CreateNewStatement (P) DBWaitTime (P) JDBC_Connection_Url (S) JDBC_Connection_Username (S)</p> <p>Where: P: Phase Sensor S: State Sensor E: Event Sensor</p>	<p>Producer Name: JDBC</p> <p>The Producer Name is based on the leading component of the noun path.</p> <p>Event Type 1 Event Type Name: JDBC_Connection State <noun type> State Event Type Path: dms/JDBC/JDBC_Connection_State dms/<leading component of noun path>/<noun type>/_State</p> <p>Fields:</p> <ul style="list-style-type: none"> ▪ ECID ▪ RID ▪ JDBC_Connection name Value will be the full path of the noun ▪ JDBC_Connection_Url Value will be that of the state sensor of this name at the time of the event ▪ JDBC_Connection_Username Value will be that of the state sensor of this name at the time of the event ▪ Event Name Value will be one of the following: <ul style="list-style-type: none"> ▪ The name of the DMS event sensor whose activation caused this JFR event instance ▪ Null if this JFR event instance was created for a state sensor update
	<p>Producer Name: JDBC</p> <p>Event Type 2 Event Type Name: JDBC_Connection CreateStatement Event Type Path: dms/JDBC/JDBC_Connection_CreateStatement</p> <p>Fields:</p> <ul style="list-style-type: none"> ▪ ECID ▪ RID ▪ JDBC_Connection name ▪ JDBC_Connection_Url ▪ JDBC_Connection_Username

Table 5–7 (Cont.) Examples of Dynamically Derived Producers and Events

DMS	JRocket Flight Recorder (JFR)
	<p data-bbox="540 262 781 283">Producer Name: JDBC</p> <p data-bbox="540 304 678 325">Event Type 3</p> <p data-bbox="540 346 1162 367">Event Type Name: JDBC_Connection CreateNewStatement</p> <p data-bbox="540 388 716 409">Event Type Path:</p> <p data-bbox="540 430 1068 451">dms/JDBC/JDBC_Connection_CreateNewStatement</p> <p data-bbox="540 472 613 493">Fields:</p> <ul data-bbox="540 504 899 682" style="list-style-type: none"><li data-bbox="540 504 646 525">▪ ECID<li data-bbox="540 541 630 562">▪ RID<li data-bbox="540 579 841 600">▪ JDBC_Connection name<li data-bbox="540 617 824 638">▪ JDBC_Connection_Url<li data-bbox="540 655 899 676">▪ JDBC_Connection_Username

Table 5–7 (Cont.) Examples of Dynamically Derived Producers and Events

DMS	JRockit Flight Recorder (JFR)
	<p>Producer Name: JDBC</p> <p>Event Type 4 Event Type Name: JDBC_Connection DBWaitTime</p> <p>Event Type Path: dms/JDBC/JDBC_Connection_DBWaitTime</p> <p>Fields:</p> <ul style="list-style-type: none"> ▪ ECID ▪ RID ▪ JDBC_Connection name ▪ JDBC_Connection_Url ▪ JDBC_Connection_Username
<p>Noun type: webcenter_lifecycle</p> <p>Noun path: /oracle/webcenter/webcenter/lifecycle</p> <p>Sensors: ProcessingTime (P) status (S) successCount (E)</p>	<p>Producer Name: webcenter</p> <p>Event Type 1 Event Type Name: webcenter_lifecycle State</p> <p>Fields:</p> <ul style="list-style-type: none"> ▪ ECID ▪ RID ▪ webcenter_lifecycle name ▪ status ▪ event name
<p>Where: P: Phase Sensor S: State Sensor E: Event Sensor</p>	<p>Producer Name: webcenter</p> <p>Event Type 2 Event Type Name: webcenter_lifecycle ProcessingTime</p> <p>Fields:</p> <ul style="list-style-type: none"> ▪ ECID ▪ RID ▪ webcenter_lifecycle name ▪ status

5.7.3 Understanding DMS Event Output

Table 5–8 describes the fields that make up a DMS event. Field elements are separated by ":" (with a few exceptions). Sample events are provided to illustrate the position of the field within an actual event string.

Table 5–8 Event Formatting Descriptions

Applicable Events	Field Number	Name	Description
All	1	Version number	The version number of the event format For example: v1:1280737384058:HTTP_REQUEST:STOP:/MyWebApp/emp
All	2	Event time	The time at which the event occurred For example: v1:1280737384058:HTTP_REQUEST:STOP:/MyWebApp/emp
All	3	Source object type	The type of object on which an action was performed to produce the event including: <ul style="list-style-type: none"> ■ NOUN ■ EVENT_SENSOR ■ STATE_SENSOR ■ PHASE_SENSOR ■ EXECUTION_CONTEXT ■ HTTP_REQUEST For example: v1:1280737384058:HTTP_REQUEST:STOP:/MyWebApp/emp
All	4	Action type	The type of action that resulted in the generation of this event. A given source object type may not necessarily produce events for every action type: <ul style="list-style-type: none"> ■ CREATE ■ UPDATE ■ DELETE ■ START ■ STOP ■ ABORT For example: v1:1280737384058:HTTP_REQUEST:STOP:/MyWebApp/emp
Nouns	5	Noun type	The name of the noun type For example: v1:1281344803506:NOUN:CREATE:JDBC_Connection:/JDBC/JDBC Data Source-0/CONNECTION_1
	6	Noun path	The full path identifying the noun to which the sensor belongs For example: v1:1281344803506:NOUN:CREATE:JDBC_Connection:/JDBC/JDBC Data Source-0/CONNECTION_1

Table 5–8 (Cont.) Event Formatting Descriptions

Applicable Events	Field Number	Name	Description
All Sensor Types	5	Noun type	The name of the noun type to which this sensor belongs For example: v1:1280503318973:STATE_SENSOR:UPDATE:JDBC_ Connection: LogicalConnection:/JDBC/JDBC Data Source-0/CONNECTION_ 1:State.ANY:LogicalConnection@13bed086
	6	Sensor name	The name of the sensor For example: v1:1280737383069:PHASE_SENSOR:STOP:JDBC_ Connection:DBWaitTime:/JDBC/JDBC Data Source-0/CONNECTION_1:1280737382950:1280737383069
	7	Noun path	The full path identifying the noun to which the sensor belongs For example: v1:1280737383069:PHASE_SENSOR:STOP:JDBC_ Connection:DBWaitTime:/JDBC/JDBC Data Source-0/CONNECTION_1:1280737382950:1280737383069
Phase Sensor Types	8	Start token	The start token of the phase. For example: v1:1280737383069:PHASE_SENSOR:STOP:JDBC_ Connection:DBWaitTime:/JDBC/JDBC Data Source-0/CONNECTION_1: 1280737382950:1280737383069
	9	Stop token	The end token of the phase. For example: v1:1280737383069:PHASE_SENSOR:STOP:JDBC_ Connection:DBWaitTime:/JDBC/JDBC Data Source-0/CONNECTION_1:1280737382950: 1280737383069

Table 5–8 (Cont.) Event Formatting Descriptions

Applicable Events	Field Number	Name	Description
State Sensor Types	8	State value type	<p>The type of value held by the state sensor including:</p> <ul style="list-style-type: none"> ■ State.DOUBLE ■ State.INTEGER ■ State.LONG ■ State.OBJECT ■ State.ANY <p>For example:</p> <pre>v1:1280503318973:STATE_SENSOR:UPDATE:JDBC_Connection:LogicalConnection:/JDBC/JDBC Data Source-0/CONNECTION_1:State.ANY:LogicalConnection@13bed086</pre>
	9	State value	<p>The value of the state represented in string form.</p> <p>For example:</p> <pre>v1:1280503318973:STATE_SENSOR:UPDATE:JDBC_Connection:LogicalConnection:/JDBC/JDBC Data Source-0/CONNECTION_1:State.ANY:LogicalConnection@13bed086</pre>
HTTP Requests	5	URI	<p>Uniform Resource Identifier (URI) identifies the resource upon which to apply the request.</p> <p>For example:</p> <pre>v1:1280737382889:HTTP_REQUEST:START:/myWebApp/showEmployees v1:1280737384058:HTTP_REQUEST:STOP:/myWebApp/showEmployees</pre>
Execution Context	5	ECID,RID	<p>The context identifier (composed of ECID and RID separated by a comma).</p> <p>For execution context events the complete substring starting at the first character after the fourth event field separator (":") records the ECID,RID identifiers - the context identifiers may contain ":" but these should not be interpreted as event field separators.</p> <p>For example:</p> <pre>v1:1280737384058:EXECUTION_CONTEXT:STOP:bc4fd0668f79d507:367c127f:12a23f2013c:-8000-00000000000000f73,0</pre>

5.7.4 Understanding DMS Event Actions

Table 5–9 shows the action types that can be performed on source object types.

Table 5–9 Actions Performed on Source Object Types

	Create	Update	Delete	Start	Stop	Abort
Noun	Yes	-	Yes	-	-	-
Event Sensor	Yes	Yes	Yes	-	-	-
Phase Sensor	Yes	-	Yes	Yes	Yes	Yes
State Sensor	Yes	Yes	Yes	-	-	-

Table 5–9 (Cont.) Actions Performed on Source Object Types

	Create	Update	Delete	Start	Stop	Abort
Execution Context	-	-	-	Yes	Yes	-
Http Request	-	-	-	Yes	Yes	-

5.8 DMS Best Practices

The use of DMS metrics can have an impact on application performance. When adding metrics, consider the following:

- Use a High Resolution Clock to increase DMS Precision

By default DMS uses the system clock for measuring time intervals during a `PhaseEvent`. The default clock reports microsecond precision in C processes such as Apache and reports millisecond precision in Java processes. Optionally, DMS supports a high resolution clock to increase the precision of performance measurements and lets you select the values for reporting time intervals. You can use a high resolution clock when you need to time phase events more accurately than is possible using the default clock or when the system's default clock does not provide the resolution needed for your requirements.

System clocks are not necessarily as accurate as their precision implies. For example, a system clock that reports time in milliseconds may not tick (change) once per millisecond. Instead, it may take up to 15ms to tick as shown in the following example:

Table 5–10 Default System Clock Time versus Actual Time (in milliseconds)

Actual Time	System Time
12:00:00.000	12:00:00.000
12:00:00.001	12:00:00.000
12:00:00.002	12:00:00.000
[...]	
12:00:00.014	12:00:00.000
12:00:00.015	12:00:00.015
12:00:00.016	12:00:00.015

[Table 5–10](#) shows a phase with a 12ms duration that runs from actual time 12:00:00.002 to 12:00:00.014 would be calculated in system time as having a duration of zero. Similarly, a phase with a 2ms duration running from 12:00:00.014 to 12:00:00.016 would be reported in system time as having a duration of 15ms.

Note: These behaviors are more evident on some operating systems than others. Use caution when analyzing individual periods of time that are shorter than the tick period of the system clock. Configuring DMS to use a higher resolution clock will cause DMS to record phase sensor activations with higher resolution, but the accuracy will still be limited by the underlying system.

- Configure DMS Clocks for Reporting Time for Java

Selecting the high resolution clock changes clocks for all applications running on the server where the clock is changed. You set the DMS clock and the reporting values globally using the `oracle.dms.clock` and `oracle.dms.clock.units` properties, which control process startup options.

For example, to use the high resolution clock with the default values, set the following property on the Java command line:

```
-Doracle.dms.clock=highres
```

Caution: If you use the high resolution clock, the default values are different from the value that Fusion Middleware Control expects (msecs). If you need the Fusion Middleware Control displays to be correct when using the high resolution clock, then you need to set the units property as follows:

```
-Doracle.dms.clock.units=msecs
```

[Table 5–11](#) shows supported values for the `oracle.dms.clock` property.

[Table 5–12](#) shows supported values for the `oracle.dms.clock.units` property.

Table 5–11 *oracle.dms.clock Property Values*

Value	Description
DEFAULT	Specifies that DMS use the default clock. With the default clock, DMS uses the Java call <code>java.lang.System.currentTimeMillis</code> to obtain times for <code>PhaseEvents</code> . The default value for the units for the default clock is <code>MSECS</code> .
HIGHRES	The Java Highres clock uses <code>System.nanoTime()</code> (no JNI required).

Table 5–12 *oracle.dms.clock.units Property Values*

Value	Description
MSECS	Specifies that the time be converted to milliseconds and reported as "msecs". A millisecond is 10^{-3} seconds. Note: This is the default value for the default clock.
USECS	Specifies that the time be converted to microseconds and reported as "usecs". A microsecond is 10^{-6} seconds.
NSECS	Specifies that the time be converted to nanoseconds and reported as "nsecs". A nanosecond is 10^{-9} seconds. Note: This is the default value for the high resolution clock.

Note the following when using the high resolution DMS clock:

- When you set the `oracle.dms.clock` and the `oracle.dms.clock.units` properties, any combination of upper and lower case characters is valid for the value that you select (case is not significant). For example, any of the following values are valid to select the high resolution clock: `highres`, `HIGHRES`, `HighRes`.
- DMS checks the property values at startup. When the clock property is set with a value not listed in [Table 5–11](#), DMS uses the default clock. If the `oracle.dms.clock` property is not set, DMS uses the default clock.

- When the clock units property is set to a value not listed in [Table 5-12](#), DMS uses the default units for the specified clock.

Oracle HTTP Server Performance Tuning

This chapter discusses the techniques for optimizing Oracle HTTP Server performance. This chapter contains the following sections:

- [Section 6.1, "About Oracle HTTP Server"](#)
- [Section 6.2, "Monitoring Oracle HTTP Server Performance"](#)
- [Section 6.3, "Basic Tuning Considerations"](#)
- [Section 6.4, "Advanced Tuning Considerations"](#)

Note: The configuration examples and recommended settings described in this chapter are for illustrative purposes only. Consult your own use case scenarios to determine which configuration options can provide performance improvements.

6.1 About Oracle HTTP Server

Oracle HTTP Server (OHS) is the Web server component for Oracle Fusion Middleware. It provides a listener for Oracle WebLogic Server and the framework for hosting static pages, dynamic pages, and applications over the Web. Oracle HTTP Server is based on the Apache 2.2.x infrastructure, and includes modules developed specifically by Oracle. The features of single sign-on, clustered deployment, and high availability enhance the operation of the Oracle HTTP Server.

For more information see *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*.

For more information on the Apache open-source software infrastructure, see the Apache Software Foundation web site at <http://www.apache.org/>.

6.2 Monitoring Oracle HTTP Server Performance

Oracle Fusion Middleware automatically and continuously measures run-time performance for Oracle HTTP Server. The performance metrics are automatically enabled; you do not need to set options or perform any extra configuration to collect them. If you encounter a problem, such as an application that is running slowly or is hanging, you can view particular metrics to find out more information about the problem.

Note: Fusion Middleware Control provides real-time data. For more information on using Fusion Middleware Control to view performance metrics for HTTP Server, see "Monitoring Oracle HTTP Server Performance" in *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*.

If you are interested in viewing historical data, consider using Grid Control. See [Section 4.8, "Oracle Enterprise Manager Cloud Control"](#).

In addition to the Fusion Middleware Control, Oracle HTTP Server also has Dynamic Monitoring Service (DMS), which collects metrics for every functional piece. You can review these metrics as needed to understand system behavior at a given point of time. This displays memory, CPU information and the minimum, maximum, and average times for the request processing at every layer in Oracle HTTP Server. The metrics also display details about load level, number of threads, number of active connections, and so on, which can help in tuning the system based on real usage.

You can use Oracle Enterprise Manager or SpyServlet to monitor the metrics. See [Chapter 4, "Monitoring Oracle Fusion Middleware"](#). Another way to view DMS metrics for OHS is shown in the following example:

1. `cd $INSTANCE_HOME/bin`
2. `./opmnctl metric op=query COMPONENT_NAME=<component_name>
dmsarg=[name=/OHS/Modules/<module_name>.c`

Examples:

```
./opmnctl metric op=query COMPONENT_NAME=ohs1  
dmsarg=[name=/OHS/Modules/mod_cgi.c
```

```
./opmnctl metric op=query COMPONENT_NAME=ohs1 dmsarg=[name=*
```

6.3 Basic Tuning Considerations

The following tuning configurations may improve the performance of the Oracle HTTP Server. Always consult your own use case scenarios to determine if these settings are applicable to your deployment.

6.3.1 Tuning Oracle HTTP Server Directives

Oracle HTTP Server uses directives in `httpd.conf`. This configuration file specifies the maximum number of HTTP requests that can be processed simultaneously, logging details, and certain limits and time outs.

More information on configuring the Oracle HTTP Server, see "Management Tools for Oracle HTTP Server" in *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*.

Oracle HTTP Server supports three different Multi-Processing Modules (MPMs) by default. The MPMs supported are:

- Worker - This uses Multi-Process-Multi-Threads model and is the default MPM on all platforms other than Microsoft Windows platforms. Multi-thread support makes it more scalable by using fewer system resources and multi-process support makes it more stable.

- WinNT - This MPM is for Windows platforms only. It consists of a parent process and a child process. The parent process is the control process, and the child process creates threads to handle requests.
- Prefork - This is Apache 1.3.x style and uses processes instead of threads. This is considered the least efficient MPM.

The directives for each MPM type are defined in the `ORACLE_INSTANCE/config/OHSCComponent/<ohsname>/httpd.conf` file. The default MPM type is Worker MPM. To use a different MPM (such as Prefork MPM), edit the `ORACLE_HOME/ohs/bin/apachectl` file.

Note: The information in this chapter is based on the use of Worker and WinNT MPMs, which use threads. The directives listed below may not be applicable if you are using the prefork MPM. If you are using Oracle HTTP Server based on Apache 1.3.x or Apache 2.2 with prefork MPM, refer to the Oracle Application Server 10g Release 3 documentation at

<http://www.oracle.com/technology/documentation/appserver10132.html>.

Table 6–1 Oracle HTTP Server Configuration Properties

Directive	Description
<p><code>ListenBackLog</code></p> <p>This directive maps to the Maximum Queue Length field on the Performance Directives screen.</p>	<p>Specifies the maximum length of the queue of pending connections. Generally no tuning is needed. Note that some operating systems do not use exactly what is specified as the backlog, but use a number based on, but normally larger than, what is set.</p> <p>Default Value: 511</p>
<p><code>MaxClients</code></p> <p>This directive maps to the Maximum Requests field on the Performance Directives screen.</p> <p>Note that this parameter is not available in <code>mod_winnt</code> (Microsoft Windows). <code>Winnt</code> uses a single process, multi-threaded model and is controlled by <code>ThreadLimit</code> directive.</p>	<p>Specifies a limit on the total number of servers running, that is, a limit on the number of clients who can simultaneously connect. If the number of client connections reaches this limit, then subsequent requests are queued in the TCP/IP system up to the limit specified with the <code>ListenBackLog</code> directive (after the queue of pending connections is full, new requests generate connection errors until a thread becomes available).</p> <p>You can configure the <code>MaxClients</code> directive in the <code>httpd.conf</code> file up to a maximum of 8000 (8K) (the default value is 150). If your system is not resource-saturated and you have a user population of more than 150 concurrent HTTP/Thread connections, you can improve your performance by increasing <code>MaxClients</code> to increase server concurrency. Increase <code>MaxClients</code> until your system becomes fully utilized (85% is a good threshold).</p> <p>Conversely, when system resources are saturated, increasing <code>MaxClients</code> does not improve performance. In this case, the <code>MaxClients</code> value could be reduced as a throttle on the number of concurrent requests on the server.</p> <p>If the server handles persistent connections, then it may require sufficient concurrent <code>httpd</code> or thread server processes to handle both active and idle connections. When you specify <code>MaxClients</code> to act as a throttle for system concurrency, you must consider that persistent idle <code>httpd</code> connections also consume <code>httpd</code>/thread processes. Specifically, the number of connections includes the currently active persistent and non-persistent connections and the idle persistent connections. A persistent, <code>KeepAlive</code>, <code>http</code> connection consumes an <code>httpd</code> child process, or thread, for the duration of the connection, even if no requests are currently being processed for the connection.</p> <p>If you have sufficient capacity, <code>KeepAlive</code> should be enabled; using persistent connections improves performance and prevents wasting CPU resources reestablishing HTTP connections. Normally, you should not change <code>KeepAlive</code> parameters.</p> <p>The maximum allowed value for <code>MaxClients</code> is 8192 (8K).</p> <p>Default Value: 150</p>
<p><code>StartServers</code></p> <p>This directive maps to the Initial Child Server Processes field on the Performance Directives screen.</p>	<p>Specifies the number of child server processes created on startup. If you expect a sudden load after restart, set this value based on the number child servers required.</p> <p>Note that the following parameters are inter-related and applicable only on UNIX platforms (<code>worker_mpm</code>):</p> <ul style="list-style-type: none"> ■ <code>MaxClients</code> ■ <code>MaxSpareThreads</code> and <code>MinSpareThreads</code> ■ <code>ServerLimit</code> and <code>StartServers</code> <p>On the Windows platform (<code>mpm_winnt</code>), as well as UNIX platforms, the following parameters are important to tune:</p> <ul style="list-style-type: none"> ■ <code>ThreadLimit</code> ■ <code>ThreadsPerChild</code> <p>Note that each child process has a set of child threads defined for them and that can actually handle the requests. Use <code>ThreadsPerChild</code> in connection with this directive.</p> <p>The values of <code>ThreadLimit</code>, <code>ServerLimit</code>, and <code>MaxClients</code> can indirectly affect this value. Read the notes for these directives and use them in conjunction with this directive.</p> <p>Default Value: 2</p>

Table 6–1 (Cont.) Oracle HTTP Server Configuration Properties

Directive	Description
<p><code>ServerLimit</code></p> <p>Note that this parameter is not available in <code>mod_winnt</code> (Microsoft Windows). <code>Winnt</code> uses a single process, multi-threaded model</p>	<p>Specifies an upper limit on the number of server (child) processes that can exist or be created. This value overrides the <code>StartServers</code> value if that value is greater than the <code>ServerLimit</code> value. This is used to control the maximum number of server processes that can be created.</p> <p>Default Value: 16</p>
<code>ThreadLimit</code>	<p>Specifies the upper limit on the number of threads that can be created under a server (child) process. This value overrides the <code>ThreadsPerChild</code> value if that value is greater than the <code>ThreadLimit</code> value. This is used to control the maximum number of threads created per process to avoid conflicts/issues.</p> <p>Default Values:</p> <ul style="list-style-type: none"> ■ Windows Multi-Processing Module (<code>mpm_winnt</code>): 1920 ■ All others: 64
<p><code>ThreadsPerChild</code></p> <p>This directive maps to the Threads Per Child Server Process field on the Performance Directives screen.</p>	<p>Sets the number of threads created by each server (child) process at startup.</p> <p>Default Value: 64 when <code>mpm_winnt</code> is used and 25 when Worker MPM is used.</p> <p>The <code>ThreadsPerChild</code> directive works with other directives, as follows:</p> <p>At startup, Oracle HTTP Server creates a parent process, which creates several child (server) processes as defined by the <code>StartServers</code> directive. Each server process creates several threads (server/worker), as specified in <code>ThreadsPerChild</code>, and a listener thread which listens for requests and transfers the control to the worker/server threads.</p> <p>After startup, based on load conditions, the number of server processes and server threads (children of server processes) in the system are controlled by <code>MinSpareThreads</code> (minimum number of idle threads in the system) and <code>MaxSpareThreads</code> (maximum number of idle threads in the system). If the number of idle threads in the system is more than <code>MaxSpareThreads</code>, Oracle HTTP Server terminates the threads and processes if there are no child threads for a process. If the number of idle threads is fewer than <code>MinSpareThreads</code>, it creates new threads and processes if the <code>ThreadsPerChild</code> value has already been reached in the running processes.</p> <p>The following directives control the limit on the above directives. Note that the directives below should be defined before the directives above for them to take effect.</p> <ul style="list-style-type: none"> ■ <code>ServerLimit</code> - Defines the upper limit on the number of servers that can be created. This affects <code>MaxClients</code> and <code>StartServers</code>. ■ <code>ThreadLimit</code> - Defines the upper limit on <code>ThreadsPerChild</code>. If <code>ThreadsPerChild</code> is greater than <code>ThreadLimit</code>, then it is automatically trimmed to the latter value. ■ <code>MaxClients</code> - Defines the upper limit on the number of server threads that can process requests simultaneously. This should be equal to the number of simultaneous connections that can be made. This value should be a multiple of <code>ThreadsPerChild</code>. If <code>MaxClients</code> is greater than <code>ServerLimit</code> multiplied by <code>ThreadsPerChild</code>, it is automatically be trimmed to the latter value.

Table 6–1 (Cont.) Oracle HTTP Server Configuration Properties

Directive	Description
<p>MaxRequestsPerChild</p> <p>This directive maps to the Max Requests Per Child Server Process field on the Performance Directives screen.</p>	<p>Specifies the number of requests each child process is allowed to process before the child process dies. The child process ends to avoid problems after prolonged use when Apache (and any other libraries it uses) leak memory or other resources. On most systems, this is not needed, but some UNIX systems have notable leaks in the libraries. For these platforms, set MaxRequestsPerChild to 10000; a setting of 0 means unlimited requests.</p> <p>This value does not include KeepAlive requests after the initial request per connection. For example, if a child process handles an initial request and 10 subsequent "keep alive" requests, it would only count as 1 request toward this limit.</p> <p>Default Value: 0</p> <p>Note: On Windows systems MaxRequestsPerChild should always be set to 0 (unlimited) since there is only one server process.</p>
<p>MaxSpareThreads</p> <p>MinSpareThreads</p> <p>These directives map to the Maximum Idle Threads and Minimum Idle Threads fields on the Performance Directives screen.</p> <p>Note that these parameters are not available in mod_winnt (Windows platform).</p>	<p>Controls the server-pool size. Rather than estimating how many server threads you need, Oracle HTTP Server dynamically adapts to the actual load. The server tries to maintain enough server threads to handle the current load, plus a few additional server threads to handle transient load increases such as multiple simultaneous requests from a single browser.</p> <p>The server does this by periodically checking how many server threads are waiting for a request. If there are fewer than MinSpareThreads, it creates a new spare. If there are more than MaxSpareThreads, some of the spares are removed.</p> <p>Default Values:</p> <p>MaxSpareThreads: 75</p> <p>MinSpareThreads: 25</p>
<p>Timeout</p> <p>This directive maps to the Request Timeout field on the Performance Directives screen.</p>	<p>The number of seconds before incoming receives and outgoing sends time out.</p> <p>Default Value: 300</p>
<p>KeepAlive</p> <p>This directive maps to the Multiple Requests Per Connection field on the Performance Directives screen.</p>	<p>Whether or not to allow persistent connections (more than one request per connection). Set to Off to deactivate.</p> <p>Default Value: On</p>

Table 6–1 (Cont.) Oracle HTTP Server Configuration Properties

Directive	Description
MaxKeepAliveRequests	The maximum number of requests to allow during a persistent connection. Set to 0 to allow an unlimited amount. If you have long client sessions, consider increasing this value. Default Value: 100
KeepAliveTimeout	Number of seconds to wait for the next request from the same client on the same connection. Default Value: 5 seconds
This directive maps to the Allow With Connection Timeout (seconds) field, which is located under the Multiple Requests Per Connection field, on the Performance Directives screen.	
limit ulimit	Number of objects that a program uses to read or write to an open file or open network sockets. A lack of available file descriptors can impact operating system performance. Tuning the file descriptor limit can be accomplished by configuring the hard limit (<code>ulimit</code>) in a shell script which starts the OHS. Once the hard limit has been set the OHS will then adjust the soft limit (<code>limit</code>) to match. Note that configuring file descriptor limits is platform specific. Refer to your operating system documentation for more information.

6.3.2 Reducing Httpd Process Availability with Persistent Connections

If your browser supports persistent connections, you can support them on the server using the `KeepAlive` directives in the Oracle HTTP Server. Persistent Connections can improve performance by reducing the work load on the server. With Persistent Connections enabled, the server does not have to repeat the work to set up the connections with a client.

The default settings for the `KeepAlive` directives are:

```
KeepAlive on
MaxKeepAliveRequests 100
KeepAliveTimeOut 5
```

These settings allow enough requests per connection and time between requests to reap the benefits of the persistent connections, while minimizing the drawbacks. You should consider the size and behavior of your own user population when setting these values. For example, if you have a large user population and the users make small infrequent requests, you may want to reduce the `keepAlive` directive default settings, or even set `KeepAlive` to off. If you have a small population of users that return to your site frequently, you may want to increase the settings.

`KeepAlive` option should be used judiciously along with `MaxClients` directive. `KeepAlive` option would tie a worker thread to an established connection until it times out or the number of requests reaches the limit specified by `MaxKeepAliveRequests`. This means that the connections or users in the `ListenBacklog` queue would be starving for a worker until the worker is relinquished by the `keep-alive` user. The starvation for resources happens on the `KeepAlive` user load with user population consistently higher than that specified in the `MaxClients`.

Note: The `MaxClients` property is applicable only to UNIX platforms. On Windows, the same functionality is achieved through the `ThreadLimit` and `ThreadsPerChild` parameters.

Increasing `MaxClients` may impact performance in the following ways:

- A high number of `MaxClients` can overload the system resources and may lead to poor performance.
- For a high user population with fewer requests, consider increasing the `MaxClients` to support `KeepAlive` connections to avoid starvation. Note that this can impact overall performance if the user concurrency increases. System performance is impacted by increased concurrency and can possibly cause the system to fail.

`MaxClients` should always be set to a value where the system would be stable or performing optimally (~85% CPU).

Typically for high user population with less frequent requests, consider turning the `KeepAlive` option off or reduce it to a very low value to avoid starvation.

Disabling the `KeepAlive` connection may impact performance in the following ways:

- Connection establishment for every request has a cost.
- If the frequency of creating and closing connections is higher, then some system resources are used. The TCP connection has a `time_wait` interval before it can close the socket connection and open file descriptors for every connection. The default `time_wait` value is 60 seconds and each connection can take 60 seconds to close, even after it is relinquished by the server.

WARNING: To avoid potential performance issues, values for any parameters should be set only after considering the nature of the workload and the system capacity.

6.3.3 Logging Options for Oracle HTTP Server

This section discusses types of logging, log levels, and the performance implications for using logging.

6.3.3.1 Access Logging

Access logs are generally enabled to track who accessed what. The `access_log` file, available in the `ORACLE_INSTANCE/diagnostics/logs/OHS/ohsname` directory, contains an entry for each request that is processed. This file grows as time passes and can consume disk space. Depending on the nature of the workload, the `access_log` has little impact on performance. If you notice that performance is becoming an issue, the file can be disabled if some other proxy or load balancer is used and gives the same information.

6.3.3.2 Configuring the `HostNameLookups` Directive

By default, the `HostNameLookups` directive is set to `Off`. The server writes the IP addresses of incoming requests to the log files. When `HostNameLookups` is set to `On`, the server queries the DNS system on the Internet to find the host name associated with the IP address of each request, then writes the host names to the log. Depending on the server load and the network connectivity to your DNS server, the performance

impact of the DNS HostNameLookup may be high. When possible, consider logging only IP addresses. On UNIX systems, you can resolve IP addresses to host names off-line, with the `logresolve` utility found in the `ORACLE_HOME/Apache/Apache/bin/` directory.

6.3.3.3 Error logging

The server notes unusual activity in an error log. The `ohsname.log` file, available in `ORACLE_INSTANCE/diagnostics/logs/OHS/ohsname` directory, contains errors, warnings, system information, and notifications (depending on the log-level setting).

The `httpd.conf` file contains the error log configuration for OHS. The logging mode is defined by the "OraLogMode" directive. The default is "odl-text", which produces the Oracle diagnostic logging format in a text file. Alternatively, change this to "odl-xml" to produce the Oracle diagnostic logging format in an XML file.

For Oracle diagnostic-style logging, "OraLogSeverity" directive is used for setting the log level.

For Apache-style logging, the `ErrorLog` and `LogLevel` directives identify the log file and the level of detail of the messages recorded. The default debug level is `Warn`.

Excessive logging can have some performance cost and may also fill disk space. The log level control should be used based on need. For requests that use dynamic resources, for example, requests that use `mod_ossso` or `mod_plsql`, there is a performance cost associated with setting higher debugging levels, such as the `debug` level.

6.4 Advanced Tuning Considerations

This section provides advanced tuning recommendations which may or may not apply to your environment. Review the following recommendations to determine if the changes would improve your HTTP Server performance.

6.4.1 Tuning Oracle HTTP Server Security

This section covers the following topics:

- [Tuning Oracle HTTP Server Secure Sockets Layer \(SSL\)](#)
- [Tuning Oracle HTTP Server Port Tunneling](#)

6.4.1.1 Tuning Oracle HTTP Server Secure Sockets Layer (SSL)

Secure Sockets Layer (SSL) is a protocol developed by Netscape Communications Corporation that provides authentication and encrypted communication over the Internet. Conceptually, SSL resides between the application layer and the transport layer on the protocol stack. While SSL is technically an application-independent protocol, it has become a standard for providing security over HTTP, and all major web browsers support SSL.

SSL can become a bottleneck in both the responsiveness and the scalability of a web-based application. Where SSL is required, the performance challenges of the protocol should be carefully considered. Session management, in particular session creation and initialization, is generally the most costly part of using the SSL protocol, in terms of performance.

This section covers the following SSL performance-related information:

- [Section 6.4.1.1.1, "Caching SSL on Oracle HTTP Server"](#)

- [Section 6.4.1.1.2, "Using SSL Application Level Data Encryption"](#)
- [Section 6.4.1.1.3, "Tuning SSL Performance"](#)

See Also: *Oracle Fusion Middleware Application Security Guide*

6.4.1.1.1 Caching SSL on Oracle HTTP Server When an SSL connection is initialized, a session-based handshake between client and server occurs that involves the negotiation of a cipher suite, the exchange of a private key for data encryption, and server and, optionally, client, authentication through digitally-signed certificates.

After the SSL session state has been initiated between a client and a server, the server can avoid the session creation handshake in subsequent SSL requests by saving and reusing the session state. The Oracle HTTP Server caches a client's SSL session information by default. With session caching, only the first connection to the server incurs high latency.

The `SSLSessionCacheTimeout` directive in `ssl.conf` determines how long the server keeps a saved SSL session (the default is 300 seconds). Session state is discarded if it is not used after the specified time period, and any subsequent SSL request must establish a new SSL session and begin the handshake again. The `SSLSessionCache` directive specifies the location for saved SSL session information (the default location is the following directory):

```
$ORACLE_INSTANCE/diagnostics/logs/$COMPONENT_ TYPE/$COMPONENT_NAME
```

Note that multiple Oracle HTTP Server processes can use a saved session cache file.

Saving SSL session state can significantly improve performance for applications using SSL. For example, in a simple test to connect and disconnect to an SSL-enabled server, the elapsed time for 5 connections was 11.4 seconds without SSL session caching. With SSL session caching enabled, the elapsed time for 5 round trips was 1.9 seconds.

The reuse of saved SSL session state has some performance costs. When SSL session state is stored to disk, reuse of the saved state normally requires locating and retrieving the relevant state from disk. This cost can be reduced when using HTTP persistent connections. Oracle HTTP Server uses persistent HTTP connections by default, assuming they are supported on the client side. In HTTP over SSL as implemented by Oracle HTTP Server, SSL session state is kept in memory while the associated HTTP connection is persisted, a process which essentially eliminates the performance impacts associated with SSL session reuse (conceptually, the SSL connection is kept open along with the HTTP connection). For more information see [Section 6.3.2, "Reducing Httpd Process Availability with Persistent Connections"](#).

6.4.1.1.2 Using SSL Application Level Data Encryption In most applications using SSL, the data encryption cost is small compared with the cost of SSL session management. Encryption costs can be significant where the volume of encrypted data is large, and in such cases the data encryption algorithm and key size chosen for an SSL session can be significant. In general there is a trade-off between security level and performance.

Oracle HTTP Server negotiates a cipher suite with a client based on the `SSLCipherSuite` attribute specified in `ssl.conf`. OHS 11g uses 128 bit Encryption algorithm by default and no longer supports lower encryption. Note that the previous release [10.1.3x] used 64 bit encryption for Windows. For UNIX, the 10.x releases had 128 bit encryption used by default.

See Also: *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server* for information on using supported cipher suites.

6.4.1.1.3 Tuning SSL Performance The following recommendations can assist you with determining performance requirements when working with Oracle HTTP Server and SSL.

1. The SSL handshake is an inherently resource intensive process in terms of both CPU usage and response time. Thus, use SSL only where needed. Determine the parts of the application that require the security, and the level of security required, and protect only those parts at the requisite security level. Attempt to minimize the need for the SSL handshake by using SSL sparingly, and by reusing session state as much as possible. For example, if a page contains a small amount of sensitive data and several non-sensitive graphic images, use SSL to transfer the sensitive data only, use normal HTTP to transfer the images. If the application requires server authentication only, do not use client authentication. If the performance goals of an application cannot be met by this method alone, additional hardware may be required.
2. Design the application to use SSL efficiently. Group secure operations to take advantage of SSL session reuse and SSL connection reuse.
3. Use persistent connections, if possible, to minimize cost of SSL session reuse.
4. Tune the session cache timeout value (the `SSLSessionCacheTimeout` directive in `ssl.conf`). A trade-off exists between the cost of maintaining an SSL session cache and the cost of establishing a new SSL session. As a rule, any secured business process, or conceptual grouping of SSL exchanges, should be completed without incurring session creation more than once. The default value for the `SSLSessionCacheTimeout` attribute is 300 seconds. It is a good idea to test an application's usability to help tune this setting.
5. If large volumes of data are being protected through SSL, pay close attention to the cipher suite being used. The `SSLCipherSuite` directive specified in `ssl.conf` controls the cipher suite. If lower levels of security are acceptable, use a less-secure protocol using a smaller key size (this may improve performance significantly). Finally, test the application using each available cipher suite for the specified security level to find the optimal suite.
6. If SSL remains a bottleneck to the performance and scalability of your application, after taking the preceding considerations into account, consider deploying multiple Oracle HTTP Server instances over a hardware cluster or consider the use of SSL accelerator cards.

6.4.1.2 Tuning Oracle HTTP Server Port Tunneling

When OracleAS Port Tunneling is configured, every request processed passes through the OracleAS Port Tunneling infrastructure. Thus, using OracleAS Port Tunneling can have an impact on the overall Oracle HTTP Server request handling performance and scalability.

With the exception of the number of OracleAS Port Tunneling processes to run, the performance of OracleAS Port Tunneling is self-tuning. The only performance control available is to start more OracleAS Port Tunneling processes; this increases the number of available connections and the scalability of the system.

The number of OracleAS Port Tunneling processes is based on the degree of availability required, and the number of anticipated connections. This number cannot

be automatically determined because for each additional process a new port must be opened through the firewall between the DMZ and the intranet. You cannot start more processes than you have open ports, and you do not want less processes than open ports, since in this case ports would not have any process bound to them.

To measure the OracleAS Port Tunneling performance, determine the request time for servlet requests that pass through the OracleAS Port Tunneling infrastructure. The response time running with OracleAS Port Tunneling should be compared with a system without OracleAS Port Tunneling to determine whether your performance requirements can be met using OracleAS Port Tunneling.

See Also: *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server* for information on configuring OracleAS Port Tunneling

6.4.2 Tuning Oracle HTTP Server

The following tips can enable you to avoid or debug potential Oracle HTTP Server performance problems:

- [Analyzing Static Versus Dynamic Requests](#)
- [Managing PL/SQL Requests](#)
- [Limiting the Number of Enabled Modules](#)
- [Monitoring Oracle HTTP Server Performance](#)

6.4.2.1 Analyzing Static Versus Dynamic Requests

It is important to understand where your server is spending resources so you can focus your tuning efforts in the areas where the most stands to be gained. In configuring your system, it can be useful to know what percentage of the incoming requests are static and what percentage are dynamic.

Generally, you want to concentrate your tuning effort on dynamic pages because dynamic pages can be costly to generate. Also, by monitoring and tuning your application, you may find that much of the dynamically generated content, such as catalog data, can be cached, sparing significant resource usage.

6.4.2.2 Managing PL/SQL Requests

You can get unrepresentative results when data outliers appear. This can sometimes occur at start-up. To simulate a simple example, assume that you ran a PL/SQL "Hello, World" application for about 30 seconds. Examining the results, you can see that the work was all done in `mod_plsql.c`:

```
/ohs_server/ohs_module/mod_plsql.c
handle.maxTime:      859330
handle.minTime:      17099
handle.avg:          19531
handle.active:       0
handle.time:         24023499
handle.completed:    1230
```

Note that `handle.maxTime` is much higher than `handle.avg` for this module. This is probably because when the first request is received, a database connection must be opened. Later requests can make use of the established connection. In this case, to obtain a better estimate of the average service time for a PL/SQL module, that does

not include the database connection open time which causes the `handle.maxTime` to be very large, recalculate the average as in the following:

```
(time - maxTime)/(completed -1)
```

For example:

```
(24023499 - 859330)/(1230 - 1) = 18847.98
```

6.4.2.3 Limiting the Number of Enabled Modules

Oracle HTTP Server, which is now based on Apache 2.2, has a slight change in architecture in the way the requests are handled, compared to the previous release of Oracle HTTP Server, which was based on Apache 1.3.

In the new architecture, Oracle HTTP Server invokes the service function of each module that is loaded (in the order of definition in `httpd.conf` file) until the request is serviced. This indicates that there is some cost associated with invoking the service function of each module, to know if the service is accepted or declined.

Because of this change in architecture, consider placing the most frequently hit modules above the others in the `httpd.conf` file.

For the static page requests, which are directly deployed to Oracle HTTP Server and served by the default handler, the request has to go through all the modules before the default handler is invoked. This process can impact performance of the request so consider enabling only the modules that are required by the deployed application. Example, if "mod_plsql" is never used by the deployed application, disable it to maintain performance.

In addition, there are a few modules that register their hooks to do some work during the URL translation phase, which would add to the cost of request processing time. Example: `mod_security`, when enabled, has a cost of about 10% on CPU Cost per Transaction for the specweb benchmark. Again, enable only those modules that are required by your deployed applications to save CPU time.

6.4.2.4 Tuning the File Descriptor Limit

A lack of available file descriptors can cause a wide variety of symptoms which are not always easily traced back to the operating system's file descriptor limit. Tuning the file descriptor limit can be accomplished by configuring the operating system's hard limit for the user who starts the OHS. Once configured, the OHS will adjust the soft limit to match the operating system limit.

Configuring file descriptor limits is platform-specific. Refer to your operating system documentation for more information. The following code example shows the command for Linux:

```
APACHECTL_ULIMIT=ulimit -S -n `ulimit -H -n`
```

Note that this limit must be reconfigured after applying a patch set.

Oracle Metadata Service (MDS) Performance Tuning

This chapter provides tuning tips for Oracle Metadata Service (MDS).

- [Section 7.1, "About Oracle Metadata Services \(MDS\)"](#)
- [Section 7.2, "Monitoring Oracle Metadata Service Performance"](#)
- [Section 7.3, "Basic Tuning Considerations"](#)
- [Section 7.4, "Advanced Tuning Considerations"](#)

7.1 About Oracle Metadata Services (MDS)

Oracle Metadata Services (MDS) is an application server and Oracle relational database that keeps metadata in these areas: a file-based repository data, dictionary tables (accessed by built-in functions) and a metadata registry. One of the primary uses of MDS is to store customizations and persisted personalization for Oracle applications. Oracle Metadata Services (MDS) is used by components such as Oracle WebCenter Portal: Framework and Oracle Application Development Framework (ADF) to manage metadata. Examples of metadata objects managed by MDS are: JSP pages and page fragments, ADF page definitions and task flows, and customized variants of those objects.

Note: Most of the Oracle Metadata Service configuration parameters are immutable and cannot be changed at run time unless otherwise specified.

7.2 Monitoring Oracle Metadata Service Performance

MDS uses DMS sensors to provide tuning and diagnostic information which can be viewed using Enterprise Manager. This information is useful, for example, to see if the MDS caches are large enough.

Information on DMS metrics can be found in the Fusion Middleware Control Console. Click **Help** at the top of the page to get more information. In most cases, the Help window displays a help topic about the current page. Click **Contents** in the Help window to browse the list of help topics, or click **Search** to search for a particular word or phrase.

7.3 Basic Tuning Considerations

Tuning is the adjustment of parameters to improve performance. The default MDS configuration must be tuned in almost all deployments. Please review the requirements and recommendations in this section carefully.

7.3.1 Tuning the Database Repository

For optimal performance of MDS APIs, the database schema for the MDS repository must be monitored and tuned by the database administrator. This section lists some recommended actions to tune the database repository:

- [Collecting Schema Statistics](#)
- [Increasing Redo Log Size](#)
- [Reclaiming Disk Space](#)
- [Monitoring the Database Performance](#)

For additional information on tuning the database, see "Optimizing Instance Performance" in *Oracle Database Performance Tuning Guide*.

7.3.1.1 Collecting Schema Statistics

While MDS provides database indexes, they may not be used as expected due to a lack of schema statistics. If performance is an issue with MDS operations such as accessing or updating metadata in database repository, the database administrator must ensure that the statistics are available and current.

The following example shows one way that the Oracle database schema statistics can be collected:

```
execute dbms_stats.gather_schema_stats(ownname => '<username>',  
estimate_percent => dbms_stats.auto_sample_size, method_opt=> 'for all  
columns size auto', cascade=>true);
```

If the performance does not improve after statistics collection, then try to flush the database shared pool to clear out the existing SQL plans by using the following command:

```
alter system flush shared_pool;
```

In general, the database should be configured with automatic statistics recollection. For additional information on gathering statistics, see "Automatic Performance Statistics" in *Oracle Database Performance Tuning Guide*.

7.3.1.2 Increasing Redo Log Size

The size of the redo log files can influence performance because the behavior of the database writer and archiver processes depend on the redo log sizes. Generally, larger redo log files provide better performance. Undersized log files increase checkpoint activity and can reduce performance.

Consider having at least 3 redo log groups with 2G of size each. Redo log files should be placed on a disk separate from data files to improve I/O performance.

For more information see "Sizing Redo Log Files" in *Oracle Database Performance Tuning Guide*.

7.3.1.3 Reclaiming Disk Space

While manual and auto-purge operations delete the metadata content from the repository, the database may not immediately reclaim the space held by tables and indexes. This may result in the disk space consumed by MDS schema growing. Database administrators can manually rebuild the indexes and shrink the tables to increase performance and to reclaim disk space.

For more information see "Reclaiming Unused Space" in *Oracle Database Performance Tuning Guide*.

7.3.1.4 Monitoring the Database Performance

Database administrators must monitor the database (for example, by generating automatic workload repository (AWR) reports for Oracle database) to observe lock contention, I/O usage and take appropriate action to address the issues.

For more information see:

- "Generating Automatic Workload Repository Reports" in *Oracle Database Performance Tuning Guide*
- "Monitoring Performance" in *Oracle Database Administrator's Guide*.

7.3.2 Tuning Cache Configuration

MDS uses a cache to store metadata objects and related objects (such as XML content) in memory. MDS Cache is a shared cache that is accessible to all users of the application (on the same JVM). If a metadata object is requested repeatedly, with the same customizations, that object may be retrieved more quickly from the cache (a "warm" read). If the metadata object is not found in the cache (a "cold" read), then MDS may cache that object to facilitate subsequent read operations depending on the cache configuration, the type of metadata object and the frequency of access.

Cache can be configured or changed post deployment through MBeans. This element maps to the `MaximumCacheSize` attribute of the `MDSAppConfig` MBean. For more information see "Changing MDS Configuration Attributes for Deployed Applications" in *Oracle Fusion Middleware Administrator's Guide*.

Note: MDS Metrics, visible in Enterprise Manager, are useful for tuning the MDS cache. In particular, "IOs Per MO Content Get" or "IOs Per Metadata Object Get" should be less than 1. If not, consider increasing the size of the MDS cache. For more information on viewing DMS metric information, see [Section 7.2, "Monitoring Oracle Metadata Service Performance"](#).

Having a correctly sized cache can significantly improve throughput for repeated reading of metadata objects. The optimal cache size depends on the number of metadata objects used and the individual sizes of these objects. Prior to packaging the Enterprise ARchive (EAR) file, you can manually update the cache-config in `adf-config.xml`, by adding the following entry:

```
<mds-config>
  <cache-config>
    <max-size-kb>200000</max-size-kb>
  </cache-config>
</mds-config>
```

Note: MDS cache grows in size as metadata objects are accessed until it hits `max-size-kb`. After that, objects are removed from the cache to make room as needed on a least recently used (LRU) basis to make room for new objects. Unless time-to-live (TTL) is set, the MDS cache continues to occupy the `max-size-kb` of memory.

7.3.2.1 Enabling Document Cache

In addition to the main MDS cache, MDS uses a document cache in conjunction with each metadata store to store thumbnail information about metadata documents (base document and customization documents) in memory. The entry for each document is small (<100 bytes) and the cache size limit is specified in terms of the number of document entries. MDS calculates an appropriate default size limit for the document cache based on the configured maximum size of the MDS Cache, as follows:

- If MDS cache is disabled, MDS defaults to having no document cache.
- If MDS cache is enabled, MDS defaults the document cache size to one document entry per KB of document cache configured.
- If `cache-config` is not specified, MDS defaults to 10000 document entries.
- If MDS cache is set to a very small value, MDS uses a minimum size of 500 for document cache.

In general, the defaults should be sufficient in most cases. However, insufficient document cache size may impact performance. Prior to packaging the Enterprise ARchive (EAR) file, you can explicitly set document cache size by adding this entry to `adf-config.xml`:

```
<metadata-store-usage id="db1">
  <metadata-store ...>
    <property name = .../>
  </metadata-store>
  <document-cache max-entries="10000"/>
</metadata-store-usage>
```

Note: Document cache is cleared when it exceeds the `document-cache max-entries` value. To avoid performance issues, consider increasing the document cache size if you receive a notification like the following for example:

```
NOTIFICATION: Document cache DBMetadataStore : MDS
Repository connection = <> exceeds its maximum number of
entries <NNNN>, so the cache is cleared.
```

The DMS metric "IOs Per Document Get" (visible in Enterprise Manager, see [Section 7.2](#)) should be less than 1. If not, consider increasing the document cache size.

7.3.3 Purging Document Version History

MDS keeps document version history in the database's metadata store. As version history accumulates, it requires more disk space and degrades read/write performance. Assuming the document versions are not part of an active label, there are two ways to purge version history:

- [Auto Purge](#)

- [Manual Purge](#)

Note: Purging version history manually may impact performance depending on the number of metadata updates that have been made since the last purge.

7.3.3.1 Auto Purge

The auto-purge interval can be configured or changed post deployment through MBeans. This element maps to the `AutoPurgeTimeToLive` attribute of the `MDSAppConfig` MBean. If your application uses the database store for MDS, you can set auto-purge by adding this entry in `adf-config.xml` prior to packaging the EAR:

```
<persistence-config>
  <auto-purge seconds-to-live="T" />
</persistence-config>
```

In the example above, the auto-purge interval removes versions that are older than the specified time *T* (in seconds). For more information, see "Changing MDS Configuration Attributes for Deployed Applications" in *Oracle Fusion Middleware Administrator's Guide*.

Tip: Adjust the auto-purge interval based on document versions created in your application. Purging can take longer based on number of versions created. See also "Advanced Systems Administration for Oracle WebCenter Portal" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.

7.3.3.2 Manual Purge

When you suspect that the database is running out of space or performance is becoming slower, you can manually purge existing version history using `WLST` command or through Oracle Enterprise Manager. Manual purging may impact performance, so plan to purge in a maintenance window or when the system is not busy.

For more information about manually purging version history, see "Purging Metadata Version History" in *Oracle Fusion Middleware Administrator's Guide*.

7.3.4 Using Database Polling Interval for Change Detection

MDS employs a polling thread which queries the database to gauge if the data in the MDS in-memory cache is out of sync with data in the database. This can happen when metadata is updated in another JVM. If it is out of sync, MDS clears any out of date-cached data so subsequent operations see the latest versions of the metadata. MDS invalidates the document cache, as well as MDS cache, so subsequent operations have the latest version of the metadata.

The polling interval can be configured or changed post deployment through MBeans. The element maps to the `ExternalChangeDetection` and `ExternalChangeDetectionInterval` attributes of the `MDSAppConfig` MBean. Prior to packaging the Enterprise ARchive (EAR) file, you can configure the polling interval by adding this entry in `adf-config.xml`:

```
<mds-config>
  <persistence-config>
    <external-change-detection enabled="true" polling-interval-secs="T" />
  </persistence-config>
```

</mds-config>

In the example above, 'T' specifies the polling interval in seconds. The minimum value is 1. Lower values cause metadata updates, that are made in other JVMs, to be seen more quickly. It is important to note, however, that a lower value can also create increased middle tier and database CPU consumption due to the frequent queries. By default, polling is enabled ('true') and the default value of 30 seconds should be suitable for most purposes. For more information, see "Changing MDS Configuration Attributes for Deployed Applications" in *Oracle Fusion Middleware Administrator's Guide* .

Note: When setting the polling interval, consider the following: if you poll too frequently, the database is queried for out-of-date versions; too infrequently, and those versions may stack up and polling can take longer to process.

7.4 Advanced Tuning Considerations

After you have performed the modifications recommended in the previous section, you can make additional changes that are specific to your deployment. Consider carefully whether the recommendations in this section are appropriate for your environment.

7.4.1 Analyzing Performance Impact from Customization

MDS customization may impact performance at run time. The impact from customization depends on many factors including:

- The type of customization that has been created (shared or user level)
- The percentage of metadata objects in the system which is customized. The lower this percentage the lower the impact of customization.
- The number of configured customization layers, and the efficiency of the customization classes.

There are two main types of customization:

- **Shared Customizations:** these are layers of customization corresponding to customization classes whose `getCacheHint` method returns `ALL_USERS` or `MULTI_USER`, meaning the layer applies to all or multiple users. Shared customizations are cached in the (shared) MDS cache.
- **User Level Customizations (also known as Personalizations):** these are layers of customization corresponding to customization classes whose `getCacheHint` method returns `SINGLE_USER`, meaning the layer applies to just one user. User customizations are generally cached on the user's session (`HttpSession`) until the user logs out.

For more information about customization concepts, writing customization classes, and configuring customization classes, see "Customizing Applications with MDS" in *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

Part III

Oracle Fusion Middleware Server Components

This part describes configuring Oracle Fusion Middleware server components to improve performance. It contains the following chapters:

- [Chapter 8, "Oracle Application Development Framework Performance Tuning"](#)
- [Chapter 9, "Oracle TopLink \(EclipseLink\) JPA Performance Tuning"](#)
- [Chapter 10, "Oracle Web Cache Performance Tuning"](#)

Oracle Application Development Framework Performance Tuning

This chapter provides basic guidelines on how to maximize the performance and scalability of the Oracle Application Development Framework (ADF). This chapter covers design, configuration, and deployment performance considerations in the following sections:

- [Section 8.1, "About Oracle ADF"](#)
- [Section 8.2, "Basic Tuning Considerations"](#)
- [Section 8.3, "Advanced Tuning Considerations"](#)

This chapter assumes that you are familiar with building ADF applications. To learn about ADF, see the following guides:

- *Oracle Application Development Framework Developer's Guide*
- *Oracle Fusion Middleware Web User Interface Developer's Guide for Oracle Application Development Framework*

8.1 About Oracle ADF

Oracle Application Development Framework (Oracle ADF) is an end-to-end application framework that builds on Java Platform, Enterprise Edition (Java EE) standards and open-source technologies to simplify and accelerate implementing service-oriented applications. Oracle ADF is suitable for enterprise developers who want to create applications that search, display, create, modify, and validate data using web, wireless, desktop, or web services interfaces.

For more information see "Introduction to Oracle ADF" in *Oracle Application Development Framework Developer's Guide*.

8.2 Basic Tuning Considerations

Before building, configuring, and deploying ADF applications, review the following tuning recommendations to achieve optimal performance:

- [Oracle ADF Faces Configuration and Profiling](#)
- [Tuning ADF Faces](#)
- [Tuning ADF Faces Component Attributes](#)
- [Tuning Table and Tree Components](#)
- [Tuning Data Visualization Tool \(DVT\) Components](#)

■

8.2.1 Oracle ADF Faces Configuration and Profiling

This section discusses the configuration and profiling concepts of the ADF Faces. Configuration options for Oracle ADF Faces are set in the `web.xml` file. Most of parameters have default values that can be tuned for performance. [Table 8–1](#) describes some of these configuration options.

Table 8–1 ADF Configuration Options

Parameter	Description
Compress View State <code>org.apache.myfaces.trinidad.COMPRESS_VIEW_STATE</code>	Controls whether or not the page state is compressed. Latency can be reduced if the size of the data is compressed. The default is <code>False</code> , but to optimize performance, consider setting this parameter to <code>True</code> .
Debug Mode <code>org.apache.myfaces.trinidad.resource.DEBUG</code>	Controls whether output should be enhanced for debugging or not. This parameter should be removed or set to <code>False</code> .
Check File Modification <code>org.apache.myfaces.trinidad.resource.CHECK_FILE_MODIFICATION</code>	Controls whether ADF faces check for modification date of JSP pages and discard any saved state if the file is changed. This parameter should be removed or set to <code>False</code> .
Client State Method <code>oracle.adf.view.rich.CLIENT_STATE_METHOD</code>	Specifies which type of saving (<code>all</code> or <code>token</code>) should be used when client-side state saving is enabled. The default value is <code>token</code> . Consider using <code>token</code> to optimize performance.
Log Level <code>oracle.adf.view.rich.LOGGER_LEVEL</code>	Sets the log level on the client side. The default value is <code>OFF</code> . This parameter should be removed or set to <code>OFF</code> .
Assert Enabled <code>oracle.adf.view.rich.ASSERT_ENABLED</code>	Specifies whether to process assertions on the client side. The default value is <code>OFF</code> . This parameter should be removed or set to <code>OFF</code> .

Note: When you are profiling or measuring client response time using the Firefox browser, ensure that the Firebug plug-in is disabled. While this plug-in is very useful for getting information about the page and for debugging JavaScript code on the page, it can impact the total response time.

For more information on disabling the Firefox Firebug plug-in, see the Firefox Support Home Page at <http://support.mozilla.com/en-US/kb/>.

8.2.2 Tuning ADF Faces

[Table 8–2](#) provides configuration recommendations that may improve performance of ADF Faces:

Table 8–2 Configuration Parameters for ADF Faces

Configuration Recommendation	Description
Avoid inline JavaScript in pages.	<p>Inline JavaScript can increase response payload size, will never be cached in browser, and can block browser rendering. Instead of using inline JavaScript, consider putting all scripts in .js files in JavaScript libraries and add scripts to the page using af:resource tag.</p> <p>NOTE: Consider using af:resource rather than trh:script when possible.</p>
Avoid self-signed SSL certificates. Opt for officially signed security certificates instead.	<p>In some browsers, using a self-signed SSL certificate prevents resources from being cached. The resources, such as JavaScript libraries, images, and style sheets, must be reloaded each time a user accesses the application, which can impact network performance.</p> <p>An added benefit of using officially-signed SSL certificates is the additional security assurances they can provide over self-signed certificates.</p>
Configure the JSP timeout parameter.	<p>Using the JavaServer Pages (JSP) timeout parameter causes infrequently used pages to be flushed from the cache by the following setting in web.xml:</p> <pre><servlet> <servlet-name> oraclejsp <init-param> <param-name> jsp_timeout </param-name> <param-value> 600 </param-value> </init-param> </servlet-name> </servlet></pre> <p>NOTE: Set this parameter based on your own use case scenarios.</p>
Create a single toolbar item with a drop-down popup.	<p>When the browser size is small because of the screen resolution, the menubar/toolbar overflow logic becomes expensive in Internet Explorer 7 and 8. It especially has problems with laying out DOM structures with command items.</p> <p>Create a single toolbar item with a drop-down popup and put all the input fields inside it. This popup should have deferred child creation and contentDelivery="lazy".</p>
Use shared popups in collection-based components (like table).	<p>Stamping popups in a table can impact performance, so consider using a shared popup in collection-based components (like table) to reduce the overhead.</p>
Do not use hover popups on navigation links.	<p>A hover popup on a navigation link causes the navigation to wait for the hover to be fetched first.</p> <p>Move hover popups on navigation links in a table to a separate icon within the table cell</p>

Table 8–2 (Cont.) Configuration Parameters for ADF Faces

Configuration Recommendation	Description
Use partial page navigation.	<p>Partial Page Navigation is a feature of the ADF Faces framework that enables navigating from one ADF Faces page to another without a full page transition in the browser. The new page is sent to the client using Partial Page Rendering (PPR)/Ajax channel.</p> <p>The main advantage of partial page navigation over traditional full page navigation is improved performance: the browser no longer re-interprets and re-executes Javascript libraries, and does not spend time for cleanup/initialization of the full page. The performance benefit from this optimization is very big; it should be enabled whenever possible.</p> <p>Some known limitations of this feature are:</p> <ul style="list-style-type: none"> ■ For the document's "metaContainer" facet (the HEAD section), only scripts are brought over with the new page. Any other content, such as icon links or style rules can be ignored. ■ Applications cannot use anchor (hash) URLs for their own purposes. <p>For more information, see "Using Partial Page Navigation" in the <i>Oracle Fusion Middleware Web User Interface Developer's Guide for Oracle Application Development Framework</i>.</p>
Use page templates.	<p>Page templates enable developers to build reusable, data-bound templates that can be used as a shell for any page. A developer can build one or more templates that provide structure and consistency for other developers building web pages. The templates have both static areas on them that cannot be changed when they are used and dynamic areas on them where the developer can place content specific to the page they are building.</p> <p>There are some important considerations when using templates:</p> <ul style="list-style-type: none"> ■ Since templates are present in every application page, they have to be optimized so that common performance impacts are avoided. Adding round corners to the template, for example, can impact the performance for every page. ■ When building complex templates, avoid building them in multiple pieces and including them in the top-level template using <code><f:subview></code> tag. From a performance perspective, doing so can impact memory usage on the server side. (<code><f:subview></code> introduces another level into the ID scoping hierarchy, which results in longer IDs. Long IDs have a negative impact on performance. Developers are advised to avoid using <code><f:subview></code> unless it is required. It is not necessary to use <code><f:subview></code> around <code><jsp:include></code> if you can ensure that all IDs are unique. For example, if you are using <code><jsp:include></code>, break a large page into multiple pieces for easier editing. And whenever possible, avoid using <code><f:subview></code>. If you are including content developed by someone else, use <code><f:subview></code> if you do not know which IDs the developer used. In addition, you do not have to put <code><f:subview></code> at the top of a region definition. ■ Avoid long IDs in all cases, especially on pageTemplates, subviews, subforms, and on tables or within tables. Long IDs can have a performance impact on the server side, network traffic, and client processing. <p>For more information, see the topic "Using Page Templates" in <i>Oracle Fusion Middleware Web User Interface Developer's Guide for Oracle Application Development Framework</i>.</p>

Table 8–2 (Cont.) Configuration Parameters for ADF Faces

Configuration Recommendation	Description
Enable ADF faces geometry management.	<p>ADF Rich Client supports geometry management of the browser layout where parent components are in the UI explicitly. The children components are sized to stretch and fill up available space in the browser. While this feature makes the UI look better, it has a cost. The impact is on the client side where the browser must spend time resizing the components. The components that have geometry management by default are:</p> <ul style="list-style-type: none"> ■ PanelGridLayout ■ PanelAccordion ■ PanelStretchLayout ■ PanelTabbed ■ BreadCrumbs ■ NavigationPane ■ PanelSplitter ■ Toolbar ■ Toolbox ■ Table ■ Train <p>Notes:</p> <ul style="list-style-type: none"> ■ Consider using panelGridLayout to flatten HTML hierarchy and improve performance when designing page layouts for non-stamped component. panelGridLayout uses JavaScript to accomplish page layout like PSL, but it does not have the overhead of nesting JS layout components. To avoid performance degradation, do not use PanelGridLayout in stamped components ■ When using geometry management, try minimizing the number of child components that are under a parent geometry managed component. ■ The cost of geometry management is directly related to the complexity of child components. ■ The performance cost of geometry management can be smaller (as perceived by the user) for the pages with table or other data stamped components when table data streaming is used. The client-side geometry management can be executed while the browser is waiting for the data response from the server. <p>For more information, see the topic "Geometry Management and Component Stretching" in <i>Oracle Fusion Middleware Web User Interface Developer's Guide for Oracle Application Development Framework</i>.</p>
Use the ADF rich client overflow feature.	<p>ADF Rich Client supports overflow feature. This feature moves the child components to the non-visible overflow area if they cannot fit the page. The components that have built-in support for overflow are: PanelTabbed, BreadCrumbs, NavigationPane, PanelAccordion, Toolbar, and Train. Toolbar should be contained in a Toolbox to handle the overflow.</p> <p>While several optimizations methods are in place to reduce the cost of overflow, it is necessary to pay special attention to the number of child components and complexity of each of them in the overflow component. Consider setting a big enough initial size of the overflow component so that overflow does not happen in most cases.</p> <p>To prevent performance degradation, however, consider adding a drop-down with the command items to avoid these overflow areas.</p>

Table 8–2 (Cont.) Configuration Parameters for ADF Faces

Configuration Recommendation	Description
Use ADF Rich Client Partial Page Rendering (PPR).	<p data-bbox="428 289 1357 422">ADF Rich Client is based on Asynchronous JavaScript and XML (Ajax) development technique. Ajax is a web development technique for creating interactive web applications, where web pages feel more responsive by exchanging small amounts of data with the server behind the scenes, without the whole web page being reloaded. The effect is to improve a web page's interactivity, speed, and usability.</p> <p data-bbox="428 436 1357 569">With ADF Faces, the feature that delivers the Ajax partial page refresh behavior is called partial page rendering (PPR). PPR enables small areas of a page to be refreshed without having to redraw the entire page. For example, an output component can display what a user has chosen or entered in an input component or a command link or button can cause another component on the page to be refreshed.</p> <p data-bbox="428 583 1235 609">Two main Ajax patterns are implemented with partial page rendering (PPR):</p> <ul data-bbox="428 623 743 688" style="list-style-type: none"> <li data-bbox="428 623 743 651">■ native component refresh <li data-bbox="428 659 743 688">■ cross-component refresh <p data-bbox="428 703 1341 753">While the framework builds in native component refresh, cross-component refresh has to be done by developers in certain cases.</p> <p data-bbox="428 768 1341 978">Cross-Component refresh is a concept where different components on the page have a semantic relationship to one another. For example, if a user selects one choice component on a page that prompts a country selection. The outcome of this selection requires another component on the page to be changed and updated to show the valid states of the country. This relationship requires a cross-component refresh. This can be accomplished either declarative or programmatically. For more information, refer to <i>Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework</i>.</p> <p data-bbox="428 993 1341 1098">It is important to note the differences between partial targets and partial triggers when using cross-component refresh. Partial triggers are used when a component wants to listen for changes on another component. Partial targets are used when a component that changes wants to make sure other components on the page are re-rendered.</p> <p data-bbox="428 1113 1357 1297">Consider a typical situation in which a page includes an <code>af:inputText</code> component, an <code>af:commandButton</code> component, and an <code>af:outputText</code> component. When the user enters a value for the <code>af:inputText</code>, then clicks the <code>af:commandButton</code>, the input value is reflected in the <code>af:outputText</code>. Without PPR, clicking the <code>af:commandButton</code> triggers a full-page refresh. Using PPR, you can limit the scale of the refresh to only those components you want to refresh, in this case the <code>af:outputText</code> component. To achieve this, you would do two things:</p> <ul data-bbox="428 1312 1341 1535" style="list-style-type: none"> <li data-bbox="428 1312 1341 1388">■ Set up the <code>af:commandButton</code> for partial submit by setting the <code>partialSubmit</code> attribute to <code>true</code>. Doing this causes the command component to start firing partial page requests each time it is clicked. <li data-bbox="428 1402 1341 1535">■ Define which components are to be refreshed when the partial submit takes place, in this example the <code>af:outputText</code> component, by setting the <code>partialTriggers</code> attribute for each of them to the id of the component triggering the refresh. In this example, this means setting the <code>partialTriggers</code> attribute of the <code>af:outputText</code> component to give the id of the <code>af:commandButton</code> component. <p data-bbox="428 1549 1281 1600">The steps above achieve PPR using a command button to trigger the partial page refresh.</p> <p data-bbox="428 1614 1341 1747">The main reason why partial page rendering can significantly boost the performance is that full page refresh does not happen and the framework artifacts (such as ADF Rich Client JS library, and style sheets) are not reloaded and only a small part of page is refreshed. In several cases, this means no extra data is fetched or no geometry management.</p> <p data-bbox="428 1761 1357 1837">The ADF Rich Client has shown that partial page rendering results in the best client-side performance. Besides the impact on the client side, server-side processing can be faster and can have better server-side throughput and scalability.</p> <p data-bbox="428 1852 1341 1927">For more information, see the topic "Rerendering Partial Page Content" in <i>Oracle Fusion Middleware Web User Interface Developer's Guide for Oracle Application Development Framework</i>.</p>

Table 8–2 (Cont.) Configuration Parameters for ADF Faces

Configuration Recommendation	Description
Use ADF rich client navigation.	<p>ADF Rich Client has an extensive support for navigation. One of the common use cases is tabbed navigation. This functionality is currently supported by components like <code>navigationPane</code> which can bind to <code>xmlMenuModel</code> to easily define navigation.</p> <p>There is one drawback in this approach, however. Using ADF rich client navigation results in a full page refresh every time the user switches the tab. One option is to use <code>panelTabbed</code> instead. <code>panelTabbed</code> has built-in support for partial page rendering of the tabbed content without requiring any developer work. However, <code>panelTabbed</code> cannot bind to any navigational model and the content has to be available from within the page, so it has limited applicability.</p>
Cache resources.	<p>Developers are strongly encouraged to ensure that any resources that can be cached (images, CSS, JavaScript) have their cache headers specified appropriately. Also, client requests for missing resources on the server result in addition round trips to the server. To avoid this, make sure all the resources are present on the server.</p> <p>Consider using the <code>ResourceServlet</code> to configure <code>web.xml</code> to enable resource caching:</p> <pre data-bbox="508 720 1036 947"> <servlet-mapping> <servlet-name>resources</servlet-name> <url-pattern>/js/*</url-pattern> </servlet-mapping> <servlet-mapping> <servlet-name>resources</servlet-name> <url-pattern>/images/*</url-pattern> </servlet-mapping> </pre>
Reduce the size of state token cache.	<p>This property is defined in <code>web.xml</code> <code>org.apache.myfaces.trinidad.CLIENT_STATE_MAX_TOKENS</code> in "token"-based client-side state saving and determines how many tokens should be preserved at any one time. The default value is 15. When this value is exceeded, state will be "forgotten" for the least recently viewed pages, which can impact users that actively use the Back button or that have multiple windows open simultaneously.</p> <p>In order to reduce live memory per session, consider reducing this value to 2. Reducing the state token cache to 2 means one Back button click is supported. For applications without support for a Back button, this value should be set to 1.</p> <p>CAUTION: Setting this parameter to 1 is not recommended for Oracle WebCenter applications as it may cause the JSP error <code>IllegalStateException</code>. For example, WebCenter Spaces and WebCenter Framework Portal applications launch dialog windows, implying that the previous state of the JSF view tree has to be stored. If the number is set to 1, then the system can store the current main view page and then view the dialog page when viewing a mail, for example. The dialog page is now trying to open another dialog page (for reply or forward), but the server cannot save the state of the previous page and returns the <code>IllegalStateException</code>.</p>
Use <code>af:resource</code> component to deliver css and js scripts.	<p>A common developer task is to define custom styles inside a regular page or template page. Since most browsers use progressive scanning of the page, a late introduction of styles forces the browser to recompute the page. This impacts the page layout performance. For better performance use <code>af:resource</code> component to deliver css and js scripts.</p> <p>To get a component (or static CDATA content) to display in the "head", use the "metaContainer" facet.</p> <p>See the <i>Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework</i> for details about <code>af:facetRef</code>.</p>

Table 8–2 (Cont.) Configuration Parameters for ADF Faces

Configuration Recommendation	Description
Optimize custom JavaScript code.	<p>ADF Rich Client uses JavaScript on the client side. The framework itself provides most of the functionality needed. However, you may have to write custom JavaScript code. To get the best performance, consider bundling the JavaScript code into one JS lib (one JavaScript file) and deliver it to the client. The easiest approach is to use the ADF tag: <code><af:resource type="javascript" source=" " /></code>.</p> <p>If most pages require custom JavaScript code, the tag should be included in the application template. Otherwise, including it in particular pages can result in better performance. If custom the JavaScript code lib file becomes too big, then consider splitting it into meaningful pieces and include only the pieces needed by the page. Overall, this approach is faster since the browser cache is used and the html content of the page is smaller.</p>
Disable debug output mode.	<p>The debug-output element in the <code>trinidad-config.xml</code> file specifies whether output should be more verbose to help with debugging. When set to <code>TRUE</code>, the output debugging mechanism in Trinidad produces pretty-printed, commented HTML content. To improve performance by reducing the output size, you should disable the debug output mode in production environments.</p> <p>Set the debug-output element to <code>FALSE</code>, or if necessary, remove it completely from the <code>trinidad-config.xml</code> file.</p>
Disable test automation.	<p>Enabling test automation parameter <code>oracle.adf.view.rich.automation.ENABLED</code> generates a client component for every component on the page which can negatively impact performance.</p> <p>Set the <code>oracle.adf.view.rich.automation.ENABLED</code> parameter value to <code>FALSE</code> (the default value) in the <code>web.xml</code> file to improve performance.</p>
Disable animation.	<p>ADF Rich Client framework has client side animation enabled by default. Animation is introduced to provide an enhanced user experience. Some of the components, like popup table, have animation set for some of the operations. While using animation can improve the user experience, it can increase the response time when an action is executed. If speed is the biggest concern, then animation can be disabled by setting the flag in <code>trinidad-config.xml</code></p>
Disable client-side assertions.	<p>Assertions on client-side code base can have a significant impact on client-side performance. Set the parameter value to <code>FALSE</code> (the default value) to disable client-side assertions. Also ensure that the <code>oracle.adf.view.rich.ASSERT_ENABLED</code> is not explicitly set to <code>TRUE</code> in the <code>web.xml</code> file.</p>
Disable JavaScript Profiler.	<p>When the JavaScript <code>oracle.adf.view.rich.profiler.ENABLED</code> profiler is enabled, an extra round-trip occurs on every page in order to fetch the profiler data. Disable the profiler in the <code>web.xml</code> file to avoid this extra round-trip. Default is <code>DISABLED</code>.</p>
Disable resource debug mode.	<p>When resource debug mode is enabled, the HTTP response headers do not tell the browser (or WebCache) that resources (JS libraries, CSS style sheets, or images) can be cached.</p> <p>Disable the <code>org.apache.myfaces.trinidad.resource.DEBUG</code> parameter in the <code>web.xml</code> file to ensure that caching is enabled.</p>
Disable timestamp checking.	<p>The <code>org.apache.myfaces.trinidad.CHECK_FILE_MODIFICATION</code> parameter controls whether jsp or jsp files are checked for modifications each time they are accessed.</p> <p>Ensure that the parameter value <code>org.apache.myfaces.trinidad.CHECK_FILE_MODIFICATION</code> is set to <code>FALSE</code> (the default value) in the <code>web.xml</code> file.</p>
Disable checking for CSS file modifications.	<p>The <code>org.apache.myfaces.trinidad.CHECK_FILE_MODIFICATION</code> parameter controls when CSS file modification checks are made. To aid in performance, this configuration option defaults to <code>false</code> - do not check for css file modifications. Set this to <code>TRUE</code> if you want the skinning css file changes to be reflected without stopping or starting the server.</p>

Table 8–2 (Cont.) Configuration Parameters for ADF Faces

Configuration Recommendation	Description
Enable content compression.	<p>By default, style classes that are rendered are compressed to reduce page size. In production environments, make sure you remove the <code>DISABLE_CONTENT_COMPRESSION</code> parameter from the <code>web.xml</code> file or set it to <code>FALSE</code>.</p> <p>For debugging, turn off the style class content compression. You can do this by setting the <code>DISABLE_CONTENT_COMPRESSION</code> property to <code>TRUE</code>.</p>
Enable JavaScript obfuscation.	<p>ADF Faces supports a runtime option for providing a non-obfuscated version of the JavaScript library. The obfuscated version is supplied by default, but the non-obfuscated version is supplied for development use. Obfuscation reduces the overall size of the JavaScript library by about 50%.</p> <p>To provide an obfuscated ADF Faces build, set the <code>org.apache.myfaces.trinidad.DEBUG_JAVASCRIPT</code> parameter to <code>FALSE</code> in the <code>web.xml</code> file.</p>
Enable library partitioning.	<p>In the Oracle 11g Release, library partitioning is on by default. In previous versions library partitioning was off by default. Ensure that the library partitioning is on by validating the <code>oracle.adf.view.rich.libraryPartitioning.DISABLED</code> property is set to <code>false</code> in the <code>web.xml</code> file.</p>

8.2.3 Tuning ADF Faces Component Attributes

[Table 8–3](#) provides configuration recommendations for ADF Faces Component Attributes:

Table 8–3 ADF Faces Component Attributes

Configuration Recommendation	Description
Use the <code>childCreation</code> property for <code>panelTabbed</code> components to control when the contents of the <code>ShowDetailItem</code> children are created.	<p>By default, the child components of <code>showDetailItems</code> inside a <code>panelTabbed</code> are created even if they are not disclosed. This can cause unnecessary memory and CPU usage.</p> <p>To control when the contents of the <code>showDetailItems</code> children are created, use the <code>childCreation</code> property. Using this property can improve server-side performance by postponing construction of the components until they are likely to be disclosed.</p> <p>The <code>childCreation</code> property options are:</p> <ul style="list-style-type: none"> ■ <code>immediate</code> All <code>showDetailItem</code> children will be populated immediately ■ <code>lazy</code> <code>showDetailItem</code> children will be populated when they are likely to be disclosed and kept in memory thereafter ■ <code>lazyuncached</code> <code>showDetailItem</code> children will be populated when they are likely to be disclosed and may be removed when undisclosed
Use the <code>immediate</code> data delivery attribute with ADF Rich Client components.	<p>ADF Rich Client components have an <code>immediate</code> attribute. If a component has its <code>immediate</code> attribute set to <code>TRUE</code> (<code>immediate="true"</code>), then the validation, conversion, and events associated with the component are processed during the <code>applyRequestValues</code> phase. These are some cases where setting <code>immediate</code> to <code>TRUE</code> can lead to better performance.</p> <ul style="list-style-type: none"> ■ The <code>commandNavigationItem</code> in the <code>navigationPane</code> can use the <code>immediate</code> attribute set to <code>TRUE</code> to avoid processing the data from the current screen while navigating to the new page. ■ If the input component value has to be validated before the other values, <code>immediate</code> should be set to <code>TRUE</code>. In case of an error it be detected earlier in the cycle and additional processing be avoided. <p>ADF Rich Client is built on top of JSF and uses standard JSF lifecycle. See "Understanding the JSF and ADF Faces Lifecycles" in <i>Oracle Fusion Middleware Web User Interface Developer's Guide for Oracle Application Development Framework</i>.</p> <p>There are some important issues associated with the <code>immediate</code> attribute. Refer to "Using the Immediate Attribute" in <i>Oracle Fusion Middleware Web User Interface Developer's Guide for Oracle Application Development Framework</i> for more information.</p> <p>Note that this is an advanced feature. Most of the performance improvements can be achieved using the <code>af:subForm</code> component. Refer to <i>Oracle Fusion Middleware Web User Interface Developer's Guide for Oracle Application Development Framework</i> for <code>af:subForm</code> details.</p>
Use the "visible" and "rendered" attributes.	<p>All ADF Faces Rich Client display components have two properties that dictate how the component is displayed on the page:</p> <ul style="list-style-type: none"> ■ The <code>visible</code> property specifies simply whether the component is to be displayed on the page, or is to be hidden. ■ The <code>rendered</code> property specifies whether the component shall exist in the client page at all. <p>The EL expression is commonly used to control these properties. For better performance, consider setting the component to "not rendered" instead of "not visible", assuming there is no client interaction with the component. Making a component "not rendered" can improve server performance and client response time since the component does not have client side representation.</p>

Table 8–3 (Cont.) ADF Faces Component Attributes

Configuration Recommendation	Description
Use client-side events.	<p data-bbox="678 260 1442 369">ADF Rich Client framework provides the client-side event model based on component-level events rather than DOM level. The client-side event model is a very useful feature that can speed up the application. Review the following performance considerations:</p> <ul data-bbox="678 380 1442 636" style="list-style-type: none"> <li data-bbox="678 380 1442 516">■ Consider using client-side events for relatively simple event handling that can be done on the client side. This improves client side performance by reducing the number of server round trips. Also, it can increase server-side throughput and scalability since requests do not have to be handled by the server. <li data-bbox="678 527 1442 636">■ By default, the events generated on the client by the client components are propagated to the server. If a client-side event handler is provided, consider canceling the event at the end of processing so that the event does not propagate to the server.
Limit character length when using "id" attribute.	<p data-bbox="678 653 1442 762">The "id" attribute should not be longer than 7 characters in length. This is particularly important for naming containers. A long id can impact performance as the amount of HTML that must be sent down to the client is impacted by the length of the ids.</p>

Table 8–3 (Cont.) ADF Faces Component Attributes

Configuration Recommendation	Description
Use client-side components.	<p>ADF Rich Client framework has client-side components that play a role in client-side event handling and component behavior. The <code>clientComponent</code> attribute is used to configure when (or if) a client-side component should be generated. Setting <code>clientComponent</code> attribute to <code>TRUE</code> has a performance impact, so determine if its necessary to generate client-side components.</p> <p>For more information, see "Client-side Components" in <i>Oracle Fusion Middleware Web User Interface Developer's Guide for Oracle Application Development Framework</i>.</p>
Optimize popups by setting the <code>childCreation</code> attribute on <code>af:popup</code> to <code>deferred</code> for a server-side performance enhancement	<p>Setting <code>childCreation</code> to <code>deferred</code> postpones construction of the components under the popup until the content is delivered. A deferred setting can therefore reduce the footprint of server-side state in some cases.</p> <p>CAUTION: This approach CANNOT be used if any of the following tags are present inside the popup:</p> <ul style="list-style-type: none"> ■ <code>f:attribute</code> ■ <code>af:setPropertyListener</code> ■ <code>af:clientListener</code> ■ <code>af:serverListener</code> <p>It also CANNOT be used if you need to refer to any child components of the popup before the popup is displayed. Setting <code>childCreation="deferred"</code> will postpone creating any child components of the popup and you cannot refer to them until after the popup is shown.</p>
Use the lazy data delivery attribute to optimize popups that include regions.	<p>If setting the <code>childCreation</code> attribute to <code>deferred</code> cannot be used, or you have regions within the popup, consider using the lazy data delivery attribute to optimize the popup.</p> <p>For this scenario, perform the following:</p> <ol style="list-style-type: none"> 1. Edit the <code>pageDef</code> for the page or fragment containing the region in the popup. 2. Locate the <code><taskflow></code> binding entry for the task flow in question. 3. Set the activation attribute to conditional (this is one of the options from the pull down list in the property inspector). 4. Set the active attribute to an EL expression that initially resolves to a boolean false. 5. In a <code>popupFetch</code> <code>setPropertyListener</code>, update the underlying value tested by the EL used in Step 4 so that it resolves to true.

8.2.4 Tuning Table and Tree Components

Table, Tree, and TreeTable are some of the most complex, and frequently used, components. Since these components can include large sets of data, they can be the common source of performance problems. [Table 8–4](#) provides some performance recommendations.

Table 8–4 Table and Tree Component Configurations

Configuration Recommendation	Description
Use <code>editingMode="clickToEdit"</code> .	<p>When using <code>editingMode="editAll"</code> all content of the editable values holders and their client components is sent. This can significantly increase the HTTP payload and the Document Object Model (DOM) content on the client.</p> <p>Consider switching to <code>editingMode="clickToEdit"</code> to reduce the amount of transmitted data and potentially improve user interaction.</p>
Set the <code>RowCountTreshold</code> to optimize the <code>SELECT COUNT</code> execution for ADF table rendering.	<p>By default table components render in "known row count" mode which means it will size scrollbar based on total number of rows. This mode provides rich UI, but also requires <code>SELECT COUNT</code> execution for ADF table rendering to get total the number of rows. In some cases <code>SELECT COUNT</code> can be slow and hard to optimize on the database.</p> <p>If ADF-Business Components (BC) is used, then <code>RowCountTreshold</code> can be set to -1 which will avoid <code>SELECT COUNT</code> but the table will render in "unknown row count" mode.</p> <p>The other option is to set <code>RowCountTreshold</code> to positive <i>N</i> which causes an optimized <code>SELECT COUNT</code> to execute and the table will have a richer rendering as long as user scrolls within <i>N</i> rows.</p>
Reduce <code>fetchSize</code> when possible.	<p>A larger fetch size attribute on <code>af:table</code> implies that more data needs to be processed, fetched from the server, and displayed on the client. This can also increase the amount of DOM displayed on the client.</p>
Modify table fetch size.	<p>Tables have a fetch size which defines the number of rows to be sent to the client in one round-trip. To get the best performance, keep this number low while still allowing enough rows to fulfill the initial table view port. This ensures the best performance while eliminating extra server requests.</p> <p>In addition, consider keeping the table fetch size and iterator range size in sync. By default, the table fetch size is set to the EL expression <code>{bindings.<name>.rangeSize}</code> and should be equal to the iterator size.</p> <p>For more information see "Using Tables and Trees" in <i>Oracle Fusion Middleware Web User Interface Developer's Guide for Oracle Application Development Framework</i>.</p>
Disable column stretching.	<p>Columns in the table and <code>treeTable</code> components can be stretched so that there is no unused space between the end of the last column and the edge of the table or <code>treeTable</code> component. This feature is turned off by default due to potential performance impacts. Turning this feature on may have a performance impact on the client rendering time, so use caution when enabling this feature with complex tables.</p>
Use header rows and frozen columns only when necessary.	<p>The table component provides features that enable you to set the row Header and frozen columns. These options can provide a well-designed interface which can lead to a good user experience. However, they can impact client-side performance. To get the best performance for table components, use these options only when they are needed.</p>
Disable popups that cannot be displayed by the user.	<p>Most cells have no attachments initially and only one popup can be displayed by the user. Therefore, popups that cannot be displayed by the user should have <code>renderer="false"</code>. This will cut down the unnecessary DOM/client components sent to the browser. Similarly the DOM has a <code>panelGroupLayout</code> with a number of cells which are empty. There is no need to send DOM for empty cells.</p>

8.2.4.1 Specifying a Data Delivery Method

Data for Table, Tree, and other stamped components can be delivered immediately or lazily. The table below describes tuning scenarios when one option is better for performance than the other. In many cases you must consider what is being tuned to determine which option to use.

By default, **lazy** delivery is used. This means that data is not delivered in the initial response from the server. Rather, after the initial page is rendered, the client asks the server for the data and gets it as a response to the second request.

In the case of **immediate** delivery, data can be in line with the response to the page request. It is important to note that data delivery is per component and not per page. This means that these two can be mixed on the same page.

When choosing between these two options, consider the following:

Lazy Delivery (default)	<p>Lazy delivery should be used for tables, or other stamped components, which are known to have slow fetch time. The examples are stamped components are the ones based on data controls using web services calls or other data controls with slow data fetch. Lazy delivery can also be used on pages where content is not immediately visible unless the user scrolls down to it. In this case the time to deliver the visible context to the client will be shorter, and the user perceives better performance.</p> <p>Lazy delivery is implemented using data streaming technique. The advantage of this approach is that the server has the ability to execute data fetches in parallel and stream data back to the client as soon as the data is available. The technique performs very well for a page with two tables, one that returns data very quickly and one that returns data very slowly. Users see the data for the fast table as soon as the data is available.</p> <p>NOTE: When using lazy delivery, consider using <code>whenAvailable</code> to first determine if the data is available, then render if it is not available.</p> <p>Executing data fetches in parallel also speeds up the total time to fetch data. This gives an advantage to lazy loading in cases of multiple, and possibly slow, data fetches. While streaming is the default mechanisms to deliver data in lazy mode, parallel execution of data controls is not. In order to enable parallel execution, open the page definition and change <code>RenderHint</code> on the iterator to <i>background</i>.</p> <p>In certain situations, the advantage of parallel execution is faster response time. Parallel execution could potentially use more resources due to multiple threads executing request in parallel and possibly more database connections will be opened.</p> <p>Consider using parallel execution only when there are multiple slow components on the page and the stamped components belong to different data control frames (such as isolated taskflows). Since parallel execution synchronizes on the data control frame level, when there is a single data control frame parallel execution may not improve performance.</p>
Immediate Delivery	<p>Immediate delivery (<code>contentDelivery="immediate"</code>) should be used if table data control is fast, or if it returns a small set of data. In these cases the response time be faster than using lazy delivery.</p> <p>Another advantage of immediate delivery is less server resource usage, compared to lazy delivery. Immediate delivery sends only one request to the server, which results in lower CPU and memory usage on the server for the given user interaction.</p> <p>There are some important issues associated with the <code>immediate</code> attribute. Refer to "Using the Immediate Attribute" in <i>Oracle Fusion Middleware Web User Interface Developer's Guide for Oracle Application Development Framework</i> for more information.</p>

8.2.4.2 Performance Considerations for autoSuggest

`autoSuggest` is a feature that can be enabled for `inputText`, `inputListOfValues`, and `inputComboboxListOfValues` components. When the user types characters in the input field, the component displays a list of suggested items. The feature performs a query in the database table to filter the results. In order to speed up database processing, a database index should be created on the column for which `autosuggest` is enabled.

This improves the component's response times especially when the database table has a large number of rows.

Note: In order to prevent large number of `autoSuggest` items being returned, consider using the `maxSuggestedItems` property on `autoSuggestBehavior` with a value appropriate for your usage.

8.2.5 Tuning Data Visualization Tool (DVT) Components

Data Visualization Tool (DVT) components are built on top of ADF Rich Client components. DVT components include graphs, gauges, Gantt charts, pivot tables and maps. [Table 8-5](#) provides some configuration recommendations for DVT components:

Table 8-5 DVT Component Configurations

Configuration Recommendation	Description
Modify the <code>RangeSize</code> attribute.	<p>The <code>RangeSize</code> attribute defines the number of rows to return simultaneously. A <code>RangeSize</code> value of -1 causes the iterator to return all the rows. Using a lower value may improve performance, but it may be harder to manage the data. Also, any data beyond <code>rangeSize</code> will not be available in the view.</p>
Use horizontal text instead of vertical text.	<p>By default, pivot tables use horizontal text for column headers. However, there is an option to use vertical text as well. Vertical text can be used by specifying a CSS style for the header format such as:</p> <pre>writing-mode:tb-rl;filter:flipV flipH;</pre> <p>While vertical text can look better in some cases, it has a performance impact when the Firefox browser is used.</p> <p>The problem is that vertical text is not native to Firefox as it is in Internet Explorer. To show vertical text, the pivot table uses images produced by <code>GaugeServlet</code>. These images cannot be cached as the text is dynamic and depends on the binding value. Due to this, every rendering of the pivot table incurs extra round-trips to the server to fetch the images, which impact network traffic, server memory, and CPU.</p> <p>To have the best performance, consider using horizontal text instead of vertical text.</p>
Avoid using dynamic resizing.	<p>The <code>dynamicResize="DYNAMIC_SIZE"</code> property on DVT graph and gauge provides the ability to size a component based on available area on the browser. Since the available area is not always known before geometry management calculates it, a second request is needed to fetch the correct graph or gauge size.</p> <p>The additional fetch can impact performance. If possible, set the <code>dynamic</code> resize property on the component and avoid using <code>dynamic</code> <code>resize</code> and explicitly set the correct area size.</p>
Use HTML5, instead of Flash, as the default image format.	<p>As of Oracle Fusion Middleware 11g Release 1 (11.1.1.7.0), graph and gauge supports HTML5 output format. A new <code>web.xml</code> context parameter has been introduced to change the default output format to HTML5. The context parameter, <code>oracle.adf.view.rich.dvt.DEFAULT_IMAGE_FORMAT</code>, will be automatically added to new applications.</p> <p>Applications that have this context parameter, but have not yet specified an image format, will automatically default to HTML5. Valid values are HTML5 and FLASH. For the applications that do not have this context parameter, the default rendering for Graph and Gauge would be Flash.</p> <p>To maintain optimal performance, use HTML5 for all modern browsers.</p>

8.3 Advanced Tuning Considerations

After you have performed the tuning modifications recommended in the previous section, you can make additional changes that are specific to your ADF Server deployment. Consider carefully whether the recommendations in this section are appropriate for your environment.

- [Tuning ADF Server Performance](#)
- [Tuning Groovy Usage](#)

8.3.1 Tuning ADF Server Performance

Oracle ADF Server components consist of the non-UI components within ADF. These include the ADF implementations of the model layer (ADFm), business services layer (ADFbc), and controller layer (ADFc). As the server components are highly configurable, it is important to choose the combination of configurations that best suits the available resources with the specified application performance and functionality.

8.3.1.1 HTTP Session Timeout Tuning

For ADF applications with a significant user community, the amount of memory held by sessions waiting to expire can negatively impact performance when the default HTTP session timeout of 45 minutes is used. The memory being held can be higher than what is physically available, causing the server to not be able to handle the load. For large numbers of users, such as those using a public facing website, the session timeout should be as short as possible.

To improve performance, consider modifying the default session timeout value (in minutes) in the `web.xml` file. Use a session timeout value that works with your use case scenario. The example below shows a session timeout of 10 minutes:

```
<session-config>
  <session-timeout>
    10
  </session-timeout>
</session-config>
```

8.3.1.2 View Objects Tuning


View objects (VOs) provide many tuning options to enable a developer to tailor the View Object to the application's specific needs. View Objects should be configured to use the minimal feature set required to fulfill the functional requirement. The *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework* provides detailed information on tuning View Objects. The following sections provide some tips pertaining to View Object performance.

8.3.1.2.1 Creating View Objects To maximize View Object performance, the View Object should match the intended usage. For instance, data retrieved for a list of values pick-list is typically read-only, so a read-only View Object should be used to query this data. Tailoring the View Object to the specific needs of the application can improve performance, memory usage, CPU usage, and network usage.

View Object Type	Description
Read-only View Objects	<p>Consider using a read-only View Object if the View Object does not have to insert or update data. There are two options for read-only View Objects:</p> <ul style="list-style-type: none"> ■ Non-updatable EO-based View Objects ■ Expert-mode View Objects <p>Non-updatable EO-based View Objects offer the advantage of a customizable select list at runtime which retrieve attributes needed in the UI, data reads from local cache (instead of re-executing a database query), and data consistency with other updatable View Objects based on the same EO.</p> <p>Expert-mode View Objects have the ability to perform SQL operations not supported by EOs and avoid the small performance impact from coordinating View Object and EO rows. EO-based View Objects can be marked non-updatable by deselecting the "updatable" option in the selected EO for the View Object, which can also be done by adding the parameter <code>ReadOnly="true"</code> on the <code>EntityUsage</code> attribute in the View Object XML definition.</p>
Insert-only View Objects	<p>For View Objects that are used only for inserting records, you can prevent unnecessary select queries from being executed when using the View Object. To do this, set the option <code>No Rows</code> in the <code>Retrieve from the Database</code> group box in the View Objects Overview tab. This sets <code>MaxFetchSize</code> to 0 (zero) for the View Object definition.</p>
Runtime-created View Objects	<p>View Objects can be created at runtime using the <code>createViewObjectFromQueryStmt()</code> API on the AM. However, avoid using runtime-created View Objects unless absolutely necessary due to potential performance impacts and complexity of tuning.</p>

8.3.1.2.2 Configuring View Object Data Fetching View Object performance is largely dependent on how the view object is configured to fetch data. If the fetch options are not tuned correctly for the application, then the view object may fetch an excessive amount of data or may take too many round-trips to the database. Fetch options can be configured through the **Retrieve from the Database** group box in the View Object dialog [Figure 8-1](#).

Figure 8–1 View Object Dialog



Fetch Option	Description
Fetch Mode	<p>The default fetch option is the All Rows option, which is retrieved as needed (FetchMode="FETCH_AS_NEEDED") or all at once (FetchMode="FETCH_ALL"), depending on which option is appropriate. The FETCH_AS_NEEDED option ensures that an executeQuery() operation on the view object initially retrieves only as many rows as necessary to fill the first page of a display. The number of rows is set based on the view object's range size.</p> <p>CAUTION: Avoid using the FetchMode="FETCH_ALL" unless required. This option can impact performance.</p>
Fetch Size	<p>In conjunction with the fetch mode option, the Batches field controls the number of records fetched simultaneously from the database (FetchSize in the View Object, XML). The default value is 1, which may impact performance unless only 1 row is fetched. The suggested configuration is to set this value to $n+1$ where n is the number of rows to be displayed in the user interface.</p> <p>Note that for DVT objects, Fetch Size should be $n+1$ where n is either rangeSize or the likely maximum rowset size if rangeSize is -1.</p>

Fetch Option	Description
Max Fetch Size	<p>The default max fetch size for a View Object is -1, which means that there is no limit to the number of rows the View Object can fetch. Setting a max fetch size of 0 (zero) makes the View Object insert-only. In cases where the result set should only contain <i>n</i> rows of data, the option Only Up to Row Number should be selected and set or call <code>setMaxFetchSize(N)</code> to set this programmatically. To set this manually, add the parameter <code>MaxFetchSize</code> to the View Object XML.</p> <p>For View Objects whose WHERE clause expects to retrieve a single row, set the option At Most One Row. This option ensures that the view object knows not to expect any more rows and skips its normal test for that situation. In this case no select query is issued and no rows are fetched.</p> <p>Max fetch size can also be used to limit the impact from a non-selective query that may return hundreds (or thousands) of rows. In such cases, specifying the max fetch size limits the number of rows that can be fetched and stored into memory.</p>
Forward-Only Mode	<p>If a data set is only traversed going forward, then forward-only mode can help performance when iterating through the data set. This can be configured by programmatically calling <code>setForwardOnly(true)</code> on the View Object. Setting forward-only can also prevent caching previous sets of rows as the data set is traversed.</p>

8.3.1.2.3 Additional View Object Configurations Table 8–6 provides additional tuning considerations when using the View Object:

Table 8–6 Additional View Object Configurations

Configuration Recommendation	Description
Optimize large data sets.	View Objects provide a mechanism to page through large data sets so that a user can jump to a specific page in the results. This is configured by calling <code>setRangeSize(N)</code> followed by <code>setAccessMode(RowSet.RANGE_PAGING)</code> on the View Object where <i>N</i> is the number of rows contained within 1 page. When navigating to a specific page in the data set, the application can call <code>scrollToRangePage(P)</code> on the View Object to navigate to page <i>P</i> . Range paging fetches and caches only the current page of rows in the View Object row cache at the cost of another query execution to retrieve each page of data. Range paging is not appropriate where it is beneficial to have all fetched rows in the View Object row cache (for example, when the application must read all rows in a data set for an LOV or page back and forth in records of a small data set).
Get an estimated row count using <code>getEstimatedRowCount</code> .	Use <code>getEstimatedRowCount</code> instead of <code>getRowCount</code> to prevent any impacts to performance. <code>getRowCount</code> can impact performance because it fetches all rows to memory and then count them.
Disable "spillover" configurations when possible.	You can use the data source as "virtual memory" when the JVM container runs out of memory. By default this is disabled and can be enabled (if needed) by setting <code>jbo.use.pers.coll=true</code> . Keep this option disabled (if possible) to avoid a potential performance impact.
Review SQL style configuration.	If the generic SQL92 SQL style is used to connect to generic SQL92-compliant database, then some View Object tuning options do not apply. The View Object fetch size is one such tuning option. When SQL92 SQL style is used, the fetch size defaults to 10 rows, regardless of what is configured for the View Object. The SQL style is set when defining the database connection. By default when defining an Oracle database connection, the SQL style can be <code>Oracle</code> . To manually override the SQL style, pass the parameter <code>-Djbo.SQLBuilder="SQL92"</code> to the JVM at startup.

Table 8–6 (Cont.) Additional View Object Configurations

Configuration Recommendation	Description
Use bind variables for view object queries.	<p>If the query associated with the View Object contains values that may change from execution to execution, consider using bind variables. This may help to avoid re-parsing the query on the database. Bind variables can be added to the View Object in the Query section of the View Object definition.</p> <p>Remove unused list bindings to improve performance.</p>
Use query optimizer hints for view object queries.	<p>The View Object can pass hints to the database to influence which execution plan to use for the associated query. The optimizer hints can be specified in the Retrieve from the Database group box.</p> <p>CAUTION: To prevent performance issues, do not use <code>ENDSWITH/CONTAINS</code> on required/selectively required view criteria items.</p>
Use dynamic SQL generation.	<p>View Objects can be configured to dynamically generate SQL statements at runtime instead of defining the SQL at design time. A View Object instance, configured with generating SQL statements dynamically, can avoid re-querying a database. This is especially true during page navigation if a subset of all attributes with the same key Entity Object list is used in the subsequent page navigation. Performance can be improved by activating a superset of all the required attributes to eliminate a subsequent query execution.</p>
Passivate transient expressions.	<p>Consider passivating <code>TransientExpressions</code> to maintain performance. This is controlled by a flag on EO/VO level or individual expression level. From the View Object dialog screen, select Including All Transient Values when enabling Passivate.</p>

8.3.1.3 Batch Processing

Batch processing enables multiple inserts, updates, and deletes to be processed together when sending the operations to the database. Enabling this feature is done on the Entity Object (EO) by either selecting the "Use Update Batching" check box in the Tuning section of the EO's General tab, or by directly modifying the EO's XML file and adding the parameter `BatchThreshold` with the specified batch size to the `Entity` attribute.

The `BatchThreshold` value is the threshold at which a group of operations can be batched instead of performing each operation one at a time. If the threshold is not exceeded, then rows may be affected one at a time. On the other hand, more rows than specified by the threshold can be batched into a single batch.

Note that the `BatchThreshold` configuration for the EO is not compatible if an attribute in the EO exists with the configuration to refresh after insert (`RetrievedOnInsert="true"`) or update (`RetrievedOnUpdate="true"`).

8.3.1.4 RangeSize Tuning

This parameter controls the number of records ADFm requests from the BC layer simultaneously. The default `RangeSize` is 25 records. Consider setting this value to the number of records to be displayed in the UI simultaneously for the View Object so that the number of round-trips between the model and BC layers is reduced to one. This is configured in the `Iterator` attribute of the corresponding page's page definition XML.

8.3.1.5 Application Module Design Considerations

Designing an application's module granularity is an important consideration that can significantly impact performance and scalability. It is important to note that each root application module generally holds its own database connection. If a user session consumes multiple root application modules, then that user session can potentially

hold multiple database connections simultaneously. This can occur even if the connections are not actively being used, due to the general affinity maintained between an application module and a user session. To reduce the possibility that a user can hold multiple connections at once, consider the following options:

- Design larger application modules to encompass all of the functionality that a user needs.
- Nest smaller application modules under a single root application module so that the same database connection can be shared among the nested application modules.
- Use lazy loading for application modules. In the Application Module tuning section, customize runtime instantiation behavior to use lazy loading. Lazy loading can also be set JVM-wide by adding the following JVM argument:

```
-Djbo.load.components.lazily=true
```

8.3.1.6 Application Module Pooling

Application module (AM) pooling enables multiple users to share several application module instances. The configurations for the AM pool vary depending on the expected usage of the application.

Most of the AM pool parameters can be set through Oracle JDeveloper. The configurations are saved in `bc4j.xcfg`, which can be manually edited if needed. Parameters can also be set at the system level by specifying these as JVM parameters (`-Dproperty=value`). The `bc4j.xcfg` configuration takes precedence over the JVM configuration; this enables a generic system-level configuration to be overridden by an application-specific exception.

Table 8–7 Application Module (AM) Pool Tuning

Configuration Recommendation	Description
Use the lazy data delivery method.	<p>By default AM loads all the associated definitions when the instance of AM is created. AM can have a large number of VOs associated with it. This can make loading of AM slower and passivation slower.</p> <p>To prevent possible performance issues, set <code>LoadComponentsLazily="true"</code>.</p>
Optimize the number of AM pools in the application.	<p>Parameters applied at the system level are applied per AM pool. If the application uses more than 1 AM pool, then system-level values for the number of AM instances must be multiplied by the number of AM pools to realize the actual limits specified on the system as a whole.</p> <p>For example, if an application uses 4 separate AM pools to service the application, and a system-level configuration is used to limit the max AM pool size to 100, then this can result in a maximum of 400 AM instances (4 pools * 100 max pool size).</p> <p>If the intent is to limit the entire application to a max pool size of 100, then the system-level configuration should specify a max pool size of 25 (100 max pool size / 4 pools). Finer granularity for configuring each AM pool can be achieved by configuring each pool separately through JDev or directly in <code>bc4j.xcfg</code>.</p>
Optimize the number of database connections.	<p>By default AM instances retain their database connections even when checked back into the AM pool. There are many performance benefits to maintain this association. To maintain performance, consider configuring more AM instances than the maximum number of specified database connections.</p> <p>NOTE: If you have an AM pool that needs to be used as root pool, consider tuning at the specific AM pool level. For pools that are infrequently used, consider tuning pool sizes on the pool level so that top-level application parameters are not used.</p>

8.3.1.6.1 General AM Pool Configurations The following guidelines can be used as a general starting point when tuning AM and AM pool behavior. More specific tuning for memory or CPU usage can be found in .

Table 8–8 AM Pool Tuning Parameters

Parameter	Description
Initial Pool Size <code>jbo.ampool.initpoolsize</code>	<p>Specifies the number of application module instances to create when the pool is initialized (default is zero). Setting a nonzero initial pool size increases the time to initialize the application, but improves subsequent performance for operations requiring an AM instance.</p> <p>Configure this value to 10% more than the anticipated number of concurrent AM instances required to service all users.</p>
Maximum Pool Size <code>jbo.ampool.maxpoolsize</code>	<p>Specifies the maximum number of application module instances that the pool can allocate (default is 4096). The pool can never create more application module instances than this limit imposes. A general guideline is to configure this to 20% more than the initial pool size to allow for some additional growth.</p>
Minimum Available Size <code>jbo.ampool.minavailablesize</code>	<p>The minimum number of available application module instances that the pool monitor should leave in the pool during a resource cleanup operation, when the server is under light load.</p> <p>Set to 0 (zero) if you want the pool to shrink to contain no instances when all instances have been idle for longer than the idle time-out after a resource cleanup.</p> <p>The default is 5 instances.</p> <p>While application module pool tuning allows different values for the <code>jbo.ampool.minavailablesize</code> <code>jbo.ampool.maxavailablesize</code> parameters, in most cases it is fine to set these minimum and maximum tuning properties to the same value.</p>
Maximum Available Size <code>jbo.ampool.maxavailablesize</code>	<p>The ideal maximum number of available application module instances in the pool when the server is under load.</p> <p>When the pool monitor wakes up to do resource cleanup, it will try to remove available application module instances to bring the total number of available instances down to this ideal maximum. Instances that have been not been used for a period longer than the idle instance time-out will always get cleaned up at this time, then additional available instances will be removed if necessary to bring the number of available instances down to this size.</p> <p>The default maximum available size is 25 instances. Configure this to leave the maximum number of available instances desired after a resource cleanup. A lower value generally results in more application module instances being removed from the pool on a cleanup.</p> <p>While application module pool tuning allows different values for the <code>jbo.ampool.maxavailablesize</code> <code>jbo.ampool.minavailablesize</code> parameters, in most cases it is fine to set these minimum and maximum tuning properties to the same value.</p>
Referenced Pool Size <code>jbo.recyclethreshold</code>	<p>Specifies the maximum number of application module instances in the pool that attempt to preserve session affinity for the next request made by the session that used them last before releasing them to the pool in managed-state mode (default is 10).</p> <p>The referenced pool size should always be less than or equal to the maximum pool size. This enables the configured number of available instances to try and remain "loyal" to the affinity they have with the most recent session that released them in managed state mode.</p> <p>Configure this value to the expected number of concurrent users that perform multiple operations with short think times. If there are no users expected to use the application with short think times, then this can be configured to 0 (zero) to eliminate affinity.</p>

Table 8–8 (Cont.) AM Pool Tuning Parameters

Parameter	Description
Maximum Instance Time to Live <code>jbo.ampool.timetolive</code>	<p>The number of milliseconds after which to consider an connection instance in the pool as a candidate for removal during the next resource cleanup regardless of whether it would bring the number of instances in the pool below <code>minavailablesize</code>.</p> <p>The default is 3600000 milliseconds of total time to live (which is 3600 seconds, or one hour). A lower value reduces the time an application module instance can exist before it must be removed at the next resource cleanup. The default value is sufficient for most applications. A higher value increases the time an application module instance can exist before it must be removed at the next cleanup.</p>
Idle Instance Timeout <code>jbo.ampool.maxinactiveage</code>	<p>The number of milliseconds after which to consider an inactive application module instance in the pool as a candidate for removal during the next resource cleanup.</p> <p>The default is 600000 milliseconds of idle time (which is 600 seconds, or ten minutes). A lower value results in more application module instances being marked as a candidate for removal at the next resource cleanup. A higher value results in fewer application module instances being marked as a candidate for removal at the next resource cleanup.</p>
Pool Polling Interval <code>jbo.ampool.monitorsleepinterval</code>	<p>The length of time in milliseconds between pool resource cleanup.</p> <p>While the number of application module instances in the pool will never exceed the maximum pool size, available instances which are candidates for getting removed from the pool do not get "cleaned up" until the next time the application module pool monitor wakes up to do its job.</p> <p>The default is to have the application module pool monitor wake up every 600000 milliseconds (which is 600 seconds, or ten minutes). Configuring a lower interval results in inactive application module instances being removed more frequently to save memory. Configuring a higher interval results in less frequent resource cleanups.</p>
Failover <code>jbo.dofailover</code>	<p>Specifies whether to disable or enable failover. By default, failover is disabled. To enable failover, set the parameter to <code>true</code>.</p> <p>NOTE: When enabling application module state passivation, a failure can occur when Oracle WebLogic Server is configured to forcibly release connection back into the pool. A failure of this type produces a <code>SQLException</code> (Connection has already been closed) that is saved to the server log. The exception is not reported through the user interface.</p> <p>To ensure that state passivation occurs and changes are saved, set an appropriate value for the <code>weblogic-application.xml</code> deployment descriptor parameter <code>inactive-connection-timeout-seconds</code> on the <code><connection-check-params> pool-params</code> element.</p> <p>Setting the deployment descriptor parameter to several minutes, in most cases, should avoid forcing the inactive connection timeout and the resulting passivation failure. Adjust the setting as needed for your environment.</p>

Table 8–8 (Cont.) AM Pool Tuning Parameters

Parameter	Description
Locking Mode <code>jbo.locking.mode</code>	Specifies the locking mode (<i>optimistic</i> or <i>pessimistic</i>). The default is <i>pessimistic</i> , which means that a pending transaction state can be created on the database with row-level locks. With <i>pessimistic</i> locking mode, each time an AM is recycled, a rollback is issued in the JDBC connection. Web applications should set the locking mode to <i>optimistic</i> to avoid creating the row-level locks.
Database Connection Pooling <code>jbo.doconnectionpooling</code>	Specifies whether the AM instance can be disconnected from the database connection when the AM instance is returned to the AM pool. This enables an application to size the AM pool larger than the database connection pool. The default is <i>false</i> , which means that an AM instance can retain its database connection when the AM instance is returned to the AM pool. When set to <i>true</i> , the AM can release the database connection back to the database connection pool when the AM instance is returned to the AM pool. Note that before an AM is disconnected from the database connection, a rollback can be issued on that database connection to revert any pending database state.
Transaction Disconnect Level <code>jbo.txn.disconnect_level</code>	When used in conjunction with <code>jbo.doconnectionpooling=true</code> , specifies BC4J behavior for maintaining JDBC ResultSets. By default <code>jbo.txn.disconnect_level</code> is 0, and passivation can be used to close any open ResultSets when the database connection is disconnected from the AM instance. Configuring <code>jbo.txn.disconnect_level</code> to 1 can prevent this behavior to avoid the passivation costs for this situation.

For parameters that can be configured for memory-constrained systems, see [Table 8–9](#).

Table 8–9 AM Pool Sizing Configurations - Memory Considerations

Parameter	Description
Initial Pool Size <code>jbo.ampool.initpoolsize</code>	Set this to a low value to conserve memory at the cost of slower performance when additional AM instances are required. The default value of 0 (zero) does not create any AM instances when the AM pool is initialized.
Maximum Pool Size <code>jbo.ampool.maxpoolsize</code>	Configure this to prevent the number of AM instance from exceeding the determined value. However, if this is set too low, then some users may see an error accessing the application if no AM instances are available.
Minimum Available Pool Size <code>jbo.ampool.minavailablesize</code>	Set to 0 (zero) to shrink the pool to contain no instances when all instances have been idle for longer than the idle time out after a resource cleanup. However, a setting of 1 is commonly used to avoid the costs of re-creating the AM pool.
Maximum Available Pool Size <code>jbo.ampool.maxavailablesize</code>	Configure this to leave the maximum number of available instances specified after a resource cleanup.

For parameters that can be configured to reduce the load on the CPU to some extent through a few parameters, see [Table 8–10](#).

Table 8–10 AM Pool Sizing Configurations - CPU Considerations

Parameter	Description
Initial Pool Size <code>jbo.ampool.initpoolsize</code>	Set this value to the number of AM instances you want the application pool to start with. Creating AM instances during initialization takes the CPU processing costs of creating AM instances during the initialization instead of on-demand when additional AM instances are required.
Session Recycle Threshold <code>jbo.recyclethreshold</code>	Configure this value to maintain the AM instance's affinity to a user's session. Maintaining this affinity as much as possible save the CPU processing cost of needing to switch an AM instance from one user session to another.

8.3.1.6.2 AM Pool Resource Cleanup Configurations These parameters affect the frequency and characteristics for AM pool resource cleanups.

For memory-constrained systems, configure the AM pool to clean up more AM instances more frequently so that the memory consumed by the AM instance can be freed for other purposes.

The AM pool can be configured to reduce the need for CPU processing by allowing more AM instances to exist in the pool for longer periods of time. This generally comes at the cost of consuming more memory.

CAUTION: Reducing the number of available AM instances and increasing the frequency of cleanups can result in higher CPU usage and longer response times.

Table 8–11 AM Pool Resource Cleanup Configurations - Memory Considerations

Parameter	Description
Minimum Available Size jbo.ampool.minavailablesize	A setting of 0 (zero) shrinks the pool to contain no instances when all instances have been idle for longer than the idle time out. However, a setting of 1 is commonly used to avoid the costs of re-creating the AM pool
Maximum Available Size jbo.ampool.maxavailablesize	A lower value generally results in more AM instances being removed from the pool on a cleanup. Setting these to a higher value leaves more idle instances in the pool, so that AM instances do not have to be recreated at a later time. However, the values should not be set excessively high to keep more AM instances than can be required at maximum load.
Time to Live jbo.ampool.timetolive	A lower value reduces the time an AM instance can exist before it must be removed at the next resource cleanup. A higher value increases the time an AM instance can exist before it must be removed at the next resource cleanup. CAUTION: A setting of -1 implies unlimited time to live. This setting should only be used for well tested applications with no memory leaks.
Maximum Inactive Age jbo.ampool.maxinactiveage	A low value results in more AM instances being marked as a candidate for removal at the next resource cleanup. A higher value results in fewer AM instances being marked as a candidate for removal at the next resource cleanup.
Monitor Sleep Interval jbo.ampool.monitorsleepinterval	This controls how frequent resource cleanups can be triggered. Configuring a lower interval results in inactive AM instances being removed more frequently to save memory. Configuring a higher interval results in less frequent resource cleanups.

8.3.1.7 ADFc: Region Usage

Table 8–12 ADFc Region Configurations

Configuration Recommendation	Description
Add regions to a page only when the specific functionality is required.	Adding regions to a page can be a powerful addition to the application. However, regions can be a resource-intensive component on the page. For better performance, consider using regions only when the specific functionality is required.
Disable <code>detect-metadata-changes</code> in <code>adf-config</code> .	Disable <code>detect-metadata-changes</code> by setting it to <code>false</code> in <code>adf-config</code> . Detecting metadata changes is only needed for applications that are not using ADFc DT@RT operations to customize unbounded task flows. If the application does customize unbounded flow, the property can be enabled and disabled as needed using RT APIs.
Consider nesting "hidden" popups inside of the 1 "visual" root component when possible.	Regions with multiple root components can have a negative impact on performance. Even if only 1 "visual" root component exists, and the rest are "hidden" popups, there may still be an impact to performance. For example, if you do not want stretching, consider wrapping everything in a <code>panelGroupLayout layout="vertical" or "scroll"</code> (as applicable). If you do want stretching, consider putting the "visual" root in the "center" facet and then put the popups in the "bottom" facet and set <code>bottomHeight="0px"</code> on the <code>panelStretchLayout</code> .

8.3.1.8 Defer Task Flow Execution

By default, task flows are activated when the page is loaded, even when the task flow is not initially rendered. This causes unnecessary overhead if the task flow is never displayed.

8.3.1.9 Configuring the Task Flow Inside Switcher

By default, task flows under switchers are activated when the page is loaded, not when the switcher facet is displayed. To avoid this, use conditional activation and set "active" to an expression language (EL) expression that returns 'true' when the facet is displayed. Consider using the same condition for switcher to drive task flow conditional activation.

Example of switcher:

```
<af:switcher id="s1" defaultFacet="1" facetName="#{pageFlowScope.facet}">
  <f:facet name="1">
    <af:region value="#{bindings.TF1.regionModel}" id="r1"/>
  </f:facet>
  <f:facet name="2">
    <af:region value="#{bindings.TF2.regionModel}" id="r2"/>
  </f:facet>
</af:switcher>
```

Associated binding:

```
<taskFlow id="tTF1" taskFlowId="-----" active="#{pageFlowScope.facet=='1'}"
activation="conditional" xmlns="http://example.host.com/adf/controller/binding"/>
<taskFlow id="tTF2" taskFlowId="-----" active="#{pageFlowScope.facet=='2'}"
activation="conditional" xmlns="http://example.host.com/adf/controller/binding"/>
```

8.3.1.10 Reusing Static Data

If the application contains static data that can be reused across the application, the cache data can be collected using a shared application module. More information on creating and using shared application modules can be found in "Sharing Application Module View Instances" in *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

8.3.1.11 Conditional Validations

For resource-intensive validations on entity attributes, consider using preconditions to selectively apply the validations only when needed. The cost of validation must be weighted against the cost of the precondition to determine if the precondition is beneficial to the performance. More information on specifying preconditions for validation can be found in "How to Set Preconditions for Validation" in *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

8.3.2 Tuning Groovy Usage

If you have applications that are using Groovy to evaluate expressions, then you may need to tune the expressions to improve performance. Groovy script parsing can impact performance, so its best to use it only when required.

For example, if you are using groovy to evaluate a common expression like this:

```
Date currentDate = null;
.....
    ExprEval expression = new ExprEval("adf.currentDate", ExprEval.EXPR_STYLE_
GROOVY);
    currentDate = (Date)expression.evaluate(null);
```

You might improve performance by making the following changes:

- If date and time are important, you might improve performance by using the following:

```
currentDate = new Date(System.currentTimeMillis());
```

- If the current time is not required, consider using the following:

```
Calendar cal = Calendar.getInstance();
cal.set(Calendar.HOUR, 0);
cal.set(Calendar.HOUR_OF_DAY, 0);
cal.set(Calendar.MINUTE, 0);
cal.set(Calendar.SECOND, 0);
cal.set(Calendar.MILLISECOND, 0);
currentDate = new java.sql.Date(cal.getTimeInMillis());
```

Oracle TopLink (EclipseLink) JPA Performance Tuning

This chapter describes some of the available performance tuning features for EclipseLink, an open-source persistence framework used with Oracle TopLink. The chapter includes the following topics:

- [Section 9.1, "About Oracle TopLink and EclipseLink"](#)
- [Section 9.3, "Basic Tuning Considerations"](#)
- [Section 9.4, "Advanced Tuning Considerations"](#)

Note: For more information on performance tuning in these areas, see the following:

- EclipseLink Performance Tuning at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Advanced_JPA_Development/Performance
 - Performance Monitoring and Profiling at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Advanced_JPA_Development/Performance/Performance_Profiling
 - Introduction to Optimization at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Advanced_JPA_Development/Performance#Identifying_General_Performance_Optimization
-
-

9.1 About Oracle TopLink and EclipseLink

Oracle TopLink includes the open source EclipseLink as the Java Persistence API (JPA) implementation. Oracle TopLink extends EclipseLink with advanced integration into the Oracle Application Server.

The Java Persistence API (JPA) is a specification for persistence in Java EE and Java SE applications. In JPA, a persistent class is referred to as an entity. An entity is a plain old Java object (POJO) class that is mapped to the database and configured for usage through JPA using annotations, persistence XML, or both. This chapter focuses on tuning JPA in the context of EJB3.0 and a Java EE environment.

The information in this chapter assumes that you are familiar with the basic functionality of EclipseLink. Before you begin tuning, consider reviewing the introductory information found at the following:

- The EclipseLink JPA User's Guide at <http://wiki.eclipse.org/EclipseLink/UserGuide/JPA>
- "Considering JPA Entity Architecture" at <http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Introduction/Architecture>
- Introduction to EclipseLink Queries at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Querying
- Introduction to Cache at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Caching
- Introduction to Mapping and Configuration at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Mapping

For more information on Oracle TopLink, see the TopLink page on OTN <http://www.oracle.com/technology/products/ias/toplink/index.html>.

[Note that as of Oracle TopLink Release 11g, the older Toplink APIs have been deprecated. For more information, see the TopLink Release Notes at <http://www.oracle.com/technology/products/ias/toplink/doc/11110/relnotes/toplink-relnotes.html#CHDGAEDJ>]

Note: This chapter serves as a 'quick start' guide to performance tuning JPA in the context of a Java EE environment. While the chapter provides common performance tuning considerations and related documentation resources, it is not meant to be comprehensive list of areas to tune.

9.2 Monitoring TopLink Performance

This section lists a few features in EclipseLink that can help you analyze your JPA application performance:

- For monitoring performance, see "Performance Monitoring" in the [EclipseLink User's Guide](#). Note that this tool is intended to profile and monitor information in a multithreaded server environment.
- For profiling performance, see "Measuring EclipseLink Performance with the EclipseLink Profiler" in the [EclipseLink User's Guide](#). Note that this tool is intended for use with single-threaded finite use cases.
- For debugging performance issues and testing, you can view the SQL generated from EclipseLink. To view the SQL, increase the logging level to "FINE" by using the EclipseLink JPA extensions for logging.

For best performance, restore the logging levels to the default levels when you are done profiling or debugging.

9.3 Basic Tuning Considerations

The following tuning recommendations will apply to most deployments. Always consult your use case scenarios to determine which tuning parameters are appropriate for your environment.

- [Using Efficient SQL Statements and Queries](#)
- [Tuning Cache Configuration](#)
- [Integrating Oracle Toplink with Coherence](#)

9.3.1 Using Efficient SQL Statements and Queries

This section covers using efficient SQL statements and SQL querying. [Table 9–1](#) and [Table 9–2](#) show tuning parameters and performance recommendations related to SQL statements and querying.

Table 9–1 EJB/JPA Using Efficient SQL Statements and Querying

Tuning Parameter	Description	Performance Notes
Parameterized SQL Binding	<p>Using parameterized SQL and prepared statement caching, you can improve performance by reducing the number of times the database SQL engine parses and prepares SQL for a frequently called query. EclipseLink enables parameterized SQL by default. However, not all databases and JDBC drivers support these options. Note that the Oracle JDBC driver bundled with Oracle Application Server does support this option. The persistence property in persistence.xml "eclipselink.jdbc.bind-parameters" is used to configure this.</p> <p>See Also: "Caching" at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Caching and "Querying" at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Querying</p> <p>Default Value: PERSISTENCE_UNIT_DEFAULT (which is true by default)</p>	<p>Leave parameterized SQL binding enabled for selected databases and JDBC drivers that support these options.</p>
JDBC Statement Caching	<p>Statement caching is used to lower the performance impact of repeated cursor creation and repeated statement parsing and creation; this can improve performance for applications using a database.</p> <p>Note: For Java EE applications, use the data source's statement caching (and do not use EclipseLink Statement Caching for EJB3.0/JPA, for example: <code>eclipselink.jdbc.cache-statements="true"</code>).</p> <p>Set this option in an Oracle Weblogic data-source by setting Statement Cached Type and Statement Cached Size configuration options.</p> <p>See also "Increasing Performance with the Statement Cache" in <i>Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server</i>.</p> <p>Default Value: The Oracle Weblogic Server data source default statement cache size is 10 statements per connection.</p>	<p>You should always enable statement caching if your JDBC driver supports this option. The Oracle JDBC driver supports this option.</p>

Table 9–1 (Cont.) EJB/JPA Using Efficient SQL Statements and Querying

Tuning Parameter	Description	Performance Notes
Fetch Size	<p>The JDBC fetch size gives the JDBC driver a hint as to the number of rows that should be fetched from the database when more rows are needed.</p> <p>For large queries that return a large number of objects, you can configure the row fetch size used in the query to improve performance by reducing the number database hits required to satisfy the selection criteria.</p> <p>Most JDBC drivers use a default fetch size of 10. If you are reading 1000 objects, increasing the fetch size to 256 can significantly reduce the time required to fetch the query's results.</p> <p>Note: The default value means use the JDBC driver default value, which is typically 10 rows for the Oracle JDBC driver.</p> <p>To configure this, use query hint "eclipselink.jdbc.fetch-size".</p> <p>Default Value: 0</p>	<p>The optimal fetch size is not always obvious. Usually, a fetch size of one half or one quarter of the total expected result size is optimal. Note that if you are unsure of the result set size, incorrectly setting a fetch size too large or too small can decrease performance.</p>
Batch Writing	<p>Batch writing can improve database performance by sending groups of INSERT, UPDATE, and DELETE statements to the database in a single transaction, rather than individually.</p> <p>The persistence property in persistence.xml "eclipselink.jdbc.batch-writing"="JDBC" is used to configure this.</p> <p>Default Value: Off</p>	<p>Enable for the persistence unit.</p>
Change Tracking	<p>This is an optimization feature that lets you tune the way EclipseLink detects changes in an Entity.</p> <p>Default Value: AttributeLevel if using weaving (Java EE default), otherwise Deferred.</p>	<p>Leave at default AttributeLevel for best performance.</p>
Weaving	<p>Can disable through persistence.xml properties "eclipselink.weaving"</p> <p>Default Value: On</p>	<p>Leave on for best performance.</p>

Table 9–1 (Cont.) EJB/JPA Using Efficient SQL Statements and Querying

Tuning Parameter	Description	Performance Notes
Read Only	<p>Setting an EJB3.0 JPA Entity to read-only ensures that the entity cannot be modified and enables EclipseLink to optimize unit of work performance.</p> <p>Set through query hint "eclipselink.read-only".</p> <p>Can also be set at entity level using @ReadOnly class annotation.</p> <p>Default Value: False</p>	For optimal performance use read-only on any query where the resulting objects are not changed.
firstResult and maxRows	<p>These are JPA query properties that are used for paging large queries. Typically, these properties can be used when the entire result set of a query returning a large number of rows is not needed. For example, when a user scans the result set (a page at a time) looking for a particular result and then discards the rest of the data after the record is found.</p>	Use on queries that can have a large result set and only a subset of the objects is needed.
Sequence number pre-allocation	<p>Sequence number pre-allocation enables a batch of ids to be queried from the database simultaneously in order to avoid accessing the database for an id on every insert.</p> <p>Default Value: 50</p>	Always use sequence number pre-allocation for best performance for inserts. SEQUENCE or TABLE sequencing should be used for optimal performance, not IDENTITY which does not allow pre-allocation.

9.3.1.1 Entity Relationships Query Parameter Tuning

Table 9–2 shows the Entity relationship query parameters for performance tuning.

Table 9–2 EJB3.0 Entity Relationship Query Performance Options

Tuning Parameter	Description	Performance Notes
Batch Fetching	<p>The eclipselink.batch hint supplies EclipseLink with batching information so subsequent queries of related objects can be optimized in batches instead of being retrieved one-by-one or in one large joined read.</p> <p>Batch fetching has three types: JOIN, EXISTS and IN. The type is set through the query hint "eclipselink.batch.type"</p> <p>Note that batching is only allowed on queries that have a single object in their select clause. The query hint to configure this is "eclipselink.batch". Batch fetching can also be set using the @BatchFetch annotation.</p> <p>Default Value: Off</p>	<p>Use for queries of tables with columns mappings to table data you need. You should only use either batch fetching or joining if you know that you are going to access all of the data; if you do not intend to access the relationships, then just let indirection defer their loading.</p> <p>Batch fetching is more efficient than joining because it avoids reading duplicate data; therefore for best performance for queries where batch fetching is supported, consider using batch fetching instead of join reading.</p>
Join Fetching	<p>Join fetching is a query optimization feature that enables a single query for a class to return the data to build the instances of that class and its related objects.</p> <p>Use this feature to improve query performance by reducing database access. By default, relationships are not join-read: each relationship is fetched separately when accessed if you are using lazy-loading, or as a separate database query if you are not using lazy-loading.</p> <p>You can specify the use of join in JPQL (<code>JOIN FETCH</code>), or you can set it multi-level in a query hint, "eclipselink.join-fetch". It also can be set in the mapping annotation @JoinFetch.</p> <p>Joining is part of the JPA specification, whereas batch fetching is not. And, joining works on queries that not work with batch fetching. For example, joining works on queries with multiple objects in the select clause, queries with a single result, and for cursors and first/max results, whereas batch fetching does not.</p> <p>See Also: "Join Fetch" at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Querying/Query_Hints#Join_Fetch</p> <p>Default Value: Not Used</p>	<p>Use for queries of tables with columns mappings to table data you need. You should only use either batch fetching or joining if you know that you are going to access all of the data; if you do not intend to access the relationships, then just let indirection defer their loading. For the best performance of selects, where batch fetching is not supported, a join is recommended</p>

Table 9–2 (Cont.) EJB3.0 Entity Relationship Query Performance Options

Tuning Parameter	Description	Performance Notes
Lazy loading	<p>Without lazy loading on, when EclipseLink retrieves a persistent object, it retrieves all of the dependent objects to which it refers. When you configure lazy reading (also known as indirection, lazy loading, or just-in-time reading) for an attribute mapped with a relationship mapping, EclipseLink uses an indirection object as a place holder for the referenced object.</p> <p>EclipseLink defers reading the dependent object until you access that specific attribute. This can result in a significant performance improvement, especially if the application is interested only in the contents of the retrieved object, rather than the objects to which it is related.</p> <p>See Also: "Lazy Loading" at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Mapping/Basic_Mappings/Lazy_Basics</p> <p>Default Value: On for collection mapping (ToMany mappings, @OneToMany, @ManyToMany)</p> <p>Default Value: Off for reference (ToOne mappings, @OneToOne, @ManyToOne)</p> <p>(Note that setting lazy loading On for @OneToOne, @ManyToOne requires weaving, which is On by default for Java EE.)</p>	<p>Use lazy loading for all mappings. Using lazy loading and querying the referenced objects using batch fetching or Join is more efficient than Eager loading.</p> <p>You may also consider using optimized loading with LoadGroups which allows a query to force instantiation of relationships.</p>

9.3.2 Tuning Cache Configuration

This section describes tuning the default internal cache that is provided by EclipseLink. Oracle Toplink/EclipseLink can also be integrated with Oracle Coherence. For information on configuring and tuning an EclipseLink Entity Cache using Oracle Coherence, see [Section 9.3.3, "Integrating Oracle Toplink with Coherence"](#).

The default settings for EJB3.0/JPA used with the EclipseLink persistence manager and cache are no locking, no cache refresh, and cache-usage DoNotCheckCache. To ensure that your application uses the cache and does not read stale data from the cache (when you do not have exclusive access), you must configure these and other isolation related settings appropriately. [Table 9–3](#) shows the cache configuration options.

For more information on cache configuration, see "Caching" at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Caching.

Note: By default, EclipseLink assumes that your application has exclusive access to the data it is using (that is, there are no external, non-EclipseLink, applications modifying the data). If your application does not have exclusive access to the data, then you must change some of the defaults from [Table 9–3](#).

Table 9–3 EJB3.0 JPA Entities and Cache Configuration Options

Tuning Parameter	Description	Performance Notes
Object Cache	<p>EclipseLink sessions provide an object cache. EJB3.0 JPA applications that use the EclipseLink persistence manager create EclipseLink sessions that by default use this cache. This cache, known as the session cache, retains information about objects that are read from or written to the database, and is a key element for improving the performance of an EclipseLink application.</p> <p>Typically, a server session's object cache is shared by all client sessions acquired from it. Isolated sessions provide their own session cache isolated from the shared object cache.</p> <p>The annotation type <code>@Cacheable</code> specifies whether an entity should be cached. Caching is enabled when the value of the persistence.xml caching element is <code>ENABLE_SELECTIVE</code> or <code>DISABLE_SELECTIVE</code>. The value of the <code>Cacheable</code> annotation is inherited by subclasses; it can be overridden by specifying <code>Cacheable</code> on a subclass.</p> <p><code>Cacheable(false)</code> means that the entity and its state must not be cached by the provider.</p> <p>Default Value: Enabled (shared is True)</p>	<p>Generally it is recommended that you leave caching enabled. If you have an object that is always read from the database, as in a pessimistic locked object, then the cache for that entity should be disabled. Also, consider disabling the cache for infrequently accessed entities</p>
Query Result Set Cache	<p>In addition to the object cache in EclipseLink, EclipseLink also supports a query cache:</p> <ul style="list-style-type: none"> ■ The object cache indexes objects by their primary key, allowing primary key queries to obtain cache hits. By using the object cache, queries that access the data source can avoid the cost of building the objects and their relationships if the object is already present. ■ The query cache is distinct from the object cache. The query cache is indexed by the query and the query parameters - not the object's primary key. This enables any query executed with the same parameters to obtain a query cache hit and return the same result set. <p>The query hints for a query cache are:</p> <pre>"eclipselink.query-cache"</pre> <pre>"eclipselink.query-cache.size"</pre> <pre>"eclipselink.query-cache.invalidation"</pre> <p>See Also: "Caching" at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Caching and "EclipseLink JPA Query Hints" at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Querying/Query_Hints</p> <p>Default Value: Not Used</p>	<p>Use for frequently executed non-primary key queries with infrequently changing result sets. Use with a cache invalidation time out to refresh as needed.</p>

Table 9–3 (Cont.) EJB3.0 JPA Entities and Cache Configuration Options

Tuning Parameter	Description	Performance Notes
Cache Size	<p>Cache size can be configured through persistence properties: "eclipselink.cache.size.<entity>" "eclipselink.cache.size.default" "eclipselink.cache.type.default"</p> <p>See Also: "Configuring Persistence Units Using persistence.xml" at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Configuration/JPA/persistence.xml and "Class PersistenceUnitProperties" at http://www.eclipse.org/eclipselink/api/2.3/org/eclipse/persistence/config/PersistenceUnitProperties.html</p> <p>Default Value: Type SoftWeak, Size 100 (per Entity).</p>	<p>Set the cache size relative to how much memory you have available, how many instances of the class you have, the frequency the entities are accessed, and how much caching you want based on your tolerance for stale data.</p> <p>Consider creating larger cache sizes for entities that have many instances that are frequently accessed and stale data is not a big issue.</p> <p>Consider using smaller cache sizes or no cache for frequently updated entities that must always have fresh data, or infrequently accessed entities.</p>
Locking	<p>Oracle supports the locking policies shown in Table 9–4: no locking, optimistic, pessimistic, and read-only.</p> <p>Locking is set through JPA @Version annotation, eclipselink.read-only</p> <p>How to Use EclipseLink Locking at http://wiki.eclipse.org/EclipseLink/Examples/JPA/Locking</p> <p>Default Value: No Locking</p>	<p>For entities that can be updated concurrently, consider using the locking policy to prevent a user from writing over another users changes. To optimize performance for read-only entities, consider defining the entity as read-only or use a read-only query hint.</p>

Table 9–3 (Cont.) EJB3.0 JPA Entities and Cache Configuration Options

Tuning Parameter	Description	Performance Notes
Cache Usage	<p>By default, all query types search the database first and then synchronize with the cache. Unless refresh has been set on the query, the cached objects can be returned without being refreshed from the database. You can specify whether a given query runs against the in-memory cache, the database, or both.</p> <p>To get performance gains by avoiding the database lookup for objects already in the cache, you can configure that the search attempts to retrieve the required object from the cache first, and then search the data source only if the object is not in the cache. For a query that looks for a single object based on a primary key, this is done by setting the query hint "eclipselink.cache-usage" to <code>CheckCacheByExactPrimaryKey</code>.</p> <p>Default Value: <code>DoNotCheckCache</code></p>	<p>For faster performance on primary key queries, where the data is typically in the cache and does not require a lot of refreshing, it is recommended to check the cache first on these queries (using <code>CheckCacheByExactPrimaryKey</code>).</p> <p>This avoids the default behavior of retrieving the object from the database first and then for objects already in the cache, returning the cached values (not updated from the database access, unless refresh has been set on the query).</p>
Isolation	<p>There is not a single tuning parameter that sets a particular database transaction isolation level in a JPA application that uses EclipseLink.</p> <p>In a typical EJB3.0 JPA application, a variety of factors affect when database transaction isolation levels apply and to what extent a particular database transaction isolation can be achieved, including the following:</p> <ul style="list-style-type: none"> ■ Locking mode ■ Use of the Session Cache ■ External Applications ■ Database Login method <code>setTransactionIsolation</code> <p>See Also: "Shared and Isolated Cache" at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Caching/Shared_and_Isolated</p>	

Table 9–3 (Cont.) EJB3.0 JPA Entities and Cache Configuration Options

Tuning Parameter	Description	Performance Notes
Cache Refreshing	<p>By default, EclipseLink caches objects read from a data source. Subsequent queries for these objects access the cache and thus improve performance by reducing data source access and avoiding the cost of rebuilding object's and their relationships. Even if a query accesses the data source, if the objects corresponding to the records returned are in the cache, EclipseLink uses the cached objects. This default caching policy can lead to stale data in the application.</p> <p>Refreshing can be enabled at the entity level (<code>alwaysRefresh</code> or <code>refreshOnlyIfNewer</code> and <code>expiry</code>) and at the query level (with the <code>eclipselink.refresh</code> query hint). You can also force queries to go to the database with (<code>disableHits</code>). Using an appropriate locking policy is the only way to ensure that stale or conflicting data does not get committed to the database.</p> <p>For more information see: Section 9.3.2.1, "Cache Refreshing Scenarios"</p> <p>See Also: "Caching Overview" at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Caching/Caching_Overview</p> <p>Default Value: No Cache Refreshing</p>	<p>Try to avoid entity level cache refresh and instead, consider configuring the following:</p> <ul style="list-style-type: none"> ■ cache refresh on a query-by-query basis ■ cache expiration ■ isolated caching

9.3.2.1 Cache Refreshing Scenarios

There are a few scenarios to consider for data refreshing in the cache, all with performance implications:

- In the case where you never want cached data and always want fresh data, consider using an isolated cache (`Shared=False`). This is the case when certain data in the application changes so frequently that it is desirable to always refresh the data, instead of only refreshing the data when a conflict is detected.
- In the case when you want to avoid stale data, but getting stale data is not a major issue, then using a cache expiry policy would be the recommended solution. In this case you should also use optimistic locking, which automatically refresh stale objects when a locking error occurs. If using optimistic locking, you could also enable the entity `@Cache` attributes `alwaysRefresh` and `refreshOnlyIfNewer` to allow queries that access the database to refresh any stale objects returned, and avoid refreshing invalid objects when unchanged. You may also want to enable refreshing on certain query operations when you know you want refreshed data, or even provide the option of refreshing something from the client that would call a refreshing query.
- In the case when you are not concerned about stale data, you should use optimistic locking; this automatically refresh stale objects in the cache on locking errors.

9.3.2.2 Locking Modes

The locking modes, as shown in [Table 9–4](#), along with EclipseLink cache-usage and query refreshing options, ensures data consistency for EJB entities using JPA. The different combinations have both functional and performance implications, but often

the functional requirements for up-to-date data and data consistency lead to the settings for these options, even when it may be at the expense of performance.

For more information, see "Locking" at

http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Mapping/Locking.

Table 9–4 Locking Mode Policies

Locking Option	Description	Performance Notes
No Locking	The application does not prevent users overwriting each other's changes. This is the default locking mode. Use this mode if the Entity is never updated concurrently or concurrent reads and updates to the same rows with read-committed semantics is sufficient. Default Value: No Locking	In general, no locking is faster, but may not meet your needs for data consistency.
Optimistic	All users have read access to the data. When a user attempts to make a change, the application checks to ensure the data has not changed since the user read the data. See Also: "Optimistic Locking" at http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Mapping/Locking/Optimistic_Locking	If infrequent concurrent updates to the same rows are expected, then optimistic locking may provide the best performance while providing data consistency guarantees.
Pessimistic	The first user who accesses the data with the purpose of updating it locks the data until completing the update.	If frequent concurrent updates to the same rows are expected, pessimistic locking may be faster than optimistic locking that is getting a lot of concurrent access exceptions and retries. When using pessimistic locking at the entity level, it is recommended that you use it with an isolated cache (Shared=False) for best performance.
Read Only	Setting an EJB3.0 JPA Entity to read-only ensures that the entity cannot be modified and enables EclipseLink to optimize unit of work performance. Set at the entity level using @ReadOnly class annotation. Can also be set at the query level through query hint "eclipselink.read-only".	Defining an entity as read-only can perform better than an entity that is not defined as read-only, yet does no inserts, updates, or deletes, since it enables EclipseLink to optimize the unit of work performance. Always use read-only for all read-only operations

9.3.3 Integrating Oracle Toplink with Coherence

Oracle Toplink can be integrated with Oracle Coherence. This integration is provided through the Oracle TopLink Grid feature. With TopLink Grid, there are several types of integration with EclipseLink JPA features.

For example:

- Replace the default EclipseLink L2 cache with Coherence. This provides support for very large L2 caches that span cluster nodes. EclipseLink's default L2 cache improves performance for multi-threaded and Java EE server hosted applications running in a single JVM, and requires configuring special cache coordination features if used across a cluster.

- Configure entities to execute queries in the Coherence data grid instead of the database. This allows clustered application deployments to scale beyond database-bound operations.

For more information on using EclipseLink JPA with a Coherence Cache, see "JPA on the Grid" Approach at

<http://www.oracle.com/technology/products/ias/toplink/doc/11110/grid/tlgug003.htm>

For more information on Oracle Toplink integration with Oracle Coherence, see "Oracle TopLink Integration with Coherence Grid Guide" at

<http://www.oracle.com/technology/products/ias/toplink/doc/11110/grid/toc.htm>

9.4 Advanced Tuning Considerations

The following tuning recommendations may or may not apply to your deployment. Consult your use case scenarios to determine if using these parameters is appropriate for your environment:

- [Configuring Mappings](#)
- [Configuring Data Partitioning](#)

9.4.1 Configuring Mappings

EclipseLink can transform data between an object representation and a representation specific to a data source. This transformation is called mapping and it is the core of a EclipseLink project.

A mapping corresponds to a single data member of a domain object. It associates the object data member with its data source representation and defines the means of performing the two-way conversion between object and data source.

For information on Mapping see, "Configuring Mappings" at

http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Basic_JPA_Development/Mapping.

9.4.2 Configuring Data Partitioning

EclipseLink allows you to configure data partitioning using the `@Partitioned` annotation. Partitioning enables an application to scale information across multiple databases; including clustered databases. For more information on using `@Partitioned` and other partitioning policy annotations, see "Data Partitioning" at

http://wiki.eclipse.org/EclipseLink/UserGuide/JPA/Advanced_JPA_Development/Data_Partitioning.

Oracle Web Cache Performance Tuning

This chapter provides guidelines for improving the performance of Oracle Web Cache.

- [Section 10.1, "About Oracle Web Cache"](#)
- [Section 10.2, "Performance Considerations"](#)
- [Section 10.3, "Basic Tuning Considerations"](#)
- [Section 10.4, "Advanced Tuning Considerations"](#)

10.1 About Oracle Web Cache

Oracle Web Cache is a content-aware server accelerator, or a reverse proxy, for the Web tier. Oracle Web Cache is the primary caching mechanism provided with Oracle Fusion Middleware. Caching improves the performance, scalability, and availability of Web sites that run on Oracle Fusion Middleware by storing frequently accessed URLs in memory. It can also improve the performance, scalability, and availability of Web sites that run on any Web server or application server, such as Oracle HTTP Server and Oracle WebLogic Server.

For more information, see the *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache*.

10.2 Performance Considerations

Effective Oracle Web Cache performance tuning starts with a good understanding of its usage and general performance issues. Before you begin tuning Oracle Adaptive Access Manager, review this section as well as the recommendations discussed in [Top Performance Areas](#):

- [Optimizing Hardware Resources](#)
- [Optimizing Platform Connections](#)

10.2.1 Optimizing Hardware Resources

- [Hardware Resources](#)
- [Memory Configuration](#)

10.2.1.1 Hardware Resources

Oracle Web Cache performs best with one very powerful CPU or two CPUs. Because Oracle Web Cache is an in-memory cache, it is rarely limited by CPU cycles. Additional CPUs do not increase performance significantly. However, the speed of the

processors is critical—use the fastest CPUs you can afford. Use more CPUs if Web Cache is sharing the system with other Oracle application server components or other applications.

Note that Oracle Web Cache is limited by the available addressable memory. Additional memory can increase performance and scalability. For information about the amount of memory needed, see [Section 10.2.1.2, "Memory Configuration"](#).

Oracle Web Cache has two processes: one for the administration server and one for the cache server.

- The administration server process is used for configuring and monitoring Oracle Web Cache. This process consumes very little CPU time. However, when viewing the statistics pages in Oracle Web Cache Manager, the administration server process must query the cache server process to obtain the relevant metrics. Accessing the statistics pages frequently, or setting a high refresh rate on a statistics page can affect cache server performance.
- The cache server process uses three threads: one to manage the front-end activities, a second to manage the back-end activities, and a third to process requests.

For a cost-effective way to use Oracle Web Cache, run it on a fast two-CPU dedicated computer with lots of memory. See the *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache* for information about various deployment scenarios.

For a Web site with more than one Oracle Web Cache instance, consider installing each instance on a separate two-CPU node, either as part of a cache cluster or as a standalone instance. When Oracle Web Cache instances are on separate nodes, you are less likely to encounter operating system limitations, particularly in network throughput. For example, two caches on two separate two-CPU nodes are less likely to encounter operating system limitations than two caches on one four-CPU node.

Of course, if other resources are competing with Oracle Web Cache for CPU usage, you should take the requirements of those resources into account when determining the number of CPUs needed. Although a separate node for Oracle Web Cache is optimal, you can also derive a significant performance benefit from Oracle Web Cache running on the same node as the rest of the application Web server.

10.2.1.2 Memory Configuration

To avoid swapping documents in and out of the cache, configure enough memory for the cache. Generally, the amount of memory (maximum cache size) for Oracle Web Cache should be set to at least 512 MB. Your application's memory requirements can vary based upon factors such as document size, number of documents, the number of HTTP headers returned, and whether ESI is present. To get a close approximation on the maximum amount of memory required, you may apply the formula provided below. If your application uses ESI then all templates and document fragments must be accounted for when figuring the TotalDocs and the AvgDocSize.

Estimated Cache size in bytes = $1.25 * (\text{TotalDocs} * ((\text{AvgDocSize}/8192+1) * 8192 + 16384))$

- 0.25 accounts for the run time memory usage. The Web Cache action limit is set to 5% below than the maximum Web Cache size by default. Web Cache also allocates 5% of the total cache size to optimize access misses that cannot be cached.
- TotalDocs refers to the total number of documents you intend to place in Web Cache.
- The AvgDocSize is self-explained.
- Remember to convert the estimated cache size is returned in bytes by the formula.

The memory formula presented above was verified against actual memory usage measurements and it showed very close results as can be seen in the table below:

Number of Cached	Doc Size In	Measured Cache	Formula Generated Results
Docs	Bytes	Size in MB	Size in MB
3300.00	102400	499.61	499.51
5525.00	51200	499.27	499.08
11050.00	51200	998.54	998.17
6600.00	102400	999.22	999.02
13200.00	102400	1998.44	1998.05
22100.00	51200	1997.07	1996.34
3300.00	102400	499.61	499.51

10.2.1.2.1 Configuring WebCache Memory The cache is empty when Oracle Web Cache starts. For monitoring to be valid, ensure that the cache is fully populated. That is, ensure that the cache has received enough requests so that a representative number of documents are cached.

The Oracle Web Cache Statistics page (Monitoring > Web Cache Statistics) provides information about the current memory use, the maximum memory use and the total documents currently resident in Oracle Web Cache. Note the following metrics in the Cache Overview table:

- Size of Documents in Cache shows the current logical size of the cache, which is the size of the valid documents in the cache. For example, if the cache contains two documents, one 3 KB and one 50 KB, the Size of Documents in Cache is 53 KB, the total of the two sizes.
- Configured Maximum Cache Size indicates the maximum cache size as specified in the Resource Limits page.
- Current Allocated Memory displays the physical size of the cache, which is the amount of data memory allocated by Oracle Web Cache for cache storage and operation. This number is always smaller than the process size shown by operating system statistics because the Oracle Web Cache process, like any user process, consumes memory in other ways, such as instruction storage, stack data, thread, and library data.
- Current Action Limit is 95% of the Configured Maximum Cache Size. This number is usually larger than the Current Allocated Memory.

If the Current Allocated Memory is greater than the Current Action Limit, Oracle Web Cache begins to use allocated but unused memory, and may begin garbage collection to free more memory. During garbage collection, Oracle Web Cache removes the less popular and less valid documents from the cache in favor of the more popular and more valid documents to obtain space for new HTTP responses without exceeding the maximum cache size.

If the Current Allocated Memory is close to or greater than the Current Action Limit, increase the maximum cache size to avoid swapping documents in and out of the cache. For more information, see "Specifying Properties for an Oracle Web Cache System Component" in *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache*.

10.2.2 Optimizing Platform Connections

- [UNIX Connections](#)
- [Windows Connections](#)

10.2.2.1 UNIX Connections

On most UNIX platforms, each client connection requires a separate file descriptor. The Oracle Web Cache server attempts to reserve the maximum number of file descriptors when it starts. If you have root privileges, you can increase this number. For example, for the LINUX Operating System you can increase the maximum number of file descriptors by modifying Oracle Web Cache users file descriptors limits in `/etc/security/limits.conf`.

For example to allow the user "WC_USER" to have 4092 connections, in the `/etc/security/limits.conf` file add the following entries:

```
WC_User soft nofile 4092
WC_User hard nofile 4092
```

Ensure that there are adequate file descriptors available to any process on the host by increasing the `fs.file-max` parameter in the `/etc/sysctl.conf` file.

On Solaris Operating System you can increase the maximum number of file descriptors by setting the `rlim_fd_max` parameter. If `webcached` is not run as `root`, the Oracle Web Cache server logs an error message and fails to start.

10.2.2.2 Windows Connections

On Windows, only available kernel resources limit the number of file handles as well as socket handles - the size of paged and non-paged pools. However, the number of TCP ports the system can open restricts the number of active TCP/IP connections.

For more information on establishing connections, see "Set Resource Limits and Network Thresholds" in *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache*.

10.3 Basic Tuning Considerations

This section provides the basic tuning considerations for Oracle Web Cache. It contains the following tuning recommendations:

- [Optimizing Network Connections](#)
- [Increasing Cache Hit Rates](#)
- [Optimizing Response Time](#)

10.3.1 Optimizing Network Connections

- [Network Bandwidth](#)
- [Network Connections](#)
- [Network-Related Parameters](#)

10.3.1.1 Network Bandwidth

When you use Oracle Web Cache, ensure that each system has sufficient network bandwidth to accommodate the throughput load. Otherwise, the network may be

saturated but Oracle Web Cache has additional capacity. For example, if an application generates 100 megabits of data or more per second, 10/100 Megabit Ethernet can be saturated.

If the network is saturated, consider using Gigabit Ethernet rather than 10/100 Megabit Ethernet. Gigabit Ethernet provides the most efficient deployment scenario to avoid network collisions, retransmissions, and bandwidth starvation. Additionally, consider using two separate network cards: one for incoming client requests and one for requests from the cache to the application Web server.

Use network-monitoring utilities that show network bandwidth usage. If the network is under utilized and throughput is less than expected, check whether the CPUs are saturated.

10.3.1.2 Network Connections

It is important to specify a reasonable number for the maximum connection limit for the Oracle Web Cache server. If you set a number that is too high, performance can be affected, resulting in slower response time. If you set a number that is too low, fewer requests can be satisfied. Strike a balance between response time and the number of requests processed concurrently.

To help determine a reasonable number, consider the following factors:

- The maximum number of clients that you intend to serve concurrently at any given time.
- The average size of a document and the average number of requests per document.
- Network bandwidth. The amount of data that can be transferred at any one time is limited by the network bandwidth.
- The percentage of cache misses. Cache misses are forwarded to the application Web server. Those requests consume additional network bandwidth, resulting in longer response times; especially if a large percentage of requests are cache misses.
- How quickly a document is processed. Use a network monitoring utility, such as `ttcp` or LoadRunner to determine how quickly your system processes a document.
- The cache cluster member capacity, if you have a cache cluster environment. The capacity reflects the number of incoming connections from other cache cluster members. Set the cluster member capacity using the Clustering page (Properties > Clustering) of Oracle Web Cache Manager.

WARNING: Do not set the values listed above to an arbitrarily high value. Oracle Web Cache sets aside some resources such as memory for each connection. Altering these values can adversely affect performance.

Use various tools, such as those available with the operating system and with Oracle Web Cache, to help determine the maximum number of connections. For example, the `netstat-a` command enables you to determine the number of established connections; the `ttcp` utility enables you to determine how fast a document is processed. The Oracle Web Cache Manager provides statistics on hits and misses.

For detailed instructions on how to set the maximum number of incoming connections, see "Specifying Properties for an Oracle Web Cache System Component" in *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache*.

10.3.1.3 Network-Related Parameters

Besides the number of network connections, other network-related parameters for Oracle Web Cache, the application Web server, and the operating system can affect response time. In most situations, the default settings are sufficient.

If response time is slow, you should tune Oracle Web Cache, the application Web server, and operating system parameters that affect connections, as explained in this section.

For Oracle Web Cache, check the values of the following settings:

- **Keep-Alive Timeout**

The amount of time a network connection is left open after Oracle Web Cache sends a response to a browser. Keep-Alive enables an HTTP client to send multiple requests to Oracle Web Cache using the same network connection. By default, the connection is left open for five seconds, which is typically enough time for the browser to send subsequent requests to Oracle Web Cache using the same connection.

If the network between the browser and Oracle Web Cache is slow, consider increasing the timeout, experiment with 10 seconds then 20 seconds and perhaps up to 30 seconds.

If you receive the following error, either increase the maximum incoming connections for Oracle Web Cache or lower the Keep-Alive Timeout:

```
11313: The cache server reached the maximum number of allowed incoming connections. Listening is temporarily suspended.
```

With a heavy load, such as during stress-testing, if clients continuously send one request and then disconnect, set the Keep-Alive Timeout to 0. With this value, Oracle Web Cache closes the connection as soon as the request is completed, to free up resources.

Set the Keep-Alive Timeout value in the Network Timeouts page (Properties > Network Timeouts).

- **Origin Server Timeout**

The amount of time for the application Web server to generate a response to Oracle Web Cache. If the application Web server or proxy server is unable to generate a response within that time, Oracle Web Cache sends a network apology page to the browser.

Usually, this value should be equal to the response time of the slowest document served by the application Web Server. If the value is too low, long-running requests can timeout before the response is complete. If the value is too high and the application Web server hangs for some reason, it can take longer for Oracle Web Cache to failover to another application Web server.

Set this value in the Network Timeouts page (Properties > Network Timeouts).

For the application Web server, check the values of the following settings in the application Web server's configuration file (`httpd.conf`). (These particular parameter names are specific to the Oracle HTTP Server.)

- **KeepAlive**

Whether to allow persistent connections. Persistent connections allow a client to send multiple sequential requests through the same connection.

Make sure `KeepAlive` is enabled. This can improve performance because the connection is set up only once and is kept open for subsequent requests from the same client.

- `KeepAliveTimeout`: The time a connection is left open to wait for the next request from the same client. If requests are primarily from Oracle Web Cache, you can set this value fairly high. A reasonable value is 30 seconds.
- `MaxKeepAliveRequests`: The maximum number of requests to allow during a persistent connection. Set to 0 to allow an unlimited number of requests.
- `MaxClients`: The maximum number of clients that can connect to the application Web server simultaneously.

If `KeepAlive` is enabled for the application Web server, you may require more concurrent `httpd` server processes, and you may have to set the `MaxClients` directive to a higher value.

If client requests have a short response time, you may be able to improve performance by setting `MaxClients` to a lower value. However, when the `MaxClients` value is reached, no additional processes can be created, causing other requests to fail. The `MaxClients` limit on the application Web server should be greater than or equal to the application Web server capacity as set through the Oracle Web Cache Manager.

For the operating system, check the TCP time-wait setting. This setting controls the amount of time that the operating system holds a port, not allowing new connections to use the same port.

On the Linux operating system, validate the value of `/proc/sys/net/ipv4/tcp_fin_timeout`. On the Solaris Operating System, check the `tcp_time_wait_interval` setting, using the following command:

```
ndd -get /dev/tcp tcp_time_wait_interval.
```

On Windows, check the value of `TcpTimeWaitDelay` in the following key in the registry:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters
```

This setting is usually only an issue during stress testing, if you continuously open more TCP/IP connections from one client computer. In this situation, lower the TCP time-wait setting. In real world deployments, this is rarely an issue because it is unlikely that a single client can generate a huge number of connections.

10.3.2 Increasing Cache Hit Rates

A *cache hit* is a web browser request that can be satisfied from documents stored in the cache. A *cache miss* is a web browser request that cannot be satisfied from documents stored in the cache and must be forwarded to the application web server.

If the ratio of cache hits to cache misses is low, consider the following ways to raise the cache hit rate:

- Use cookies and URL parameters to increase cache hit rates.

Oracle Web Cache can cache different versions of a document with the same URL, based on request cookies or headers. To use this feature, applications may need to implement simple changes, such as creating a cookie or header that differentiates the documents.

Some applications contain insignificant URL parameters, which can lead to different URLs representing the same content. If the documents are cached under

their full URLs, the cache hit/miss ratio becomes very low. You can configure Oracle Web Cache to ignore the non-differentiating URL parameter values, so that a single document is cached for different URLs, greatly increasing cache hit rates.

Sometimes the content for a set of documents is nearly identical. For example, the documents may contain hyperlinks composed of the same URL parameters with different session-specific values, or they may include some personalized strings in the document text, such as welcome greetings or shopping cart totals. You can configure Oracle Web Cache to store a single copy of the document with placeholders for the embedded URL parameters or the personalized strings, and to dynamically substitute the correct values for the placeholders when serving the document to clients.

For more information on multiple version documents, sessions, ignoring URL parameter values, and simple personalization, see "Getting Started with Administering Oracle Web Cache" in *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache*.

- Use redirection to cache entry documents.

For some popular site entry documents, such as "/", that typically require session establishment, session establishment effectively makes the document non-cacheable to all new users without a session. To cache these documents while preserving session establishment, you can either:

- Create a blank document that provides session establishment for all initial requests and redirects to the actual popular document. Subsequent redirected requests to the popular document can specify the session, enabling the popular document to be served from the cache.
- Use a JavaScript that sets a session cookie for the popular documents.

Note: For more information on configuring caching rules for documents requiring session establishment, see "Caching and Compressing Content" in *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache*.

- Use partial page caching where possible.

Many Web documents, such as pages generated by OracleAS Portal, are composed of fragments with unique caching properties. For these pages, full-page caching is not feasible. However, Oracle Web Cache provides partial page caching using Edge Side Includes (ESI). With ESI, you can divide each Web page into a template and multiple fragments that can, in turn, be further divided into templates and lower level fragments. Each fragment or template is stored and managed independently; a full page is assembled from the underlying fragments upon request. Fragments can be shared among different templates, so that common fragments are not duplicated to waste cache space. Sharing can also greatly reduce the number of updates required when fragments expire.

Depending on the application, updating a fragment can be cheaper than updating a full page. In addition, each template or fragment can have its own unique caching policies such as expiration, validation, and invalidation, so that each fragment in a full Web page can be cached if possible, even when some fragments are not cached or are cached for a much shorter period of time.

- Use ESI variables for improved cache hit/miss ratio for personalized pages.

Personalized information often appears in Web pages, making them unique for each user. For example, many Web pages contain tens or hundreds of hyperlinks embedding application session IDs. To resolve this, create your ESI pages with variables. Because variables can resolve to different pieces of request information or response information, the uniqueness of templates and fragments can be significantly reduced. This, in turn, results in better cache hit/miss ratios.

10.3.3 Optimizing Response Time

If you have not configured the application Web server or the cache correctly, response time may be slower than anticipated. This section summarizes much of the information presented in this chapter.

If the application Web server is responding more slowly than expected or if the application Web server is not responding to requests from the cache because it has reached its capacity, check the application Web server and Oracle Web Cache settings.

First, check the following:

- **Caching rules:** Ensure that you are caching the appropriate objects. Are there popular objects that you should cache but are not caching? Use the Popular Requests page (Monitoring > Popular Requests) to see a list of the most popular requests and to check that those objects are being cached.
- **Priority rankings of the caching rules:** Give frequently accessed non-cacheable documents a higher priority than cacheable documents. Give frequently accessed cacheable documents the lowest priority. Note that parsing of caching rules may be resource-intensive if a large number of rules are defined.
- **Compression:** If the network is a bottleneck for the client, compressing documents as they are cached can relieve some of the congestion on the network because compressed documents are smaller.

Then, check the following:

The application Web server configuration, particularly the `MaxClients`, `KeepAlive`, `KeepAliveTimeout`, and `MaxKeepAliveRequests` settings.

The `MaxClients` limit on the application Web server should be greater than or equal to the application Web server capacity as set through the Oracle Web Cache Manager.

The application Web server capacity as set using the Origin Servers page (Origin Servers, Sites, and Load Balancing > Origin Servers) of the Oracle Web Cache Manager. See the *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache* for information about setting application Web server capacity.

Then, if the application Web server is still busier than anticipated, it may mean that the cache cannot process the requests and is routing more requests to the application Web server. Check the following Oracle Web Cache settings in the Oracle Web Cache Manager:

- **The number of cache connections.** Check Maximum Incoming Connections in the Resource Limits page (Properties > Resource Limits).
- **The memory size for the cache.** Check Maximum Cache Size in the Resource Limits page (Properties > Resource Limits).
- **The cache cluster capacity.** In a cache cluster, if cluster capacity is too low, a cache may not receive a response for owned content from a peer cache in the specified interval. As a result, the request is sent to the application Web server. Check Capacity in the Clustering page (Properties > Clustering). See the *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache* for more information.

If the settings for the application Web server and Oracle Web Cache are set correctly, but the response times are still higher than expected, check system resources, especially:

- Network bandwidth
- CPU usage

10.4 Advanced Tuning Considerations

The following Oracle Web Cache tuning considerations are provided as a guide. Always consult your own use case scenarios to determine if these configurations should be used in your deployment.

10.4.1 Optimizing Performance with Oracle ADF

Consider the following configuration options for optimizing Oracle Web Cache performance with Oracle ADF Rich Client Applications:

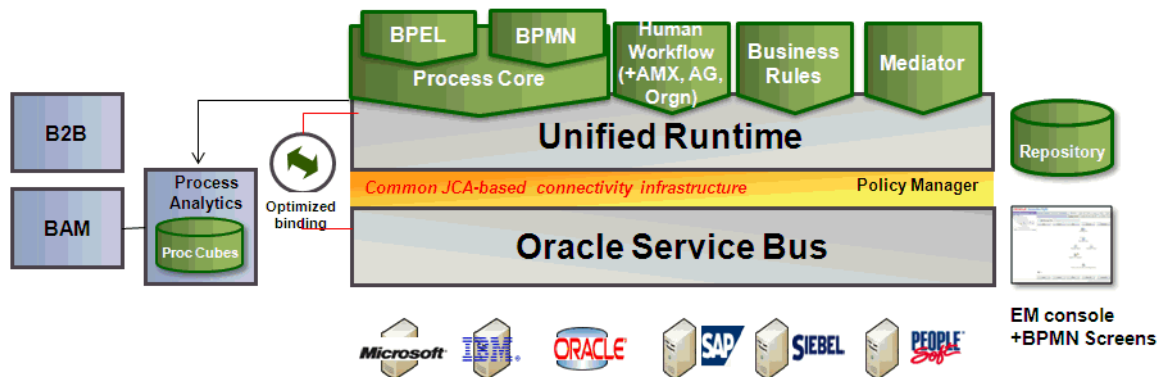
- After you configure the Maximum Cache Size setting in the Resource Limits page of Oracle Web Cache Manager, use a simulated load or an actual load to monitor the cache to see how much memory is actually used. Verify that any additional memory usage does not result in the host swapping memory to disk, as this may impact performance.
- Personalization and compression rules for all sites include the following:
 - Images should be cached but not compressed
 - CSS files should be both cached and compressed for all request types
 - JS files should be both cached and compressed for all request types
 - HTML files should be both cached and compressed
 - SWF files should be cached but not compressed
 - Add a rule to compress but not cache .jspx files for all GET and POSTS
 - Add a rule to compress but not cache \.jspx.*\$ files for all GET and POSTS
 - Add a rule to compress but not cache adw\.jspx for all request types
 - Add a rule not to compress and not cache profiling.js for all request types

For more detail on setting cache and compression rules, see "Caching and Compressing Content," in *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache*.

Part IV

SOA Suite Components

This part describes configuring Oracle Service-Oriented Architecture (SOA) Suite components to improve performance. Oracle SOA Suite is a component of Oracle Fusion Middleware. Oracle SOA Suite provides a complete set of service infrastructure components for designing, deploying, and managing SOA composite applications. The image below shows the Oracle SOA Platform.



Oracle SOA Suite enables services to be created, managed, and orchestrated into SOA composite applications. Composites enable you to easily assemble multiple technology components into one SOA composite application. SOA composite applications consist of:

- **Service components:** Service components are the basic building blocks of SOA composite applications. Service components implement a part of the overall business logic of the SOA composite application. BPEL Process, Oracle Mediator, Human task flow and decision services are examples of the service components.
- **Binding components:** Binding components connect SOA composite applications to external services, applications, and technologies. Binding components are organized into two groups:
 - **Services:** Provide the outside world with an entry point to the SOA composite application. The WSDL file of the service advertises its capabilities to external applications. The service bindings define how a SOA composite service can be invoked (for example, through SOAP).
 - **References:** Enable messages to be sent from the SOA composite application to external services (for example, the same functionality that partner links provide for BPEL processes, but at the higher SOA composite application level).

The SOA Suite Components are documented in the following chapters:

- [Chapter 11, "General Tuning for SOA Suite Components"](#)
- [Chapter 12, "Oracle Business Rules Performance Tuning"](#)
- [Chapter 13, "Oracle BPEL Process Manager Performance Tuning"](#)
- [Chapter 15, "Oracle Mediator Performance Tuning"](#)
- [Chapter 16, "Oracle Business Process Management Performance Tuning"](#)
- [Chapter 17, "Oracle Human Workflow Performance Tuning"](#)
- [Chapter 18, "Oracle Adapters Performance Tuning"](#)
- [Chapter 14, "Oracle Business Activity Monitoring Performance Tuning"](#)
- [Chapter 19, "User Messaging Service Performance Tuning"](#)
- [Chapter 20, "Oracle B2B Performance Tuning"](#)

General Tuning for SOA Suite Components

This chapter describes tuning configurations that can apply to multiple SOA Suite applications.

- [Section 11.1, "About SOA Suite Configuration Properties"](#)
- [Section 11.2, "SOA Infrastructure Configurations"](#)
- [Section 11.3, "Modifying SOA Configuration Parameters"](#)
- [Section 11.4, "JVM Tuning Parameters"](#)
- [Section 11.5, "Database Settings"](#)

For more information on any of the SOA Suite Applications, see [Section IV, "SOA Suite Components"](#) for a list of the application-specific documentation provided in this guide.

Note: Additional SOA tuning recommendations can be found in "Managing Large Documents and Large Number of Instances" in the *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

11.1 About SOA Suite Configuration Properties

Refer to the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite* for more information on configuring the SOA Applications.

11.2 SOA Infrastructure Configurations

SOA Infrastructure configuration parameters impact the entire SOA Infrastructure. The following configurations are modified through the SOA-INFRA component:

- Viewing and setting the SOA Infrastructure audit level
- Capturing the state of the SOA composite application instance
- Enabling the payload validation of incoming messages
- Specifying the callback server and server URLs
- Setting UDDI registry properties
- Viewing the data source JNDI locations
- Setting the non-fatal connection retry count
- Setting Web service binding properties

For more information on SOA configuration, see "Configuring SOA Infrastructure Properties" in *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

11.2.1 Audit Level

The Audit Level property enables you to select the level of information to be collected by the message tracking infrastructure. This information is collected in the instance data store (database) associated with the SOA Infrastructure. This setting has no impact on what is written to log files.

Value	Description
Off	No composite instance tracking and payload tracking information is collected. Note that no logging and display of instances in Oracle Enterprise Manager Fusion Middleware Control Console can result in a slight performance increase for processing instances. Instances are created, but are not displayed.
Production	Composite instance tracking is collected, but the Oracle Mediator service engine does not collect payload details and the process service engine does not collect payload details for assign activities (payload details for other activities are collected). This level is optimal for most normal production operations.
Development	Enables both composite instance tracking and payload detail tracking. However, this setting may impact performance. This level is useful largely for testing and debugging purposes.

11.2.2 Instance Tracking Audit Trail Threshold

This parameter is used to limit the audit trail size while it is being built. The default value is 1MB. If the audit trail exceeds the `instanceTrackingAuditTrailThreshold` size (1MB by default), then an exception is thrown, and the audit trail is not fully built. The value is in Bytes, so the default value is 1024*1024. This parameter can improve performance, as it prevents huge audit trails to potentially consume a lot or all the memory available on the SOA server where the audit trail is built. So in many way it acts as a safety valve. However in some cases users might want to increase the default value, if they get an exception while retrieving audit trails from Enterprise Manager, that states the " `instanceTrackingAuditTrailThreshold`" has been exceeded.

11.2.3 Logging Level

For more information on setting the logging levels for your applications, see "Configuring Log Files" in *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

11.3 Modifying SOA Configuration Parameters

SOA soa-infra level configurations can be set through Oracle Enterprise Manager.

For more information, see "Getting Started with Administering Oracle SOA Suite and Oracle BPM Suite" in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

11.4 JVM Tuning Parameters

JVM parameters can have an impact on SOA performance. The major factors that impact a SOA component's performance relate to the heap size. For more information on tuning the JVM for performance, see [Section 2.4, "Tuning Java Virtual Machines \(JVMs\)"](#).

11.5 Database Settings

Tuning your database configurations may be useful with the SOA Suite of applications. Configurations and specific settings may vary for different use cases. See your database-specific administration manuals for more information on tuning database properties.

For additional basic database tuning guidelines, see [Section 2.6, "Tuning Database Parameters"](#).

11.5.1 Configuring Data Sources for SOA

SOA obtains database connections using an application server managed data source. You can use the WebLogic Server Console to configure SOA data source. For more information on using the WebLogic Server Console, see the *Oracle Fusion Middleware Administrator's Guide*.

Consider the following data source configurations when performance is an issue:

- When configuring the data source, ensure that the connection pool has enough free connections.
- Statement caching can eliminate potential performance impacts caused by repeated cursor creation and repeated statement parsing and creation. Statement caching also reduces the performance impact of communication between the application server and the database server
- Disable unnecessary connection testing and profiling.

For more information, see "Tuning JDBC Stores" in *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*.

11.5.2 Managing Tables and Indexes

Consider using hash partitioning on your tables and indexes if your data does not easily lend itself to range partitioning, but you would like to partition for performance and manageability reasons. Hash partitioning provides a method of evenly distributing data across a specified number of partitions. Rows are mapped into partitions based on a hash value of the partitioning key. Creating and using hash partitions gives you a highly tunable method of data placement, because you can influence availability and performance by spreading these evenly sized partitions across I/O devices (striping).

Caution: Do not use partitioning on any date-based index that is likely to be used in range queries by Enterprise Manager or SOA. Partitioning date-based indexes, such as STATE_TYPE_DATE and COMPOSITE_INSTANCE_CREATED, for example, could significantly impact performance.

To improve performance, consider using hash partitioning on the following tables and indexes. Note that the performance improvement of each partitioning will vary considerably depending on the type of composite.

Partitioned Table Name	Partition Type	Number
CUBE_INSTANCE	Partitioned and Reverse key index CI_CREATION_DATE	Not Applicable
CUBE_SCOPE	Partition by hash (CIKEY)	Partitions = 200
MEDIATOR_CASE_INSTANCE	Partition by hash ("ID")	Partitions = 200
XML_DOCUMENT	Partition by hash (document_id)	Partitions = 200

Hash Partitioned Indexes

BRDECISIONINSTANCE_INDX3
MEDIATOR_INSTANCE_INDEX2
MEDIATOR_INSTANCE_INDEX5
MEDIATOR_INSTANCE_INDEX6
MEDIATOR_INSTANCE_INDEX1
MEDIATOR_INSTANCE_INDEX3
MEDIATOR_CASE_INSTANCE_INDEX2
MEDIATOR_CASE_DETAIL_INDEX1
REFERENCE_INSTANCE_CO_ID
CI_NAME_REV_STATE
DOC_DLV_MSG_GUID_INDEX
REFERENCE_INSTANCE_ECID
DMEDIATOR_CASE_INSTANCE_INDEX3
REFERENCE_INSTANCE_ID
MEDIATOR_CASE_DETAIL_INDEX1
CI_ECID
COMPOSITE_INSTANCE_ID
MEDIATOR_CASE_DETAIL_INDEX2

11.5.3 Weblogic Server Performance Tuning

For complete performance tuning of Weblogic Server, refer to *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*.

Oracle Business Rules Performance Tuning

Oracle Business Rules technology enables automation of business rules; it also enables extraction of business rules from procedural logic such as Java code or BPEL processes.

The chapter includes the following sections:

- [Section 12.1, "About Oracle Business Rules"](#)
- [Section 12.2, "Basic Tuning Considerations"](#)

12.1 About Oracle Business Rules

Oracle Business Rules provides high performance and easy to use implementation of Business Rules technology. It provides easy to use authoring environment as well as a very high performance inference capable rules engine. Oracle Business Rules is part of the Oracle Fusion Middleware stack and will be a core component of many Oracle products including both middleware and applications.

12.2 Basic Tuning Considerations

In most cases, writing of Rules should not require a focus on performance. However, as in any technology, there are tips and tricks that can be used to maximize performance when needed. Most of the considerations are focused on the initial configuration of the data model.

- [Section 12.2.1, "Use Java Beans"](#)
- [Section 12.2.2, "Assert Child Facts instead of Multiple Dereferences"](#)
- [Section 12.2.3, "Avoid Side Affects in Rule Conditions"](#)
- [Section 12.2.4, "Avoid Expensive Operations in Rule Conditions"](#)
- [Section 12.2.5, "Consider Pattern Ordering"](#)
- [Section 12.2.6, "Consider the Ordering of Tests in Rule Conditions"](#)
- [Section 12.2.7, "Enable assertXPath Support"](#)

12.2.1 Use Java Beans

The rule engine is most efficient when the facts it is reasoning on are Java Beans (or RL classes) and the associated tests involve bean properties. The beans should expose get and set methods (if set is allowed) for each bean property. If application data is not directly available in Java Beans, flatten the data to a collection of Java Beans that will be asserted as facts (and used in the rules).

12.2.2 Assert Child Facts instead of Multiple Dereferences

Expressions like `Account.Contact.Address` involve more than one object dereference. In a rule condition, this is not as efficient as expressions with single dereferences. It is a best practice to flatten fact types as much as possible. If the fact type has a hierarchical structure, consider using `assertXPath` or other means to assert object hierarchy; that is for the preceding example, assert both `Account` and `Contact` as Fact Types.

12.2.3 Avoid Side Affects in Rule Conditions

Methods or functions that have side affects such as changing a value or state should not be used in a rule condition. Due to the optimizations performed when the rule engine builds the Rete network, and the Rete network operations that are performed as facts are asserted, modified (and re-asserted), or retracted, the tests in a rule condition may be evaluated a greater or lesser number of times than would occur in a procedural program. If a method or function has side effects, those side effects may be performed an unexpected number of times.

12.2.4 Avoid Expensive Operations in Rule Conditions

Expensive operations should be avoided in rule conditions. Expensive operations would include any operation that involves I/O (disk or network) or even intensive computations. In general, consider avoiding I/O or DBMS access from the rules engine directly. These operations should be done external to the rules engine. For other expensive operations or calculations, consider performing the computations and assert the results as a Java or RL fact. These facts are used in the rule conditions instead of the expensive operations.

12.2.5 Consider Pattern Ordering

Reordering rule patterns can improve the performance of rule evaluation in time, memory use, or both. There are two main guidelines for ordering fact clauses (patterns) within a rule condition.

- If a fact is not expected to change (or will not change frequently) during rule evaluation, place its fact clause before fact clauses that change more frequently. That is, order the fact clauses by expected rate of change from least to greatest. Ordering fact clauses in this way can improve the performance (time) of rule evaluation.
- If a fact clause (including any tests that involve only that fact) is expected to match fewer facts than other fact clauses in the rule condition, place that fact clause before the others. That is, order the fact clauses from most restrictive (matches fewest facts) to least restrictive. This can reduce the amount of memory used during rule evaluation. It may also improve the performance.

Sometimes these two guidelines conflict and it may require some experimentation to arrive at the best ordering.

12.2.6 Consider the Ordering of Tests in Rule Conditions

Similar to the recommendations for fact clauses, the tests in a rule condition should be ordered such that a test that will be more restrictive is placed before a test that is less restrictive. This can reduce the amount of computation required for facts that do not satisfy the rule condition. If the degree of restrictiveness is not known, or estimated to be equal for a collection of tests, then the simpler tests should be placed before more expensive tests.

12.2.7 Enable assertXPath Support

Note: AssertXPath is only relevant for 10g rules usage or for 10g projects migrated to 11g. In 11g Release 1 and later, AssertXPath was replaced by AssertTree. In 11g, no XLink facts are asserted as the linkage is now handled in a different manner. AssertTree has the same performance, by default, as the special check box for 10g rules.

Most of the work done by the rules engine is done during assert, retract, or modify operations. In particular, the `assertXPath` method, though very convenient, may have a performance impact. The power of this method is not only that it asserts the whole hierarchy in one call, but also asserts some XLink facts for children facts to link back to parent facts. However, if these features are not needed, and you need to assert only a few levels as facts, it is better to turn off the "Supports XPath" for the relevant fact types and then use a function to do custom asserts.

Instead of using `assertXPath` the following example uses a function to assert `ExpenseReport` and `ExpenseLineItems`:

```
function assertAllObjectsFromList(java.util.List objList)
{
  java.util.Iterator iter = objList.iterator();
  while (iter.hasNext())
  {
    assert(iter.next());
  }
}

function assertExpenseReport (demo.ExpenseReport expenseReport)
{
  assert(expenseReport);
  assertAllObjectsFromList(expenseReport.getExpenseLineItem());
}
```

To improve performance of `assertXPath`, select the "Enable improved `assertXPath` support for performance" check box in the Dictionary Properties page in Rule Author. Taking advantage of this will require that the following conditions are met:

- `assertXPath` is only invoked with an XPath expression of `"/*"`. Any other XPath expression will result in an `RLIllegalArgumentException`.
- XLink facts should not be used in rule conditions as the XLink facts will not be asserted.

Oracle BPEL Process Manager Performance Tuning

Oracle Business Process Execution Language (BPEL) Process Manager provides several property settings that can be configured to optimize performance at the composite, fabric, application and server levels. This chapter describes these property settings and provides recommendations on how to use them.

This chapter contains the following sections:

- [Section 13.1, "About BPEL Process Manager"](#)
- [Section 13.2, "Basic Tuning Considerations"](#)
- [Section 13.3, "Advanced Tuning Considerations"](#)

13.1 About BPEL Process Manager

BPEL is the standard for assembling a set of discrete services into an end-to-end process flow, radically reducing the cost and complexity of process integration initiatives. Oracle BPEL Process Manager offers a comprehensive and easy-to-use infrastructure for creating, deploying and managing BPEL business processes.

For more information, see "Configuring BPEL Process Service Components and Engines" in *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite* and "Using the BPEL Process Service Component" in the *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

13.2 Basic Tuning Considerations

This section describes the performance tuning properties at the BPEL engine level. They can be configured using Oracle Enterprise Manager. For more information, see "Configuring BPEL Process Service Engine Properties" in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

Note: The configuration examples and recommended settings described in this chapter are for illustrative purposes only. Consult your own use case scenarios to determine which configuration options can provide performance improvements.

13.2.1 Tuning Audit Levels

The following properties can be set to audit levels.

13.2.1.1 Audit Level

The `auditLevel` property sets the audit trail level. This configuration property is applicable to both durable and transient processes. This property controls the amount of audit events that are logged by a process. Audit events result in more database inserts into the `audit_trail` table which may impact performance. Audit information is used only for viewing the state of the process from Oracle Enterprise Manager Console.

Use the `off` value if you do not want to store any audit information. Always choose the audit level according to your business requirements and use cases. For more information on setting the audit level, see "Introduction to the Order of Precedence for Audit Level Settings" in *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

Value	Description
Inherit	Inherits the audit level from infrastructure level.
off	No audit events (activity execution information) are persisted and no logging is performed; this can result in a slight performance boost for processing instances.
Minimal	All events are logged; however, no audit details (variable content) are logged.
Error	Logs only serious problems that require immediate attention from the administrator and are not caused by a bug in the product. Using this level can help performance.
Production	All events are logged. The audit details for assign activities are not logged; the details for all other activities are logged.
Development	All events are logged; all audit details for all activities are logged.

13.2.1.2 Audit Trail Threshold

The `audittrailthreshold` property sets the maximum size (in kilobytes) of an audit trail details string before it is stored separately from the audit trail. If an audit trail details string is larger than the threshold setting, it is not immediately loaded when the audit trail is initially retrieved; a link is displayed with the size of the details string. Strings larger than the threshold setting are stored in the `audit_details` table, instead of the `audit_trail` table.

The details string typically contains the contents of a BPEL variable. In cases where the variable is very large, performance can be severely impacted by logging it to the audit trail.

The default value is 50000 (50 kilobytes).

13.2.1.3 Audit Store Policy

This property specifies the strategy to persist the BPEL audit data.

Value	Description
<code>syncSingleWrite</code> (default)	AuditTrail and dehydration are persisted to DB in one transaction.
<code>syncMultipleWrite</code>	AuditTrail and dehydration are persisted in the same thread but separate transactions.

Value	Description
async	<p>AuditTrail and dehydration are persisted by separate threads and separate transactions.</p> <p>CAUTION: If you set <code>auditStorePolicy</code> to <code>async</code>, audit messages write to the database in a separate thread and transaction, which increases transaction speed. However, if a node crash occurs, audit trace functionality may be lost.</p> <p>Oracle recommends setting <code>auditStorePolicy</code> to <code>async</code> only if enhanced performance is critical.</p> <p>For more information, see "Oracle BPEL Process Manager High Availability Architecture and Failover Considerations" in the <i>Oracle Fusion Middleware High Availability Guide</i>.</p>

By default, audit messages are stored as part of the main BPEL transaction. A BPEL instance holds on to the audit messages until the flow reaches dehydration. In some use cases, for example when you have a large loop, and there is no dehydration point in the loop, a large number of audit logs are accumulated. This could lead to an out-of-memory issue and BPEL main transaction can experience timeout errors. You may consider using `syncMultipleWrite` or `async` to store the audit message separately from the main transaction.

When you use `syncMultipleWrite` and `async` `auditStorePolicy`, there are a few other properties that need to be considered. Please see the sections below.

13.2.1.4 Audit Flush ByteThreshold

This property controls how often the engine should flush the audit events, basically after adding an event to the current batch, the engine checks to see if the current batch byte size is greater than this value or not.

Consider tuning this property when `async` or `syncMultipleWrite` audit strategies are used. This size needs to be tuned based on the application.

13.2.1.5 Audit Flush EventThreshold

This property controls how often the engine should flush the audit events, basically when it reaches this limit of the number of events, the engine would trigger the store call.

Consider tuning this property when `async` or `syncMultipleWrite` audit strategies are used. This size needs to be tuned based on the application.

13.2.2 Tuning Database Persistence for BPEL

The `oneWayDeliveryPolicy` property controls database persistence of messages entering Oracle BPEL Server. By default, incoming requests are saved in the delivery service database table `dlv_message`. These requests are later acquired by Oracle BPEL Server worker threads and delivered to the targeted BPEL process. This property persists delivery messages and is applicable to durable processes.

This property can be set during creation of a BPEL process service component. For more information, see "How to Add a BPEL Process Service Component" in *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

Note: One-way invocation messages are stored in the delivery cache until delivered. If the rate at which one-way messages arrive is much higher than the rate at which Oracle BPEL Server delivers them, or if the server fails, messages may be lost.

The `oneWayDeliveryPolicy` is from the Oracle 10g configuration property `deliveryPersistencePolicy`. The configuration property name in 11g is `bpel.config.oneWayDeliveryPolicy`.

Value	Description
<code>async.persist (Default)</code>	Delivery messages are persisted in the database. With this setting, reliability is obtained with some performance impact on the database. In some cases, overall system performance can be impacted.
<code>async.cache</code>	<p>Incoming delivery messages are kept only in the in-memory cache. If performance is preferred over reliability, this setting should be considered.</p> <p>CAUTION: If you set <code>OneWayDeliveryPolicy</code> to <code>async.cache</code>, invoke messages do not persist to the database; they are in memory. If the SOA node that holds these messages fails, those messages will be lost.</p> <p>Use the <code>async.cache</code> option only when performance is critical and when you have an external means to track messages and redeliver them.</p> <p>For more information, see "Oracle BPEL Process Manager High Availability Architecture and Failover Considerations" in the <i>Oracle Fusion Middleware High Availability Guide</i>.</p>
<code>sync</code>	Directs Oracle BPEL Server to bypass the scheduling of messages in the invoke queue, and invokes the BPEL instance synchronously. In some cases this setting can improve database performance.

13.2.3 Tuning Invoke Messages

The `MaximumNumberOfInvokeMessagesInCache` property specifies the number of invoke messages that can be kept in the in-memory cache. Once the engine hits this limit, the message is pushed to dispatcher in-memory cache. The saved messages can be recovered using a recovery job. Use value -1 to disable.

The default value is 100000 messages.

13.2.4 Tuning Processed Requests List

The `StatsLastN` property sets the size of the most-recently processed request list. After each request is finished, statistics for the request are kept in a request list. A value less than or equal to 0 disables statistics gathering. To optimize performance, consider disabling statistics collection if you do not need them.

This property is applicable to both durable and transient processes.

The default value is -1.

13.2.5 Tuning XML Document Persistence

The `largedocumentthreshold` property sets the large XML document persistence threshold. This is the maximum size (in kilobytes) of a BPEL variable before it is stored in a separate table from the rest of the instance scope data.

This property is applicable to both durable and transient processes.

Large XML documents impact the performance of the entire Oracle BPEL Server if they are constantly read in and written out whenever processing on an instance must be performed.

The default value is 10000 (100 kilobytes).

13.2.6 Validating XML

The `validateXML` property validates incoming and outgoing XML documents. If set to True, the Oracle BPEL Process Manager applies schema validation for incoming and outgoing XML documents. Nonschema-compliant payload data is intercepted and displayed as a fault.

This setting is independent of the SOA composite application and SOA Infrastructure payload validation level settings. If payload validation is enabled at both the service engine and SOA Infrastructure levels, data is checked twice: once when it enters the SOA Infrastructure, and again when it enters the service engine.

CAUTION: Enabling XML payload validation can impact performance.

This property is applicable to both durable and transient processes.

The default value is False.

13.2.7 Tuning Wait Time

The `SyncMaxWaitTime` property sets the maximum time the process result receiver waits for a result before returning. Results from asynchronous BPEL processes are retrieved synchronously by a receiver that waits for a result from Oracle BPEL Server.

The default value is 45 seconds.

The `SyncMaxWaitTime` property can be found on the BPEL Service Engine Properties page by clicking the **More BPEL Configuration Properties** link at the bottom of the screen. For more information, see "Configuring BPEL Process Service Engine Properties" in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

13.2.8 Tuning Instance Key Block Size

The `InstanceKeyBlockSize` property controls the instance ID range size. Oracle BPEL Server creates instance keys (a range of process instance IDs) in batches using the value specified. After creating this range of in-memory IDs, the next range is updated and saved in the `ci_id_range` table.

For example, if `instanceKeyBlockSize` is set to 100, Oracle BPEL Server creates a range of instance keys in-memory (100 keys, which are later inserted into the `cube_instance` table as `cikey`). To maintain optimal performance, ensure that the block size is larger than the number of updates to the `ci_id_range` table.

The default value is 10000.

13.2.9 Tuning Automatic Recovery Attempts

The `MaxRecoveryAttempt` parameter allows you to configure the number of automatic recovery attempts to submit in the same recoverable instance. The value you provide specifies the maximum number of times invoke and callback messages are recovered. Once the number of recovery attempts on a message exceeds the specified value, a message is marked as nonrecoverable.

When a BPEL instance makes a call to another server using `invokeMessage`, and that call fails due to a server down, validation error, or security exception, the `invokeMessage` is placed in a recovery queue and BPEL attempts to retry those messages. When there are many messages, and a majority of them are being sent to the same target, the target can become overloaded. Setting the appropriate value of `MaxRecoveryAttempt` will prevent excessive load on servers that are targeted from BPEL web service calls.

For more information on configuring the number of automatic recovery attempts to submit in the same recoverable instance, see "Configuring Automatic Recovery Attempts for Invoke and Callback Messages" in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

13.3 Advanced Tuning Considerations

The following BPEL tuning considerations may not be applicable to all BPEL deployments. Always consult your own use case scenarios to determine if these configurations should be used in your deployment.

13.3.1 Tuning BPEL Properties Set Inside a Composite

The BPEL properties set inside a composite affect the behavior of the component containing the BPEL process only. Each BPEL process can be created as a component of a composite. Component properties include `inMemoryOptimization`, `completionPersistPolicy`, and `auditLevel`.

For more information on configuring these properties, see the "Deployment Descriptor Properties" appendix of the *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

13.3.2 Identifying Tables Impacted By Instance Data Growth

Instance data occupies space in Oracle BPEL Process Manager schema tables. Data growth from auditing and dehydration can have a significant impact on database performance and throughput. See [Section 13.2.1, "Tuning Audit Levels"](#) for audit configuration and [Section 13.3.1, "Tuning BPEL Properties Set Inside a Composite"](#) for dehydration configuration. The table below describes the tables that are impacted by instance data growth. A brief description is provided of each table.

Table 13–1 Oracle BPEL Process Manager Tables Impacted by Instance Data Growth

Table Name	Table Description
<code>audit_trail</code>	Stores the audit trail for instances. The audit trail viewed in Oracle BPEL Control is created from an XML document. As an instance is processed, each activity writes events to the audit trail as XML.

Table 13–1 (Cont.) Oracle BPEL Process Manager Tables Impacted by Instance Data

Table Name	Table Description
audit_details	Stores audit details that can be logged through the API. Activities such as an assign activity log the variables as audit details by default. Audit details are separated from the audit_trail table due to their large size. If the size of a detail is larger than the value specified for this property, it is placed in this table. Otherwise, it is placed in the audit_trail table.
cube_instance	Stores process instance metadata (for example, the instance creation date, current state, title, and process identifier)
cube_scope	Stores the scope data for an instance (for example, all variables declared in the BPEL flow and some internal objects that help route logic throughout the flow).
dlv_message	Stores incoming (invocation) and callback messages upon receipt. This table only stores the metadata for a message (for example, current state, process identifier, and receive date).
dlv_subscription	Stores delivery subscriptions for an instance. Whenever an instance expects a message from a partner (for example, the receive or onMessage activity) a subscription is written out for that specific receive activity.
document_ci_ref	Stores cube instance references to data stored in the xml_document table.
document_dlv_msg_ref	Stores references to dlv_message documents stored in the xml_document table.
wftask	Stores tasks created for an instance. The TaskManager process keeps its current state in this table.
work_item	Stores activities created by an instance. All activities in a BPEL flow have a work_item table. This table includes the metadata for the activity (current state, label, and expiration date (used by wait activities)).
xml_document	Stores all large objects in the system (for example, dlv_message documents). This table stores the data as binary large objects (BLOBs). Separating the document storage from the metadata enables the metadata to change frequently without being impacted by the size of the documents.
Headers_properties	Stores headers and properties information.

Oracle Business Activity Monitoring Performance Tuning

This chapter describes how to tune the Oracle Business Activity Monitoring (BAM) dashboard application for optimal performance. Oracle BAM provides the tools for monitoring business services and processes in the enterprise.

This chapter discusses useful parameters that can be modified to enhance the overall performance of BAM:

- [Section 14.1, "About Oracle Business Activity Monitoring"](#)
- [Section 14.2, "Basic Tuning Considerations"](#)

14.1 About Oracle Business Activity Monitoring

Oracle Business Activity Monitoring (BAM) provides the tools for monitoring business services and processes in the enterprise. It allows correlating of market indicators to the actual business process and to changing business processes quickly or taking corrective actions if the business environment changes. Oracle BAM also provides the necessary tools and run-time services for creating dashboards that display real-time data inflow and define rules to send alerts under specified conditions.

For more information see *Oracle Fusion Middleware User's Guide for Oracle Business Activity Monitoring*.

14.2 Basic Tuning Considerations

The following sections provide Oracle BAM tuning considerations that can be used to address performance issues:

- [BAM Server Tuning](#)
- [BAM Dashboard Tuning](#)
- [BAM Database Tuning](#)
- [Internet Browser Tuning](#)
- [Enterprise Message Source Tuning](#)

14.2.1 BAM Server Tuning

The following tuning configurations can be used to improve performance of the BAM Server:

14.2.1.1 Set the ViewSetSharing and ElementCountLimit Parameters

The `ViewSetSharing` parameter can be set to `TRUE` or `FALSE` in the BAM server configuration file. This parameter enables view set sharing when possible. Typically a particular view set can be shared with other users if they are trying to access the same dashboard, if the view sets are not dissimilar due to factors like row level security or prompts/parameters tied to filters.

Consider setting the `ViewSetSharing` parameter to `TRUE` so that Active Data Cache (ADC) can reuse the same viewset and snapshot and avoid creating more viewsets. This reduces the BAM server resource usage and improves user response time.

If this parameter is turned on, it does not always guarantee that ADC can reuse the existing viewset. If there have been too many changes to the underlying snapshot for the existing viewset, ADC may choose to create new viewset instead.

The `ReportCache` parameter used to determine if there have been too many changes is `ElementsCountLimit`. This defines the number of changes to the snapshot used by Report Cache to do the determination. In cases where the active data comes in at a fast rate, try to set this parameter to a large number so that ADC can use view sharing at the expense of more server CPU usage. The default value of `ElementsCountLimit` is 50.

14.2.1.2 Enable the Async Servlet

During periods of higher active data rates, the browser uses more memory. To prevent potential impacts to performance, consider providing more memory on the client machine. To do this, set the `UseAsynchServlet=TRUE` for the BAM dashboard application.

The BAM dashboard application uses the Async servlet feature so that the BAM server does not bind a specific thread to a specific user request. This provides for better server-side system resource usage.

This parameter can be turned off by adding `UseAsynchServlet=FALSE` in the server configuration file. During debugging, consider turning it off to make the process easier.

Otherwise this should always be turned on, which is the default.

See "Creating the Dashboard View" in *Oracle Fusion Middleware User's Guide for Oracle Business Activity Monitoring*.

14.2.2 BAM Dashboard Tuning

This section provides information on tuning the BAM dashboard for performance.

14.2.2.1 Tune the Active Data Retrieval Interval

The Active Data Retrieval Interval parameter controls the rate in milliseconds at which the Oracle BAM Active Data Cache (ADC) pushes events to the Oracle BAM Report Server. This is one of the factors that can affect the frequency of viewing active events on the dashboard page. Increasing this interval reduces the load on the Oracle BAM Server. Note that larger intervals increase the likelihood of multiple updates in the dashboard collapsing into a single update.

The default `ADCPushInterval` value is 1 second. You can override the default `ADCPushInterval` value within a particular report using the Active Data Retrieval Interval property in Active Studio.

For more information on using Active Studio, see "Getting Started With Oracle BAM Active Studio" in *Oracle Fusion Middleware User's Guide for Oracle Business Activity Monitoring*.

14.2.3 BAM Database Tuning

To achieve the best performance for Oracle Business Activity Monitoring, consider maintaining a database on its own hardware dedicated to the Oracle Business Activity Monitoring system. General database administration practices, as described in the *Oracle Database Performance Tuning Guide*, also apply to a database dedicated to Oracle Business Activity Monitoring.

For more information on general database configurations, see [Section 2.6, "Tuning Database Parameters"](#).

14.2.4 Internet Browser Tuning

This section provides performance tuning configurations for Internet browsers:

14.2.4.1 Set iActiveDataScriptsCleanupFactor

BAM sends active data in <script> blocks to the browser over a persistent connection. In some cases, the browser does not free up the memory used by the <script> blocks. This can impact dashboard performance over time.

The `iActiveDataScriptsCleanupFactor` parameter provides a solution for these memory issues. A periodic browser refresh is forced after receiving the specified number of characters. The issue may become apparent when active data is being sent to the dashboard at a fast pace. You may need to increase this value further for particularly high rates of data such as when active data is coming to the dashboard at a rate of 25 events per second or greater. Ultimately the value you set depends on factors like your data, number of views, number of viewsets, `ADCPushinterval`, and so on). You can monitor the browser's memory consumption to help determine an appropriate value.

If performance continues to be an issue, consider increasing the value for this parameter. For example, set the value to 2 or 3 times the default value if active data is predicted to increase. The default value for this parameter is 1048576 bytes. The default value often prevents frequent reconnects and prevents CPU/memory on the client machine from creeping up too high.

14.2.4.2 Set Browser Cache Settings

If you are using Microsoft Internet Explorer, consider setting the Browsing History Settings to "Automatic." See the Microsoft Internet Explorer online help for more information.

14.2.5 Enterprise Message Source Tuning

BAM Enterprise Message Source (EMS) provides inbound JMS connectivity to BAM. After setup, a BAM EMS instance can monitor JMS queues/topics and read data from them. Each EMS instance is configured to publish data to a single Data Object in BAM Server. The Enterprise Message Source supports four types of operations: Insert, Update, Upsert, or Delete. Two types of JMS messages are supported: `MapMessage` and `TextMessage`.

14.2.5.1 Message Batching

The EMS batching process clubs messages into one single message before it is sent to BAM EMS. This feature enables the sender to send all messages in one batch over JMS. The batching process can improve network performance by limiting the number of round trips from the sender to JMS server to BAM EMS.

Oracle Mediator Performance Tuning

This chapter describes how to tune Oracle Mediator for optimal performance. It contains the following topics:

- [Section 15.1, "About Oracle Mediator"](#)
- [Section 15.2, "Basic Tuning Considerations"](#)
- [Section 15.3, "Tuning Event Delivery Network \(EDN\)"](#)

15.1 About Oracle Mediator

Mediator is a component of Oracle SOA offering that provides mediation capabilities like selective routing, transformation and validation capabilities, along with various message exchange patterns, like synchronous, asynchronous and event publishing or subscription. Oracle Mediator provides the framework to mediate between various providers and consumers of services and events. The Mediator service engine runs with the SOA Service Infrastructure Java EE application.

See Also: For details about the SOA Suite, see *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

For details about Oracle Mediator, see "Administering Oracle Mediator Service Components and Engines" in *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

15.2 Basic Tuning Considerations

In most business environments, customer data resides in disparate sources including business partners, legacy applications, enterprise applications, databases, and custom applications. The challenge of integrating this data efficiently can be met by using Oracle Mediator to deliver real-time data access to all applications that update or have a common interest in the same data.

Note: Before you begin tuning Oracle Mediator properties, be sure that you have read and understand the Oracle Mediator chapters in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

This section provides details about setting common Oracle Mediator properties such as:

- [Tuning metricsLevel](#)
- [Using Domain-Value Maps](#)
- [Deploying Deferred Routing Rules](#)
- [Tuning Error and Retry Parameters](#)
- [Setting the Audit Level](#)
- [Using Resequencer for Messages](#)

15.2.1 Tuning metricsLevel

This property controls DMS metrics tracking level. By default, DMS metrics collections is enabled. If you do not need to collect DMS metrics data, consider setting the `metricsLevel` to Disabled to improve performance.

15.2.2 Using Domain-Value Maps

When performance is an issue, consider using domain-value maps instead of database lookup within XSL transformations to minimize file I/O.

For more information on using domain value maps, see "Working with Domain Value Maps" in *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

15.2.3 Deploying Deferred Routing Rules

The following performance configuration parameters can be used for tuning components with parallel routing rules deployed:

- `DeferredWorkerThreadCount`: Specifies the number of deferred dispatchers for processing messages in parallel. For higher loads consider increasing this parameter to have more number of outbound threads for deferred processing as each parallel rule is processed by one of the `DeferredWorkerThreads`. Default value is 4 threads.
- `DeferredMaxRowsRetrieved`: When Mediator routing rule type is set to 'Parallel', `DeferredMaxRowsRetrieved` sets the number of maximum rows (maximum number of messages for parallel routing rule processing) that are retrieved from Mediator store table (which stores messages for parallel routing rule for processing.) Note that each message retrieved in this batch is processed by one worker thread at a time. Default value is 200 rows.
- `DeferredLockerThreadSleep`: For processing parallel routing rules, Oracle Mediator has a daemon locker thread that retrieves and locks messages from Mediator store database. The thread polls the database to retrieve messages for parallel processing. When no messages are available, the locker thread "sleeps" for the amount of time specified in the `DeferredLockerThreadSleep` and prevents round trips to database. Default value is 2 seconds. Consider increasing this value to improve performance. Some use case scenarios can benefit from a 'sleep' of 3600 seconds (60 minutes.)

During the specified time, no messages are available for parallel routing in either of the following cases:

- There are no Mediator components with parallel routing rules deployed.
- Mediator component(s) with parallel routing rule is deployed, but there are no continuous incoming messages for such components.

Note: You can specify Oracle Mediator component Priority through JDeveloper Mediator designer. This property is used to set priority among Oracle Mediator components with parallel routing rules.

For more information, see "Creating Mediator Routing Rules" in *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

15.2.4 Tuning Error and Retry Parameters

Consider increasing the `ErrorLockerThreadSleep` parameter value when you do not want to reduce the number of database trips.

The `ErrorLockerThreadSleep` parameter specifies the idle time between two successive iterations for retrieving errored out messages when there is no errored out message from parallel processing. The time is measured in seconds. Default value is 5 seconds. Consider increasing this value to improve performance. Some use case scenarios can benefit from an idle time of 3600 seconds (60 minutes.)

For more information on routing, see "Creating Mediator Routing Rules" in *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

15.2.5 Setting the Audit Level

The `auditLevel` property sets the audit trail logging level. This configuration property is applicable to all the Mediator components. This property controls the amount of audit events that are logged by a Mediator component. Audit events result in more database inserts into the `audit_trail` table which may impact performance. Audit information is used only for viewing the state of the Mediator component from Oracle Enterprise Manager Console.

Use the off value if you do not want to store any audit information. This value can improve performance in some use cases. Always choose the audit level according to your business requirements and use cases. For more information on setting the audit level, see "Understanding the Order of Precedence for Audit Level Settings" in *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

Value	Description
Inherit	Inherits the audit level from infrastructure level.
off	No audit events (flow execution information) are persisted and no logging is performed; this can result in a slight performance boost for processing instances.
Production	All events are logged. For each audit event, the payload details are not persisted.
Development	All audit events are logged. For each audit event, the payload details are also persisted.

15.2.6 Using Resequencer for Messages

A Resequencer is used to rearrange a stream of related but out-of-sequence messages back into order. It sequences the incoming messages that arrive in a random order and then send them to the target services in an orderly manner.

For more information about Resequencers, refer to "Administering Oracle User Messaging Service" in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

You can fine tune Resequencer by setting the value of the following properties in the Mediator Service Engine Properties page:

- `ResequencerWorkerThreadCount`: Specifies the worker thread count. Default is 4.
- `ResequencerMaxGroupsLocked`: Specifies the maximum number of groups locked in each iteration. Default is 4.
- `ResequencerLockerThreadSleep`: Specifies the sleep interval for the locker threads in seconds. Default is 10.

15.3 Tuning Event Delivery Network (EDN)

The Event Delivery Network (EDN) delivers events published by Oracle Mediator, Oracle BPEL Process Manager components, and external publishers such as Oracle Application Development Framework entity objects.

To improve performance of the Event Delivery Network, consider increasing the thread count (default is 3.) This property can be modified through WLST. For more information, see [Section 11.3, "Modifying SOA Configuration Parameters"](#).

Oracle Business Process Management Performance Tuning

The Oracle Business Process Management (BPM) Suite provides a seamless integration of all stages of the application development life cycle from design-time and implementation to run-time and application management.

This chapter contains the following sections:

- [Section 16.1, "About Oracle Business Process Management"](#)
- [Section 16.2, "Basic Tuning Considerations"](#)
- [Section 16.3, "Tuning Oracle Workspace and Worklist Applications"](#)
- [Section 16.4, "Tuning Process Analytics"](#)

16.1 About Oracle Business Process Management

The Oracle BPM Suite provides an integrated environment for developing, administering, and using business applications centered around business processes. BPM is layered on the Oracle SOA Suite and shares many of the same product components, including Business Rules, Human Workflow, and Oracle Adapter Framework for Integration.

For more information on using BPM, see the *Oracle Fusion Middleware User's Guide for Oracle Business Process Management*.

For more information on tuning Oracle BPM with your other Oracle Fusion Middleware components, see [Chapter 2, "Top Performance Areas"](#).

16.2 Basic Tuning Considerations

This section describes the following basic BPM performance tuning properties:

- [Audit Level](#)
- [LargeDocumentThreshold](#)
- [Dispatcher System Threads](#)
- [Dispatcher Engine Threads](#)
- [Dispatcher Invoke Threads](#)

Note: The configuration examples and recommended settings described in this chapter are for illustrative purposes only. Consult your own use case scenarios to determine which configuration options can provide performance improvements.

16.2.1 Audit Level

The `auditLevel` property sets the audit trail logging level. This configuration property is applicable to both durable and transient processes. This property controls the amount of audit events that are logged by a process. Audit events result in more database inserts into the `audit_trail` table which may impact performance. Audit information is used only for viewing the state of the process from Oracle Enterprise Manager Console.

Use the `off` value if you do not want to store any audit information. Always choose the audit level according to your business requirements and use cases. For more information on setting the audit level, see "Understanding the Order of Precedence for Audit Level Settings" in *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

Value	Description
Inherit	Inherits the audit level from infrastructure level.
off	No audit events (activity execution information) are persisted and no logging is performed; this can result in a slight performance boost for processing instances.
Minimal	All events are logged; however, no audit details (variable content) are logged.
Error	Logs only serious problems that require immediate attention from the administrator and are not caused by a bug in the product. Using this level can help performance.
Production	All events are logged. The audit details for assign activities are not logged; the details for all other activities are logged.
Development	All events are logged; all audit details for all activities are logged.

16.2.2 LargeDocumentThreshold

The `largedocumentthreshold` property sets the large XML document persistence threshold. This is the maximum size (in kilobytes) of a BPMN Data Object before it is stored in a separate location from the rest of the instance scope data.

This property is applicable to both durable and transient processes.

Large XML documents impact the performance of the entire Oracle BPM Runtime if they are constantly read in and written out whenever processing on an instance must be performed.

The default value is 10000 (100 kilobytes).

16.2.3 Dispatcher System Threads

The `dspSystemThreads` property specifies the total number of threads allocated to process system dispatcher messages. System dispatcher messages are general clean-up

tasks that are typically processed quickly by the server (for example, releasing stateful message beans back to the pool). Typically, only a small number of threads are required to handle the number of system dispatch messages generated during run time.

The minimum number of threads for this thread pool is 1 and it cannot be set to 0 or a negative number.

The default value is 2. Any value less than 1 thread is changed to the default.

16.2.4 Dispatcher Engine Threads

The `dspEngineThreads` property specifies the total number of threads allocated to process engine dispatcher messages. Engine dispatcher messages are generated whenever an activity must be processed asynchronously. If the majority of processes deployed are durable with a large number of dehydration points (mid-process receive, `onMessage`, `onAlarm`, and wait activities), greater performance may be achieved by increasing the number of engine threads. Note that higher thread counts can cause greater CPU utilization due to higher context switching costs.

The minimum number of threads for this thread pool is 1 and it cannot be set to 0 or a negative number.

The default value is 30 threads. Any value less than 1 thread is changed to the default.

16.2.5 Dispatcher Invoke Threads

The `dspInvokeThreads` property specifies the total number of threads allocated to process invocation dispatcher messages. Invocation dispatcher messages are generated for each payload received and are meant to instantiate a new instance. If the majority of requests processed by the engine are instance invocations (as opposed to instance callbacks), greater performance may be achieved by increasing the number of invocation threads. Higher thread counts may cause greater CPU utilization due to higher context switching costs.

The minimum number of threads for this thread pool is 1 and it cannot be set to 0 or a negative number.

The default value is 20 threads. Any value less than 1 thread is changed to the default.

16.3 Tuning Oracle Workspace and Worklist Applications

The following settings can be used to tune Oracle Workspace and Worklist applications:

Parameter	Description
HTTP Session Timeout	<p>To manage over resource usage, adjust the session timeout value, in minutes, in the web.xml file.</p> <p>The following is a sample snippet of web.xml:</p> <pre><session-config> <session-timeout> 5 </session-timeout> </session-config></pre> <p>NOTE: If you must modify this property, post deployment, you must edit web.xml manually. See "Editing web.xml Properties" in <i>Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal</i>.</p>
ADF Client State Token	<p>Through this setting, you can control the number of pages users can navigate using the browser Back button without losing information. To reduce CPU and memory usage, you can decrease the value in the web.xml file.</p> <p>The following is a sample snippet of web.xml:</p> <pre><context-param> <param-name> org.apache.myfaces.trinidad.CLIENT_STATE_ MAX_TOKENS </param-name> <param-value> 3 </param-value> </context-param></pre> <p>NOTE: If you must modify this property, post deployment, you must edit web.xml manually. See "Editing web.xml Properties" in <i>Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal</i>.</p>
Compress_View_State Token	<p>This setting controls whether or not the page state is compressed. Zipping greatly reduced the memory being taken up by page state in the session object.</p> <p>The following is a snippet of the web.xml:</p> <pre><param-name>org.apache.myfaces.trinidad.COMPRESS_VIEW_ STATE</param-name> <param-value>true</param-value></pre>
DISABLE_CONTENT_COMPRESSION	<p>By default, style classes that are rendered are compressed to reduce page size. In production environments, make sure you remove the DISABLE_CONTENT_COMPRESSION parameter from the web.xml file or set it to FALSE.</p> <p>The following is a snippet of the web.xml:</p> <pre><param-name>org.apache.myfaces.trinidad.DISABLE_ CONTENT_COMPRESSION</param-name> <param-value>>false</param-value></pre>

16.4 Tuning Process Analytics

Tuning Process Analytics includes the following:

16.4.1 Process Measurement

Process Analytics uses measurement events to sample the process and publish measurements to registered consumers. These measurements can be disabled using the BPMN Configuration "Disable Sensors". Specific consumers for these measurements can be disabled by setting the BPMN Configuration "Disable Actions". For more information, see the *Oracle Fusion Middleware Administrator's Guide*.

Note: Only data that is useful should be published. The process design specifies what data (dimensions, measure, counters) should be published and at what point(s). If data is being generated that is not useful, then it could be adding unnecessary load to the system.

Measurement events are published on the JMS Topic: MeasurementTopic, and consumed by registered Action MDBs. In order to tune JMS for Measurements, consider changing the following, as needed, in a high volume environment:

- MeasurementTopic
 - Bytes Max 800 MB
 - Message Max 1000000
- MeasurementTopicConnectionFactory
 - Send Timeout 240000
- BPMJMSServer
 - MessageBuffer Size 100000

Note that the BPMJMSServer uses a Paging File and JMSFileStore.

16.4.2 Tuning Process Cubes

Process Cubes perform periodic aggregations to compute workload information. The frequency of these computations is determined by the `CubeUpdateFrequency` parameter of `BPMNConfig` mbean and can be changed from the Oracle Enterprise Manager console. In a high volume environment, consider changing this parameter to an appropriately higher value such as 12 hours, for example, to conserve computing resources.

Note: The creation of workload snapshots can impact performance. Consider using the properties in Oracle Fusion Middleware Control to tune the frequency and time to live (TTL) for workload snapshots. For more information on using Fusion Middleware Control, see the *Oracle Fusion Middleware Administrator's Guide*.

Process Cube Aggregator uses the `BPM_CUBE_AUDITINSTANCE` table to compute workload and performance information. Unwanted records from the `BPM_CUBE_AUDITINSTANCE` table get purged as part of the SOA Purge script. Additionally, consider running the following delete script periodically to purge the unwanted records from `BPM_CUBE_AUDITINSTANCE` table for improving the performance of Process Cube computations.

```
DELETE FROM BPM_CUBE_AUDITINSTANCE A
WHERE EXISTS
(SELECT 1 FROM BPM_CUBE_AUDITINSTANCE B
```

```
WHERE A.COMPONENTINSTANCEID = B.COMPONENTINSTANCEID AND
B.OPERATION=' INSTANCE_CREATED' AND
B.ACTIVITYSTATUS=' PROCESSED' )
```

After deleting unwanted records as described above, you should also consider "shrinking space". Deleted space is not immediately reclaimed, therefore run the following commands:

```
alter table BPM_CUBE_AUDITINSTANCE enable row movement;
```

```
alter table BPM_CUBE_AUDITINSTANCE shrink space;
```

```
ALTER TABLE BPM_CUBE_AUDITINSTANCE MODIFY LOB (<lob_column>) (SHRINK SPACE);
```

16.4.2.1 Tuning Workload Table Record Proliferation

1. Use the "Attribute" business indicators instead of "Dimension" business indicators when there is high cardinality. Workload records get generated in large numbers if a dimension business indicator has high cardinality. Before 11.1.1.7.0 release, its best to avoid using such business indicator. If the value is numeric, please consider making them as measure business indicator

From 11.1.1.7.0 release onwards, "Attribute" type business indicators have been introduced. It is recommended that all business indicators with high cardinality, which are neither dimensions nor measures, should be ideally created as an "Attribute" type business indicator. If you notice high memory usage (by workload calculation) and generation of large number of workload records for each interval, then re-evaluate all the dimension business indicators in the composite and determine if any of them should be actually created as an "Attribute" business indicator (keeping the same name). Once the change is done, re-deploy the updated composite with "Keep instances running" option.

2. Workload timer does the purge and workload calculation as part of the same transaction. If the purge times out, for example, `txn timeout`, `connection resource not available` and more, workload timer stops purging from the next snapshot onwards so that it does not affect the workload calculation.
 - The timer based purge can be resumed by either updating the `BPMNConfig CubeUpdateFrequency` parameter or by the restarting server.
 - If the environment has high number of records in `bpm_cube_workload` table, then it is best to update the `bpm_cube_workload` table manually. The only impact of this action is that the workload dashboards in workspace do not show any data till a fresh snapshot of workload calculation is populated.

Oracle Human Workflow Performance Tuning

This chapter describes how to tune Oracle Human Workflow for optimal performance. You can tune Oracle Human Workflow in these areas:

- [Section 17.1, "About Oracle Human Workflow"](#)
- [Section 17.2, "Basic Tuning Considerations"](#)
- [Section 17.3, "Improving Server Performance"](#)
- [Section 17.4, "Completing Workflows Faster"](#)
- [Section 17.5, "Tuning Identity Provider"](#)
- [Section 17.6, "Tuning the Database"](#)

17.1 About Oracle Human Workflow

Oracle Human Workflow is a service engine running in Oracle SOA Service Infrastructure that allows the execution of interactive human driven processes. A human workflow provides the human interaction support such as approve, reject, and reassign actions within a process or outside of any process. The Human Workflow service consists of a number of services that handle various aspects of human interaction with a business process.

For more information, see "Using the Human Workflow Service Component" in *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

See also the Oracle Human Workflow web site at <http://www.oracle.com/technology/products/soa/hw/index.html>.

17.2 Basic Tuning Considerations

This section discusses the various options available to address performance issues:

- [Minimize Client Response Time](#)
- [Choose the Right Workflow Service Client](#)
- [Narrow Qualifying Tasks Using Precise Filters](#)
- [Retrieve Subset of Qualifying Tasks \(Paging\)](#)
- [Fetch Only the Information That Is Needed for a Qualifying Task](#)
- [Reduce the Number of Return Query Columns](#)
- [Use the Aggregate API for Charting Task Statistics](#)
- [Use the Count API Methods for Counting the Number of Tasks](#)

- [Create Indexes On Demand for Flexfields](#)
- [Use the doesTaskExist Method](#)

17.2.1 Minimize Client Response Time

Since workflow client applications are interactive, it is important to have good response time at the client. Some of the factors that affect the response time include service call performance impacts, querying time to determine the set of qualifying tasks for the request, and the amount of additional information to be retrieved for each qualifying task.

17.2.2 Choose the Right Workflow Service Client

Workflow services support two major types of clients: SOAP and EJB clients. EJB clients can be further separated into local EJB clients and remote EJB clients.

If the client application is based on .Net technologies, then only the SOAP workflow services can be used. However, if the client application is based on Java EE technology, then consider which client should be used based on your use case scenarios. The options are listed below:

- Remote client - This is the best option in terms of performance in most cases. If the client is running in the same JVM as the workflow services (soa-infra application), the API calls are optimized so that there is no remote method invocation (RMI) involved. If the client is on a different JVM, then RMI is used, which can impact performance due to the serialization and de-serialization of data between the API methods.
- SOAP client - While this option is preferred for standardization (based on web services), there are additional performance considerations when compared to the remote method invocation (RMI) used in the remote client. Additional processing is performed by the web-services technology stack which causes the marshalling and unmarshalling of API method arguments between XML.

For more information, see *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

17.2.3 Narrow Qualifying Tasks Using Precise Filters

Using precise filters is one of the most important factors in improving response time. When a task list is retrieved, the query should be as precise as possible so the maximum filtering can be done at the database level.

For example, when the inbox view is requested for a user, the tasks are filtered mainly based on whether they are assigned to the current user or to the groups the user belongs to. By specifying additional predicate filters on the inbox view, the overall response time for the query can be reduced since lesser number of tasks qualify.

Alternatively, you can define views by specifying predicate filters and the overall response time for such views is reduced since lesser number of tasks qualify. All predicates passed to the query APIs (or defined in the views) are directly pushed to the database level SQL queries. With this information, the database optimizer can use the best indexes to create an optimal execution plan. The additional filters can be based on task attributes or promoted flex fields. For example, instead of listing all PO approval tasks, views can be defined to present tasks to the user based on priority, date, category, or amount range.

Example: To retrieve all assigned tasks for a user with priority = 1, you can use the following API call:

```
Predicate pred = new Predicate(TableConstants.WFTASK_STATE_COLUMN,
Predicate.OP_EQ,
IWorkflowConstants.TASK_STATE_ASSIGNED);
pred.addClause(Predicate.AND,
TableConstants.WFTASK_PRIORITY_COLUMN,
Predicate.OP_EQ,

1);
List tasks = querySvc.queryTasks(ctx,
queryColumns,
null,
ITaskQueryService.AssignmentFilter.MY ITaskQueryService.AssignmentFilter.MY,
null,
pred,
null,
startRow,
endRow);
```

17.2.4 Retrieve Subset of Qualifying Tasks (Paging)

Once the task list has been narrowed down to meet a specific criteria as discussed in the previous section, the next level of filtering is based on how many tasks are to be presented to the user. You want to avoid fetching too many rows, which not only increases the query time but also increases the application process time and the amount of data returned to client. The query API has paging parameters that control the number of qualifying rows returned to the user and the start row.

For example, in the `queryTasks` method:

```
List tasks = querySvc.queryTasks(ctx,
queryColumns,
null,
ITaskQueryService.AssignmentFilter.MY,
null,
pred,
null,
startRow,
endRow);
```

Consider setting the `startRow` and `endRow` parameters to values that may limit the number of return matching records.

17.2.5 Fetch Only the Information That Is Needed for a Qualifying Task

When using the `queryTask` service, consider reducing the amount of optional information retrieved for each task returned in the list. This may reduce the performance impacts from additional SQL query and Java logic.

For example, in the following `queryTasks` method, only the group actions information is retrieved. You can also retrieve attachment and payload information directly in the listing, but you may encounter performance impacts.

```
List<ITaskQueryService.OptionalInfo> optionalInfo
= new ArrayList<ITaskQueryService.OptionalInfo>();
optionalInfo.add(ITaskQueryService.OptionalInfo.GROUP_ACTIONS);
// optionalInfo.add(ITaskQueryService.OptionalInfo.ATTACHMENTS);
// optionalInfo.add(ITaskQueryService.OptionalInfo.PAYLOAD);
List tasks = querySvc.queryTasks(ctx,
queryColumns,
```

```
optionalInfo,  
ITaskQueryService.AssignmentFilter.MY,  
null,  
pred,  
null,  
startRow,  
endRow);
```

In rare cases where the entire payload is needed, then the payload information can be requested. Typically only some of the payload fields are needed for displaying the task list. For example, for PO Tasks, the PO amount may be a column that must be displayed. Rather than fetching the payload as additional information and then retrieving the amount using an xpath expression and displaying it in the listing, consider mapping the amount column from the payload to a flex field. The flex field can then be directly retrieved during SQL querying which may significantly reduce the processing time.

Similarly, for attachments where the name of the attachment is to be displayed in the listing and the document itself is stored in an external repository, consider capturing the attachment name in the payload and mapping it to a flex field, so that processing time is optimized. While constructing the listing information, the link to the attachment can be constructed by fetching the appropriate flex field.

17.2.6 Reduce the Number of Return Query Columns

When using the `queryTask` service, consider reducing the number of query columns to improve the SQL time. Also, try to use the common columns as they are most likely indexed and the SQL can execute faster.

For example, in the following `queryTasks` method, only the `TASKNUMBER` and `TITLE` columns are returned:

```
List queryColumns = new ArrayList();  
queryColumns.add("TASKNUMBER");  
queryColumns.add("TITLE");  
...  
List tasks = querySvc.queryTasks(ctx,  
null,  
ITaskQueryService.AssignmentFilter.MY,  
null,  
pred,  
null,  
startRow,  
endRow);
```

17.2.7 Use the Aggregate API for Charting Task Statistics

Sometimes it is necessary to display charts or statistics to summarize task information. Rather than fetching all the tasks using the query API, and computing the statistics at the client layer, consider using the new aggregate APIs to compute the statistics at the database level.

For example, the following call illustrates the use of the API to get summarized statistics based on state for tasks assigned to a user:

```
List taskCounts = querySvc.queryAggregatedTasks(ctx,  
Column.getColumn(WFTaskConstants.STATE_COLUMN),  
ITaskQueryService.AssignmentFilter.MY,  
keyWordFilter,  
filterPredicate,  
false,
```



```
false);
```

17.2.8 Use the Count API Methods for Counting the Number of Tasks

Sometimes it is only necessary to count how many tasks exist that match certain criteria. Rather than calling the `queryTasks` API method, and determining the size of the returned list, call the `countTasks` API method, which returns only the number of matching tasks. The performance impact of returning a count of tasks is much lower than returning a list of task objects.

For example, the following call illustrates the use of the API to get the total number of tasks assigned to a user:

```
int numberOfTasks = querySvc.countTasks(ctx,
ITaskQueryService.AssignmentFilter.MY,
keyWordFilter,
filterPredicate);
```

17.2.9 Create Indexes On Demand for Flexfields

The workflow schema table `WFTASK` contains several flexfield attribute columns that can be used for storing task payload values in the workflow schema. Because there are numerous columns, and their use is optional, the installed schema does not contain indexes for these columns. In certain use-cases, for example, where certain mapped flexfield columns are frequently used in query predicates, performance can be improved if you create indexes on these columns.

For example, to create an index on the `TEXTATTRIBUTE1` column, the following SQL command should be run:

```
create index WFTASKTEXTATTRIBUTE1_I on WFTASK(TEXTATTRIBUTE1);
```

Note: The exact indexes required depend on the flexfield attribute columns being used, and the nature of the queries being executed. After creating the indexes, the statistics for the `WFTASK` table should be re-computed and flushed.

17.2.10 Use the doesTaskExist Method

Sometimes it is necessary to check whether any tasks exist that match particular query criteria. Rather than calling the `countTasks` method, and checking if the number returned is zero, consider using `doesTaskExist`. The `doesTaskExist` method performs an optimized query that simply checks if any rows exist that match the specified criteria. This method may achieve better results than calling the `countTasks` method.

For example, the following call illustrates the use of the API method to determine if a user owns any task instances:

```
boolean userOwnsTask = querySvc.doesTaskExist(ctx,
ITaskQueryService.AssignmentFilter.OWNER,null,null);
```

17.3 Improving Server Performance

Server performance essentially determines the scalability of the system under heavily loaded conditions. [Section 17.2.1, "Minimize Client Response Time"](#) lists several ways in which client response times can be minimized by fetching the right of amount of information and reducing the potential performance impact associated with querying. These techniques also reduce the database and service logic performance impacts at

the server and can improve server performance. In addition, a few other configuration changes can be made to improve server performance:

- [Archive Completed Instances Periodically](#)
- [Select the Appropriate Workflow Callback Functionality](#)
- [Minimize Performance Impacts from Notification](#)
- [Deploy Clustered Nodes](#)

17.3.1 Archive Completed Instances Periodically

The database scalability of a system is largely dependent on the amount of data in the system. Since business processes and workflows are temporal in nature, once they are processed, they are not queried frequently. Having numerous completed instances in the system can slow the system. Consider using an archival scheme to periodically move completed instances to another system that can be used to query historical data. Archival should be done carefully to avoid orphan task instances.

17.3.2 Select the Appropriate Workflow Callback Functionality

The workflow callback functionality can be used to query or update external systems after any significant workflow event, such as assignment or completion of task. While this functionality is very useful, it has to be implemented correctly to avoid impacting performance.

When performance is critical, ensure that there are sufficient resources to update the external system after the task is completed instead of after every workflow event. For example, instead of using a callback, the service can be invoked once after the completion of the task. If a callback cannot be avoided, then consider using a Java callback instead of a BPEL callback. Java callbacks do not have the performance impact associated with a BPEL callback since the callback method is executed in the same thread. In contrast, a BPEL callback may impact performance when sending a message to the BPEL engine, which in turn must be correlated so that it is delivered to the correct process instance. The workflow service has to be called by the BPEL engine after the invocation of the service.

17.3.3 Minimize Performance Impacts from Notification

Notifications are useful for alerting users that they have a task to execute. In environments where most approvals happen through email, actionable notifications are especially useful. This also implies that there is not much load in terms of worklist usage. However if most users interact through the Worklist, and notifications serve a secondary purpose, then notifications should be used judiciously. Consider minimizing the notification to just alert a user when a task is assigned instead of sending out notifications for each workflow event. Also, if the task content is also mailed in the notification there may be an impact to performance. To minimize the impact, consider making the notifications secure in which case only a link to the task is sent in the notification and not the task content itself.

17.3.4 Deploy Clustered Nodes

All workflow instances and state information are stored in the dehydration database. Workflow services are stateless which means they can be used concurrently on a cluster of nodes. When performance is critical and a highly scalable system is needed, a clustered environment can be used for supporting workflow. For more information on clustered architecture, see [Section 30.2, "Using Clusters with Oracle Fusion](#)

Middleware".

17.4 Completing Workflows Faster

The time it takes for a workflow to complete depends on the routing type specified for the workflow. The workflow functionality provides some options that can be used to improve the amount of time it takes to complete workflows. Some of these options are discussed in this section:

- [Use Workflow Reports to Monitor Progress](#)
- [Specify Escalation Rules](#)
- [Specify User and Group Rules for Automated Assignment](#)
- [Use Task Views to Prioritize Work](#)

17.4.1 Use Workflow Reports to Monitor Progress

Several workflow reports (and corresponding views) are available that can make monitoring and proactively fixing problems easier. A few of these reports are listed below:

- The Unattended Tasks Report provides a list of group tasks that need attention since they have not yet been acquired by any user to work on.
- The Task Cycle Time Report gives an idea of how much time it takes for a particular type of workflow to complete.
- The Task Productivity Report indicates the inflow and outflow of tasks for different users.
- The Assignee Time Distribution Report provides a detailed drill-down of the time spent by each user during the task life cycle (including the idle time when the task was waiting to be picked up by a user.)

All of these reports can be used effectively to fix problems. By checking unattended tasks report, you can assign tasks that have been in the queue for a long time to specific users. By monitoring cycle time and other statistics, you can add staff to groups that are overloaded or take a longer time to complete. Thus reports can be used effectively to ensure workflows complete faster.

17.4.2 Specify Escalation Rules

To ensure that tasks do not get stuck at any user, you can specify escalation rules. For example, you can move a task to a manager if a certain amount of time passes without any action being taken on the task. Custom escalation rules can also be plugged in if the task must be escalated to some other user based on alternative routing logic. By specifying proper escalation rules, you can reduce workflow completion times.

17.4.3 Specify User and Group Rules for Automated Assignment

Instead of manually reassigning tasks to other users or members of a group, you can use user and group rules to perform automated reassignment. This ensures that workflows get timely attention. For example, a user can set up a user rule such that workflows of a specific type and matching a certain filter criteria are automatically reassigned to another user in a specified time window. Similarly, a group rule can be used to automatically reassign workflows to a member of the group based on different routing criteria such as round robin or most productive. Thus rules can help

significantly reduce workflow waiting time, which results in faster workflow completion.

17.4.4 Use Task Views to Prioritize Work

A user's inbox can contain tasks of various types with various due dates. The user has to manually sift through the tasks or sort them to find out which one he or she should work on next. Instead, by creating task views where tasks are filtered based on due dates or priority, users can get their work prioritized automatically so they can focus on completing their tasks instead of wasting their time on deciding which tasks to work on. This also results in faster completion of workflows.

17.5 Tuning Identity Provider

The workflow service uses information from the identity provider in constructing the SQL query to determine the tasks qualifying for a user based on his or her role/group membership. The identity provider is also queried for determining role information to determine privileges of a user when fetching the details of a task and determining what actions can the user perform on a task. There are a few ways to speed up requests made to the identity provider.

- Set the search base in the identity configuration file to node(s) as specific as possible. Ideally you should populate workflow-related groups under a single node to minimize traversal for search and lookup. This is not always possible; for example, you may need to use existing groups and grant membership to groups located in other nodes. If it is possible to specify filters that can narrow down the nodes to be searched, then you should specify them in the identity configuration file.
- Index all critical attributes such as dn and cn in the identity provider. This ensures that when a search or a lookup is done, only a subset of the nodes are traversed instead of a full tree traversal.
- Use an identity provider that supports caching. Not all LDAP providers support caching but Oracle Internet Directory supports caching which can make lookup and search queries faster.

17.6 Tuning the Database

The Human Workflow schema is shipped with several indexes defined on the most important columns for all the tables. Based on the type of request, different SQL queries are generated to fetch the task list for a user. The database optimizer evaluates the cost of different plan alternatives (for example, full table scan, access table by index) and decides on a plan that is lower in cost. For the optimizer to work correctly, the index statistics should be current at all times. As with any database usage, it is important to make sure the database statistics are updated at regular intervals and other tunable parameters such as memory, table space, and partitions are used effectively to get maximum performance.

For more information on tuning the database, see [Section 2.6, "Tuning Database Parameters"](#).

Oracle Adapters Performance Tuning

This chapter describes how to tune Oracle Adapters for optimal performance. Oracle Adapters, a component of the Oracle SOA Suite of Applications, provide an integrated view of data and allow multiple applications to be integrated.

This chapter contains the following sections:

- [Section 18.1, "About Oracle Adapters"](#)
- [Section 18.2, "Oracle JCA Adapters for Files/FTP"](#)
- [Section 18.3, "Oracle JCA Adapter for Database Tuning"](#)
- [Section 18.4, "Oracle Socket Adapter Tuning"](#)
- [Section 18.5, "Oracle SOA JMS Adapter Tuning"](#)
- [Section 18.6, "Oracle AQ Adapter Tuning"](#)
- [Section 18.7, "Oracle MQ Adapter Tuning"](#)

18.1 About Oracle Adapters

Oracle technology adapters integrate Oracle Application Server and Oracle Fusion Middleware components such as Oracle BPEL Process Manager (Oracle BPEL PM) or Oracle Mediator components to file systems, FTP servers, database queues (advanced queues, or AQ), Java Message Services (JMS), database tables, and message queues (MQ Series).

For more information on Oracle Adapters, see *Oracle Fusion Middleware User's Guide for Technology Adapters*.

18.2 Oracle JCA Adapters for Files/FTP

This section describes the various features available for scalability and performance tuning of Oracle File and FTP Adapters. The Oracle File and FTP Adapters provide knobs to throttle the inbound and outbound operations. The Oracle File and FTP Adapters also provide knobs that can be used to tune the performance of outbound operations. The Oracle File and FTP Adapters knobs are described in the following sections:

- [Inbound Throttling Best Practices](#)
- [Outbound Throttling Best Practices](#)
- [Outbound Performance Best Practices](#)

Note: For composites with Oracle File and FTP Adapters, which are designed to consume very large number of concurrent messages, you must set the number of open files parameter for your operating system to a larger value. For example, to set the number of open files parameter to 8192 for Linux, use the `ulimit -n 8192` command

18.2.1 Inbound Throttling Best Practices

The Oracle File and FTP Adapters provide parameters that can be used to throttle the inbound operations. The table below describes the inbound throttling practices:

Parameter	Type	Values	Description
MaxRaiseSize	JCA	<pre><property name="MaxRaiseSize " value="100" /></pre> <p>Default: 10000 (ten thousand)</p>	<p>This parameter defines the maximum number of files that the inbound adapter would submit for processing on each polling cycle. For example, if your inbound directory has 1000 files and the <code>MaxRaiseSize</code> is set to 100, the adapter can increase to 100 files on each polling cycle.</p> <p>Defined in the Inbound JCA File.</p>
SingleThreadModel	JCA	<pre><property name="SingleThread Model" value="true" /></pre> <p>Default: False (In this case, the global in-memory queue is used).</p>	<p>If the value is <code>true</code>, the poller lists, translates, or publishes files in the same thread. In other words, it does not use the global in-memory queue for publishing.</p> <p>Defined in the Inbound JCA File.</p>
ThreadCount	JCA	<pre><property name="ThreadCount " value="10" /></pre> <p>Default: -1 (In this case, the adapter uses the global thread pool and in-memory queue)</p>	<p>This parameter enables the Oracle File and FTP Adapters to create their own processor threads rather than depending on the global pool of processor worker threads for processing the enqueued files. This parameter partitions the in-memory queue and each composite application receives its own in-memory queue.</p> <p>If the <code>ThreadCount</code> is set to 0, then the threading behavior is the same as that of the <code>SingleThreadModel</code>. If the <code>ThreadCount</code> is set to -1, then the global thread pool is activated, which is the same as the Default Threading Model. The maximum value that can be set for <code>ThreadCount</code> is 40.</p> <p>Defined in the Inbound JCA File.</p>
PublishSize			<p>This parameter enables you to break the file content into a specified number batches. The debatching process can improve performance when the message size is big as the receiver of the message. Debatching can create more instances with less data.</p>

18.2.2 Outbound Throttling Best Practices

The Oracle File and FTP Adapters provide parameters that can be used to throttle the outbound operations. The table below describes the outbound throttling practices:

Parameter	Type	Value	Description
ConcurrentThreshold	JCA	<pre><property name="ConcurrentThre shold" value="100"/></pre> <p>Default: 20 (In this case, not more than 20 translations occur for a particular outbound scenario.)</p>	<p>This parameter specifies the maximum number of translation activities that are allowed to start in parallel for a particular outbound scenario. The translation step during the outbound operation is CPU intensive and must be monitored as it might cause other applications or threads to starve. The maximum value is 100.</p> <p>Defined in the Outbound JCA File.</p>

18.2.3 Outbound Performance Best Practices

The Oracle File and FTP Adapters provide parameters that can be used to tune the performance of outbound operations. The table below describes the outbound performance parameters:

Parameter	Type	Value	Description
UseStaging	JCA	<pre><property name="UseStaging" value="true"/></pre> <p>Default: True</p>	<p>If the parameter is set to true, then the outbound Oracle File or FTP Adapter writes translated data to a staging file and later streams the staging file to the target file. If the parameter is set to false, then the outbound Oracle File or FTP Adapter does not use an intermediate staging file.</p> <p>Defined in Outbound JCA File.</p>
serializeTranslation	Endpoint Property	<pre><reference name="PurchaseOrderOut "> <interface.wsdl interface="..." /> <binding.jca config="PurchaseOrderO ut_ftp.jca" /> <property name="serializeTransla tion" type="xs:string" many="false" source="" override="may">true</p roperty> </reference></pre> <p>Defaults:</p> <ul style="list-style-type: none"> ■ True (If the value of UseStaging is set to True) ■ False (If the value of UseStaging is set to False) 	<p>If True, then the translation step is serialized using a semaphore. The number of permits for semaphore (monitoring the translation step) comes from ConcurrentThreshold parameter (listed in the preceding table). The default value of True is used because the translation step is CPU intensive and you do not want to starve other applications or threads.</p> <p>If False, then the translation step occurs outside the semaphore.</p> <p>Defined in Binding property for reference in composite.xml.</p>

Parameter	Type	Value	Description
inMemoryTranslation	Binding Property	<pre><reference name="PurchaseOrderOut "> <interface.wsd1 interface="..." /> <binding.jca config="PurchaseOrderO ut_ftp.jca" /> <property name="inMemoryTranslat ion" type="xs:string" many="false" source="override="may ">false</property> </reference></pre> <p>Default: False</p>	<p>This parameter is applicable only if UseStaging is False.</p> <p>If True, then the translation step occurs in-memory (an in-memory byte array is created.)</p> <p>If False, then the adapter creates an output stream to the target file (FTP, FTPS, and SFTP included) and allows the translator to translate and write directly to the stream.</p> <p>Defined in Binding property for reference in composite.xml.</p>

18.3 Oracle JCA Adapter for Database Tuning

The Oracle Database Adapter is pre-configured with many performance optimizations. You can, however, make some changes to reduce the number of round trips to the database, as described in the following sections:

- [JCA Adapter Basic Tuning Considerations](#)
- [Existence Checking](#)

Note: The tuning considerations in this chapter are listed for example only. Tuning parameters are specific to each deployment. Review your current usage and performance issues to determine which tuning considerations can improve performance.

18.3.1 JCA Adapter Basic Tuning Considerations

Adapter performance is directly related to the number of round-trips to the database, and the network cost of each trip. If performance becomes an issue, and making modifications is appropriate for your deployment, consider tuning the following parameters:

- Use Indexes

Indexes can improve performance of selects, updates and deletes. Index all queried fields, such as the primary key and the MarkReadField of the LogicalDeletePollingStrategy, when polling. For MarkReadField specify a non-null MarkUnreadValue. Caution: An index on a column containing many nulls may revert to full table scans.
- Disable OptimizeMerge

The OptimizeMerge parameter allows the detection of XML elements for which no value was specified. The related columns are excluded from inserts and updates. Disabling this parameter generally improves performance, but there is one case where it could have a negative effect. If multiple rows are being passed in as a single XML, and each row has different columns set (user entered with many optional fields), there is no benefit from batch writing, as each insert or update is different.
- Increase MaxRaiseSize

The `MaxRaiseSize` parameter indicates the maximum number of XML records that can be raised at a time to the BPEL engine. For example, if you set `MaxRaiseSize = 10`, then 10 database records are raised simultaneously. On an inbound read, for example, you can set `MaxRaiseSize = 0` (unbounded) which means that if you read 1000 rows, you can create one XML with 1000 elements. These elements are passed through a single Oracle BPEL Process Manager instance. A merge on the outbound side can then take all 1000 in one group and write them all at once with batch writing. Use the `MaxRaiseSize` parameter for publishing large payloads.

- Increase `MaxTransactionSize`

This property controls the number of records processed per transaction by each thread. If set to a large value such as 1000, turning on the `UseBatchDestroy` option could have a negative impact on performance. Setting a large `MaxTransactionSize` and a small `MaxRaiseSize` could also have negative impact on performance. Consider maintaining up to a 10:1 ratio in a synchronous scenario. Ideally, you should consider increasing `MaxRaiseSize` until it is a 1:1 ratio.

- Enable `UseBatchDestroy`

This property controls how the processed records are updated (ex: Deleted for `DeletePollingStrategy`, MarkedProcessed for `LogicalDeleteStrategy`). If set, only one update/delete is executed for all the rows that are part of that transaction. The number of rows in a transaction is controlled by the `MaxTransactionSize` option. Note that this may not always offer an improvement because, by default, batch writing is used, which also ends up in a single round trip to the database.

- Enable Batch Reading

Batch reading of one-to-many and one-to-one relationships is on by default. You can also use joined reading for one-to-one relationships instead, which may offer a slight improvement.

- Disable Delete Polling Strategy

Avoid the delete polling strategy because it must individually delete each row. The sequencing polling strategy can destroy 1000 rows with a single update to a helper table. Note that a `LogicalDelete` is also better than `Delete`, as updates are typically faster than deletes. To maintain performance, however, ensure that you have indexed the table. If you have not indexed, you can keep the total number of rows small by using deletes. In some instances deletes may be faster as the cost of a full table scan is negligible.

- Use Distributed Polling

Distributed polling enables you to configure polling for scalability. For more information, see "Scalability" in *Oracle Fusion Middleware User's Guide for Technology Adapters*.

- Use Synchronous Processes

On BPEL you can configure Database Adapter processes to be synchronous. You can also create sequential routing rules in Mediator. This can improve throughput in database-to-database scenarios, as there is less instance processing impact.

- Use Insert

The insert operation is the most performant because it uses no existence check and has no extra performance impact associated with it. There are no reads, only writes. If you know that you are inserting most of the time, use insert, and catch a

unique key constraint SQL exception inside your BPEL process, which can then perform a merge or update instead.

To monitor performance, you can enable debug logging and then watch the SQL for various inputs.

- **Disable Merge**

Merge executes one extra SELECT per related table. The SELECT is used to determine whether each row should be inserted or updated. If the row is updated, the update performed is minimal. If no rows have changed, nothing is updated.

- **Use Connection Pooling**

The adapter should also point to a tuned data source connection pool. Tuning the connection pool is important because creating and tearing down database connections can impact performance.

- **Use Attribute Filtering**

On the Attribute Filtering page of the Adapter Configuration Wizard you can choose which fields to map to the XML and vice versa. You can improve performance by deselecting columns that are not needed for your particular business case, especially large columns like LOBs.

- **Use Native Sequencing**

If you are using the XSL functions to assign primary keys to records, consider using the built-in native sequencing support in the adapter. Sequencing support obtains and caches 50 keys at a time by default. Caching improves performance by reducing the number of round trips. The chunk size can be controlled incrementally by modifying the `sequencePreallocationSize` connector property.

- **Do not use primary or foreign keys on the database**

Using primary and foreign keys can impact performance. Avoid using them when possible. If foreign keys are defined on the tables, make sure an index is created on the foreign key.

- **JDBC Driver Class**

The default JDBC driver class used to create the physical database connections in the connection pool is `oracle.jdbc.xa.client.OracleXADataSource`. Changing the driver to `oracle.jdbc.OracleDriver` may provide some performance improvement.

For more information on tuning the JDBC drivers, see "Third Party JDBC Driver and Database Connection Configuration" in *Oracle Fusion Middleware User's Guide for Technology Adapters*.

18.3.2 Existence Checking

One method of performance optimization for merge is to eliminate check database existence checking. The existence check is marginally better if the row is new, because only the primary key is returned, not the entire row. Due to the nature of merge, however, if the existence check passes, the entire row must be read to calculate what changed. Therefore, for every row to be updated, you see one extra round trip to the database during merge.

Use check cache on the root descriptor/table and any child tables if A is master and B is a privately owned child. If A does not exist, B cannot exist. And if A exists, all of its child tables are loaded as part of reading A.

Note: One way to prevent merge from performing an existence check for every record, when you know that an insert is required, is to set the primary key to null.

18.3.3 Throttling

It is possible to configure a speed limit on DbAdapter performance to protect down-stream components from message bursts. Consider leaving burst records unprocessed on the source database until SOA can process them efficiently. As of Oracle Adapters release 11.1.1.6.0 you can set the inbound DbAdapter property `RowsPerPollingInterval`. It acts as a limit on the number of records which can be processed in one polling interval. The default is unlimited.

The following sections describe the configuration options for `RowsPerPollingInterval`:

18.3.3.1 Formula

The formula for maximum rows per second is:

$$\text{Number of active nodes in SOA cluster} \times \text{NumberOfThreads} \times \text{RowsPerPollingInterval} / \text{PollingInterval}$$

18.3.3.2 RowsPerPollingInterval and MaxTransactionSize

`MaxTransactionSize` can be thought of as `RowsPerDatabaseTransaction` or `DatabaseFetchSize`. It does not affect how many rows can be processed in one polling interval period.

The one exception is the following configuration:

```
-distributed polling checked, usesSkipLocking="false"
```

In this one case `RowsPerPollingInterval` will default to `MaxTransactionSize` instead of unlimited.

If `RowsPerPollingInterval` is set to lower than `MaxTransactionSize` or `MaxRaiseSize`, they will be effectively lowered to `RowsPerPollingInterval`.

18.3.3.3 Configuration

There is no UI support for `RowsPerPollingInterval`. Instead find the `db.jca` file for the inbound polling service and add the property manually. Add it to the same section as the properties `MaxRaiseSize`, `MaxTransactionSize`, and `PollingInterval`, in any order.

18.4 Oracle Socket Adapter Tuning

This section describes performance tuning for Oracle Socket Adapter. Performance can be optimized for the Oracle Socket Adapter using Connection Pool if the socket server you are connecting to does not close the socket with each interaction. Connection pool lets you use a socket connection repeatedly, avoiding the overload of creating a new socket for each interaction.

Note: The Connection Pool feature is applicable to outbound interactions only. For more information on Socket Adapters, see "Oracle JCA Adapter for Sockets" in *Oracle Fusion Middleware User's Guide for Technology Adapters*

In order to enable the connection pool feature for the Oracle Socket Adapter, the `KeepAlive` connection factory property must be set to `True`. This connection property can be modified using the Connection Pool tab of Oracle WebLogic Server Administration Console.

For instructions on modifying the Oracle Socket Adapter connection pooling, see "Configuring Oracle Socket Adapter Connection Pooling" in *Oracle Fusion Middleware User's Guide for Technology Adapters*.

18.5 Oracle SOA JMS Adapter Tuning

This section describes some of the properties that can be set for the Oracle SOA JMS Adapter to optimize performance. See "Introduction to the Oracle JMS Adapter" in the *Oracle Fusion Middleware User's Guide for Technology Adapters* for more information.

18.5.1 `adapter.jms.receive.threads` Property

To improve performance, the `adapter.jms.receive.threads` property can be tuned for an adapter service. The default value is 1, but multiple inbound threads can be used to improve performance. When specified, the value of `adapter.jms.receive.threads` is used to spawn multiple inbound poller threads.

For example:

```
<service name="dequeue" ui:wSDLLocation="dequeue.wsdl">
<interface.wSDL
interface="http://xmlns.oracle.com/pcbpel/adapter/jms/textmessageusingqueues/textm
essageusingqueues/dequeue%2F#wsdl.interface(Consume_Message_ptt)"/>
<binding.jca config="dequeue_jms.jca">
<property name="adapter.jms.receive.threads" type="xs:string"
many="false">10</property>
</binding.jca">
</service>
```

18.6 Oracle AQ Adapter Tuning

This section describes Oracle AQ Adapter tuning configurations.

18.6.1 `adapter.aq.dequeue.threads` Property

To improve dequeue performance 'adapter.aq.dequeue.threads' property can be set for an adapter service. Default value is 1 but multiple inbound threads can be used to improve performance. The value of property 'adapter.aq.dequeue.threads' is used to spawn multiple inbound poller threads.

For example:

```
<service name="dequeue" ui:wSDLLocation="dequeue.wsdl">
<interface.wSDL
interface="http://xmlns.oracle.com/pcbpel/adapter/aq/raw/raw/dequeue/#wsdl.interfa
ce(Dequeue_ptt)"/>
<binding.jca config="dequeue_aq.jca">
```

```
<property name="adapter.aq.dequeue.threads" type="xs:string"
many="false">10</property>
</binding.jca>
</service>
```

18.7 Oracle MQ Adapter Tuning

The Oracle MQ Series Adapter supports the scalability feature for inbound operations only. Oracle MQ Series Adapter provides the parameter to control the number of threads that dequeue the messages from the inbound queue. You must specify the following property in the.jca file:

```
InboundThreadCount='N'
```

In the example above *N* is the number of threads that you want to span to dequeue the messages from the inbound queue.

User Messaging Service Performance Tuning

This chapter describes tips for tuning the User Messaging Service. It contains the following sections:

- [Section 19.1, "About Oracle User Messaging Services"](#)
- [Section 19.2, "Basic Tuning Considerations"](#)
- [Section 19.3, "Database Tuning for Optimal Throughput"](#)

19.1 About Oracle User Messaging Services

Oracle User Messaging Service enables users to receive notifications sent from SOA applications that are developed and deployed to the Oracle WebLogic Server using Oracle JDeveloper.

At the application level, there is notification activity for a specific delivery channel (such as SMS or E-Mail). For example, when you build a SOA application that sends e-mail notification, you drag and drop an Email Activity component from the JDeveloper Component Palette to the appropriate location within a workflow. The application connects then sends notifications.

For more information on Oracle User Messaging Service, see *Oracle WebLogic Communication Services Administrator's Guide*, *Oracle WebLogic Communication Services Developer's Guide*, and the *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

19.2 Basic Tuning Considerations

Depending on your User Messaging usage and performance issues, you may consider tuning the following:

- [SMPP Driver Performance Tuning](#)
- [Email Driver Polling Frequency](#)

19.2.1 SMPP Driver Performance Tuning

Short Messaging Peer-Peer Protocol (SMPP) messaging drivers can be configured using Enterprise Manager. One of the key parameters for optimizing SMPP performance is `WindowSize`. This is especially important when the SMPP driver is connected to a remote SMSC and there is high network latency between the two elements. Configuring the `WindowSize` parameter enables the SMPP driver to send several requests to the Short Messaging Service Center (SMSC) before waiting for an acknowledgment. Without windowing (i.e., a `WindowSize` of 1), the driver must wait

for a synchronous acknowledgment from the SMSC before sending the next message. With windowing, more messages can be sent per network round-trip, allowing a higher overall throughput.

To take advantage of an increased `WindowSize`, the number of MDB threads for the driver must be correspondingly increased. The two values should be matched so that driver threads can process and send messages before waiting for the requests to be acknowledged. Increasing the two values may improve performance, but only up to the point at which network latency no longer dominates the sending rate. Also, the maximum allowed value for the `WindowSize` is normally defined as a service policy by the SMSC operator.

For more information, see "Configuring Oracle User Messaging Service" in *Oracle WebLogic Communication Services Administrator's Guide*.

19.2.2 Email Driver Polling Frequency

For Email drivers, the "CheckMailFreq" configuration parameter defines how frequently the driver checks for incoming emails. For example, a value of "30" means the driver checks the configured inbox every 30 seconds. This parameter can influence performance; checking more frequently enables the driver to keep up with a higher incoming email load, but can impact performance due to frequent IMAP or POP3 operations. Default value is 30 seconds.

19.3 Database Tuning for Optimal Throughput

User Messaging Service stores messaging state such as sent and received messages and delivery status information in the database. Therefore, database and data source tuning may have an effect on messaging throughput. The connection pool size for the data sources can be tuned for higher load levels, but the defaults are sufficient for most cases.

For general database tuning considerations, see [Section 2.6, "Tuning Database Parameters"](#).

Oracle B2B Performance Tuning

This chapter describes tips for tuning Oracle B2B performance. It contains the following sections:

- [Section 20.1, "About Oracle B2B"](#)
- [Section 20.2, "Basic Tuning Considerations"](#)

20.1 About Oracle B2B

Oracle B2B (Business to Business) is an e-commerce gateway that enables the secure and reliable exchange of business documents between an enterprise and its trading partners. Oracle B2B supports business-to-business document standards, security, transports, messaging services, and trading partner management. With Oracle B2B used as a binding component within an Oracle SOA Suite composite application, end-to-end business processes can be implemented.

For more information about Oracle SOA Suite, see *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

20.2 Basic Tuning Considerations

The following sections describe basic tuning configurations that you should also consider while tuning:

- [Tuning Data Storage Configurations for B2B](#)
- [Tuning MDS Cache Size](#)
- [Tuning Number of Threads](#)
- [Tuning the JMS Multiple Out Queues Setting](#)

20.2.1 Tuning Data Storage Configurations for B2B

Tuning the B2B_DATA_STORAGE table attributes can improve performance in some deployments. Review your use case scenarios and then consider the following to improve performance:

- If the B2B_DATA_STORAGE table logging attribute is set to NOLOGGING, consider changing it to LOGGING.

When performing DML operations using NOLOGGING option, the unrecoverable System Change Number (SCN) in the control files are recorded and this can lead to "control file contention" under heavy data load.

- Consider migrating LOB columns in the B2B_DATA_STORAGE table to SecureFiles
- Partition the table B2B_DATA_STORAGE by HASH(ID) with 128 partitions
- Partition the PK index by GLocal HASH(ID) with 64 partitions

20.2.2 Tuning MDS Cache Size

Changing the value of the Metadata Service (MDS) instance cache size can improve performance. A ratio of 5:1 is recommended for the `xxx-to-mdsCache` values. For example, if the `xxx` size is 1024, maintain `mdsCache` at 200 MB.

These settings can be modified using Oracle Enterprise Manager Fusion Middleware Control. For more information, see "Configuring Oracle B2B" in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

20.2.3 Tuning Number of Threads

Changing the value of `b2b.inboundThreadCount` and `b2b.outboundThreadCount` can improve Oracle B2B message processing. The recommended value depends on your system. For a 2 GB computer, for example, a setting of 3 to 5 is recommended. The `b2b.inboundThreadSleepTime` and `b2b.outboundThreadSleepTime` properties put a thread to sleep after message processing. A setting between 10 and 1000 (milliseconds) is recommended.

These settings can be modified using Oracle Enterprise Manager Fusion Middleware Control. For more information, see "Configuring Oracle B2B" in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

20.2.4 Tuning the JMS Multiple Out Queues Setting

The JMS Out Queue component is the element that enables B2B to receive data from a JMS queue. To maximize performance, consider enabling the Multiple JMSOUTQUEUES and create the corresponding listening channels in B2B.

Oracle Service Bus Performance Tuning

This chapter describes tips for tuning Oracle Service Bus performance. It contains the following sections:

- [Section 21.1, "About Oracle Service Bus"](#)
- [Section 21.2, "Monitoring Oracle Service Bus"](#)
- [Section 21.3, "Basic Tuning Considerations"](#)
- [Section 21.4, "Advanced Tuning Considerations"](#)

21.1 About Oracle Service Bus

Within a SOA framework, Oracle Service Bus (OSB) provides connectivity, routing, mediation, management and also some process orchestration capabilities. The design philosophy for OSB is to be a high performance and stateless (non-persistent state) intermediary between two or more applications. However, given the diversity in scale and functionality of SOA implementations, OSB applications are subject to large variety of usage patterns, message sizes and QOS requirements.

In most SOA deployments, OSB is part of a larger system where it plays the role of an intermediary between two or more applications (servers). A typical OSB configuration involves a client invoking an OSB proxy which may make one or more service callouts to intermediate back-end services and then route the request to the destination back end system before routing the response back to the client.

It is necessary, therefore, to understand that OSB is part of a larger system and the objective of tuning is the optimization of the overall system performance. This involves not only tuning OSB as a standalone application, but also using OSB to implement flow-control patterns such as throttling, request-buffering, caching, prioritization and parallelism.

For more information about Oracle Service Bus, see the *Oracle Fusion Middleware Administrator's Guide for Oracle Service Bus*.

21.2 Monitoring Oracle Service Bus

Though the out-of-the-box monitoring sub-system has a very low overhead, and scales well to numerous services and multiple nodes in a cluster, when dealing with thousands of services or a large scale cluster deployment, being selective about enabling monitoring can help reduce network traffic. When a business or proxy service is created, monitoring is disabled by default for that particular service. For more information, see "Configuring Operational Settings for Proxy Services" or "Configuring

Operational Settings for Business Services" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Service Bus*.

To enable or disable monitoring of all services that have individually been enabled or disabled for monitoring, use the "Enable Monitoring" option on the **Operations Global Settings** page. For more information, see "Enabling Global Settings" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Service Bus*.

21.3 Basic Tuning Considerations

Depending on your OSB usage and performance issues, you may consider tuning the following:

- [Tuning JVM Memory](#)
- [Tuning WebLogic Server for OSB](#)

21.3.1 Tuning JVM Memory

As with all other Oracle Fusion Middleware components, OSB performance can be impacted by JVM parameters. The two primary JVM tuning parameters to consider when optimizing OSB performance are heap size and garbage collection. For more information on tuning the JVM for performance, see [Section 2.4, "Tuning Java Virtual Machines \(JVMs\)"](#).

21.3.2 Tuning WebLogic Server for OSB

To optimize OSB, consider tuning the following WebLogic Server parameters:

21.3.2.1 Domain Mode

For production environments, create a domain in "Production" mode to maximize performance. The parameter is:

```
-Dweblogic.ProductionModeEnabled=true
```

To enable Weblogic server production mode through Weblogic Administration Console, see *Oracle Fusion Middleware Understanding Domain Configuration for Oracle WebLogic Server*.

21.3.2.2 WebLogic Server Logging Levels

For OSB performance testing and production environments, consider using the lowest acceptable logging level, such as "ERROR" or "WARNING" whenever possible. For more information, see [Section 2.10, "Setting Logging Levels"](#)

21.3.2.3 HTTP Access Logging

To optimize OSB performance, consider turning off the HTTP access logging. For more information, see [Section 6.3.3.1, "Access Logging"](#).

21.3.2.4 JMS Tuning

Ensure that the right persistence level is set for the Java Message Service (JMS) destinations. Consider the following scenarios:

- For non-persistent JMS scenarios:
 - Explicitly turn off persistence at the JMS server level by un-checking the "Store Enabled" flag from the Advanced section of the General tab for the JMS server on

the WebLogic Server console. It is also possible to override the persistence mode at the JMS destination level.

- For persistent JMS scenarios:

There are two choices: file store and JDBC store. Typically operations on a File Store perform better than JDBC store. If there are multiple JMS servers involved, create each store on a separate disk to lower I/O contention.

For more information on JMS Server Tunings, see "Tuning WebLogic JMS" in the *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*.

21.3.2.5 Connection Backlog Buffering

You can tune the number of connection requests that a WebLogic Server instance will accept before refusing additional requests. The Accept Backlog parameter specifies how many Transmission Control Protocol (TCP) connections can be buffered in a wait queue. This fixed-size queue is populated with requests for connections that the TCP stack has received, but the application has not accepted yet. This parameter should be tuned when dealing a large number of concurrent clients. For more information, see "Tuning Connection Backlog Buffering" in *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*.

21.3.3 Tuning OSB Operational Settings

This section discusses the following Oracle Service Bus operational settings:

- [OSB Tracing](#)
- [Cache Tuning for Proxy Service Runtime Data](#)
- [Initialize Router Cache for Proxy Service Runtime Data](#)

21.3.3.1 OSB Tracing

Oracle Service Bus has the option to trace messages without having to shutdown the server. This is an extremely useful feature both in a development and production environment for debugging, diagnosing and troubleshooting problems involving message flows in one or more proxy services.

Tracing is disabled by default but can be enabled on a per service basis. When tracing is enabled, the entire message context is also printed including headers and message body. It is important to realize its impact for large message sizes and high throughput scenarios.

For more information, see "How to Enable or Disable Tracing" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Service Bus*.

21.3.3.2 Cache Tuning for Proxy Service Runtime Data

OSB caches proxy service runtime meta-data using a two-level cache with static and dynamic sections. The cache introduces a performance tradeoff between memory consumption and compilation cost. Note that caching proxy services may help throughput but could impact memory usage.

The static section is an upper-bound Least Recently Used (LRU) cache that is never garbage collected. When a proxy service is bumped from the static section, it is demoted to the dynamic section where the cache can be garbage collected when there is memory pressure.

The number of proxy services in the static portion of the cache can be tuned by setting its size using the system property

`com.bea.wli.sb.pipeline.RouterRuntimeCache.size`. The default value is 100. This can be increased to a desired value provided there is sufficient memory for runtime data processing for large number of proxy services.

This property value can be set in the `setDomainEnv.sh` file as an extra java argument as follows:

```
-Dcom.bea.wli.sb.pipeline.RouterRuntimeCache.size={size}
```

Example:

```
EXTRA_JAVA_PROPERTIES="-Dcom.bea.wli.sb.pipeline.RouterRuntimeCache.size=3000  
${EXTRA_JAVA_PROPERTIES}"
```

To preload the router runtime cache with the compiled runtime of proxy services, set the system property, `com.bea.wli.sb.pipeline.RouterRuntimeCache.preload`, to `true`.

```
EXTRA_JAVA_PROPERTIES="-Dcom.bea.wli.sb.pipeline.RouterRuntimeCache.preload=true
```

Enabling this property may improve performance by decreasing the time it takes to make the initial calls to the proxy server. It can also preload the cache when a configuration session is committed.

If the number of proxy services exceeds 100, then the `com.bea.wli.sb.pipeline.RouterRuntimeCache.size` property should be used to set the appropriate cache size as the cache size is finite.

21.3.3.3 Initialize Router Cache for Proxy Service Runtime Data

The OSB system property `com.bea.wli.sb.pipeline.initializeRouterCache` initializes the `RouterRuntimeCache` for proxy services at server startup. Setting this property to `true` processes the first request to the proxy service faster since the runtime metadata model required for proxy service execution is already built at server startup. It is important to note, however, that this property can lead to longer server startup times due to the proxy service model initialization that is done.

To initialize the `RouterRuntimeCache` at server startup, set the system property as follows:

```
com.bea.wli.sb.pipeline.initializeRouterCache=true
```

21.4 Advanced Tuning Considerations

After you have performed the modifications recommended in the previous section, you can make additional changes that are specific to your deployment. Consider carefully whether the recommendations in this section are appropriate for your environment.

21.4.1 Transport Tuning (Oracle WebLogic Server and Oracle Service Bus)

Latency and throughput of poller-based transports depends on the frequency with which a source is polled and the number of files and messages read per polling sweep.

The following are the main transport configurations to tune:

21.4.1.1 Polling Interval

Consider using a smaller polling interval for high throughput scenarios where the message size is not very large and the CPU is not saturated. The primary polling interval defaults are listed below with links to additional information:

Polling Intervals	Default Interval	Additional Information
File Transport	60 seconds	"File Transport Configuration Page" in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Service Bus</i>
FTP Transports	60 seconds	"FTP Transport Configuration Page" in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Service Bus</i>
MQ Transport	1000 milliseconds	"MQ Transport Configuration Page" in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Service Bus</i>
SFTP Transport	60 seconds	"SFTP Transport Configuration Page" in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Service Bus</i>
JCA Transport	60 seconds	"JCA Transport Configuration Page" in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Service Bus</i> See also Section 18.3.1, "JCA Adapter Basic Tuning Considerations"

21.4.1.2 Read Limit

The read limit determines the number of files or messages that are read per polling sweep. This defaults to 10 for the File and FTP transports. It can be set to 0 to specify no limit. Set this value to the desired concurrency. For more information, see "File Transport Configuration Page" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Service Bus*.

Note: Setting the Read Limit to a high value and the Polling Interval to a small value may result in a large number of messages being simultaneously read into memory. This can lead to an OOM (out-of-memory error) if the message size is large.

21.4.2 Design Time Considerations for Proxy Applications

Consider the following design configurations for proxy applications based on your OSB usage and use case scenarios:

- Avoid creating many OSB context variables that are used just once within another XQuery

Context variables created using an Assign action are converted to XmlBeans and then reverted to the native XQuery format for the next XQuery. Multiple "Assign" actions can be collapsed into a single Assign action using a FLWOR expression. Intermediate values can be created using "let" statements. Avoiding redundant context variable creation eliminates overheads associated with internal data

format conversions. This benefit has to be balanced against visibility of the code and reuse of the variables.

- Transforming contents of a context variable such as `$body`.

Use a Replace action to complete the transformation in a single step. If the entire content of `$body` is to be replaced, leave the XPath field blank and select "Replace node contents". This is faster than pointing to the child node of `$body` (e.g. `$body/Order`) and selecting "Replace entire node". Leaving the XPath field blank eliminates an extra XQuery evaluation.
- Use `$body/*[1]` to represent the contents of `$body` as an input to a Transformation (XQuery / XSLT) resource.

OSB treats "`$body/*[1]`" as a special XPath that can be evaluated without invoking the XQuery engine. This is faster than specifying an absolute path pointing to the child of `$body`. A general XPath like "`$body/Order`" must be evaluated by the XQuery engine before the primary transformation resource is executed.
- Enable Streaming for pure Content-Based Routing scenarios.

Read-only scenarios such as Content-Based Routing can derive better performance from enabling streaming. OSB leverages the partial parsing capabilities of the XQuery engine when streaming is used in conjunction with indexed XPaths. Thus, the payload is parsed and processed only to the field referred to in the XPath. Other than partial parsing, an additional benefit for read-only scenarios is that streaming eliminates the overhead associated with parsing and serialization of XmlBeans.

The gains from streaming can be negated if the payload is accessed a large number of times for reading multiple fields. If all fields read are located in a single subsection of the XML document, a hybrid approach provides the best performance. See [Section 21.4.3, "Design Considerations for XQuery Tuning"](#) for additional details.

The output of a transformation is stored in a compressed buffer format either in memory or on disk. Therefore, streaming should be avoided when running out of memory is not a concern.
- Set the appropriate QOS level and transaction settings.

Do not set XA or Exactly-Once unless the reliability level required is once and only once and its possible to use the setting (it is not possible if the client is a HTTP client). If OSB initiates a transaction, it is possible to replace XA with LLR to achieve the same level of reliability.

OSB can invoke a back end HTTP service asynchronously if the QOS is "Best-Effort". Asynchronous invocation allows OSB to scale better with long running back-end services. It also allows Publish over HTTP to be truly fire-and-forget.
- Disable or delete all log actions.

Log actions add an I/O overhead. Logging also involves an XQuery evaluation which can be expensive. Writing to a single device (resource or directory) can also result in lock contentions.

21.4.3 Design Considerations for XQuery Tuning

OSB uses XQuery and XPath extensively for various actions like Assign, Replace, and Routing Table. The following XML structure (`$body`) is used to explain XQuery and XPath tuning concepts:


```

<soap-env:Body>
  <Order>
    <CtrlArea>
      <CustName>Mary</CustName>
    </CtrlArea>
    <ItemList>
      <Item name="ACE_Car" >20000 </Item>
      <Item name=" Ext_Warranty" >1500</Item>
      ... a large number of items
    </ItemList>
    <Summary>
      <Total>70000</Total>
      <Status>Shipped</Status>
      <Shipping>My Shipping Firm </Shipping>
    </Summary>
  </Order>
</soap-env:Body>

```

- Avoid the use of double front slashes ("/") in XPaths.

`$body//CustName` while returning the same value as `$body/Order/CtrlArea/CustName` will perform a lot worse than the latter expression. "/" implies all occurrences of a node irrespective of the location in an XML tree. Thus, the entire depth and breadth of the XML tree has to be searched for the pattern specified after a "/". Use "/" only if the exact location of a node is not known at design time.

- Index XPaths where applicable.

An XPath can be indexed by simply adding "[1]" after each node of the path. XQuery is a declarative language and an XPath can return more than one node; it can return an array of nodes. `$body/Order/CtrlArea/CustName` implies returning all instances `Order` under `$body` and all instances of `CtrlArea` under `Order`. Therefore, the entire document has to be read in order to correctly process the above XPath. If you know that there is a single instance of `Order` under `$body` and a single instance of `CtrlArea` under `Order`, we could rewrite the above XPath as `$body/Order[1]/CtrlArea[1]/CustName[1]`.

The second XPath implies returning the first instances of the child nodes. Thus, only the top part of the document needs to be processed by the XQuery engine resulting in better performance. Indexing is key to processing only what is needed.

Note: Indexing should not be used when the expected return value is an array of nodes. For example, `$body/Order[1]/ItemList[1]/Item` returns all "Item" nodes, but `$body/Order[1]/ItemList[1]/Item[1]` only returns the first item node. Another example is an XPath used to split a document in a "for" action.

- Extract frequently used parts of a large XML document as intermediate variables within a FLWOR expression

An intermediate variable can be used to store the common context for multiple values. Sample XPaths with common context:

```

$body/Order[1]/Summary[1]/Total, $body/Order[1]/Summary[1]/Status,
$body/Order[1]/Summary[1]/Shipping

```

The above XPaths can be changed to use an intermediate variable:

```

let $summary := $body/Order[1]/Summary[1]

```

```
$summary/Total, $ summary/Status, $summary/Shipping
```

Using intermediate variables consumes more memory but reduces redundant XPath processing.

- Using a Hybrid Approach for read-only scenarios with Streaming

The gains from streaming can be negated if the payload is accessed a large number of times for reading multiple fields. If all fields read are located in a single subsection of the XML document, a hybrid approach provides the best performance. The hybrid approach includes enabling streaming at the proxy level and Assigning the relevant subsection to a context variable, The individual fields can then be accessed from this context variable.

The fields "Total" and "Status" can be retrieved using three Assign actions:

```
Assign "$body/Order[1]/Summary[1]" to "foo"  
Assign "$foo/Total" to "total"  
Assign "$foo/Status" to "total"
```

Oracle Business Intelligence Performance Tuning

This chapter describes tips for tuning Oracle Business Intelligence performance. It contains the following sections:

- [Section 22.1, "About Oracle Business Intelligence"](#)
- [Section 22.2, "Oracle BI Server Query Performance Tuning"](#)
- [Section 22.3, "Oracle BI Server Query Cache Performance Tuning"](#)
- [Section 22.4, "Oracle BI Web Client Performance Tuning"](#)

22.1 About Oracle Business Intelligence

Oracle Business Intelligence (BI) Enterprise Edition (or Oracle Business Intelligence) provides a full range of business intelligence capabilities that collects up-to-date data from the organization, presents the data in easy-to-understand formats (such as tables and graphs), and delivers the data quickly to the members of the organization.

These capabilities enable the organization to make better decisions, take informed actions, and implement more-efficient business processes.

22.2 Oracle BI Server Query Performance Tuning

This section describes some important considerations for improving query performance with the Oracle BI Server.

For detailed information on BI performance tuning, see "Managing Performance Tuning and Query Caching" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

The following list summarizes methods that you can use to improve query performance:

- **Tuning and indexing underlying databases:** For Oracle BI Server database queries to return quickly, the underlying databases must be configured, tuned, and indexed correctly. Note that different database products have different tuning considerations.

If there are queries that return slowly from the underlying databases, then you can capture the SQL statements for the queries in the query log and provide them to the database administrator (DBA) for analysis. See "Managing the Query Log" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about configuring query logging on the system.

- **Aggregate tables:** It is extremely important to use aggregate tables to improve query performance. Aggregate tables contain precalculated summarizations of data. It is much faster to retrieve an answer from an aggregate table than to recompute the answer from thousands of rows of detail.

The Oracle BI Server uses aggregate tables automatically, if they have been properly specified in the repository. See *Oracle Fusion Middleware Metadata Repository Builder's Guide for Oracle Business Intelligence Enterprise Edition* for examples of setting up aggregate navigation.

- **Query caching:** The Oracle BI Server can store query results for reuse by subsequent queries. Query caching can dramatically improve the apparent performance of the system for users, particularly for commonly used dashboards, but it does not improve performance for most ad-hoc analysis.

See "About the Oracle BI Server Query Cache" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about query caching concepts and setup.

- **Setting parameters in Fusion Middleware Control:** You can set various performance configuration parameters using Fusion Middleware Control to improve system performance. See "Setting Performance Parameters in Fusion Middleware Control" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.
- **Setting parameters in NQSConfig.INI:** The NQSConfig.INI file contains additional configuration and tuning parameters for the Oracle BI Server, including parameters to configure disk space for temporary storage, set virtual table page sizes, and several other advanced configuration settings. See "NQSConfig.INI File Configuration Settings" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.

22.3 Oracle BI Server Query Cache Performance Tuning

You can configure the Oracle BI Server to maintain a local, disk-based cache of query result sets (query cache). The query cache allows the Oracle BI Server to satisfy many subsequent query requests without having to access back-end data sources (such as Oracle or DB2). This reduction in communication costs can dramatically decrease query response time. See "About the Oracle BI Server Query Cache" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

22.4 Oracle BI Web Client Performance Tuning

You can improve the performance of the Oracle BI web client (UI) by configuring your Web server to serve up all static files, as well as enabling compression for both static and dynamic resources. BI 11g ships with WebLogic Server (WLS) serving as the default HTTP server for the BI web client. By allowing the Oracle HTTP Server (OHS) to proxy requests to WLS instead, you may see an improvement in BI Web Client performance. See "Improving Oracle BI Web Client Performance" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Part V

Identity Management Suite Components

This part describes configuring Oracle Identity Management Suite components to improve performance. The Oracle Identity Management products enable you to configure and manage the identities of users, devices, and services across diverse servers, to delegate administration of these identities, and to provide end users with self-service privileges. These products also enable you to configure single sign-on across applications and to process users' credentials to ensure that only users with valid credentials can log into and access online resources.

It contains the following chapters:

- [Chapter 23, "Oracle Internet Directory Performance Tuning"](#)
- [Chapter 25, "Oracle Virtual Directory Performance Tuning"](#)
- [Chapter 26, "Oracle Identity Federation Performance Tuning"](#)
- [Chapter 24, "Oracle Unified Directory Performance Tuning"](#)
- [Chapter 27, "Oracle Fusion Middleware Security Performance Tuning"](#)

Oracle Internet Directory Performance Tuning

This chapter provides guidelines for tuning and sizing an Oracle Internet Directory installation. It contains these topics:

- [Section 23.1, "About Oracle Internet Directory"](#)
- [Section 23.2, "Monitoring Oracle Internet Directory Performance"](#)
- [Section 23.3, "Basic Tuning Considerations"](#)
- [Section 23.4, "Advanced Tuning Considerations"](#)
- [Section 23.5, "Specific Use Cases That Require Additional Tuning"](#)

23.1 About Oracle Internet Directory

Oracle Internet Directory is Oracle's Lightweight Directory Application Protocol (LDAP) version 3 Directory Server. Oracle Internet Directory is highly scalable, available, and manageable. It has a multi-threaded, multi-process, multi-instance process architecture with Oracle Database as the directory store. This unique physical architecture enables Oracle Internet Directory to be deployed on several hardware architectures including Symmetric Multi-Processor (SMP), Non-Uniform Memory Access (NUMA) and Cluster hardware. Oracle Internet Directory's physical architecture enables linear performance scalability with hardware resources and numerous high availability configurations.

For more information see *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

Note: Oracle Internet Directory's out of box configuration is not optimal for most production or test deployments. You must follow at least the steps listed in [Section 23.3, "Basic Tuning Considerations"](#) to achieve optimal performance and availability.

See Also:

- [Section 23.2.2, "Obtaining Recommendations by Using the Tuning and Sizing Wizard."](#)
- "Troubleshooting Oracle Internet Directory" in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*

Many of the recommendations in this chapter require changes to Oracle Internet Directory system configuration attributes and replication configuration attributes.

See Also:

- The "Managing System Configuration Attributes" chapter of *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*
- The "Managing Replication Configuration Attributes" chapter of *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*
- The "Attribute Reference" chapter of *Oracle Fusion Middleware Reference for Oracle Identity Management*

for more information about Oracle Internet Directory configuration attributes.

23.2 Monitoring Oracle Internet Directory Performance

To identify performance bottlenecks, you can monitor real-time performance metrics for the Oracle Internet Directory database. For more information on how to monitor other Oracle Fusion Middleware components, see [Chapter 4, "Monitoring Oracle Fusion Middleware"](#).

23.2.1 Monitoring Performance on UNIX and Windows Systems

Knowledge of the following tools is recommended for Linux, Solaris, and other UNIX-like operating systems:

Tool	Description
top	Displays the top CPU consumers on a system
vmstat	Shows running statistics on various parts of the system including the Virtual Memory Manager
mpstat	Shows an output similar to vmstat but split across various CPUs in the system. This is available on Solaris only.
iostat	Shows the disk I/O statistics from various disk controllers
sar	Collect, report, or save system activity information.

Knowledge of the following tools is recommended for Microsoft Windows:

Tool	Description
Windows Performance Monitor	Provides a customized view of the events in the system
Windows Task Manager	Provides a high level output (like top on UNIX) of the major things happening in the system.

Knowledge of the following tools is recommended for the Oracle Database:

- `utlbstat.sql` and `utlestat.sql`, or `statspack`
- Automatic Workload Repository (AWR) and Automatic Database Diagnostic Monitor (ADDM) reports
- Collecting optimizer statistics using `DBMS_STATS` package

See Also:

- *Oracle Database Reference* in the Oracle Database Documentation Library for information about `utl1bstat.sql` and `utlestat.sql`
- *Oracle Database Performance Tuning Guide* for information about stats package
- *Oracle Database Concepts* in the Oracle Database Documentation Library for information about the ANALYZE function in the DBMS_STATS package

In addition to the operating system tools, the LDAP applications being used in a customer environment must be able to provide latency and throughput measurement.

In addition, the Database Statistics Collection Tool (`oidstats.sql`), located at `$ORACLE_HOME/ldap/admin`, is provided to analyze the various database 'ods' schema objects to estimate the statistics. See [Section 23.2.3, "Updating Database Statistics by Using `oidstats.sql`"](#).

23.2.2 Obtaining Recommendations by Using the Tuning and Sizing Wizard

Oracle Enterprise Manager Fusion Middleware Control provides a convenient tool for tuning and sizing Oracle Internet Directory.

Use the wizard to obtain tuning and sizing recommendations for your system. You can select **Tuning**, **Sizing**, or **Both**. If you select **Sizing** or **Both**, you can select **Basic** or **Advanced**

Tuning

1. From the Oracle Internet Directory menu, select **Administration**, then **Tuning and Sizing**.
2. Click the **Create** icon to invoke the wizard.
3. On the Type Selection page, change the report name, then select **Tuning**.
4. The wizard presents the following pages: Hardware, Features, Load, Data Characteristics, and Garbage Collection.

On each page, specify values for the text fields (or use defaults) and Select **Yes** or **No** for each question. Some choices might be greyed out, depending upon your previous choices. Most fields have tool tips that appear when you move the cursor over the field.

Click **Next** to go to the next page or **Back** to return to the previous page. Click **Cancel** to close the wizard.

5. On the Review page, review the data you entered. Click **Back** to change your specifications or click **Finish** to view the report.
6. The report appears on the bottom right section of the page.
To download the report, click **Download Report**. To delete the report, click **Delete**.

Sizing

1. From the Oracle Internet Directory menu, change the report name, then select **Administration**, then **Tuning and Sizing**.
2. Click the **Create** icon to invoke the wizard.
3. On the Type Selection page, select **Sizing**.

4. Select **Basic** or **Advanced**.
5. On the Sizing page, specify values for the text fields (or use defaults) and Select Yes or No for each question. Some choices might be greyed out, depending upon your previous choices.
6. Click **Next**.
7. On the Review page, review the data you entered. Click **Back** to change your specifications or click **Finish** to view the report.
8. The report appears on the bottom right section of the page.
To download the report, click **Download Report**. To delete the report, click **Delete**.

Both

1. From the Oracle Internet Directory menu, change the report name, then select **Administration**, then **Tuning and Sizing**.
2. Click the **Create** icon to invoke the wizard.
3. On the Type Selection page, select **Both**.
4. Select **Basic** or **Advanced**.
5. Click **Next**.
6. The wizard presents the following pages: Sizing, Hardware, Features, Load, Data Characteristics, and Garbage Collection.

On each page, specify values for the text fields (or use defaults) and Select **Yes** or **No** for each question. Some choices might be greyed out, depending upon your previous choices.

Click **Next** to go to the next page or **Back** to return to the previous page. Click **Cancel** to close the wizard.
7. On the Review page, review the data you entered. Click **Back** to change your specifications or click **Finish** to view the report.
8. The report appears on the bottom right section of the page.
To download the report, click **Download Report**. To delete the report, click **Delete**.

23.2.3 Updating Database Statistics by Using oidstats.sql

Database statistics are updated automatically, OIDMON runs `oidstats.sql` for every configured number of updates to the database. By default, for every 5000 entries added OIDMON runs the `oidstats.sql`. This frequency can be changed using `ldapmodify` command as shown below

```
$ORACLE_HOME/bin/ldapmodify -p <oidPort> -h <oidHost> -D cn=orcladmin -w  
<adminPassword> << eof  
dn: cn=configset,cn=oidmon,cn=subconfigsubentry  
changetype: modify  
replace: orclstatsperiodicity  
orclstatsperiodicity: <desired_number>  
eof
```

See Also: The `oidstats.sql` command-line tool reference in *Oracle Fusion Middleware Reference for Oracle Identity Management*

23.2.4 Setting Performance-Related Replication Configuration Attributes

To set the replication attributes, you can use either the Replication Wizard in Oracle Enterprise Manager Fusion Middleware Control or the command line.

The attributes `orclthreadspersupplier`, `orclchangeretrycount`, and `orclconflresolution` are replication configuration set attributes.

See Also:

- "Configure Replication Attributes by Using Fusion Middleware Control" in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*
- "Configuring Attributes of the Replication Configuration Set by Using `ldapmodify`" in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*

for information about

The attributes `orclhiqschedule` and `orclupdateschedule` are replication agreement entry attributes.

See Also:

- "Viewing or Modifying an LDAP-Based Replication Setup by Using the Fusion Middleware Control Replication Wizard" in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*
- "Configuring Replication Agreement Attributes by Using `ldapmodify`" in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*

See Also:

- "Setting Up a One-Way, Two-Way, or Multimaster LDAP-Based Replication Agreement by Using the Replication Wizard in Fusion Middleware Control" in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory* or information on setting replication attributes by using the Replication Wizard.
- "Configuring Attributes of the Replication Configuration Set by Using `ldapmodify`" in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

23.2.5 Managing System Configuration Attributes

You can set most performance-related system configuration attributes from Oracle Enterprise Manager Fusion Middleware Control or from the command line. You can also use the Data Browser in Oracle Directory Services Manager to modify system configuration attributes.

For information on setting system configuration attributes for Oracle Internet Directory, see "Managing System Configuration Attributes" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*:

- "Managing System Configuration Attributes by Using Fusion Middleware Control"

- "Managing System Configuration Attributes by Using WLST"
- "Managing System Configuration Attributes by Using LDAP Tools"
- "Managing System Configuration Attributes by Using ODSM Data Browser"

23.2.6 Setting Garbage Collection Configuration Attributes

The attributes `orclpurgetargetage` and `orclpurgeinterval` reside in the changelog purging configuration entry. You can change them with `ldapmodify` or Oracle Directory Services Manager.

23.2.6.1 Modifying Changelog Purging Attributes by Using `ldapmodify`

The following example is an LDIF file used to configure change log purging.

See Also: "Change Log Purging" in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory* for a description of change log purging.

This example configures time-based purging for 120 hours (5 days). Use an LDIF file similar to this:

```
dn: cn=changelog_purgeconfig,cn=purgeconfig,cn=subconfigsubentry
changetype:modify
replace: orclpurgetargetage
orclpurgetargetage: 240
```

To apply the LDIF file `mod.ldif`, type:

```
ldapmodify -D "cn=orcladmin" -q -p port -h host -D dn -q -f mod.ldif
```

See Also: "Configuring Time-Based Change Log Purging" in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

23.2.6.2 Modifying Changelog Purging in Oracle Directory Services Manager

You can modify `orclpurgetargetage` and `orclpurgeinterval` by using the data browser in Oracle Directory Services Manager. You cannot navigate to the changelog purging configuration entry directly in the data tree, but you can get to it by using an advanced search as follows:

1. On the Data Browser tab, click **Advanced**.
2. Expand **Garbage Collection** in the left pane, then select **changelog purgeconfig**. The Garbage Collector Window appears in the right pane.
3. In the right pane, enter the changes you want to make to the **Purge Target Age** and **Purge Interval**.
4. Choose **Apply**.

23.3 Basic Tuning Considerations

Tuning is the adjustment of parameters to improve directory performance. The default Oracle Internet Directory configuration must be tuned in almost all deployments. Please review the requirements and recommendations in this section carefully.

23.3.1 Database Parameters

The suggested minimum values for Oracle Database instance parameters are described in [Table 23-1](#):

Table 23-1 Minimum Values for Oracle Database Instance Parameters

Parameter	Value	Notes
sga_target and sga_max_size	1700M for 32-bit systems	Applicable when SGA Auto Tuning using <code>sga_target</code> and <code>sga_max_size</code> is being used. Especially important for <code>bulkdelete</code> performance. A higher value may be required if the directory size exceeds 1 million entries or a high rate of I/O is observed. In case of 64-bit systems, one can go up to 60-70% of the RAM available for the Oracle Database on the box.
db_cache_size	1200M for 32-bit systems.	Applicable when SGA Auto Tuning using <code>sga_target</code> and <code>sga_max_size</code> is not being used. (SGA auto tuning using <code>sga_target</code> and <code>sga_max_size</code> is recommended instead of this parameter.) A higher value may be required if the directory size exceeds 1 million entries or a high rate of I/O is observed. In case of 64-bit systems, one can go up to 60-70% of the RAM available for the Oracle Database on the box.
shared_pool_size	300M	Applicable when SGA Auto Tuning using <code>sga_target</code> and <code>sga_maxsize</code> is not being used
session_cached_cursors	500	
processes	500	
memory_target	0	Consider setting this to 0 and set only <code>sga_target</code> and <code>pga_aggregate_target</code> as described.
pga_aggregate_target	1 - 4 GB	Set this to 1-4 GB, if sufficient RAM is available. Increase the size if you see a pattern of increased PGA usage. This is especially useful when running large bulk loads.
job_queue_processes	1 or more.	Tune this parameter only if you are using Oracle Database Advanced Replication-based multimaster replication. NOTE: This parameter should NOT exceed the number of CPU cores available on the database server.
_b_tree_bitmap_plans	FALSE	The default value of TRUE for this parameter could cause unwanted bitmap operations in the execution plan of SQL which may impact performance.
NLS_SORT	BINARY	Consider setting <code>NLS_SORT</code> to <code>BINARY</code> , otherwise sort will do full table scan and performance can be impacted.
PL_SQL_CODE_TYPE	NATIVE	Consider setting <code>PL_SQL_CODE_TYPE</code> to <code>NATIVE</code> .

See the *Oracle Database Performance Tuning Guide* for information on setting Oracle Database instance parameters.

23.3.2 LDAP Server Attributes

The recommendations in this section are summarized in [Table 23-2](#).

- Tune the number of processes and threads for the Oracle Internet Directory server instance that services LDAP application traffic. This has a major impact on overall performance. See the recommended setting for `orclmaxcc` in [Table 23–2](#).
- Disable change log generation if you are not deploying either replication or Oracle Directory Integration Platform. Set the attribute `orclgeneratechangelog` to 0.
- Skip referrals in LDAP searches if you have no referral entries in the directory. Set `orclskiprefinsql` to 1. This can have a major impact on performance.
- Close idle LDAP connections after a period of time instead of leaving them open. This prevents the unnecessary buildup of connections. For example, you can set `orclldapconntimeout` to 60 minutes.

As of 10g (10.1.4.0.1), you can only set this for users who are not configured for operation statistics tracking. Connections by users configured for statistics collection do not time out as per this setting.

See Also: "Configuring a User for Statistics Collection by Using Fusion Middleware Control" in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

- If no clients require detailed MatchDN information when the Base DN of an LDAP search operation is not present in the directory, disable it. Change `orclmatchdnenabled` to 0.

The following values are appropriate for most deployments:

Table 23–2 LDAP Server Attributes to Tune

Attribute	Default	Recommended Value	Notes
<code>orclserverprocs</code>	1	Number of CPU cores	When setting the value, use only CPU cores and not CPU threads.
<code>orclmaxcc</code>	2	10	Server restart required.
<code>orclskiprefinsql</code>	0	1	This change is highly recommended. Do not change if you have LDAP referral entries. LDAP referral entries are not common. Server restart required.
<code>orclgeneratechangelog</code>	1	0	Disable change log generation only if you do not deploy either replication or Oracle Directory Integration Platform.
<code>orclldapconntimeout</code>	0 (no timeout)	Varies, 60 minutes is reasonable	Users configured for statistics tracking do not time out.
<code>orclmatchdnenabled</code>	1	0	Disable only if no application needs detailed MatchDN information when base DN of a search is not present.

For information about configuring `orclserverprocs`, `orclldapconntimeout`, and `orclmatchdnenabled` with Oracle Enterprise Manager Fusion Middleware Control, see "Attributes of the Instance-Specific Configuration Entry" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

For information about configuring `orclskiprefinsql` or `orclmatchdnenabled` with Oracle Enterprise Manager Fusion Middleware Control, see "Configuring Shared

Properties" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

For information about configuring these attributes, as well as `orclgeneratechangelog`, from the command line, see "Setting System Configuration Attributes by Using `ldapmodify`" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

23.3.3 Database Statistics

If you use LDAP commands to add a large number of entries to Oracle Internet Directory, it can affect directory performance. If this occurs, update the database statistics. See [Section 23.2.3, "Updating Database Statistics by Using `oidstats.sql`."](#)

Typically, you only need to do this when you add entries in bulk for the first time after installing Oracle Internet Directory. You do not need to do it again because the database statistics are updated nightly automatically. If, however, you suddenly experience slow LDAP operations, without a corresponding change in data footprint, consider running `oidstats.sql` once to see if that improves performance. The impact may be due to changes in database SQL execution plans, which `oidstats.sql` can help to improve.

See Also: *Oracle Database Performance Tuning Guide* for information about SQL tuning.

You do not need to update database statistics if you use the `bulkload` tool to add the entries. The `bulkload` command automatically updates the database statistics.

23.3.4 Low-Priority Tuning Considerations

This section describes attributes that can sometimes improve performance, but are considered low-priority.

23.3.4.1 Controlling the Number of Entries to be Returned by a Search

- The attribute `orclsizeLimit` controls the maximum number of entries to be returned by a search. The default value is 10000. Setting it very high impacts server performance. It also plays a role in limiting the maximum number of changelogs the replication server can process at a time.
- The multivalued attributes of `orclsimplemodchglogattributes` control the attributes for which change logs contain only changes, not lists of all values. Limiting the return to changes only can sometimes improve performance.

For more information, see "Change Logs in Directory Replication" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

23.3.4.2 Enabling the Group Cache

The instance-specific subentry attribute `orclenablegroupcache` controls whether privilege groups and ACL groups are cached. Using this cache can improve the performance of access control evaluation for users.

Use the group cache when a privilege group membership does not change frequently. If a privilege group membership does change frequently, then it is best to turn off the group cache. It is important to note that computing a group cache may affect performance. The default is 1 (enabled). Change to 0 (zero) to disable.

See "Setting System Configuration Attributes by Using ldapmodify" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

23.3.4.3 Timeout for Write Operations

When an LDAP client initiates an operation, then does not respond to the server for a configured number of seconds, the server closes the connection. The number of seconds is controlled by the `orclnwrtimeout` attribute of the instance-specific configuration entry. The default is 30 seconds.

You can modify `orclnwrtimeout` by using Fusion Middleware Control or the command line. See "Attributes of the Instance-Specific Configuration Entry" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

23.4 Advanced Tuning Considerations

After you have performed the modifications recommended in the previous section, you can make additional changes that are specific to your deployment. Consider carefully whether the recommendations in this section are appropriate for your environment.

- [Replication or Oracle Directory Integration Platform](#)
- [Replication Server Configuration](#)
- [Garbage Collection Configuration](#)
- [Oracle Internet Directory with Cluster Configuration](#)
- [Password Policies and Verifier Profiles](#)
- [Server Entry Cache](#)
- [Result Set Cache](#)
- [Tuning Security Event Tracking](#)
- [Optimizing Searches](#)

23.4.1 Replication or Oracle Directory Integration Platform

When you deploy Oracle Internet Directory with the Oracle Directory Integration Platform or with replication, you can improve performance by having a dedicated LDAP server instance for those two servers. This allows the default Oracle Internet Directory LDAP instance to serve the LDAP application traffic and the second instance to serve LDAP requests from the replication and Oracle Directory Integration Platform servers.

1. Create an additional server instance, as described in the chapter "Managing Oracle Internet Directory Instances" in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.
2. Set `orclmaxcc` to 10 and `orclserverprocs` to 1 in the new instance configuration.
3. Restart the server, as described in the chapter "Managing Oracle Internet Directory Instances" in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.
4. Set the SSL and non-SSL ports used by the new instance and configure the replication and Oracle Directory Integration Platform to point to them.

To configure `orclmaxcc` and `orclserverprocs`, see "Attributes of the Instance-Specific Configuration Entry" in the *Oracle Fusion Middleware Administrator's Guide for Oracle*

Internet Directory and "Setting System Configuration Attributes by Using ldapmodify" in the Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory.

Note: In an Oracle Internet Directory Cluster configuration (rack-mounted or multi-box), the replication server must be started on one hardware node only. The LDAP server instance dedicated to replication must be started on the same node. The Oracle Directory Integration Platform server can be on a different node.

23.4.2 Replication Server Configuration

The following recommendations can be useful when replication traffic is heavy. Be sure you understand the trade-offs before making these changes. The recommended values are summarized in [Table 23–3](#).

- If you are deploying a single master with read-only replica consumers, you may reduce performance impacts by turning off conflict resolution. To do so, change the value of `orclconflresolution` to 0.
- If the supplier is a bottleneck, increase `orclthreadspersupplier` on the supplier. You can also increase `orclthreadspersupplier` at the consumer if it is a bottleneck, but be aware that increased parallelism causes race conditions in the application of changelogs, resulting in more human intervention queue (HIQ) changes.
- Decrease `orclchangeretrycount` so that new changelogs get more resources. If there are conflicts, however, this increases the human intervention queue (HIQ) changes.
- Change `orclupdateschedule` to 0 to make the server process changelogs immediately, instead of at the default, 60-second intervals. Do this on both the supplier and consumer.
- Increase the `orclhiqschedule` to a higher value. For example, if accessing the human intervention queue (HIQ) four times a day is sufficient and appropriate for your deployment, set the `orclhiqschedule` to 21600 seconds (6 hours).

[Table 23–3](#) summarizes these recommendations.

Table 23–3 Replication Attributes

Attribute	Default	Recommended Value	Notes
<code>orclthreadspersupplier</code>	<code>transport=1 apply=5</code>	Set transport threads to 1 and apply threads to 10 or greater	Most useful if the supplier is the bottleneck.
<code>orclchangeretrycount</code>	10	4	Provides more resources to changelogs but might increase HIQ.
<code>orclupdateschedule</code>	60 seconds	0	Causes changelogs to be processed immediately
<code>orclhiqschedule</code>	600 seconds	21600 seconds	Provides more resources to process new changes.
<code>orclconflresolution</code>	1	0	Change only if you are deploying a single master with read-only replica consumers.

See [Section 23.2.4, "Setting Performance-Related Replication Configuration Attributes"](#) for information on setting these replication attributes.

23.4.3 Garbage Collection Configuration

By default, Oracle Internet Directory runs database jobs to purge change logs, server manageability statistics, and other data beginning at midnight, with each job starting 15 minutes after the previous one. You can change this configuration to suite your deployment needs by modifying the parameters shown in [Table 23–4](#).

Table 23–4 Garbage Collection Configuration Parameters

Parameter	Value	Notes
orclpurgetargetage	Less than 10days (240 hours)	Only if there is no requirement to retain change logs
orclpurgeinterval	6–12 hours	

You can modify these attributes by using `ldapmodify` or Oracle Directory Services Manager. See [Section 23.2.6, "Setting Garbage Collection Configuration Attributes."](#)

23.4.4 Oracle Internet Directory with Cluster Configuration

As described in [Section 23.4.2, "Replication Server Configuration"](#), you can have a dedicated LDAP server for Oracle Directory Integration Platform and replication, in addition to the default server. In an Oracle Internet Directory Cluster, start the default LDAP instance on all Oracle Internet Directory nodes, but start the dedicated instance only on the node where Oracle Directory Integration Platform and replication are running.

To maintain optimal performance, consider carefully which database instance Oracle Internet Directory should connect to:

- You can configure the Oracle Internet Directory for load balancing between Oracle Database instances in the cluster, or failover mode.
- Use a dedicated IP address and Oracle Internet Directory instance for notifications in a cluster. The `orclcachenotifyip` configuration attribute can be added to the `ldiffile` to indicate the port number and IP address to use. Note, however, that if `orclcachenotifyip` is configured for an Oracle Internet Directory instance, the IP address must be local to the node where that instance is running.

For more information, see "Configuring IP Addresses for Notifications in a Cluster" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*

- If you use a dedicated LDAP server instance for replication and Oracle Directory Integration Platform, you can configure the connection strings of that instance for failover. You would use the following in `tnsnames.ora`:

```
(FAILOVER=ON) (LOAD_BALANCE=OFF)
```

- When performing a bulk operation, such as `bulkload`, connect the tool to just one Oracle Database instance for the entire operation.
- Configure Oracle Internet Directory instances as follows:
 - One Oracle Internet Directory instance on each of the nodes to service LDAP application traffic

- An instance of the Oracle Internet Directory replication server and Oracle Directory Integration Platform server on one node

23.4.5 Password Policies and Verifier Profiles

Oracle Internet Directory has password policies and password verifier profiles enabled out of box. If Oracle Internet Directory is not required to enforce password policies in a given deployment, then the password policies can be disabled. The password verifier profiles enabled out of box control the generation of certain password verifiers required by Oracle products like Enterprise User Security and Oracle Collaboration Suite. If Oracle Internet Directory is not being deployed for other Oracle products, you can disable all the password verifier profiles.

You can disable password policies and password verifiers by using Oracle Directory Services Manager or `ldapmodify`.

See Also:

- The "Managing Password Policies" chapter in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.
- The "Managing Password Verifiers" chapter in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

23.4.6 Server Entry Cache

The Oracle Internet Directory server cache enables LDAP entries to be cached on Oracle Internet Directory SHARED MEMORY, all the Oracle Internet Directory process running on the node share the cache for better performance.

23.4.6.1 Benefits of Using the Entry Cache

Oracle Internet Directory has been enhanced to support entry cache in cluster configuration starting from release 11.1.1.6.0. One of the key benefits of using the entry cache is that LDAP search operation is 4 times fast. This applies only when all or most entries can be cached.

23.4.6.2 Values for Configuring the Entry Cache

You can configure and optimize the server entry cache by setting the values shown in [Table 23–5](#).

Table 23–5 Server Entry Cache Configuration

Attribute	Default	Recommended Value	Notes
<code>orclmaxcc</code>	2	10	Restart the server after changing this attribute.
<code>orclicacheenabled</code>	1	2	
<code>orclicachemaxsize</code>	200000000 Bytes	Total size of the directory, in bytes	To determine the optimal setting for this attribute, use the number of entries in the Directory Information Tree and multiply by the average entry size. Estimate three times the size of the entries in LDIF format.
<code>orclicachemaxentries</code>	100000	Total number of entries in the DIT	

For example, if the total size of the Directory Information Tree is 300K and the total size of 300K entries in LDAP Data Interchange Files (LDIF) format is 500M, you would set `orclectacheenabled` to 1, `orclectachemaxsize` to 1,500,000,000, and `orclectachemaxentries` to 300,000.

To obtain the number of entries in the Directory Information Tree, use the following command:

```
sqlplus ods@oiddb
select count(*) from ct_dn;
```

```
oidctl connect=oiddb status -diag
```

The following example shows the `oidctl connect=oiddb status -diag` command output:

```

+-----+
| Process      | PID  | InstName | CompName | Inst# | Port | Sport |
+-----+
| oidmon       | 8192 | inst1    | oid1     | 0     |      |        |
+-----+
| oidldapd disp| 8201 | inst1    | oid1     | 1     | 5678 | 0     |
| oidldapd serv| 8205 | inst1    | oid1     | 1     | 5678 | 0     |
| oidldapd serv| 8209 | inst1    | oid1     | 1     | 5678 | 0     |
| oidldapd serv| 8213 | inst1    | oid1     | 1     | 5678 | 0     |
| oidldapd serv| 8217 | inst1    | oid1     | 1     | 5678 | 0     |
| Config DN   | cn=oid1,cn=osldapd,cn=subconfigsubentry
+-----+

```

```

+-----+
|Printing LDAP Operation in progress status ...
+-----+
OIDLDAPD_PID: 8205 WorkerID: 8 DBSID: 168 DBPID: 8245 ==> IDLE
+-----+
OIDLDAPD_PID: 8205 WorkerID: 9 DBSID: 170 DBPID: 8253 ==> IDLE
+-----+
OIDLDAPD_PID: 8205 WorkerID: 10 DBSID: 180 DBPID: 8261 ==> IDLE
+-----+
OIDLDAPD_PID: 8205 WorkerID: 11 DBSID: 189 DBPID: 8269 ==> IDLE
+-----+
OIDLDAPD_PID: 8209 WorkerID: 13 DBSID: 171 DBPID: 8249 ==> IDLE
+-----+
OIDLDAPD_PID: 8209 WorkerID: 9 DBSID: 181 DBPID: 8257 ==> IDLE
+-----+
OIDLDAPD_PID: 8209 WorkerID: 12 DBSID: 193 DBPID: 8267 ==> IDLE
+-----+
OIDLDAPD_PID: 8209 WorkerID: 10 DBSID: 199 DBPID: 8225 ==> IDLE
+-----+
OIDLDAPD_PID: 8209 WorkerID: 11 DBSID: 190 DBPID: 8227 ==> IDLE
+-----+
OIDLDAPD_PID: 8205 WorkerID: 13 DBSID: 197 DBPID: 8223 ==> IDLE
+-----+
OIDLDAPD_PID: 8205 WorkerID: 12 DBSID: 182 DBPID: 8229 ==> IDLE
+-----+

```

```

Cache Max Size : 96000001024
Max Entries configured : 2600000
Max Entries cached : 3000000
Num Entries in Cache : 3000000
Num Entries in GC : 6000000

```

```

Page size : 46062504
Entry cache Hit count : 164805647431
Entry cache Mis count : 8981856
Hash Area bytes used : 39904896
Hash Area blocks used : 22
Free pages : 944
Freed Blks : 18008766
Entry cache bytes used : 47468777224
Entry cache blocks used : 675008984
Cache memory used : 47508682120

```

The `-opdiag` option for `OIDCTL` displays the total number of LDAP operations performed in 5 second intervals (similar to the `vmstat/iostat`).

```
oidctl connect=oiddb status -opdiag
```

To configure the attributes, see "Attributes of the Instance-Specific Configuration Entry" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory* and "Setting System Configuration Attributes by Using `ldapmodify`" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

23.4.7 Result Set Cache

Result set cache is an Oracle 11g OID feature that allows complete result sets to be stored in memory. If a SQL query is executed and its result set is in the cache then almost the entire overhead of the SQL execution is avoided: this includes parse time, logical reads, physical reads and any cache contention overhead (latches for instance) that might normally be incurred. Configuring the result cache can improve performance since most LDAP applications typically look up user entry such as `mail=john.doe@acme.com` or `uid=john.doe` from a user tree. Such queries are repeated by the application every time a user logs in or uses the application. The result set of such queries may be a single entry. Performance may be affected as OID makes a trip to the database for the entry each time the query is run.

23.4.7.1 When to Use Result Set Cache

Consider using Result Set Cache only under the following conditions:

- Filter matches one or few entries.
- SQL statement causes multiple reads from disk or buffer (expensive)

23.4.7.2 Benefits of Using Result Set Cache

Benefits of using the entry cache include:

- OID evaluates the filter without making a trip to the database and therefore reduces the load on the database.
- Performance may be improved by 3 to 5 times when compared to performance when result set cache is not used.

23.4.7.3 Configuring Result Set Cache

The `OrclRSCacheAttr` attribute is used to configure the result set cache for OID. `OrclRSCacheAttr` is a multi-valued attribute that includes `cn`, `mail`, `uid`, `orclguid`. Typically these attributes are not modified for the life of the entry.

To enable result set cache, set `orclcacheenabled=2`. Result set cache can be turned off by setting `orclcacheenabled=1` or `orclcacheenabled=0`.

Note that any change to the following configuration attributes requires a restart of OID server (all the instances).

23.4.8 Tuning Security Event Tracking

The instance-specific configuration entry attributes `orcloptrackmaxtotalsize` and `orcloptracknumelemcontainers` control how much memory is used for security event tracking.

The attribute `orcloptrackmaxtotalsize` specifies the maximum number of bytes of RAM that security events tracking can use for each type of operation. If the Directory Server exceeds this limit for information collected for an operation, the server stops collecting new information and records appropriate messages in server log files. For the compare operation, the Directory Server uses twice the value of the attribute, which is the combined amount of information about users performing compare operation and users whose passwords are being compared. The default value of `orcloptrackmaxtotalsize` is 100000000 Bytes, which should be sufficient for most deployments. It can be increased to 200MB. For information about modifying `orcloptrackmaxtotalsize`, see the instance-specific configuration attribute examples in "Setting System Configuration Attributes by Using `ldapmodify`" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

The attribute `orcloptracknumelemcontainers` allows you to choose the number of in-memory cache containers to be allocated for security event tracking in the Oracle Internet Directory server. There are two subtypes for this attribute. They are `1stlevel` and `2ndlevel`. The `1stlevel` subtype is for setting the number of in-memory cache containers for storing information about users performing operations. The `2ndlevel` subtype, which is applicable only to compare operation, sets the number of in-memory cache containers for information about the users whose user password is compared and tracked when detailed compare operation statistics is programmed.

The default value of both subtypes is 256. The appropriate values for these subtypes depend on the number of users in your environment and the number of applications used to access the directory, as follows:

- In a deployment where several applications perform operations on behalf of a large number of end users, set `1stlevel` proportional to the number of applications, plus a few hundred more for end users directly accessing the directory. Then set `2ndlevel` proportional to the number of end users.
- In a deployment where end users themselves perform the operations, set `1stlevel` proportional to the number of end users, then set `2ndlevel` to a small value, such as 25.
- A typical proportional value is one fifth. Proportions between one tenth and one half are reasonable in most environments.

If your deployment requires it, set the values for `orcloptracknumelemcontainers` only when security events collection is turned on.

23.4.9 Optimizing Searches

This section contains these topics:

- [Section 23.4.9.1, "Optimizing Searches for Large Group Entries"](#)
- [Section 23.4.9.2, "Optimizing Searches for Skewed Attributes"](#)
- [Section 23.4.9.3, "Optimizing Performance of Complex Search Filters"](#)

23.4.9.1 Optimizing Searches for Large Group Entries

Searches for group entries with several thousand attribute values for either the `member` or `uniquemember` attribute can have high latency. If you find the latency unacceptably high, there are steps you can take to reduce it.

The simplest step is to reduce the number of attributes you are searching for. If you do not need to retrieve all the attributes of the group entry, specify required attributes in the search request to optimize the latency.

23.4.9.1.1 Entry Cache Enabled Configuration If you still see unacceptable latency, even with required attributes specified, then you can try to cache the large group entry in the entry cache:

```
cn=componentname,cn=osldldapd,cn=subconfigsubentry
```

Note: If you expect frequent updates to large groups, then do not use this tuning methodology. Use the Entry Cache Disabled Configuration.

23.4.9.1.2 Entry Cache Disabled Configuration. No action is required. This configuration is enabled by default.

23.4.9.2 Optimizing Searches for Skewed Attributes

To service a typical search request, the Directory Server sends a SQL statement to the Oracle Database. If a given attribute has very different response times depending on its value, then the attribute is said to be skewed. For example, if searches for `my_attribute=value1` and `my_attribute=value2` have very different response times, then `my_attribute` is said to be a skewed.

You can uniform the response times for searches for such an attribute by adding it as a value of the `orclskewedattribute` attribute, which is in the DSA configuration entry. The DN of the DSA configuration entry is

```
cn=dsaconfig,cn=configsets,cn=oracle internet directory
```

By default, the `objectclass` attribute is listed as a value in the `orclskewedattribute` attribute.

You can change the value of `orclskewedattribute` by using `orldapmodify`. See "Attributes of the Instance-Specific Configuration Entry" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory* and "Setting System Configuration Attributes by Using `ldapmodify`" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

23.4.9.3 Optimizing Performance of Complex Search Filters

When Oracle Internet Directory receives an LDAP search filter from a client application, it sends the filter to the Oracle Database as an SQL query. Sometimes client applications send filters that include terms that match a large number of entries in the directory. For example, consider the following filter:

```
(&(uid=msmith)(objectclass=inetorgperson)(orclisenabled=TRUE))
```

The terms `(objectclass=inetorgperson)` and `(orclisenabled=TRUE)` in that filter match nearly all entries. It would be very resource-intensive to execute that entire filter in the Oracle Database. To improve performance, you can specify that Oracle Internet

Directory execute a portion of that filter in its own memory, rather than in the database. To do that, you use `orclinmemfiltprocess`, an attribute in the DSA configuration entry:

```
cn=dsaconfig,cn=configsets,cn=oracle internet directory
```

When `orclinmemfiltprocess` is configured, the following events occur each time Oracle Internet Directory receives an LDAP search:

1. Oracle Internet Directory removes all the terms that are configured in the `orclinmemfiltprocess` before forming the SQL query.
2. Oracle Internet Directory sends the SQL query to Oracle Database.
3. Oracle Database sends the entries resulting from the SQL query to Oracle Internet Directory.
4. Oracle Internet Directory applies the original filter sent by the client (the terms in `orclinmemfiltprocess`) to those entries in memory.
5. Oracle Internet Directory sends the entries that match that filter to the client.

For example, suppose `orclinmemfiltprocess` is set to `(objectclass=inetorgperson) (orclisabled=TRUE)`. When Oracle Internet Directory receives the search `(&(uid=msmith) (objectclass=inetorgperson) (orclisabled=TRUE))`, it sends a filter containing only the parameter `(uid=msmith)` to the database. After Oracle Internet Directory receives entries back from the database, Oracle Internet Directory itself applies the filter `(objectclass=inetorgperson) (orclisabled=TRUE)` to those entries.

By default, `orclinmemfiltprocess` is set to the following values:

```
(objectclass=inetorgperson)
(objectclass=oblixorgperson)
(|(! (obuseraccountcontrol=*)) (obuseraccountcontrol=activated))
(| (obuseraccountcontrol=activated) (! (obuseraccountcontrol=*)) )
(objectclass=*)
(objectclass=oblixworkflowstepinstance)
(objectclass=oblixworkflowinstance)
(objectclass=orcljaznpermission)
(obapp=groupservcenter) (! (obdynamicparticipantsset=*))
(objectclass=orclfeduserinfo)
```

You can change the value of `orclinmemfiltprocess` by using `orldapmodify`. See "Attributes of the Instance-Specific Configuration Entry" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory* and "Setting System Configuration Attributes by Using `ldapmodify`" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

Under some conditions, Oracle Internet Directory ignores `orclinmemfiltprocess` and sends the entire filter to the database. It does this if the filter it receives meets the following conditions:

- It contains only one parameter, that is, one attribute-value pair.
- It contains no filter condition other than those in `orclinmemfiltprocess`

- It contains an OR condition applied to the terms that are in `orclinmemfiltprocess`
- It contains the same terms as in `orclinmemfiltprocess`, but in a different order

The following cases illustrate those conditions. In all of the following cases, `orclinmemfiltprocess` is set to `(objectclass=inetorgperson) (employeetype=Contract)`.

Examples

Case A

```
(&(manager=cn=john doe) (objectclass=inetorgperson)
(employeetype=Contract))
```

Oracle Internet Directory sends the filter `(&(manager=cn=john doe))` to the database.

Case B

```
(&(uid=rmsmith) ((objectclass=inetorgperson) (employeetype=Contract)))
```

Oracle Internet Directory sends only `(&(uid=rmsmith))` to the database, then applies the filter `(&(objectclass=inetorgperson) (employeetype=Contract))` to the entries that are returned from the database.

Case C

```
(|(uid=rmsmith) (objectclass=inetorgperson) (employeetype=Contract))
```

In this filter, the terms that match `orclinmemfiltprocess` are part of an OR condition. Oracle Internet Directory sends the filter, as is, to the database.

Case D

```
(&(uid=rmsmith) (employeetype=Contract) (objectclass=inetorgperson))
```

Even though some of the terms in this filter match `orclinmemfiltprocess`, they are in a different order, so Oracle Internet Directory sends the whole filter to the database. You could add `(employeetype=Contract) (objectclass=inetorgperson)` to `orclinmemfiltprocess` if you do not want Oracle Internet Directory to send this filter to the database.

Case E

```
(| (&(uid=rmsmith) (sn=smith) (objectclass=inetorgperson) (employeetype=Contract)))
```

In this filter, the terms that match `orclinmemfiltprocess` are part of an OR condition. Oracle Internet Directory sends the filter, as is, to the database.

Case F

```
(&( |(uid=rmsmith) (sn=smith)) (objectclass=inetorgperson) (employeetype=Contract)))
```

Even though this filter contains an OR operator, it is not applied to the terms that match `orclinmemfiltprocess`. Oracle Internet Directory sends `(&(|(uid=rmsmith) (sn=smith)))` to the directory and applies the filter `(&(manager=cn=john doe) (&(objectclass=inetorgperson) (employeetype=Contract)))` to the entries that are returned from the database.

Configuring Multiple Filters

If the application is sending multiple filters, and the terms in one filter are a superset of the terms in the other, you must configure `orclnmemfiltprocess` for both values.

For example, suppose the application is sending the following two filters:

```
(&(uid=rmsmith)(objectclass=inetorgperson)(employeetype=Contract))
(&(uid=rmsmith)(objectclass=inetorgperson)(employeetype=Contract)(departmentNumber=627))
```

where `(departmentNumber=627)` matches a lot of entries. You must configure `orclnmemfiltprocess` as follows:

```
(objectclass=inetorgperson)(employeetype=Contract)
(departmentNumber=627)
```

Optimizing Performance for Search baseDN

In the DIT, if all the users are under one baseDN, such as `cn=users,dc=acme,dc=com`, and all the LDAP search clients send base as `cn=users,dc=acme,dc=com`, then the configuration of the `orclnmemfilter` will significantly reduce database processing time. See the following example:

```
orclnmemfiltprocess,dn: cn=users,dc=acme,dc=com
```

23.5 Specific Use Cases That Require Additional Tuning

This section describes some specific use cases that require additional tuning, in addition to [Section 23.3, "Basic Tuning Considerations"](#).

23.5.1 Bulk Load Operations

If you are planning a large bulkload operation, make the following changes:

- Set the database initialization parameter `pga_aggregate_target` to 1-4GB, if sufficient RAM is available. Consider increasing the size if there is a pattern of increased PGA usage.
- Increase the database temporary tablespace before loading a large number entries. You need about 1G of temporary tablespace per million entries being loaded. You can free up the tablespace after the operation.

23.5.2 Bulk Delete Operations

If you are planning a large bulkdelete operation, perform the following tasks:

- Ensure that the database initialization parameter `sga_target` are tuned as described in [Section 23.3.1, "Database Parameters."](#)
- Set the database initialization parameter `log_buffer` to 10M. This can provide additional performance benefit.
- Ensure that you have at least three database redo log files with at least 2 GB each.
- Ensure that the undo tablespace is at least 1 GB in total size.
- Follow the recommendations about redo logs and undo tablespace in the next section, [Section 23.5.3, "High LDAP Write Operations Load."](#)

23.5.3 High LDAP Write Operations Load

If you have a high LDAP write operations load, or if you perform many `bulkdelete` operations, consider tuning the following values:

- Increase the size or number of the database redo log files so that the total size is 1000-1500 MB. Other considerations affect the total size of redo logs.
- Depending on how the disks are configured, it might be beneficial to isolate the redo log files to a dedicated set of disks.
- Increase the undo tablespace size by adding data files to this tablespace. For most deployments, 2-4 GB should suffice.
- Do not use the Oracle Internet Directory server entry cache. See [Section 23.4.6, "Server Entry Cache."](#)
- If neither Oracle Internet Directory replication nor DIP is deployed, disable change log generation. See [Section 23.4.1, "Replication or Oracle Directory Integration Platform."](#)

[Table 23–6](#) summarizes the redo log and undo tablespace recommendations provided in this section.

Table 23–6 Redo Log and Undo Tablespace Values

Attribute	Value	Notes
Redo Log	3 logs, 2GB each	Sufficient to handle many <code>bulkdelete</code> operations and large numbers of write operations.
Undo Tablespace	At least 1GB total	Many <code>bulkdelete</code> operations.
Undo Tablespace	2-4 GB	Large number of write operations.

23.5.4 Sparc T4 Hardware Tuning

On Sparc T4 hardware, set `OID_BIND_CPU=0` environment variable, and make sure `$ORACLE_HOME/bin/oiddispd` is owned by root process with `setuid` bit ON.

`OIDDISPd` creates the processor set, so, make sure the `orclserverprocs = NumberOfcores-1`.

Oracle Unified Directory Performance Tuning

This chapter provides guidelines for tuning and sizing Oracle Unified Directory. It contains these topics

- [Section 24.1, "About Oracle Unified Directory"](#)
- [Section 24.2, "Performance Considerations"](#)
- [Section 24.3, "Monitoring Unified Directory Performance"](#)
- [Section 24.4, "Basic Tuning Considerations"](#)
- [Section 24.5, "Advanced Tuning Recommendations"](#)

24.1 About Oracle Unified Directory

Oracle Unified Directory is a comprehensive next generation directory service that is designed to address large deployments, to provide high performance, to be highly extensive and to be easy to deploy, manage, and monitor.

24.2 Performance Considerations

Oracle Unified Directory aims to be high-performing and highly-scalable. Although the server can achieve impressive results with the "out-of-the-box" server configuration and default JVM settings, performance can often be improved significantly through some basic tuning.

The default settings of Oracle Unified Directory are targeted at evaluators and developers who are running equipment with limited resources. When you deploy Oracle Unified Directory in a production environment, it useful to do some initial tuning of the Java Virtual Machine (JVM) and of the server configuration to improve scalability and performance (particularly for write operations).

In addition, performance tuning strategies differ depending on whether you are running a directory server or a proxy server. This section describes some of the areas that you should consider tuning based on your server usage. Note that the specific tuning parameters and descriptions are discussed in [Section 24.4, "Basic Tuning Considerations"](#).

- **When OUD is used as an LDAP Directory Server.** When used as a Directory Server, you can maximize performance by:
 - Tuning the database cache size, preload, and file cache size appropriately.
 - Placing the database on a fast file system.
 - Using the correct database caching mode for your deployment.

- Tuning the Oracle Berkeley DB Java Edition log cleaners.
- **When OUD is used as an LDAP Proxy Server.** When used as a Proxy Server, you can maximize performance by:
 - Making sure you have a sufficient number of worker threads. Proxying requires a large number of worker threads to optimize performance.
 - Setting the heap size to an appropriate value and using the correct JVM. It is unlikely that a proxy will need more than 4GB of heap, therefore a 32-bit JVM should be used in most cases.

In addition, the following items can improve performance in specific deployment scenarios.

- **Java Version.** Use the most recent Java Runtime Environment (JRE) release available. Oracle Unified Directory is designed to work with Java SE 6 and 7.
- **Environment Variables.** The server uses the `OPENDS_JAVA_HOME` environment variable to point to your installed JRE. If you have multiple versions of Java installed on a system, set the `JAVA_HOME` environment variable to point to the root of the desired installation. In this way, the version of the JRE specified by the `JAVA_HOME` variable can be used by other applications but not by Oracle Unified Directory.

To specify a JRE installation for the server, do one of the following:

- Use the `dsjavaproperties` command to set the appropriate environment variables.
- Set the `OPENDS_JAVA_BIN` environment variable (with the JAVA binary path).
- Set the `OPENDS_JAVA_HOME` environment variable (with the JAVA installation path).

24.3 Monitoring Unified Directory Performance

This chapter provides an overview of the key monitoring tools you can use to gather performance information for Unified Directory.

Note: Oracle Unified Directory provides an extensible monitoring framework. Oracle Unified Directory performance can also be monitored by using the Enterprise Manager Grid Control plugin. For more information, see the *System Monitoring Plug-in for Oracle Unified Directory User's Guide*.

24.3.1 Examining Log Files

Oracle Unified Directory provides several types of logs: access logs, audit logs, error logs, debug logs, oud-setup logs, server.out logs, and a replication repair log. The replication repair log is read-only and its use is restricted to enabling replication conflict resolution.

To quickly determine whether performance issues are related to problems with the server or with the client, review the **access log** at `INSTANCE_DIR/OU/LOGS/ACCESS`. Access logs record information about the types of operations processed by the directory server. Access logs are provided by default.

This log contains entries of the form:

```
[09/Sep/2009:15:36:18 +0200] SEARCH RES conn=1 op=16 msgID=17
result=0 nentries=1 etime=1
```

The value of the `etime` field is the time (in milliseconds) that the server spent processing the request. Large `etimes` generally indicate an issue on the server side (which can usually be resolved by appropriate performance tuning or indexing). If you are experiencing performance problems but the `etimes` are small, the issue is more likely to be with your client application.

24.3.2 Monitoring the Server With LDAP

The following sections of the *Oracle Unified Directory Administrator's Guide* provide detailed instructions on how to monitor the server using Lightweight Directory Application Protocol (LDAP):

- Viewing Monitoring Information Using the `cn=monitor` Entry
- Monitoring Using the `manage-tasks` Command
- Monitoring the Server With JConsole

24.3.3 Monitoring the Server With SNMP

Oracle Unified Directory provides a jar file extension that contains a Simple Network Management Protocol (SNMP) connection handler for Management Information Base (MIB) 2605 support. The extension contains the SNMP connection handler, the required classes to support MIB 2605 objects and SNMP requests, and the SNMP adapter that allows an SNMP manager to access the server monitoring information.

24.4 Basic Tuning Considerations

The following sections describe basic tuning configurations that you should also consider while tuning Oracle Unified Directory:

- [Tuning Java Virtual Machine Settings](#)
- [Tuning the Server Configuration](#)

24.4.1 Tuning Java Virtual Machine Settings

As discussed in [Chapter 2.4, "Tuning Java Virtual Machines \(JVMs\)"](#), you can use the `JAVA_ARGS` environment variable to provide global configuration arguments that can be passed to the JVM, or you can use the `java.properties` file. Any argument that can be used with the `java` command can be used with both methods.

The following table describes the main JVM tunable options for Oracle Unified Directory. Note that some of these settings apply to Sun JVM only, but that similar settings could be applied to JRockit.

Parameter	Description	When using OUD as a Directory Server...	When using OUD as a Proxy Server...
<code>-server</code>	Always use the server JVM instead of the client JVM. The client VM is better optimized for processes that run for a short period of time and need to start as quickly as possible. The server VM can take longer to warm up but is faster than the client VM.	Set this parameter as described.	Set this parameter as described.

Parameter	Description	When using OUD as a Directory Server...	When using OUD as a Proxy Server...
-d32 or -d64	Select the 32-bit or 64-bit version of the JVM as follows: <ul style="list-style-type: none"> -d32 provides better performance for JVM heaps smaller than 3.5Gbytes. -XX:+UseCompressedOops should be used for JVM heaps between 3.5Gbytes and 31Gbytes. -d64: should be used for JVM heaps over 32Gbytes. 	The JVM version used will depend on heap size.	Not applicable, except for the case when global index is used, otherwise it's unlikely that a Proxy Server will need more than 4GB of heap. In most cases the 32-bit version of the JVM should be used.
-XX:+UseCompressedOops	Use this option if you use the 64-bit JVM and if the heap size is less than 32 Gbytes.	Set this parameter as described.	Not applicable a 64-bit JVM is required.
-Xms (Initial heap size) -Xmx (Maximum heap size)	This parameter sets the initial and maximum heap size available to the JVM. Increasing the heap size can improve performance, but setting it too high can have a detrimental effect in the form of longer pauses for full garbage collection runs. The initial and maximum sizes should generally be set to the same values. For maximum performance, size the heap so that the entire DB can be cached in memory. In general, you should allocate enough heap for the server runtime and the rest to the DB cache. If you use CMS as the garbage collector of the oldgen, you must take into account the <code>-XX:CMSInitiatingOccupancyFraction</code> property when you calculate the heap size so that it is coherent with the size (or percent of the heap) occupied by the dbcache. If you set the <code>CMSInitiatingOccupancyFraction</code> to 55, the dbcache percent should be set to 50. Then if you have a database on disk that is 10GB, you need at least a heap of 22GB if you want the entire database to fit into the dbcache.	Set this parameter as described. Example: -Xmx31G -Xms31G	Set this parameter as described. Example: -Xmx3500M
-XX:NewSize	The total heap space is divided into the old generation and the young generation. This parameter sets the size of the young generation. The remaining memory (old generation) must be sufficient to hold the DB cache plus some overhead.	Set this parameter as described. Example: -XX:NewSize=512M	Set this parameter as described. Example: -XX:NewSize=2G
-XX:+UseConcMarkSweepGC	Use the Concurrent Mark Sweep (CMS) garbage collector. This option allows the JVM to minimize the response time of LDAP operations, but it can have a small impact on the overall performance (throughput) of the server. Use this option if long pause times are not tolerated.	Set this parameter as described.	Set this parameter as described.

Parameter	Description	When using OUD as a Directory Server...	When using OUD as a Proxy Server...
-XX:CMSInitiatingOccupancyFraction=<percentage>	<p>Specify the level at which the CMS garbage collection is started. The default value is approximately 68%. Use this value if you want to set the percentage to something other than the default value.</p> <p>If you use CMS as the garbage collector of the oldgen, then it is important to take into account the CMS property -XX:CMSInitiatingOccupancyFraction in the computation of the heap size such it is coherent with the size (or percent of the heap) occupied by the dbcache.</p> <p>If you setup the CMSInitiatingOccupancyFraction to 55, then the dbcache percent should be setup to 50. Then if you have a db on disk that is 10GB, this means you need at least a heap of 22GB if you want all the db to fit into the dbcache.</p>	Set this parameter as described.	Set this parameter as described.
-XX:+UseCMSInitiatingOccupancyOnly	<p>Add -XX:UseCMSInitiatingOccupancyOnly if you specified -XX:CMSInitiatingOccupancyFraction=60, and the CMS collections are still starting before they reach that threshold. Remove the -XX:CMSInitiatingOccupancy=60 (using the default value of 69%), and consider adding the line: -XX:UseCMSInitiatingOccupancyOnly.</p>	Set this parameter as described.	Set this parameter as described.
-XX:+UseBiasedLocking	Improve locking performance in the server in cases where there is not expected to be a high degree of contention.	Set this parameter as described.	Set this parameter as described.
-XX:LargePageSizeInBytes=256m	<p>Use large pages for the information it stores in memory.</p> <p>NOTE: This argument applies primarily to systems using the UltraSPARC T1 processor.</p>	Set this parameter as described.	Set this parameter as described.
-XX:+CMSScavengeBeforeRemark	Causes a minor collection to occur just before the remark which reduces the remark pause time.	Set this parameter as described.	Set this parameter as described.
-XX:+UseNUMA	NOTE: This option should only be used on UltraSPARC Tx processors.	Set this parameter as described.	Set this parameter as described.

24.4.2 Tuning the Server Configuration

Various components of the server can be tuned to provide performance improvements in specific scenarios. Most performance tuning recommendations depend on several variables, including the anticipated workload, the types of data that are stored, and the hardware and resources available. The following general tuning recommendations can improve server performance in specific deployments.

24.4.2.1 Oracle Berkeley DB Java Edition Tuning Parameters

Oracle Berkeley DB Java Edition is an open source, embeddable, transactional storage engine. The architecture of Oracle Berkeley DB Java Edition supports very high performance and concurrency for both read-intensive and write-intensive workloads.

The following Berkeley DB Java Edition (JE) tuning parameters can be used to tune performance:

Parameter	Description
db-cache-percent and db-cache-size	<p>Use these properties to configure the amount of memory that the database cache uses. db-cache-size is dependant on the size of the directory and the amount of available memory.</p> <p>There are two caching options:</p> <ul style="list-style-type: none"> Fully cached in JVM heap. <p>When the directory is small and there is enough memory available, the db-cache_size should be tuned to contain the entire node. This caching mode provides deterministic response times and low disk I/O, but requires large JVM heap and can impact garbage collection.</p> Partially cached in JVM heap, remainder cached in File System cache. <p>When the directory is too large to store all of the nodes in memory, the db-cache-size should be tuned to contain only the Upper Inner Nodes (UpperIN) and the Bottom Inner Nodes (BIN). This caching mode also provides low disk I/O and can eliminate some of the issues related to large heap, but it can be slower and has slightly less deterministic response times (performance degradation is proportional to the amount of DB cached in the JVM).</p> <p>The size of the Inner Node can be computed using the DbCacheSize utility. For more information, see "Estimating the JE Cache Size".</p> <p>Before you can accurately set the db-cache-size, you must first determine the approximate size of the database after an import. For example, after doing an import into the userRoot back end, run the following command (on UNIX systems) to determine the size of the database:</p> <pre>\$ cd INSTANCE_DIR/OUJ/db \$ du -sk userRoot/ 910616 userRoot/</pre> <p>On Windows systems, use an equivalent procedure to determine the database size. Remember that the database size is not static and can increase after an initial import when modifications are made.</p> <p>Setting the JVM heap to 2 Gbytes (-Xms2g -Xmx2g), and the db-cache-percent to 50, will cause the DB cache to use 1 Gbyte of memory.</p>
je.checkpointer.highPriority	<p>If true, the checkpointer uses more resources in order to complete the checkpoint in a shorter time interval. Btree latches are held and other threads are blocked for a longer period. Log cleaner record migration is performed by cleaner threads instead of lazily during eviction and checkpoints (see CLEANER_LAZY_MIGRATION). When set to true, application response time may be longer during a checkpoint, and more cleaner threads may be required to maintain the configured log utilization.</p> <p>Setting that property to false is a way to achieve better throughput and lower response times.</p>
preload-time-limit	<p>You can configure the server to preload some of the database contents into memory on startup. For large databases, preloading the database cache avoids a long warm-up period after server startup.</p>
db-directory	<p>Ensure that the database is held on a fast file system with adequate storage. The file system should be different to the location of the access logs. By default, the database will grow to twice its original size. For example, if the database is 1 Gbyte after an import, the file system should have at least 2 Gbytes available.</p>
db-evictor-lru-only	<p>Use this property to control how the database cache retains information. Setting this value to false ensures that the internal nodes are maintained in cache, which provides better performance when the JE cache holds only a small percentage of the database contents.</p>

Parameter	Description
db-txn-durability	Use this property to configure durability for write operations. Reducing durability can increase write performance, but it can also increase the chance of data loss in the event of a JVM crash or a system crash. This property takes the following values: write-to-disk. All data are written synchronously to disk. write-to-fs. Data are written to the file system immediately but might stay in the file system before being flushed to disk. write-to-cache. Data are written to an internal buffer and flushed to the file system, then to disk when necessary.
db-log-file-max	Use this property to control the size of JE log files. Increasing the file size can improve write performance, but it can also make it harder to maintain the desired utilization percentage.
db-num-cleaner-threads and db-cleaner-min-utilization	These properties control how the cleaner works, which keeps the database size down and keeps up with high write throughput. As older records become obsolete they are cleaned. The cleaning process can impact performance if not tuned.
db-num-lock-tables	On systems with a large number of CPUs, this property can improve concurrency within the database lock manager.

24.4.2.2 Core Server Tuning Parameters

The following core server tuning parameters can be used to tune performance:

Parameter	Description
num-request-handlers	This property can be configured so that the LDAP connection handler (and the LDAPS connection handler, if it is enabled) use multiple threads for decoding client requests. Increasing the number of threads on systems with a larger number of CPUs can improve performance. As a rule of thumb, you should set this property to a quarter the number of CPUs, with a maximum of twelve. In some cases disabling the <code>keep-stats</code> property can help reduce lock contention in the connection handlers.
num-worker-threads	The default value of this property is two times the number of CPUs. This value is sufficient in most deployments.
log-file	Ensure that the access log publisher is on a fast file system, or turn it off altogether by setting the <code>enabled</code> property to <code>false</code> .

24.5 Advanced Tuning Recommendations

The following additional recommendations can improve performance in specific scenarios.

- Enable an Entry Cache.** In some cases, particularly those involving relatively small directories (for example, up to a few hundred thousand entries), it can be useful to enable an entry cache. In general the FIFO entry cache provides better results than the soft reference entry cache.

For large databases, it is recommended that you store only a specific set of the data in the cache, by using the `include-filter` property. Storing static groups in the entry cache can greatly improve the overall performance of the server. This reduces the time required to perform group membership lookup, which is necessary in evaluating ACIs, for example.

- **Disable Unused Virtual Attributes.** If the functionality needed by one or more of the virtual attributes is not required, they can be disabled for a slight performance improvement when decoding entries.
- **Disable Unused Access Logging.** If access logging is not necessary, disabling the server access logger can help improve performance.
- **Disable Unused Access Control Handlers.** If you do not need access control processing in the server, then you can disable it by setting the `enabled` configuration property to `false` for the Access Control Handler. You can set the property by using `dsconfig`.
- **Reduce Lock Contention.** On systems with large numbers of CPUs (for example, chip multi-threading (CMT) systems with several hardware threads per core), you can reduce lock contention by setting the `org.opends.server.LockManagerConcurrencyLevel` system property to be equal to the number of worker threads you intend to use.

NOTE: This property must be set as a JVM system property, because it can be required very early in the server startup process, even before accessing the server configuration.

Oracle Virtual Directory Performance Tuning

This chapter provides tuning tips for Oracle Virtual Directory. It contains the following sections:

- [Section 25.1, "About Oracle Virtual Directory"](#)
- [Section 25.2, "Basic Tuning Considerations"](#)
- [Section 25.3, "Advanced Tuning Considerations"](#)

25.1 About Oracle Virtual Directory

Oracle Virtual Directory is an LDAP Version 3-enabled service that provides an abstracted view of one or more enterprise data sources. Oracle Virtual Directory consolidates multiple data sources into a single directory view, enabling you to integrate LDAP-aware applications with diverse directory server data stores.

The information in this chapter assumes that you have reviewed the concepts and administration information in the *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*.

Note: Oracle Virtual Directory's out of box configuration may not be optimal for many production and test deployments. You are encouraged to incorporate the recommendations listed in "Basic Tuning Considerations" to achieve optimal performance and availability.

25.2 Basic Tuning Considerations

The tuning considerations in this section apply to most deployments and usage scenarios. It is highly recommended that you review these configurations and implement those that are appropriate for your use case scenarios. The tuning information is summarized in [Table 25-1](#).

Table 25–1 Basic Tuning Considerations

Configuration Attribute	Category	Default Value	Recommended Value	Notes
Threads	Listener Properties	10	10 * Number Of central processing units (CPUs) available for Oracle Virtual Directory Server	Recommendation applies only to the active LDAP Listeners. For more information, see Tuning Worker Threads .
Work Queue Capacity	Listener Properties	2048	Expected Number of Max Concurrent Clients * 2	2048 operations are executed concurrently. Some clients may send asynchronous operations as well. For more information, see Tuning Work Queue Capacity .
Max, Initial Pool Connections	LDAP Adapter Properties	10	Total Number of 'Threads' parameter values for all active Listeners that use this Adapter	Ensure that the back-end Directory Servers can handle these connections. For more information, see Tuning the LDAP Connection Pool .
Max Heap Size	System Properties	256 MB	Up to 2 GB on 32-bit systems and higher values on 64-bit systems.	Higher values protect against Out Of Memory errors. Ensure that there is sufficient RAM on the system to handle the configured value. For more information, see Tuning Heap Size .

25.2.1 Tuning the Ping Interval

Consider increasing the ping interval to 60 seconds (or more as needed) in the `opmn.xml` file.

When the system is busy, a ping from the Oracle Process Manager and Notification Server (OPMN) to Oracle Virtual Directory may fail. As a result, OPMN will restart Oracle Virtual Directory after 20 seconds (the default ping interval). To avoid this, consider increasing the ping interval to 60 seconds or more.

The ping interval can be modified in the `$ORACLE_INSTANCE/config/OPMN/opmn/opmn.xml` as shown below:

```
<process-type id="OVD" module-id="OVD">
  <module-data>
    <category id="start-options">
      <data id="java-bin" value="$ORACLE_HOME/jdk/bin/java"/>
      <data id="java-options" value="-server -Xms2056m -Xmx2056m
```

```

-Dvde.soTimeoutBackend=0 -DdisableECID=1 -Didm.oracle.home=${ORACLE_HOME}
-Dcommon.components.home=${ORACLE_HOME}/../oracle_common
-Doracle.security.jps.config=${ORACLE_INSTANCE}/config/JPS/jps-config-jse.xml"/>
  <data id="java-classpath" value="${ORACLE_HOME}/ovd/jlib/vde.jar:${ORACLE_
HOME}/jdbc/lib/ojdbc6.jar"/>
  </category>
</module-data>
  <stop timeout="120"/>
  <ping interval="60"/>
</process-type>

```

25.2.2 Tuning Worker Threads

Tune the number of worker threads based on the number of central processing units (CPU) available for Oracle Virtual Directory Server on the system.

The 'Threads' configuration parameter in the Oracle Virtual Directory Listener settings should be set to an appropriate value. The default value for Threads in the Admin Gateway listener and DSML Gateway listener should not be changed. The number of Threads for the LDAP Listeners are typically the threads that need to be tuned since it is the LDAP Listeners that take on concurrent traffic from applications. A common configuration is to have 10 threads per CPU. For example, if there are 4 central processing units on the system, then there would be 40 threads.

For more information, see "Managing Listeners" in *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*.

25.2.3 Tuning Work Queue Capacity

Tune the Work Queue Capacity based on the expected maximum number of concurrent clients to a given LDAP Listener.

The 'WorkQueueCapacity' configuration parameter in the Oracle Virtual Directory Listener settings should be set to an appropriate value. This ensures that the connection requests from LDAP clients are not rejected due to a lack of work queue capacity. Work elements are allocated on demand only, therefore a value higher than the actual estimate can be used.

The Fusion Middleware Control Performance Monitor provides a historical report which contains the maximum number of connections. Use this report to determine how to adjust the connection value based on production data.

If Oracle Virtual Directory needs to support high number of concurrent clients, then set the ulimit 'nofiles' (descriptor) parameter to the number of LDAP Clients expected. For example, in the command window where OPMN is started, set the following ulimit when 8000 concurrent clients are expected:

```
ulimit -n 8192
```

This change requires restart of OPMN and Oracle Virtual Directory to take effect.

For more information, see "Managing Listeners" in *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*.

25.2.4 Tuning the LDAP Connection Pool

Tune the size of the LDAP connection pool in Oracle Virtual Directory LDAP Adapter to be at least as high as the total number of Threads configured in the Oracle Virtual Directory Listeners that actively use the LDAP Adapter.

This ensures that in the worker threads have enough LDAP connections to process requests. The actual number of active adapters, active listeners and traffic pattern control the usage of connections. However, since connections that are idle in the LDAP Adapter connection pool are periodically closed, a higher value should not impact performance. Ensure that the back-end Directory Server is configured to handle the number of concurrent connections from Oracle Virtual Directory LDAP Adapter connection pool.

For more information, see "Configuring LDAP Adapter" in *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*.

25.2.5 Tuning Heap Size

Tune the maximum Java heap size of the JVM running Oracle Virtual Directory. This is to ensure that Oracle Virtual Directory has sufficient heap to handle the concurrent load.

For more information, see "Controlling the Maximum Heap Size Allocated to the Oracle Virtual Directory Server" in *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*.

25.3 Advanced Tuning Considerations

Depending on your Oracle Virtual Directory deployment's use case scenarios, the following tuning configurations may improve performance.

25.3.1 Tuning Database Adapters

The Database Adapter is a fully featured LDAP-to-JDBC gateway supporting translation of all LDAP operations (add, bind, delete, baseSearch, modify, wildcardSearch) into equivalent SQL prepared statement code. The Database Adapter uses JDBC class libraries to form connections to databases for the purpose of performing LDAP searches. The database libraries are generally provided by the database vendor.

Note: For improved performance, tune the database before using the Database adapter. Consult your database documentation for more information. If the database being used is an Oracle database, see *Oracle Database Performance Tuning Guide*.

For optimal performance, consider the following configuration options for the database schema against which the Oracle Virtual Directory database adapter is configured:

- Enable the Case Insensitive Search option.

Enabling (selecting) the Enable Case Insensitive Search option makes the search case insensitive for case insensitive LDAP attributes, such as uid. Oracle Virtual Directory uses UPPER in the SQL query when Enable Case Insensitive Search is enabled. If the database cannot maintain functional indexes, such as for Oracle TimesTen or MySQL databases, then you should disable the Enable Case Insensitive Search option. When the Enable Case Insensitive Search is disabled, Oracle Virtual Directory performs case sensitive searches and does not use UPPER in the SQL query. The default value for Enable Case Insensitive Search is Enable.

- In general, the mapped columns in the underlying database schema should have an index defined if the mapped LDAP attribute is used in LDAP search filters.
- In scenarios where an LDAP attribute that is used in an LDAP search filter has a matching rule of 'caseIgnoreMatch', the mapped database table column for this attribute needs a function index to be defined for optimal look-up performance.

For example, if LDAP attribute 'CN' is mapped to database schema column EMP.NAME, then a function index on UPPER(EMP.NAME) is required for optimal performance of LDAP search filters involving CN attribute.

For more information on function-based indexes, see "Using Function-based Indexes for Performance" in *Oracle Database Performance Tuning Guide*.

Table 25–2 describes some additional Database Adapter settings:

Table 25–2 Database Adapter Settings

Parameter	Value	Notes
Adapter	Default: Active	An adapter can be configured as Active or Inactive. An inactive adapter can not start during a server restart or when you try to start it. The purpose of the Inactive setting is to keep old configurations available or on stand-by without having to delete them from the configuration.
Maximum Connections	Default: 10 connections	This defines the maximum connections the Database Adapter may make with the database.
Database Connection Timeout	Default: 10 seconds	The database connection timeout adapter property controls the LDAP request to wait for a connection to become available in the cache after reaching the maximum number of connections limit. If a connection does not become available within the number of seconds defined, the LDAP request fails. If database connection timeout system property is not used, the LDAP request waits 10 seconds for a connection to become available.

25.3.2 Tuning Join Adapters

If you are using Join Adapters, join only appropriate sources. For example if a deployment requires only to link attributes in the primary source under "cn=users" branch, create a primary adapter that only exposes this branch. And then create the join rule with that adapter. This can reduce the need for Oracle Virtual Directory to try to join entries that may never have corresponding linked entries.

Tip: Always make sure that the attributes used by join rules are properly indexed.

25.3.3 Tuning Filters

If a known client search filter does not apply to certain adapters, apply the filter to all applicable "Exclude Filters" to improve performance and reduce network traffic.

25.3.4 Tuning Load Balancer Local Store Adapter

Some load balancers query an LDAP server to determine if it is up or down. If your load balancer uses this feature, consider creating a local store adapter with a separate namespace (for example `dc=loadbalancer`) that is used only for the load balancer. While the performance impact of the load-balancer is probably not noticeable, by keeping it in a separate namespace, it makes it easier to exclude the load-balancer `KeepAlive` requests from creating large log files during troubleshooting.

25.3.5 Tuning the Cache Plug-In

The CachePlug-in provides an in-memory cache for Oracle Virtual Directory. It has the ability to cache query results from any source for re-use by LDAP clients. This plug-in can improve performance for those applications where queries are highly repetitive.

To review cache operation and configuration, set VE logging level to 'Dump' to see more details. Because the cache is a normal plug-in, the cache can be configured to run anywhere within Oracle Virtual Directory. It can be executed globally, or within the context of a single adapter. It can also be restricted to specific namespaces by using the namespace filtering available in standard plug-in configuration.

25.3.5.1 Cache Hit Logic

The cache works by storing query results and making them available for later use. If a query is repeated by the same user and the same attributes or a subset of attributes are requested, the cache can return its results instead of having Oracle Virtual Directory pull the information from the source. The plug-in can also be configured to allow cache hits to be shared between users.

Sharing cache entries between users should not be used unless the pass credentials are not being passed to back-end sources and Oracle Virtual Directory is solely responsible for security enforcement. Careful consideration should be given when sharing cache hits between users as it would then be possible for one user to see something they should not, since they may have access to a cache result from a more privileged user.

25.3.5.2 Cache Plug-in Memory Management

This plug-in periodically reviews the cache and checks for expired results, or entries that have been invalidated by a previous modify transaction. In the event that the cache quota is exceeded, the plug-in attempts to trim memory by purging the queries that were least recently used (LRU).

[Table 25–3](#) describes some parameters used to tune the Memory Management Plug-in:

Table 25–3 *Memory Management Plug-in Settings*

Parameter	Value	Notes
Size	Default: 1000 entries	The maximum number of entries that may be cached at any one time.
MaxResultSize	Default: 1000 entries	The maximum number of entries that may be cached for any particular query.
Trimsize	Default: 1000 entries	When the maximum cache size is exceeded, the amount by which the cache manager must reduce the balance. Note: when necessary, trimming is done by purging expired queries first followed by queries in order of least recent use.

Table 25–3 (Cont.) Memory Management Plug-in Settings

Parameter	Value	Notes
MaximumAge	Default: 600 seconds	The maximum age in seconds for any query/entry stored in the cache.
MaintenanceInterval	Default: 60 seconds	The interval in seconds between when the cache manager checks for expired queries.
BySubject	Default: 1 (not shared)	A flag (1 or 0) indicating whether cache results are shared between subjects. A value of 1 indicates that results are not be shared between subjects.

25.3.6 Tuning LDAP Listener

Table 25–4 describes some parameters used to tune the LDAP Listener:

Table 25–4 Listener Parameters

Parameter	Value	Notes
Backlog	Default: 128 requests	Specifies the maximum number of pending connection requests that are allowed to queue up before the server starts rejecting new connection attempts. The default value is sufficient in most cases and the need to change this value is very rare.
Reuse address	Default: False	This option determines whether LDAP listener should reuse socket descriptors. If enabled, the <code>SO_REUSEADDR</code> socket option is used on the Oracle Virtual Directory server listen socket to potentially allow the reuse of socket descriptors for clients in <code>TIME_WAIT</code> state.
Keep Alive	Default: False	This option determines whether the LDAP connection should use TCP keep-alive. If enabled, the <code>SO_KEEPALIVE</code> socket option is used to indicate that TCP keepalive messages should periodically be sent to the client to verify that the associated connection is still valid.
TCP No delay	Default: True	This option determines whether the LDAP connection should use TCP no-delay. If enabled, <code>TCP_NODELAY</code> socket option is used to ensure that response messages to the client are sent immediately rather than potentially waiting to determine whether additional response messages can be sent in the same packet.

Table 25–4 (Cont.) Listener Parameters

Parameter	Value	Notes
Read Timeout	Default: 0	<p>This option enables/disables SO_TIMEOUT with the specified timeout, in milliseconds.</p> <p>With this option set to a nonzero timeout, client connection to the Oracle Virtual Directory server can remain idle only for this amount of time. If the connection is idle for a period longer than the specified timeout, the client connection is terminated.</p> <p>A timeout of zero is interpreted as an infinite timeout.</p> <p>Warning: This option is equivalent to vde.soTimeoutFrontend system property in Oracle Virtual Directory version 10g. The vde.soTimeoutFrontend system property is not supported for 11g. Users must modify the value specified in system property</p> <p>The mapping of values from 10g to 11g are:</p> <p>.Enabled to 0</p> <p>Disabled to nonzero amount of time in milliseconds</p>

25.3.7 Tuning the Server for OVD

Table 25–5 describes some basic parameters used to tune the server:

Table 25–5 Server Parameters

Parameter	Value	Notes
Anonymous Search Limit	Default: 1000	The maximum number of entries returned for an anonymous client.
Connection Timeout	Default: 120 (minutes)	<p>The Connection Timeout system property is used to prevent service outages caused by clients that do not properly close connections. The default value is 120. This value can be set in Oracle Enterprise Manager's Server Properties page.</p> <p>When you enable quota enforcement the value for inactiveConnectionTimeout can be changed. The default value here is 0 instead of 120.</p> <p>Warning: Setting to 0 disables the enforcement and client connections can not be closed regardless of how long they are inactive. The system property is not enforced on IP addresses and subjects that are exempt from the quota limit or that have disabled quota enforcement.</p>

Table 25–5 (Cont.) Server Parameters

Parameter	Value	Notes
Logging Levels	Default: Error:1 (Severe)	By default, log messages are written to the access.log file only when logging is set to NOTIFICATION:1. To maintain performance, consider keeping the default log level or use WARNING:1 (WARNING) to limit the amount of information written to the access.log file.

Oracle Identity Federation Performance Tuning

Oracle Identity Federation is a standalone, self-contained federation server that enables single sign-on and authentication in a multiple-domain identity network. It contains the following sections:

- [Section 26.1, "About Oracle Identity Federation"](#)
- [Section 26.2, "Basic Tuning Considerations"](#)
- [Section 26.3, "Advanced Tuning Considerations"](#)

Note: The configuration examples and recommended settings described in this chapter are for illustrative purposes only. Consult your own use case scenarios to determine which configuration options can provide performance improvements.

26.1 About Oracle Identity Federation

Oracle Identity Federation (OIF) is a standalone, self-contained federation server that enables single sign-on (SSO) and authentication in a multiple-domain identity network. The federation single sign-on capabilities are based on the SAML 1.x/SAML 2.0/WS-Fed protocols. The server is a Java EE Application deployed in a WebLogic Managed Server. This enables users to federate in heterogeneous environments and business associations, whether they have implemented other Oracle Identity Management products in their solution set.

For more information see *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation*.

26.2 Basic Tuning Considerations

The following tuning considerations are generally applicable to most use case scenarios and can improve Oracle Identity Federation performance.

26.2.1 Tuning Database Parameters for Identity Federation

This section provides configuration settings that can be used to tune the database.

See "Additional RDBMS Configuration" in *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation*.

26.2.1.1 Data Sources

Oracle Identity Federation uses a Java EE data source to interact with a database for various operations, such as:

- Locating a user record in the User Data Store
- Retrieving attributes from a user record in the User Data Store
- Locating, creating, or deleting an Oracle Identity Federation record from the Federation Data Store
- Locating, creating, or deleting an Oracle Identity Federation transient record from the Session or Message Data Store. (A transient record can be a user session, an artifact record, or federation protocol or session state.)

When creating a data source in the WebLogic Administration Console that can be used by Oracle Identity Federation, the maximum and minimum connection settings should be tuned for better performance. Consult your use case scenarios to determine what the connections settings should be to improve performance in your application.

26.2.1.2 RDBMS Session Cache

When Oracle Identity Federation is integrated with RDBMS for its Session Data Store, the server uses a caching mechanism to improve performance at run time. This enables the server to keep a reference to recently used session objects in memory to avoid read access to the database.

To optimize RDBMS session caching, configure the following:

- Number of session objects kept in memory at a given time
- Length of time a specific session object is kept in memory

Note: if Oracle Identity Federation is in High Availability (HA) mode with a load balancer, sticky sessions must be enabled to ensure that the cache is always reflecting accurate data.

See "Configuring RDBMS Session Cache" in *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation*.

26.2.1.3 RDBMS Compression

To decrease the amount of data to be stored in an RDBMS, Oracle Identity Federation provides the capability to compress the data before storing it to the database. There are three kinds of data that can be compressed:

- `AuthnRequest` for SSO Artifact profile: when Oracle Identity Federation acts as an IdP for Liberty 1.x protocol, the server stores the `AuthnRequest` message in the RDBMS when the artifact profile is used.
- Assertion Response for SSO Artifact profile: when Oracle Identity Federation acts as an IdP for SSO protocols, the server stores the Response message containing the Assertion in the RDBMS when the artifact profile is used. This must be enabled if attributes are contained in the assertion.
- User Session Data: Oracle Identity Federation stores some session data related to the user at run time. If several attributes are stored in the User Session (set by a custom Authentication Engine, or because the Attributes Assertion storage was enabled when Oracle Identity Federation was a service provider), then compression should be used.

See "Configuring RDBMS Data Compression" in *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation*.

26.2.2 Tuning the Oracle HTTP Server

If Oracle Identity Federation is fronted by Oracle HTTP Server (OHS), then the configuration of the HTTP Server can be tuned to increase performance. For more information on Oracle HTTP Server, see *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*.

The following parameters can be changed in the `httpd.conf` file of the OHS. For additional Oracle HTTP tuning configurations, see [Chapter 6, "Oracle HTTP Server Performance Tuning"](#). Consult your use case scenarios to determine what your settings should be.

- Timeout
- KeepAlive
- MaxKeepAliveRequests
- KeepAlive TimeOut
- MinSpareServers
- MaxSpareServers
- StartServers
- MaxClients
- MaxRequestPerChild

After modifying these parameters, save and restart OHS.

26.3 Advanced Tuning Considerations

The following tuning considerations may or may not apply to your OIF deployment. Review your use case scenarios to determine if these changes can help improve performance:

26.3.1 Tuning the LDAP Servers

This section provides configuration settings that can be used to tune LDAP such as:

- [Connection Pool Settings](#)
- [Connection Settings](#)
- [Federation Data Store Settings](#)

For the best performance, review the tuning configurations in [Chapter 2, "Top Performance Areas"](#) before tuning Oracle Identity Federation.

26.3.1.1 Connection Pool Settings

When Oracle Identity Federation is integrated with LDAP Servers as a user data store, federation data store, or authentication engine, the server keeps a pool of LDAP connections that can be re-used for subsequent requests.

Oracle Identity Federation performs the following kind of operations to the LDAP Servers:

1. User Data Store

- Locate users during assertion mappings
 - Retrieve attributes from the user record when creating an assertion
2. Authentication Engine
 - Locate user
 - Validate user credentials during authentication operations
 3. Federation Data Store, if used
 - Create a federation record
 - Locate a federation record
 - Update or delete a federation record. The LDAP Connection Pool can be configured by:
 - Setting Maximum Connections to indicate how many LDAP connections can the pool contain.
 - Setting the Connection Wait Timeout which is the time that a thread waits before re-trying to get an LDAP connection when none are available in the pool and that the pool is at maximum capacity.

See "Configuring Oracle Identity Federation" in *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation* for more information on the User and Federation Stores as well as the LDAP Authentication Engine.

26.3.1.2 Connection Settings

When Oracle Identity Federation is integrated with LDAP Servers as a user data store, federation data store, or authentication engine, the LDAP run time connections can be configured. For more information, see "Configuring Oracle Identity Federation" in *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation*.

The LDAP Connections can be configured by:

- Setting the LDAP Inactivity setting which tells Oracle Identity Federation how long an LDAP connection should be kept in a pool before being removed due to inactivity. Over time, the LDAP server may close some connections due to a long inactivity period, and if left unchecked, this can result in errors and may impact performance in Oracle Identity Federation.

See "Configuring the LDAP Inactivity Setting" in *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation*.

- Setting the LDAP Read Timeout Setting. Sometimes the LDAP server can become unresponsive, causing the thread/user to wait for a response or an error. To avoid waiting too long for an error when the server is not responding, Oracle Identity Federation sets a read timeout property on the LDAP connection. If the LDAP server does not respond before the read timeout period, an error is generated. Oracle Identity Federation closes the connection, open a new one and re-issue the LDAP command.

See "Configuring the LDAP Read Timeout Setting" in *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation*.

- Setting the High Availability (HA) LDAP Flag. When integrated with LDAP Servers that are deployed in HA mode, Oracle Identity Federation must be configured to indicate that the LDAP Servers are in HA mode.

See "Configuring High Availability LDAP Servers" in *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation*.

26.3.1.3 Federation Data Store Settings

When using Oracle Internet Directory as the Federation Data Store, Oracle Identity Federation creates, locates, updates and deletes federation records containing Account Linking Information.

Oracle Identity Federation uses specific queries when interacting with Oracle Internet Directory, and the performance can be improved by creating filters in Oracle Internet Directory. If Oracle Internet Directory is used as the Federation Data Store, it is possible to tune the LDAP Server to improve the performance of the lookup operations. Oracle Identity Federation server can be configured to use a Federation Store to persist Federated Identities records.

The Federation server uses this store to:

- Lookup a federation record through different queries
- Create a federation record
- Delete a federation

In addition to the Oracle Identity Federation-related `orclinmemfiltprocess` filter (`objectclass=orclfeduserinfo`), which is included by default, some Oracle Identity Federation environments might benefit from additional filters with the following formats:

```
(orclfedserverid=local_oif_server_id)
(orclfedproviderid=providerid_of_remote_server)
(orclfedfederationtype=n)
```

where `orclfedserverid` denotes the Oracle Identity Federation server that is making the query, `orclfedproviderid` is the identifier of a remote SAML server, and `orclfedfederationtype` is 1 or 3. Use 1 as the value for `orclfedfederationtype` when Oracle Identity Federation is an Identity Provider and the remote provider is a Service Provider. Use 3 when Oracle Identity Federation is a Service Provider and the remote provider is an Identity Provider.

A deployment can be configured to work with many remote SAML servers, so there can be several `orclfedproviderid` filters and more than one `orclfedfederationtype` filter.

For example:

```
(orclfedserverid=my_oif_server)
(orclfedproviderid=http://server.example.com:7499/fed/idp)
(orclfedproviderid=http://server2.example.com:7492/fed/idp)
(orclfedfederationtype=1)
(orclfedfederationtype=3)
```

26.3.2 Tuning SAML Protocol

The Security Assertion Markup Language (SAML) protocol involves interacting with remote servers through the use of the Simple Object Access Protocol (SOAP).

26.3.2.1 SOAP Connections

The Oracle Identity Federation server uses the SOAP protocol to send SAML Requests and to receive SAML Responses.

To optimize performance, configure the following SOAP connections:

- Total maximum number of SOAP connections that Oracle Identity Federation can open at the same time
- Maximum number of SOAP connections that Oracle Identity Federation can open at the same time to a given remote server

For more information, see "SOAP Binding" in *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation*.

26.3.2.2 XML Digital Signatures

The SAML and WS-Fed protocols of Oracle Identity Federation rely on XML Digital Signatures to ensure the authenticity of messages and that messages are not tampered with.

When possible, sign the Assertion and/or the Response to prevent any modifications. When no XML Digital Signature is present on the message, the audited message that is archived does not contain any data that proves the authenticity and integrity of the message.

Configuring Oracle Identity Federation to *not* sign Assertion and/or Response may be appropriate if:

- Performance must be improved
- SSL with SSL authentication is enabled for SOAP communications
- Disabling XML Digital Signatures is compliant with company security regulations

Note: The content of the Assertion is viewable unless SAML 2 Encryption is used. Encrypting the Assertion is optional, but XML Encryption is resource intensive and decreases performance

26.3.2.3 POST and Artifact Single Sign-On Profiles

There are two Single Sign-On profiles defined by the SAML specifications:

- POST Profile

In the POST profile, the Assertion transits through the user's browser, therefore the Assertion and/or the Response must be signed to ensure that the content has not been modified.

Note: If the performance must be improved and if using the POST profile is compliant with company security regulations, then configuring Oracle Identity Federation to use the POST profile may be an option to improve performance.

- Artifact Profile

In the Artifact profile, the Identity Provider creates a random identifier referencing the Assertion in the IdP's local store. (The Assertion is provided directly from the Identity Provider to the Service Provider.) That identifier is carried by the user's browser and presented to the Service Provider that contacts the Identity Provider to de-reference the identifier and retrieve the corresponding Assertion.

If the SOAP connection made from the SP to the IdP is encrypted using the SSL protocol with an SSL Server Certificate, then the SP authenticates the IdP and the

content of the communication has not been tampered with: in this case, the transport layer is providing the authenticity and the integrity of the message, and the XML Digital Signature on the SAML Response and Assertion can be optional.

If no XML Digital Signature is present on the message, then the audited message that is archived does not contain any data that proves the authenticity and integrity of the message.

Note: Since the Artifact profile involves additional communication flow between the Service Provider and the Identity Provider, performance may be slower when using the Artifact profile.

Oracle Fusion Middleware Security Performance Tuning

Oracle Fusion Middleware security services enable you to secure critical applications and sensitive data. This chapter describes how you can configure security services for optimal performance.

This chapter contains the following topics:

- [Section 27.1, "About Security Services"](#)
- [Section 27.2, "Basic Tuning Considerations"](#)
- [Section 27.3, "Tuning Oracle Platform Security Services"](#)
- [Section 27.4, "Oracle Web Services Security Tuning"](#)

27.1 About Security Services

Oracle Fusion Middleware provides security services through Oracle Platform Security Services (OPSS) and Oracle Web Services.

- Oracle Platform Security Services

Oracle Platform Services is a key component of Oracle Fusion Middleware. It offers an integrated suite of security services and is easily integrated with Java SE and Java EE applications that use the Java security model. Security Services includes features that implement user authentication, authorization, and delegation services that developers can integrate into their application environments. Instead of devoting resources to developing these services, application developers can focus on the presentation and business logic of their applications.

Using Oracle Platform Security for Java, applications can enforce fine-grained access control upon resource users. The three key steps are:

- Configure and invoke a login module, as appropriate. You can use provided login modules, or you can use custom login modules.
 - Authenticate the user attempting to log in, which is the role of the identity store service.
 - Authorize the user by checking permissions for any roles the user belongs to for whatever the user is attempting to accomplish, which is the role of the policy store service.
- Oracle Web Services Security

Oracle Web Services Security provides a framework of authorization and authentication for interacting with a web service using XML-based messages.

Note: The information in this chapter assumes that you have reviewed and understand the concepts and administration information for Oracle Fusion Middleware Security Services. For more information, see the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services* before tuning any security parameters.

27.2 Basic Tuning Considerations

This section offers some general guidelines on how to identify a performance bottleneck and how to approach addressing such problems.

If you discover a performance bottleneck, you should first verify that you have addressed the expected traffic load throughout your Web services deployment. If there is a system in the critical path that is at 100% CPU usage, you may simply need to add one or more computers to the cluster.

If there is a bottleneck in your deployment, it is likely to be within one of the following:

- Traffic through a slow connection with an agent
- Latency in connections to third-party queuing systems like JMS

For any of these problems, check the following potential sources:

- Problems with policy assertions that include connections to outside resources, especially the following types:
 - Database Repositories
 - LDAP Repositories
 - Secured Resources
 - Proprietary Security Systems
- Problems with database performance

If you identify one of these as the cause of a bottleneck, you may need to change how you manage your database or LDAP connections or how you secure resources.

27.3 Tuning Oracle Platform Security Services

This section provides the following basic tuning configurations for Oracle Platform Security Services (OPSS):

- [JVM Tuning Parameters](#)
- [LDAP Tuning Parameters](#)
- [Authentication Tuning Parameters](#)
- [Authorization Tuning Properties](#)
- [OPSS PDP Service Tuning Parameters](#)

27.3.1 JVM Tuning Parameters

Tuning the JVM parameters can greatly improve performance. For example, the JVM Heap size should be tuned depending upon the number of roles and permissions in the store. At run time, all roles and permissions are stored in the in-memory cache. For more JVM tuning information, see [Section 2.4, "Tuning Java Virtual Machines \(JVMs\)"](#).

27.3.2 JDK Tuning Parameters

Starting with Java Development Kit 7 (JDK 7), the default keystore size is now 2048 bits. JDK 6 and earlier had a default size of 1024 bits.

When using the Java keytool to generate keystores, the `-keysize` parameter can be used to control the keystore size. Larger keystores provide stronger security, though at the cost of decreased security performance. Consider your environment's use case scenarios to determine if increasing the keystores would negatively impact your security or performance thresholds.

For more information see the JDK 7 release notes at

<http://www.oracle.com/technetwork/java/javase/jdk7-relnotes-418459.html>

27.3.3 LDAP Tuning Parameters

This section covers Lightweight Directory Access Protocol (LDAP) tuning. Oracle supports the management of policies in file-based repositories: Oracle Internet Directory and Oracle Virtual Directory.

If you encounter increased CPU usage due to high SQL execution times, see the following chapters for basic tuning configurations for large deployments:

- Oracle Internet Directory configuration settings can impact performance. For more information, see [Chapter 23, "Oracle Internet Directory Performance Tuning"](#).
- In addition to being configured as a LDAP server, Oracle Virtual Directory can also be configured as a local storage adapter (LSA). See [Chapter 25, "Oracle Virtual Directory Performance Tuning"](#).

27.3.4 Authentication Tuning Parameters

For OPSS Authentication tuning, see "Improving the Performance of WebLogic and LDAP Authentication Providers" in the Oracle Fusion Middleware Securing Oracle WebLogic Server guide at the Oracle Technology Network

http://download.oracle.com/docs/cd/E12840_01/wls/docs103/secmanage/atn.html#wp1199087.

27.3.5 Authorization Tuning Properties

The following Java system properties can be used to optimize authorization:

Table 27–1 Authorization Properties

Java System Properties	Default Value	Valid Values	Notes
-Djps.subject.cache.key	4	3 4 5	<p>JPS uses a Subject Resolver to convert a platform subject to JpsSubject which contains user/enterprise-role information, as well as ApplicationRole information. This information is represented as principals in the subject.</p> <p>This conversion can be CPU intensive, especially if the subject's principal set has a large population. To improve performance, JPS code caches the conversion between Platform subject and JpsSubject. Note that two subjects could be confused if their contents are the same, but the case of the principals' name is different.</p> <p>The following settings can be used to configure the cache key:</p> <ul style="list-style-type: none"> ■ 3: Use the platform subject directly as the key. Note: On WLS if the <code>principalEqualCaseinsensitive</code> flag is enabled, two subjects could be confused if their contents are the same, but the case of the principals is different. ■ 4: This setting is similar to '3' but overcomes the case-sensitive issue. This is the out-of-the-box setting. ■ 5: Instead of using the whole subject as the key, this settings uses a subset of the principal set inside the subject as the key (actually use principals of <code>WLSUserImpl</code> type). <p>This setting will accelerate the cache retrieval operation if the subject has a large principal set. On a non WLS platform (such as WAS and JBOSS, this reverts back to case '4'), so this setting is for WLS only. For this case, there is also a Time To Live setting (TTL) flag which controls how long the cache is valid, as explained below.</p>

Table 27–1 (Cont.) Authorization Properties

Java System Properties	Default Value	Valid Values	Notes
<code>-Djps.subject.cache.ttl</code>	60000ms		<p>Cache's Time To Live (TTL) for case '5' (above). This system property controls how long the cache is valid. When the time expired, the cached value is dumped. The setting can be controlled by the flag of <code>-Djps.subject.cache.ttl=xxxx</code>, where 'xxx' is the duration in milliseconds.</p> <p>Consider setting the duration of this TTL setting to the same value as the value used for the group and user cache TTL in WLS LDAP authenticator.</p>
<code>-Djps.combiner.optimize</code>	False	True False	<p>This system property is used to cache the protection domains for a given subject. Setting <code>-Djps.combiner.optimize=true</code> can improve Java authorization performance.</p>
<code>-Djps.combiner.optimize.lazyeval</code>	False	True False	<p>This system property is used to evaluate a subject's protection domain when a checkPermission occurs. Setting <code>-Djps.combiner.optimize.lazyeval=true</code> can improve Java authorization performance.</p>
<code>-Djps.policystore.hybrid.mode=true</code>	True	True False	<p>This 'hybrid mode' property is used to facilitate transition from SUN <code>java.security.Policy</code> to OPSS Java Policy Provider.</p> <p>The OPSS Java Policy Provider reads from both <code>java.policy</code> and <code>system-jazn-data.xml</code>. "Hybrid" mode can be disabled by setting the system property <code>jps.policystore.hybrid.mode</code> to false when starting the WebLogic Server. Setting <code>-Djps.policystore.hybrid.mode=false</code> can reduce runtime overhead.</p>
<code>-Djps.authz</code>	No default value.	ACC SM	<p>Delegates the call to JDK API <code>AccessController.checkPermission</code> which can reduce the performance impact at run time or while debugging.</p> <p>ACC: delegate to <code>AccessController.checkPermission</code></p> <p>SM: delegate to <code>SecurityManager</code> if <code>SecurityManager</code> is set.</p>

27.3.6 OPSS PDP Service Tuning Parameters

Table 27–2 provides OPSS tuning parameters for policy store:

Table 27–2 OPSS PDP Service Tuning Parameters

Parameter	Default Value	Valid Values	Notes
oracle.security.jps.policystore.rolemember.cache.type	STATIC	STATIC, SOFT, WEAK	<p>This parameter specifies the type of role member cache. Valid only in Java EE applications.</p> <p>Valid values:</p> <ul style="list-style-type: none"> ■ STATIC: Cache objects are statically cached and can be cleaned explicitly only according to the applied cache strategy, such as FIFO. The garbage collector does not clean a cache of this type. ■ SOFT: The cleaning of a cache of this type relies on the garbage collector when there is a memory crunch. ■ WEAK: The behavior of a cache of this type is similar to a cache of type SOFT, but the garbage collector cleans it more frequently. <p>Consider maintaining the default value for the best performance.</p>
oracle.security.jps.policystore.rolemember.cache.strategy	FIFO	FIFO NONE	<p>The type of strategy used in the role member cache. Valid only in Java EE applications.</p> <p>Valid values:</p> <ul style="list-style-type: none"> ■ FIFO: The cache implements the first-in-first-out strategy. ■ NONE: All entries in the cache grow until a refresh or reboot occurs; there is no control over the size of the cache; not recommended but typically efficient when the policy footprint is very small. <p>Consider maintaining the default value for the best performance.</p>
oracle.security.jps.policystore.rolemember.cache.size	1000		<p>The size of the role member cache. The role being referred to is the enterprise role (group). You can find out the number of the groups you have in your ID store first. Then, based on your performance requirement, you can set this number to the number of the groups - full cache scenario. Or you can change to a certain percentage of the number of the groups - partial group cache scenario.</p>
oracle.security.jps.policystore.policy.lazy.load.enable	True	True False	<p>Enables or disables the policy lazy loading. If this parameter is set to false, the server initial startup time will take longer - especially in a large policy store. For faster start-up time, the recommended value is true.</p>

Table 27–2 (Cont.) OPSS PDP Service Tuning Parameters

Parameter	Default Value	Valid Values	Notes
oracle.security.jps.policystore.policy.cache.strategy	PERMISSION_FIFO	PERMISSION_FIFO NONE	The type of strategy used in the permission cache. Valid only in Java EE applications. Valid Values: <ul style="list-style-type: none"> PERMISSION_FIFO: The cache implements the first-in-first-out strategy. NONE: All entries in the cache grow until a refresh or reboot occurs; there is no control over the size of the cache; not recommended but typically efficient when the policy footprint is very small. Consider using the default value for the best performance.
oracle.security.jps.policystore.policy.cache.size	1000		The size of the permission cache. If you cache all policies, then you can set this value to the total number of grants.
oracle.security.jps.policystore.cache.updatable	True	True False	This property is used for refresh enabling. Consider maintaining the default value for the best performance.
oracle.security.jps.policystore.refresh.enable	True	True False	This property is used for refresh enabling. Consider maintaining the default value for performance.
oracle.security.jps.policystore.refresh.purge.timeout	4320000		The time, in milliseconds, after which the policy store is refreshed. Consider maintaining the default value for the best performance.
oracle.security.jps.ldap.policystore.refresh.interval	600000 (10 minutes)		The interval, in milliseconds, at which the policy store is polled for changes. Consider maintaining the default value for the best performance. This property is valid in Java EE and J2SE applications.
oracle.security.jps.policystore.rolemember.cache.warmup.enable	False	True False	This property controls the way the ApplicationRole membership cache is created. If set to True, the cache is created at server startup; otherwise, it is created on demand (lazy loading). Set to True when the number of users and groups is significantly higher than the number of application roles; set to False otherwise, that is, when the number of application roles is very high.

27.4 Oracle Web Services Security Tuning

Oracle Web Services Security provides a framework of authorization and authentication for interacting with a web service using XML-based messages. This

section provides information on factors that might affect performance of the web service.

- [Choosing the Right Policy](#)
- [Policy Manager](#)
- [Configuring the Log Assertion to Record SOAP Messages](#)
- [Monitoring the Performance of Web Services](#)

27.4.1 Choosing the Right Policy

Oracle Web Services Security supports many policies and the appropriate policies must be implemented based on the security need of the deployment. Careful consideration should be given to performance, since each additional policy can impact performance. For example Transport level security (SSL) is faster than Application level security, but transport level security can be vulnerable in multi-step transactions. Application level security has more performance implications, but provides end-to-end security.

See "Configuring Policies" in *Oracle Fusion Middleware Security and Administrator's Guide for Web Services* to determine which security policies are required for a deployment.

27.4.2 Policy Manager

There is an inherent performance impact when using the database-based policy enforcement. When database policy enforcement is chosen, careful consideration must be given to the "polling" frequency of the agent to the database.

27.4.3 Configuring the Log Assertion to Record SOAP Messages

The request and response pipelines of the default policy include a log assertion that causes policy enforcement points (PEP) to record SOAP messages to either a database or a component-specific local file. There can be potential performance impacts to the logging level. To prevent performance issues, consider using the lowest logging level that is appropriate for your deployment.

The following logging levels can be configured in the log step:

- Header - Only the SOAP header is recorded.
- Body - Only the message content (body) is recorded.
- Envelope - The entire SOAP envelope, which includes both the header and the body, is recorded. Any attachments are not recorded.
- All - The full message is recorded. This includes the SOAP header, the body, and all attachments, which might be URLs existing outside the SOAP message itself.

Note: Typically, system performance improves when log files are located in topological proximity to the enforcement component. If possible, use multiple distributed logs in a highly distributed environment.

27.4.4 Configuring Connection Pooling

When you request that a Context instance use connection pooling by using the "com.sun.jndi.ldap.connect.pool" environment property, the connection that is used might or might not be pooled. The default rule is that plain (non-SSL) connections that use simple or no authentication are allowed to be pooled. You can

change this default to include SSL connections and the DIGEST-MD5 authentication type by using system properties. To allow both plain and SSL connections to be pooled, set the "com.sun.jndi.ldap.connect.pool.protocol" system property to the string "plain ssl" as shown below:

```
"-Dcom.sun.jndi.ldap.connect.pool.protocol="plain ssl"
```

27.4.5 Monitoring the Performance of Web Services

You can monitor the performance on the following Oracle Web Services through the Web Services home page of Oracle Fusion Middleware Control:

- Endpoint Enabled Metrics such as:
 - Policy Reference Status
 - Total Violations
 - Security Violations
- Invocations Completed
- Response Time, in seconds
- Policy Violations such as:
 - Total Violations
 - Authentication Violations
 - Authorization Violations
 - Confidentiality Violations
 - Integrity Violations
- Total Faults

For general information on monitoring Oracle Fusion Middleware components, see [Chapter 4, "Monitoring Oracle Fusion Middleware"](#).

For detailed information on using Oracle Fusion Middleware Control to monitor Oracle Web Services, see "Monitoring the Performance of Web Services" in *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*.

Part VI

Oracle WebCenter Components

This part describes configuring Oracle WebCenter components to improve performance. It contains the following chapter:

- [Chapter 28, "Oracle WebCenter Portal Performance Tuning"](#)

Oracle WebCenter Portal Performance Tuning

This chapter outlines how to tune configuration properties for the operating system on which WebCenter Portal applications are installed, WebCenter Portal applications, and their back-end components.

- [Section 28.1, "About Oracle WebCenter Portal"](#)
- [Section 28.2, "Basic Tuning Considerations"](#)
- [Section 28.3, "Tuning WebCenter Portal Application Configuration"](#)
- [Section 28.4, "Tuning Back-End Component Configuration"](#)
- [Section 28.5, "Tuning Identity Store Configuration"](#)
- [Section 28.6, "Tuning Portlet Configuration"](#)

28.1 About Oracle WebCenter Portal

Oracle WebCenter Portal 11g is an integrated suite of products used to create social applications, enterprise portals, communities, composite applications, and internet or intranet Web sites on a standards-based, service-oriented architecture (SOA). Oracle WebCenter Portal combines the development of rich internet applications, a multi-channel portal framework, and a suite of horizontal Enterprise 2.0 applications, which provide content, presence, and social networking capabilities to create a highly interactive user experience. Interacting with services such as instant messaging, blogs, wikis, RSS, tags, discussion forums, activities and social networks directly within the context of a portal or an application improves user and group productivity and enhances the return on IT investments.

Oracle WebCenter Portal: Spaces is an out-of-the-box WebCenter Portal application that brings you the latest technology in terms of social networking, communication, collaboration, and personal productivity with no development effort. Through the robust set of integrated services and applications provided by Oracle WebCenter Portal's Framework, Composer, and Resource Catalog, the Spaces application enables you to deploy instant community portals, team sites and other collaborative applications.

For more information about Oracle WebCenter Portal, see *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal* and *Oracle Fusion Middleware Developer's Guide for Oracle WebCenter Portal*.

28.2 Basic Tuning Considerations

The tuning considerations in this section apply to most WebCenter Portal and WebCenter Portal: Spaces deployments and usage scenarios. It is highly recommended

that you review these configurations and implement those that are appropriate for your use case scenarios.

28.2.1 Setting System Limit

To run a WebCenter Portal application at moderate load, set the `open-files-limit` to 4096. If you encounter errors, such as **running out of file descriptors**, then increase the system limit.

For example, on Linux, you can use this command:

```
ulimit -n 8192
```

Refer to your operating system documentation to find out how to change this system limit.

28.2.2 Setting JDBC Data Source

To determine the correct setting for the JDBC data source, use the Oracle WebLogic Server Administration Console to monitor the running system database connection usage as described "Configuring JDBC Data Sources". If the "Waiting for Connection Failure" rate is noticeably higher, and the "Active Connections Current Count" is close to reaching the maximum capacity, then consider increasing capacity to avoid potential database connection contention.

However, if the "Active Connections Current Count" is routinely lower than the maximum capacity, consider reducing the capacity to save memory.

For more information, see "Configuring Connection Pool Features" in *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server*.

The following data source settings are the defaults for `mds-SpacesDS` and `WebCenterDS`. These settings can be adjusted depending on the application's usage pattern and load.

```
<jdbc-connection-pool-params>
  <initial-capacity>10</initial-capacity>
  <max-capacity>50</max-capacity>
  <capacity-increment>1</capacity-increment>
  <shrink-frequency-seconds>0</shrink-frequency-seconds>
  <highest-num-waiters>2147483647</highest-num-waiters>

<connection-creation-retry-frequency-seconds>0</connection-creation-retry-frequency-seconds>

<connection-reserve-timeout-seconds>60</connection-reserve-timeout-seconds>
  <test-frequency-seconds>0</test-frequency-seconds>
  <test-connections-on-reserve>true</test-connections-on-reserve>

<ignore-in-use-connections-enabled>true</ignore-in-use-connections-enabled>

<inactive-connection-timeout-seconds>0</inactive-connection-timeout-seconds>
  <test-table-name>SQL SELECT 1 FROM DUAL</test-table-name>
  <login-delay-seconds>0</login-delay-seconds>
  <statement-cache-size>5</statement-cache-size>
  <statement-cache-type>LRU</statement-cache-type>
  <remove-infected-connections>true</remove-infected-connections>

<seconds-to-trust-an-idle-pool-connection>60</seconds-to-trust-an-idle-pool-connection>
  <statement-timeout>-1</statement-timeout>
  <pinned-to-thread>>false</pinned-to-thread>
```

```
</jdbc-connection-pool-params>
```

To edit JDBC data source settings:

1. Login to WebLogic Server Administration Console.
2. From the Home page, select **Summary of JDBC Data Sources, Settings for mds-SpacesDS**, and then the **Connection Pool** tab.
3. Edit properties, as required.

To edit WebCenter Portal data source settings:

1. Login to WebLogic Server Administration Console.
2. From the Home page, select **Summary of JDBC Data Sources** and navigate to the **Connection Pool** tab.

See also "Tuning Data Source Connection Pools" in *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server*.

28.2.3 Setting JRockit Virtual Machine (JVM) Arguments

JVM arguments are set in the `setDomainEnv.sh` file on Unix operating systems and `setDomainEnv.cmd` on Windows operating systems. The `setDomainEnv` file is located in the `<domain_dir>/bin` directory.

See Also: [Section 2.4, "Tuning Java Virtual Machines \(JVMs\)"](#)

- **WebLogic Server production mode:** When Webcenter is installed for production deployment, the WebLogic Server is set to production mode. However, if it is installed for development and then switched to production mode for better performance, you need to include the following parameter in the startup command:

```
-Dweblogic.ProductionModeEnabled=true
```

For information on setting your domain to production mode using the Administration Console, see "Change to production mode" in the Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help.

- **Heap size:** If the server is overloaded, that is, garbage is collected or out of memory error occurs frequently, then increase the heap size as appropriate to your server's available physical memory.

For more information, see [Section 2.4.1.1, "Specifying Heap Size Values"](#) and "Set Java options for servers started by Node Manager" in the Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help.

The following parameters can be modified in the server's startup command or through the Administration Console to increase heap size:

```
jrockit vm: -Xms2048M -Xmx2048M -Xns512M
```

```
hotspot vm: -Xms2048M -Xmx2048M -XX:MaxPermSize512M
```

- **Garbage Collection**

In some usecase scenarios modifying the following garbage collection settings may improve performance:

```
-XX:+UseParallelOldGC -XX:+ParallelRefProcEnabled
```

```
-XX:SoftRefLRUPolicyMSPerMB=0
```

Note that the values of `-Xmn` and `SoftRefLRUPolicyMSPerMB(0-1000)` can also be adjusted in accordance with the respective application memory settings and requirements. For example:

```
-XX:PermSize=512m -XX:MaxPermSize=512m -Xmn512m
-Dweblogic.threadpool.MaxPoolSize=25
```

28.2.4 Using Content Compression to Reduce Downloads

If clients connect to your server using relatively slow connections, that is, using modems or VPN from remote locations, consider compressing content before it downloads to the client. While content compression increases the load on the server, the client's download experience is much improved.

Several content compression methods are available. The following steps describe how to use the `mod_deflate` module from Apache.

1. Enable `mod_deflate` module on Apache.

To do this, add the following to `httpd.conf` (`$(OH)/instances/$INSTANCE_NAME/config/OHS/$OHS_NAME`)

```
LoadModule deflate_module "${ORACLE_HOME}/ohs/modules/mod_deflate.so"
```

2. Setup the Output Filter and specify the rules for compression.

Here is a sample snippet that you can add to the `httpd.conf` (same location mentioned above). Modify the content based on your content and the compression requirements.

```
<IfModule mod_deflate.c>
SetOutputFilter DEFLATE
AddOutputFilterByType DEFLATE text/plain
AddOutputFilterByType DEFLATE text/xml
AddOutputFilterByType DEFLATE application/xhtml+xml
AddOutputFilterByType DEFLATE text/css
AddOutputFilterByType DEFLATE application/xml
AddOutputFilterByType DEFLATE image/svg+xml
AddOutputFilterByType DEFLATE application/rss+xml
AddOutputFilterByType DEFLATE application/atom+xml
AddOutputFilterByType DEFLATE application/x-javascript
AddOutputFilterByType DEFLATE text/html
SetEnvIfNoCase Request_URI \.(?:gif|jpe?g|png)$ no-gzip dont-vary
SetEnvIfNoCase Request_URI \.(?:exe|t?gz|zip|bz2|sit|rar)$ no-gzip dont-vary
SetEnvIfNoCase Request_URI \.(?:pdf|doc?x|ppt?x|xls?x)$ no-gzip dont-vary
SetEnvIfNoCase Request_URI \.avi$ no-gzip dont-vary
SetEnvIfNoCase Request_URI \.mov$ no-gzip dont-vary
SetEnvIfNoCase Request_URI \.mp3$ no-gzip dont-vary
SetEnvIfNoCase Request_URI \.mp4$ no-gzip dont-vary
</IfModule>
```

For more information about `mod_deflate`, refer to:

http://httpd.apache.org/docs/2.0/mod/mod_deflate.html

28.3 Tuning WebCenter Portal Application Configuration

This section describes parameters that enable administrators to tune performance of WebCenter Portal applications.

This section includes the following:

- [Setting Session Timeout for a Spaces Application](#)
- [Setting HTTP Session Timeout for a Framework Application](#)
- [Setting JSP Page Timeout](#)
- [Setting ADF Client State Token](#)
- [Setting ADF View State Compression](#)
- [Setting MDS Cache Size and Purge Rate](#)
- [Configuring Concurrency Management](#)

28.3.1 Setting Session Timeout for a Spaces Application

The default session timeout for the Spaces application is 45 minutes. Spaces administrators can customize the session time to suit their installation, for details see "Specifying Session Timeout Settings" in *Oracle Fusion Middleware User's Guide for Oracle WebCenter Portal: Spaces*.

28.3.2 Setting HTTP Session Timeout for a Framework Application

To manage overall resource usage for WebCenter Portal Framework applications, adjust the application's http session timeout value, in minutes, in the `web.xml` file. In general, shorter session timeout values correspond to less memory and CPU usage on the server.

If you must modify this property, post deployment, you must edit `web.xml` manually. See "Editing web.xml Properties" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.

Note: Note that this value applies only to WebCenter Portal Framework applications and not to WebCenter Portal: Spaces.

The following is a sample snippet of `web.xml`:

```
<session-config>
  <session-timeout>
    45
  </session-timeout>
</session-config>
```

28.3.3 Setting JSP Page Timeout

You can specify an integer value, in seconds, after which any JSP page will be removed from memory if it has not been requested in the `web.xml` file. This frees up resources in situations where some pages are called infrequently.

Increasing the value reduces user response time, and decreasing it reduces application memory footprint. The default value for is 600 seconds or 10 minutes. If `jsp_timeout` is not specified, it means there is no timeout.

To modify this property post deployment, you must edit `web.xml` manually. See "Editing web.xml Properties" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.

The following is a sample snippet of `web.xml`:

```
<servlet>
  <servlet-name>
    oraclejsp
    <init-param>
      <param-name>
        jsp_timeout
      </param-name>
      <param-value>
        600
      </param-value>
    </init-param>
```

28.3.4 Setting ADF Client State Token

Through this setting, you can control the number of pages users can navigate using the browser Back button without losing page state. To reduce CPU and memory usage, you can decrease the value in the `web.xml` file.

If you must modify this property, post deployment, you must edit `web.xml` manually. See "Editing `web.xml` Properties" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.

The following is a sample code snippet of `web.xml`:

```
<context-param>
  <param-name>
    org.apache.myfaces.trinidad.CLIENT_STATE_MAX_TOKENS
  </param-name>
  <param-value>
    3
  </param-value>
</context-param>
```

28.3.5 Setting ADF View State Compression

Enabling ADF View State Compression can improve performance for some, but consider the following:

- Enabling `COMPRESS_VIEW_STATE` saves heap usage, but compressing and decompressing can impact CPU resources.
- Enabling `COMPRESS_VIEW_STATE` in a clustered environment can improve performance during session replication as less data has to be sent.
- Enabling `COMPRESS_VIEW_STATE` in a single node environment should only be considered when the CPU/memory usage trade-off is acceptable.

The following is a sample code snippet of enabling the parameter in `web.xml`:

```
<context-param>
  <param-name> org.apache.myfaces.trinidad.COMPRESS_VIEW_STATE
  </param-name>
  <param-value> true </param-value>
</context-param>
```

28.3.6 Setting MDS Cache Size and Purge Rate

The default MDS cache size is 100MB. If you encounter the error message, **JOC region full**, then you can increase the MDS cache size in the `adf-config.xml` file.

Post deployment, modify these properties through the System MBeans Browser. For more information, see the section "Changing MDS Configuration Attributes for Deployed Applications" in *Oracle Fusion Middleware Administrator's Guide*.

The following is a sample snippet of `adf-config.xml`:

```
<cache-config>
<max-size-kb>150000</max-size-kb>
</cache-config>
```

Purging MDS data improves MDS queries. If your portal site changes frequently, you may want to purge old MDS data more often, by reducing the time between purges.

Consider setting the MDS `auto-purge seconds-to-live` parameter (as shown in the example below) to remove older versions of metadata automatically every hour. By default, old versions of metadata are automatically purged every hour, that is, the `auto-purge seconds-to-live` parameter is set to 3600 seconds (as shown in the example below).

Note: Each purge incurs CPU usage in the database. Do not purge too often (for example, every 5 or 10 minutes) because the database CPU impact might outweigh the performance gains from the purge.

If excessive metadata is accumulated and each purge is very expensive, reduce this interval in the `adf-config.xml` file.

By default there is no `auto-purge` entry in `adf-config.xml`. Use the following sample snippet of `adf-config.xml` to modify `auto-purge`:

```
<mdsC:adf-mds-config version="11.1.1.000">
  <mds-config xmlns="http://xmlns.oracle.com/mds/config">
    <persistence-config>
      <metadata-namespaces>
        ...
      </metadata-namespace>
      <auto-purge seconds-to-live="3600"/>
    </persistence-config>
```

To ensure the initial purge does not impact ongoing user activities, consider using the following WLST command to induce an MDS purge immediately before the bulk of the user load hits the system:

The following example shows how to purge all documents in the application repository whose versions are older than 10 seconds:

```
wls:/weblogic/serverConfig>purgeMetadata(application=' [AppName] ', server=' [ServerName] ', olderThan=10)
```

28.3.7 Configuring Concurrency Management

Concurrency management includes global settings that impact the entire WebCenter Portal and service- and resource-specific settings that only impact a particular service.

You can define deployment-specific overrides or additional configuration in the `adf-config.xml` file. For example, you can specify resource-specific (producers) values that are appropriate for a particular deployment.

The following describes the format of the global, service, and resource entries in `adf-config.xml`:

```
<concurrent:adf-service-config
  xmlns="http://xmlns.oracle.com/webcenterportal/concurrent/config">
  <global
    queueSize="SIZE"
    poolCoreSize="SIZE"
    poolMaxSize="SIZE"
    poolKeepAlivePeriod="TIMEPERIOD"
    timeoutMinPeriod="TIMEPERIOD"
    timeoutMaxPeriod="TIMEPERIOD"
    timeoutDefaultPeriod="TIMEPERIOD"
    timeoutMonitorFrequency="TIMEPERIOD"
    hangMonitorFrequency="TIMEPERIOD"
    hangAcceptableStopPeriod="TIMEPERIOD" />
  <service
    service="SERVICENAME"
    timeoutMinPeriod="TIMEPERIOD"
    timeoutMaxPeriod="TIMEPERIOD"
    timeoutDefaultPeriod="TIMEPERIOD" />
  <resource
    service="SERVICENAME"
    resource="RESOURCENAME"
    timeoutMinPeriod="TIMEPERIOD"
    timeoutMaxPeriod="TIMEPERIOD"
    timeoutDefaultPeriod="TIMEPERIOD" />
</concurrent:adf-service-config>
```

Where:

SIZE: A positive integer. For example: 20.

TIMEPERIOD: Any positive integer followed by a suffix indicating the time unit, which must be one of: `ms` for milliseconds, `s` for seconds, `m` for minutes, or `h` for hours. For example: `50ms`, `10s`, `3m`, or `1h`. The following are examples of default settings for different services. These settings are overwritten with any service-specific configurations in `connections.xml` or `adf-config.xml` files:

```
<concurrent:adf-service-config
  xmlns="http://xmlns.oracle.com/webcenter/concurrent/config">
  <service service="oracle.webcenter.community" timeoutMinPeriod="2s"
  timeoutMaxPeriod="50s" timeoutDefaultPeriod="30s"/>
  <resource service="oracle.webcenter.community"
    resource="oracle.webcenter.doclib"
    timeoutMinPeriod="2s" timeoutMaxPeriod="10s" timeoutDefaultPeriod="5s"/>
  <resource service="oracle.webcenter.community"
    resource="oracle.webcenter.collab.calendar.community"
    timeoutMinPeriod="2s" timeoutMaxPeriod="10s" timeoutDefaultPeriod="5s"/>
  <resource service="oracle.webcenter.community"
    resource="oracle.webcenter.collab.rtc"
    timeoutMinPeriod="2s" timeoutMaxPeriod="10s" timeoutDefaultPeriod="5s"/>
  <resource service="oracle.webcenter.community"
    resource="oracle.webcenter.list"
    timeoutMinPeriod="2s" timeoutMaxPeriod="10s" timeoutDefaultPeriod="5s"/>
  <resource service="oracle.webcenter.community"
    resource="oracle.webcenter.collab.tasks"
    timeoutMinPeriod="2s" timeoutMaxPeriod="10s" timeoutDefaultPeriod="5s"/>
</concurrent:adf-service-config>
```

Note: All of the attributes except `service` and `resource` are optional, and therefore, for example, the following tags are valid:

```
<global queueSize="20"/>
<resource service="foo" resource="bar" timeoutMaxPeriod="5s"/>
```

You can use the Enterprise Manager System MBean Browser to view, add, modify, and delete the concurrency configuration based on your usage pattern. To access the MBean Browser for your WebCenter Portal application, see "Accessing the System MBean Browser" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.

1. In System MBean Browser, navigate to:

Application Defined MBeans -> oracle.adf.share.config -> Server: (your server name) -> Application: (your application name) -> ADFConfig -> ADFConfig (bean) -> ADFConfig -> WebCenterConcurrentConfiguration -> Operations -> listResource

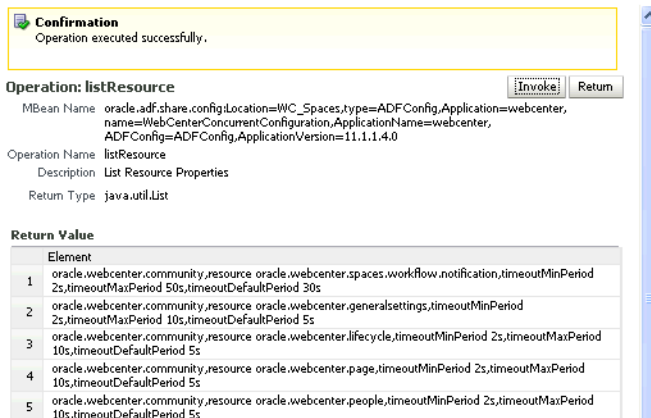
Figure 28–1 System MBean Browser - WebCenterConcurrentConfiguration

The screenshot shows the System MBean Browser interface. The left pane displays a tree view of MBeans, with the path `oracle.adf.share.config > Server: WC_Spaces > Application: webcenter > ADFConfig > ADFConfig > WebCenterConcurrentConfiguration` selected. The right pane shows the 'Operations' tab for the selected MBean, displaying a table of available operations.

Name	Description	Parameters	Return Type
1 listGlobal	List Global Properties	0	java.util.List
2 listResource	List Resource Properties	0	java.util.List
3 listService	List Service Properties	0	java.util.List
4 removeGlobalProperty	Remove Global Property	1	void
5 removeResource	Remove Resource	2	void
6 removeResourceProperty	Remove Resource Property	3	void
7 removeService	Remove Service	1	void
8 removeServiceProperty	Remove Service Property	2	void
9 setGlobal	Set the Global Properties.	2	void
10 setResource	Set the Resource description	4	void
11 setService	Set the Service description	3	void

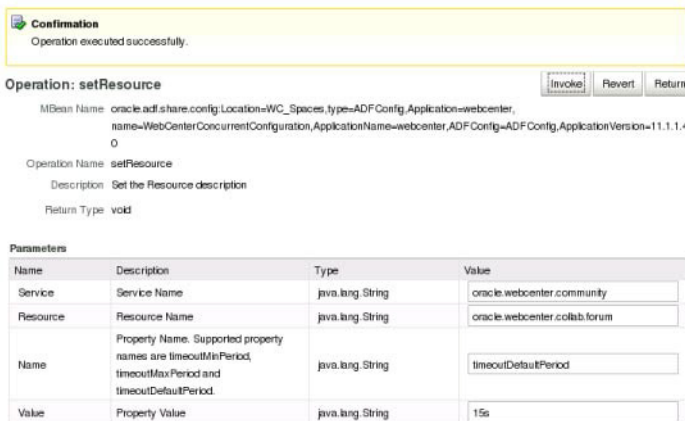
2. To view the current concurrency settings, select `listResource`, and then click **Invoke** (Figure 28–2).

Figure 28–2 System MBean Browser - listResource



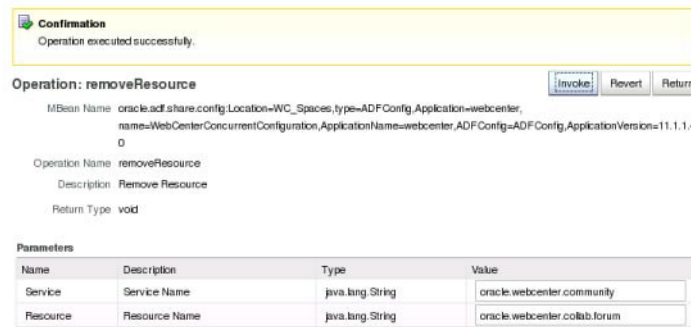
3. To change a setting, select **setResource**, enter the resource details, and then click **Invoke** (Figure 28–3).

Figure 28–3 System MBean Browser - setResource



Take care to enter the correct values for **service**, **resource**, **name** and **value**.

NOTE: If the resource parameter you are attempting to modify already has a [value] setting, you must remove the setting first by invoking the [removeResource] operation (Figure 28–4).

Figure 28–4 System MBean Browser - removeResource

4. To save changes, navigate to **Application Defined MBeans: ADFConfig:ADFConfig** -> **save**, and click **Invoke**.

28.4 Tuning Back-End Component Configuration

This section describes performance configuration for back-end services used by WebCenter Portal applications. Performance of back-ends such as BPEL servers or Oracle WebCenter Content servers, for example, should be tuned as described in guidelines for those back-ends.

This section includes the following sub sections:

- [Tuning Performance of the Announcements Service](#)
- [Tuning Performance of the Discussions Service](#)
- [Tuning Performance of the Instant Messaging and Presence \(IMP\) Service](#)
- [Tuning Performance of the Mail Service](#)
- [Tuning Performance of the Personal Events Service](#)
- [Tuning Performance of the RSS News Feed Service](#)
- [Tuning Performance of the Search Service](#)
- [Tuning Policy Store Parameters](#)

28.4.1 Tuning Performance of the Announcements Service

To manage overall resource usage for the Announcements service, you can tune the Connection Timeout property:

- Default: 10 seconds
- Minimum: 0 seconds
- Maximum: 45 seconds

Post deployment, modify the Connection Timeout property through Fusion Middleware Control or using WLST. For details, see:

- "Modifying Discussions Server Connection Details Using Fusion Middleware Control" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*

- "Modifying Discussions Server Connection Details Using WLST" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*

The following is a sample code snippet of the connections.xml to change the default timeout to 5 seconds:

```
<Reference name="Jive-7777"
className="oracle.adf.mbean.share.connection.webcenter.Announcement.
AnnouncementConnection">
    <Factory
className="oracle.adf.mbean.share.connection.webcenter.forum.ForumConnectionFactory" />
    <REFAddresses>
        <StringRefAddr addrType="connection.time.out">
            <Contents>5</Contents>
        </StringRefAddr>
    </RefAddresses>
</Reference>
```

28.4.2 Tuning Performance of the Discussions Service

To manage overall resource usage for the Discussions service, you can tune the Connection Timeout property:

- Default: 10 seconds
- Minimum: 0 seconds
- Maximum: 45 seconds

Post deployment, modify the Connection Timeout property through Fusion Middleware Control or using WLST. For details, see:

- "Modifying Discussions Server Connection Details Using Fusion Middleware Control" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*
- "Modifying Discussions Server Connection Details Using WLST" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*

The following is a sample snippet of connections.xml:

```
<Reference name="Jive-7777"
className="oracle.adf.mbean.share.connection.webcenter.forum.ForumConnection">
    <Factory
className="oracle.adf.mbean.share.connection.webcenter.forum.ForumConnectionFactory" />
    <REFAddresses>
        <StringRefAddr addrType="forum.url">
            <Contents>http://[machine]:[port]/owc_discussions_
5520</Contents>
        <StringRefAddr addrType="connection.time.out">
            <Contents>5</Contents>
        </StringRefAddr>
    </REFAddresses>
</Reference>
```

28.4.3 Tuning Performance of the Instant Messaging and Presence (IMP) Service

To manage overall resource usage for the IMP service, you can tune the Connection Timeout property:

- Default: 10 seconds

- Minimum: 0 seconds
- Maximum: 45 seconds

Post deployment, modify the Connection Timeout property through Fusion Middleware Control or using WLST. For details, see:

- "Modifying Instant Messaging and Presence Connections Details Using Fusion Middleware Control" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*
- "Modifying Instant Messaging and Presence Connections Details Using WLST" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*

The following is a sample code snippet of the connections.xml to change the default timeout to 5 seconds:

```
<Reference name="IMPService-LCS"
  className="oracle.adf.mbean.share.connection.webcenter.rtc.RtcConnection">
  <Factory
    className="oracle.adf.mbean.share.connection.webcenter.rtc.RtcConnectionFactory"/>
  <RefAddresses>
    <StringRefAddr addrType="connection.time.out">
      <Contents>5</Contents>
    </StringRefAddr>
  </RefAddresses>
</Reference>
```

28.4.4 Tuning Performance of the Mail Service

To manage overall resource usage for the Mail service, you can tune the Connection Timeout property:

- Default: 10 seconds
- Minimum: 0 seconds
- Maximum: 45 seconds

Post deployment, modify the Connection Timeout property through Fusion Middleware Control or using WLST. For details, see:

- "Modifying Mail Server Connection Details Using Fusion Middleware Control" in the *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*
- "Modifying Mail Server Connection Details Using WLST" in the *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*

The following is a sample code snippet of the connections.xml to change the default timeout to 5 seconds:

```
<Reference name="MailConnection"
  className="oracle.adf.mbean.share.connection.webcenter.mail.MailConnection">
  <StringRefAddr addrType="connection.time.out">
    <Contents>5</Contents>
  </StringRefAddr>
</Reference>
```

28.4.5 Tuning Performance of the Personal Events Service

To manage overall resource usage for the Personal Events service, you can tune the Connection Timeout property:

- Default: 10 seconds

- Minimum: 0 seconds
- Maximum: 45 seconds

You can also set a cache expiration period:

- Default: 10 seconds
- Minimum: 0 seconds
- Maximum: 45 seconds

Post deployment, modify the Connection Timeout and Cache Expiration properties through Fusion Middleware Control or using WLST. For details, see:

- "Modifying Event Server Connection Details Using Fusion Middleware Control" in the *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*
- "Modifying Event Server Connection Details Using WLST" in the *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*

The following is a sample code snippet of the connections.xml to change the default timeout to 5 seconds:

```
<Reference
name="MSExchange-my-pc"className="oracle.adf.mbean.share.connection.webcenter.calendar.PersonalEventConnection">
  <Factory
className="oracle.adf.mbean.share.connection.webcenter.calendar.PersonalEventConnectionFactory"/>
  <StringRefAddr addrType="eventservice.connection.timeout">
    <Contents>5</Contents>
  </StringRefAddr>
  <StringRefAddr addrType="eventservice.cache.expiration.time">
    <Contents>5</Contents>
  </StringRefAddr>
</RefAddresses>
</Reference>
```

28.4.6 Tuning Performance of the RSS News Feed Service

To manage overall resource usage for the RSS News Feed service, you can adjust the refresh interval and timeout in the adf-config.xml file.

If you must modify these properties, post deployment, use the System MBeans Browser.

The following is a sample snippet of adf-config.xml:

```
<rssC:adf-rss-config>
  <rssC:RefreshSecs>3600</rssC:RefreshSecs>
  <rssC:TimeoutSecs>3</rssC:TimeoutSecs>
  <rssC:Configured>true</rssC:Configured>
</rssC:adf-rss-config>
```

28.4.7 Tuning Performance of the Search Service

To manage overall resource usage and user response time for searching, you can adjust the number of saved searches displayed, the number of results displayed, and these timeout values:

- prepareTimeoutMs - Maximum time that a service is allowed to initialize a search (in ms).

- `timeoutMs` - Maximum time that a service is allowed to execute a search (in ms).
- `showAllTimeoutMs` - Maximum time that a service is allowed to display search all results (in ms).

Post deployment, modify timeout properties through Fusion Middleware Control or using WLST. For details, see:

- "Modifying Oracle SES Connection Details Using Fusion Middleware Control" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.
- "Modifying Oracle SES Connection Details Using WLST" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.

The following is a sample snippet of `adf-config.xml`:

```
<searchC:adf-search-config
xmlns="http://xmlns.oracle.com/webcenter/search/config">
  <display-properties>
    <common numSavedSearches="25" />
    <region-specific>
      <usage id="simpleSearchResultUIMetadata" numServiceRows="5" />
      <usage id="searchResultUIMetadata" numServiceRows="5" />
      <usage id="localToolbarRegion" numServiceRows="5" />
    </region-specific>
  </display-properties>
  <execution-properties prepareTimeoutMs="1000" timeoutMs="3000"
showAllTimeoutMs="20000" />
</execution-properties>
</searchC:adf-search-config>
```

28.4.8 Tuning Policy Store Parameters

If you are experiencing performance issues post login, especially in the area of permission checks, you may need to tune the policy store parameters as described in [Section 27.3.6, "OPSS PDP Service Tuning Parameters"](#). Depending on your use case scenarios, performance of WebCenter Portal and WebCenter Portal: Spaces, specifically, can be improved by modifying the following parameters:

- Set `oracle.security.jps.policystore.rolemember.cache.warmup.enable` to `True`
- Modify `oracle.security.jps.policystore.rolemember.cache.size` based on the number of active groups you expect to have in your WebCenter Portal - Spaces environment.

NOTE: This parameter should only be modified if you expect to have more than 3000 active Spaces in your WebCenter Portal: Spaces environment.

- Set `oracle.security.jps.policystore.policy.cache.size` to 5 times the number of group spaces

Note: Always refer to your own use case scenarios before modifying the policy store parameters. For more information, see the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services* before tuning any security parameters.

28.5 Tuning Identity Store Configuration

The following sections describe performance-related configurations that may be required for specific environments.

This section includes the following subsections:

- [Section 28.5.1, "Tuning the Identity Store when Using SSL"](#)
- [Section 28.5.2, "Tuning Performance when Using OVD"](#)
- [Section 28.5.3, "Tuning Performance when Using Active Directory"](#)

28.5.1 Tuning the Identity Store when Using SSL

When you configure an identity store with WebCenter Portal (using WebLogic Server providers), you can choose to configure either an SSL port or a non-SSL port. If you choose an SSL port, by default, the JNDI connections are not pooled causing increased response time and decreased performance when looking up users, groups, or other identity store entities. To address this, do the following:

1. Open the `jps-config.xml` file under `domain_home/config/fmwconfig/jps-config.xml`, locate the `idstore.ldap` service instance and add the line highlighted below:

```
<!-- JPS WLS LDAP Identity Store Service Instance -->
  <serviceInstance name="idstore.ldap" provider="idstore.ldap.provider">
    <property name="idstore.config.provider"
value="oracle.security.jps.wls.internal.idstore.WlsLdapIdStoreConfigProvider"/>
    <property name="CONNECTION_POOL_CLASS"
value="oracle.security.idm.providers.stdldap.JNDIPool"/>
    <property name="java.naming.ldap.factory.socket"
value="javax.net.ssl.SSLSocketFactory"/>
  </serviceInstance>
```

2. Restart all the servers within the domain that are connected to the identity store on an SSL port with the following JVM parameter:

```
-Dcom.sun.jndi.ldap.connect.pool.protocol=ssl
```

You can specify this by modifying `setDomainEnv.sh` or directly from the console.

3. Verify that the servers are running with this JVM parameter, and then (for *nix systems) run the following `grep` command:

```
ps -aef | grep WC_Spaces
```

and verify that the process state specifies

```
com.sun.jndi.ldap.connect.pool.protocol=ssl.
```

28.5.2 Tuning Performance when Using OVD

For OVD, the only object class against which attributes are looked up is `inetOrgPerson` (and its parent object classes). Since the Profile Gallery can display attributes not defined in `inetOrgPerson`, all the additional attributes not covered in `inetOrgPerson` would require an additional round trip to the identity store.

For best performance when using OVD in a production environment, Oracle recommends that you add the following configuration entry (in bold) to the domain-level `jps-config.xml` file:

```
<!-- JPS WLS LDAP Identity Store Service Instance -->
```

```

<serviceInstance name="idstore.ldap"
  provider="idstore.ldap.provider">
  <property name="idstore.config.provider"
value="oracle.security.jps.wls.internal.idstore.WlsLdapIdStoreConfigProvider" />
  <property name="CONNECTION_POOL_CLASS"
value="oracle.security.idm.providers.stldldap.JNDIPool" />

  <extendedProperty>
    <name>user.object.classes</name>
    <values>
      <value>top</value>
      <value>person</value>
      <value>inetorgperson</value>
      <value>organizationalperson</value>
      <value>orcluser</value>
      <value>orcluserv2</value>
      <value>ctCalUser</value>
    </values>
  </extendedProperty>
</serviceInstance>

```

28.5.3 Tuning Performance when Using Active Directory

For best performance when using Active Directory in a production environment, Oracle recommends that you add the following configuration entries (in bold) to the domain-level `jps-config.xml` file:

```

<serviceInstance provider="idstore.ldap.provider"
  name="idstore.ldap">
  <property
value="oracle.security.jps.wls.internal.idstore.WlsLdapIdStoreConfigProvider"
  name="idstore.config.provider" />
  <property value="oracle.security.idm.providers.stldldap.JNDIPool"
  name="CONNECTION_POOL_CLASS" />
  <property name="PROPERTY_ATTRIBUTE_MAPPING" value="WIRELESS_ACCT_
NUMBER=mobile:MIDDLE_NAME=middlename:MAIDEN_NAME=sn:DATE_OF_HIRE=pwdLastSet:NAME_
SUFFIX=generationqualifier:DATE_OF_BIRTH=pwdLastSet:DEFAULT_GROUP=primaryGroupID"
  />
  <property value="sAMAccountName" name="username.attr" />
  <property value="sAMAccountName" name="user.login.attr" />
</serviceInstance>

```

The People Profile Service queries for all these attributes and there is no default mapping for these attributes in the Active Directory provider. A vanilla Active Directory installation doesn't have any mapping corresponding to `DATE_OF_HIRE`, `DATE_OF_BIRTH`.

Note that the two attributes are simply a mapping to some attribute of the correct data type to reduce unnecessary LDAP server calls as Active Directory really doesn't have corresponding attributes with the same semantic meaning.

28.6 Tuning Portlet Configuration

This section describes portlet performance-related configuration. This section includes the following sub sections:

- [Tuning Performance of the Portlet Service](#)
- [Configuring Portlet Cache Size](#)

- [Enabling Java Object Cache for WSRP Producers](#)
- [Suppressing Optimistic Rendering for WSRP Portlets](#)
- [Tuning Performance of Oracle PDK-Java Producers](#)
- [Setting Portlet Container Runtime Options](#)
- [Setting WSRP Attribute for Portlet-served Resources](#)
- [Setting WSRP Attribute for Resources Not Served by the Portlet](#)
- [Configuring Portlet Timeout](#)
- [Tuning Performance of OmniPortlet](#)

28.6.1 Tuning Performance of the Portlet Service

To manage overall resource usage and user response time, you can remove unnecessary locale support, modify portlet timeout and cache size in the `adf-config.xml` file.

For the Portlet service, 28 supported locales are defined out-of-the-box. You can remove the locales that are unnecessary for your application.

If you must modify these properties, post deployment, you must edit `adf-config.xml` manually. See "Editing `adf-config.xml`" in the *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.

The following is a sample snippet of `adf-config.xml`:

```
<portletC:adf-portlet-config xmlns="http://xmlns.oracle.com/adf/portlet/config">
  <supportedLocales>
    <value>es</value>
    <value>ko</value>
    <value>ru</value>
    <value>ar</value>
    <value>fi</value>
    <value>nl</value>
    <value>sk</value>
    <value>cs</value>
    <value>fr</value>
    <value>no</value>
    <value>sv</value>
    <value>da</value>
    <value>hu</value>
    <value>pl</value>
    <value>th</value>
    <value>de</value>
    <value>it</value>
    <value>pt</value>
    <value>tr</value>
    <value>el</value>
    <value>iw</value>
    <value>pt_BR</value>
    <value>zh_CN</value>
    <value>en</value>
    <value>ja</value>
    <value>ro</value>
    <value>zh_TW</value>
  </supportedLocales>
  <defaultTimeout>20</defaultTimeout>
  <minimumTimeout>1</minimumTimeout>
</portletC:adf-portlet-config>
```

```

<maximumTimeout>300</maximumTimeout>
<parallelPoolSize>10</parallelPoolSize>
<parallelQueueSize>20</parallelQueueSize>
<cacheSettings enabled="true">
  <maxSize>10000000</maxSize>
</cacheSettings>
</portletC:adf-portlet-config>

```

28.6.2 Configuring Portlet Cache Size

You can modify the portlet cache size in the `adf-config.xml` file. The default portlet cache size is set to 10 MB.

If you must modify these properties, post deployment, you must edit `adf-config.xml` manually.

The following is a sample snippet of `adf-config.xml`:

```

<adf-portlet-config>
  ....
  <supportedLocales>
  <cacheSettings enabled="true">
<maxSize>10000000</maxSize>
  </cacheSettings>
</adf-portlet-config>

```

28.6.3 Enabling Java Object Cache for WSRP Producers

Oracle recommends that you enable the Java Object Cache (JOC) for WSRP producers so that objects written to the persistent store are cached.

The following is a sample snippet of `web.xml`:

```

<env-entry>
<env-entry-name>oracle/portal/wsrp/server/enableJavaObjectCache</env-entry-name>
<env-entry-type>java.lang.String</env-entry-type>
<env-entry-value>>false</env-entry-value>
</env-entry>

```

28.6.4 Suppressing Optimistic Rendering for WSRP Portlets

To suppress the optimistic render of WSRP portlets after a WSRP `PerformBlockingInteraction` or `HandleEvents` call, set the Portlet container runtime option (specified in `portlet.xml`) as follows:

```
com.oracle.portlet.suppressWsrpOptimisticRender=true.
```

- `true` - optimistic render always suppressed
- `false` - optimistic render may be performed

Normally, if a WSRP portlet receives a WSRP `PerformBlockingInteraction` request (processAction in JSR168/JSR286 portlets) and the portlet does not send any events as a result, the WSRP producer renders the portlet and returns the portlet's markup in the response to the `PerformBlockingInteraction` SOAP message. This markup may be cached by the consumer until the consumer's page renders, and if nothing else affecting the state of the portlet happens (such as the portlet receiving an event), the cached markup can be used by the consumer, eliminating the need for a second SOAP call to `GetMarkup`.

This assumes that the portlet's render phase is idempotent, which is always a best practice. However, if the portlet expects to receive an event, or rendering the portlet is

more costly than a second SOAP message for `GetMarkup`, the developer may use this container option to suppress the optimistic render of the portlet after a `processAction` or `handleEvent` call. The portlet still renders normally when the producer receives the WSRP `GetMarkup` request.

28.6.5 Tuning Performance of Oracle PDK-Java Producers

To manage overall resource usage for a Web producer, you can tune the `Connection Timeout` property:

- Default: 30000 ms
- Minimum: 5000 ms
- Maximum: 60000 ms

Post deployment, modify the `Connection Timeout` property through Fusion Middleware Control or using WLST. For details, see:

- "Editing Producer Registration Details Using Fusion Middleware Control" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.
- "Editing Producer Registration Details Using WLST" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.

The following is a sample snippet of `connections.xml`:

```
<webproducerconnection producerName="wc-WebClipping"
urlConnection="wc-WebClipping-urlconn" timeout="10000" establishSession="true"
mapUser="false"/>
```

28.6.6 Setting Portlet Container Runtime Options

You can use the WebCenter Portal-specific `excludedActionScopeRequestAttributes` container runtime option to specify how to store action-scoped request attributes so that they are available to portlets until a new action occurs.

Request attributes which match any of the regular expressions are not stored as action-scoped request attributes if the `javax.portlet.actionScopedRequestAttributes` container runtime option is used, in addition to any request parameters whose values match the regular expressions defined in the `com.oracle.portlet.externalScopeRequestAttributes` container runtime option.

If set to `true`, you can specify a second value of `numberOfCachedScopes` and a third value indicating the number of scopes to be cached by the portlet container.

28.6.7 Excluding Request Attributes for Portlets

The `excludedActionScopeRequestAttributes` is a multi-valued, Portlet container runtime property, where each value is a regular expression.

If using the `javax.portlet.actionScopedRequestAttributes` container runtime option with a portlet, it is possible to optimize which request attributes get stored between portlet lifecycles using the `com.oracle.portlet.excludedActionScopeRequestAttributes` container runtime option. Any request attributes which are unnecessary to store between lifecycles can be indicated to increase performance.

Consider setting the Portlet container runtime option (specified in `portlet.xml`) as follows:

```
com.oracle.portlet.excludedActionScopeRequestAttributes
```

Default values:

```
javax\portlet.*
oracle\portlet.*
com\oracle\portlet.*
```

28.6.8 Setting WSRP Attribute for Portlet-served Resources

To specify the default WSRP `requiresRewrite` flag to use when generating Resource URLs for portlet-served resources, set the Portlet container runtime option (specified in `portlet.xml`) as follows:

```
com.oracle.portlet.defaultServedResourceRequiresWsrpRewrite.
```

This setting is used for all ResourceURLs created by the portlet, unless overridden by the presence of the `oracle.portlet.server.resourceRequiresRewriting` request attribute when the ResourceURL methods `write()` or `toString()` are called. This setting is also used to specify the WSRP `requiresRewriting` flag on the served resource response, but can be overridden by the presence of the `oracle.portlet.server.resourceRequiresRewriting` request attribute when the portlet's `serveResource()` method returns.

Valid values:

- `unspecified` - (Default) The `requiresRewrite` URL flag is not given a value, and the `requiresRewriting` response flag for a `serveResource` operation is based on the MIME type of the response.
- `true` - The `requiresRewrite` URL flag and `requiresRewriting` response flag is set to `true`, indicating that the resource should be rewritten by the consumer.
- `false` - The `requiresRewrite` URL flag and `requiresRewriting` response flag is set to `false`, indicating that the resource does not necessarily need to be rewritten by the consumer, though the consumer may choose to rewrite the resource.

28.6.9 Setting WSRP Attribute for Resources Not Served by the Portlet

To specify the default WSRP `requiresRewrite` flag to use when encoding URLs for resources not served by the portlet, set the Portlet container runtime option (specified in `portlet.xml`) as follows:

```
com.oracle.portlet.defaultProxiedResourceRequiresWsrpRewrite.
```

This setting is used for all URLs returned by the `PortletResponse.encodeURL()` method, unless overridden by the presence of the `oracle.portlet.server.resourceRequiresRewriting` request attribute when the `PortletResponse.encodeURL()` method is called.

Valid values:

- `true` - (Default) The `requiresRewrite` URL flag is set to `true`, indicating that the resource should be rewritten by the consumer.
- `false` - The `requiresRewrite` URL flag is set to `false`, indicating that the resource does not necessarily need to be rewritten by the consumer.

28.6.10 Configuring Portlet Timeout

You can modify the portlet timeout value in the `adf-portlet-config` element of the `adf-config.xml` file. Default: 10 seconds, minimum: 0.1 seconds, maximum: 60 seconds.

If you must modify these properties, post deployment, you must edit `adf-config.xml` manually. See "Editing `adf-config.xml`" in the *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.

The following is a sample snippet of `adf-config.xml`:

```
<adf-portlet-config>
    . . . .
    <defaultTimeout>5</defaultTimeout>
    <minimumTimeout>2</minimumTimeout>
    <maximumTimeout>300</maximumTimeout>
</adf-portlet-config>
```

28.6.11 Tuning Performance of OmniPortlet

To manage overall resource usage for OmniPortlets, you can tune the Connection Timeout property:

- Default: 30000 ms
- Minimum: 5000 ms
- Maximum: 60000 ms

Post deployment, modify the Connection Timeout property through Fusion Middleware Control or using WLST. For details, see:

- "Editing Producer Registration Details Using Fusion Middleware Control" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.
- "Editing Producer Registration Details Using WLST" in *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.

The following is a sample snippet of `connections.xml`:

```
<webproducerconnection producerName="wc-OmniPortlet"
urlConnection="wc-OmniPortlet-urlconn" timeout="10000" establishSession="false"
mapUser="false"/>
```


Part VII

Capacity Planning, Scalability, and Availability

This part describes how to plan your site for high traffic, scalability, and availability. It contains the following chapters:

- [Chapter 29, "Capacity Planning"](#)
- [Chapter 30, "Using Clusters and High Availability Features"](#)

Capacity Planning

Capacity Planning is the process of determining what type of hardware and software configuration is required to meet application needs. Like performance planning, capacity planning is an iterative process. A good capacity management plan is based on monitoring and measuring load data over time and implementing flexible solutions to handle variances without impacting performance.

Note: The information contained in this chapter is meant to provide an overview of various techniques that can be used to develop an effective capacity management plan. The steps you take - and the plan you ultimately create - depends on your specific requirements and deployment structure.

The following sections provide an introduction to capacity planning:

- [Section 29.1, "About Capacity Planning for Oracle Fusion Middleware"](#)
- [Section 29.2, "Determining Performance Goals and Objectives"](#)
- [Section 29.3, "Measuring Your Performance Metrics"](#)
- [Section 29.4, "Identifying Bottlenecks in Your System"](#)
- [Section 29.5, "Implementing a Capacity Management Plan"](#)

29.1 About Capacity Planning for Oracle Fusion Middleware

While performance tuning can be defined as optimizing your *existing* system for better performance, capacity planning determines what your system needs (and when it needs it) to maintain performance in both steady-state and peak usage periods.

Capacity Planning involves designing your solution and testing the configuration, as well as identifying business expectations, periodic fluctuations in demand, and application constraints. You need to plan carefully, test methodically, and incorporate design principles that focus on performance. Before deploying any application into a production environment, the application should be put through a rigorous performance testing cycle. Creating an effective Capacity Management plan includes some of the same steps as performance planning:

- Step 1: [Determining Performance Goals and Objectives](#)
- Step 2: [Measuring Your Performance Metrics](#)
- Step 3: [Identifying Bottlenecks in Your System](#)
- Step 4: [Implementing a Capacity Management Plan](#)

29.1.1 Capacity Planning Factors to Consider

Before you can create a plan, you must have the data to support your deployment strategy. The following list of questions should be asked - and the information you receive should be analyzed carefully - to ensure a successful capacity management plan.

Table 29–1 Capacity Planning Factors to Consider

Capacity Planning Questions	For more information see,
What are your performance goals and objectives?	Section 29.2, "Determining Performance Goals and Objectives"
How many users need to run simultaneously (concurrently)?	Section 29.2, "Determining Performance Goals and Objectives"
Is the simulated workload adequate? (Is the workload likely to increase?)	Section 29.2, "Determining Performance Goals and Objectives"
Is the Oracle Fusion Middleware deployment configured to support clustering and other high availability factors?	Section 29.4.1, "Using Clustered Configurations"
Does the hardware meet the configuration requirements?	Section 29.5.1, "Hardware Configuration Requirements"
Do you have adequate JVMs to support your users?	Section 29.5.2, "JVM Requirements"
Is the database a limiting factor?	Section 29.5.4, "Database Configuration"

29.2 Determining Performance Goals and Objectives

The first step in creating an effective capacity management plan is to determine your network load and performance objectives. You need to understand the applications deployed and the environmental constraints placed on the system. Ideally you have information about the levels of activity that components of the application are expected to meet, such as:

- The anticipated number of users.
- The number of concurrent sessions.
- The number of SSL connections required.
- The number and size of requests.
- The amount of data and its consistency.
- Determining your target CPU utilization.

Performance objectives are limited by constraints, such as

- The configuration of hardware and software such as CPU type, disk size versus disk speed, sufficient memory.
- The ability to interoperate between domains, use legacy systems, support legacy data.
- The security requirements and use of SSL. SSL involves intensive computing operations and supporting the cryptography operations in the SSL protocol can impact the performance of the WebLogic Server.
- Development, implementation, and maintenance costs.

You can use this information to set realistic performance objectives for your application environment, such as response times, throughput, and load on specific hardware.

29.3 Measuring Your Performance Metrics

After you have determined your performance criteria in [Section 29.2, "Determining Performance Goals and Objectives"](#), take measurements of the metrics you can use to quantify your performance objectives. Benchmarking key performance indicators provides a performance baseline. See [Chapter 4, "Monitoring Oracle Fusion Middleware"](#) for information on measuring your performance metrics with Oracle Fusion Middleware applications.

29.4 Identifying Bottlenecks in Your System

Bottlenecks, or areas of marked performance degradation, should be addressed while developing your capacity management plan. If possible, profile your applications to pinpoint bottlenecks and improve application performance. Oracle provides the following profilers:

- Oracle Jrockit Mission Control provides profiling capabilities for processes using Jrockit JVM.
<http://www.oracle.com/technology/products/jrockit/missioncontrol/index.html>
- Oracle Application Diagnostics provides profiling capabilities for java processing using SUN JDK.
<http://www.oracle.com/technology/software/products/oem/htdocs/jade.html>

The objective of identifying bottlenecks is to meet your performance goals, not eliminate all bottlenecks. Resources within a system are finite. By definition, at least one resource (CPU, memory, or I/O) can be a bottleneck in the system. Planning for anticipated peak usage, for example, may help minimize the impact of bottlenecks on your performance objectives.

There are several ways to address system bottlenecks. Some common solutions include:

- [Using Clustered Configurations](#)
- [Using Connection Pooling](#)
- [Setting the Max HeapSize on JVM](#)
- [Increasing Memory or CPU](#)
- [Segregation of Network Traffic](#)
- [Segregation of Processes and Hardware Interrupt Handlers](#)

29.4.1 Using Clustered Configurations

Clustered configurations distribute work loads among multiple identical cluster member instances. This effectively multiplies the amount of resources available to the distributed process, and provides for seamless fail over for high availability.

For more information see [Chapter 30, "Using Clusters and High Availability Features"](#).

29.4.2 Using Connection Pooling

You may be able to improve performance by using existing database connections. You can limit the number of connections, timing of the sessions and other parameters by modifying the connection strings.

See [Section 2.7, "Reusing Database Connections"](#) for more information on configuring the database connection pools.

29.4.3 Setting the Max HeapSize on JVM

This is a application-specific tunable that enables a trade off between garbage collection times and the number of JVMs that can be run on the same hardware. Large heaps are used more efficiently and often result in fewer garbage collections. More JVM processes offer more fail over points.

See [Section 2.4, "Tuning Java Virtual Machines \(JVMs\)"](#) for more information.

29.4.4 Increasing Memory or CPU

Aggregating more memory and/or CPU on a single hardware resource allows localized communication between the instances sharing the same hardware. More physical memory and processing power on a single machine enables the JVMs to scale and run much larger and more powerful instances, especially 64-bit JVMs. Large JVMs tend to use the memory more efficiently, and Garbage Collections tend to occur less frequently. In some cases, adding more CPU means that the machine can have more instruction and data cache available to the processing units, which means even higher processing efficiency.

See [Section 2.2, "Securing Sufficient Hardware Resources"](#) for more information.

29.4.5 Segregation of Network Traffic

Network-intensive applications can introduce significant performance issues for other applications using network. Segregating the network traffic of time-critical applications from network-intensive applications, so that they get routed to different network interfaces, may reduce performance impacts. It is also possible to assign different routing priorities to the traffic originating from different network interfaces.

29.4.6 Segregation of Processes and Hardware Interrupt Handlers

When planning for the capacity that a specific hardware resource can handle, it is important to understand that the operating system may not be able to efficiently schedule the JVM processes as well as other system processes and hardware interrupt handlers. The JVM may experience performance impacts if it shares even a few of its CPU cores with the hardware interrupt handlers. For example, disk and network-intensive applications may induce performance impacts that are disproportionate to the load experienced by the CPU. In addition, hardware interrupts can prevent the active Java threads from reaching a "GC-safe point" efficiently. Separating frequent hardware interrupt handlers from the CPUs running the JVM process can reduce the wait for Garbage Collections to start.

It may also be beneficial to dedicate sibling CPUs on a multi-core machine to a single JVM to increase the efficiency of its CPU cache. If multiple processes have to share the CPU, the data and instruction cache can be contaminated with the data and instructions from both processes, thus reducing the amount of the cache used effectively. Assigning the processes to specific CPU cores, however, can make it impossible to use other CPU cores during peak load bursts. The capacity management

plan should include a determination on whether the CPUs should be used more efficiently for the nominal load, or should there be some extra capacity for a burst of activity.

29.5 Implementing a Capacity Management Plan

Once you have defined your performance objectives, measured your workload, and identified any bottlenecks, you must create and implement a capacity management plan. The goal of your plan should be to meet or exceed your performance objectives (especially during peak usage periods) and to allow for future workload increases. To achieve your performance objectives, you must implement your management plan and then continuously monitor the performance metrics as discussed in [Chapter 4, "Monitoring Oracle Fusion Middleware"](#).

Since no two deployments are identical, it's virtually impossible to illustrate how a capacity management plan would be implemented for all configurations. Capacity planning is an iterative process and your plan must be calibrated as changes in your workload or environment change. The following section provides key factors that should be addressed in the plan:

29.5.1 Hardware Configuration Requirements

There is no single formula for determining your hardware requirements. The process of determining what type of hardware and software configuration involves assessment of your system performance goals and an understanding of your application. Capacity planning for server hardware should focus on maximum performance requirements.

The hardware requirements you have today are likely to change. Your plan should allow for workload increases, environment changes (such as added servers or 3rd party services), software upgrades (operating systems, middleware or other applications), network connectivity and network protocols.

29.5.1.1 CPU Requirements

Your target CPU usage should not be 100%, you should determine a target CPU utilization based on your application needs, including CPU cycles for peak usage. If your CPU utilization is optimized at 100% during normal load hours, you have no capacity to handle a peak load. In applications that are latency sensitive and maintaining the ability for a fast response time is important, high CPU usage (approaching 100% utilization) can reduce response times while throughput stays constant or even increases because of work queuing up in the server. For such applications, a 70% - 80% CPU utilization is recommended. A good target for non-latency sensitive applications is about 90%.

29.5.1.2 Memory Requirements

Memory requirements are determined by the optimal heap size for the applications you are going to use, for each JVM co-located on the same hardware. Each JVM needs up to 500MB in addition to the optimal heap size; the actual impact to performance depends on the JVM brand, and on the type of application being run. For example, applications with more Java classes loaded need more space for compiled classes. 32-bit JVMs normally cannot exceed a limit of approximately 3GB on some architecture when a limit is imposed by the hardware architecture and the Operating System. It is recommended to reserve some memory for the Operating System, IO buffers and shared-memory devices.

29.5.2 JVM Requirements

The number of users/processes that a single Java Virtual Machine (JVM) can handle varies widely on the types of requests and the type of JVM you are running. As part of your performance monitoring and benchmarking procedures, you should determine how many and what kinds of processes are executed and determine if your hardware meets the requirements for your specific JVM.

29.5.3 Managed Servers

Using multiple managed servers across multiple nodes in a clustered configuration is recommended for both high performance and reliability. It is important to note, however, that having multiple managed servers may mean using more memory which can enable some applications to optimize certain operations in-memory, therefore reducing impact of disk, database and network latency.

For more information on using clustered configurations, see "Understanding Managed Servers and Managed Server Clusters" in *Oracle Fusion Middleware Administrator's Guide*.

29.5.4 Database Configuration

To maintain sustained performance, you must ensure that your existing database can scale with the increases in capacity planned for the application server tier. Tuning the database parameters and monitoring database metrics during peak usage, can help you determine if the existing database resources can scale to handle increased loads. You may need to add additional memory or upgrade the database hardware configuration. For more information on tuning an Oracle database, see the *Oracle Database Performance Tuning Guide*.

In some cases, however, you may find that the database is still not able to effectively manage increases in load, even after increasing the memory or upgrading the CPU. In these situations, consider deploying an Oracle Real Application Cluster (Oracle RAC) environment to handle the increases. Oracle RAC configurations not only provide enhanced performance, but they can also improve reliability and scalability. For more information on Oracle RAC, see *Oracle Real Application Clusters Administration and Deployment Guide*.

Using Clusters and High Availability Features

A high availability architecture is one of the key requirements for any Enterprise Deployment. Oracle Fusion Middleware has an extensive set of high availability features, which protect its components and applications from unplanned down time and minimize planned downtime.

This chapter provides an overview of the architecture, interaction, and dependencies of Oracle Fusion Middleware components, and explains how they can be deployed in a high availability architecture to maximize performance.

This chapter includes the following sections:

- [Section 30.1, "About Clusters and High Availability Features"](#)
- [Section 30.2, "Using Clusters with Oracle Fusion Middleware"](#)
- [Section 30.3, "Using High Availability Features with Oracle Fusion Middleware"](#)

Note: Using clusters and other high availability options is a complex and detailed process. This chapter is meant to introduce the concepts as they relate to Oracle Fusion Middleware. [Table 30–1](#) provides a list of Oracle Fusion Middleware guides that contain detailed high availability information.

30.1 About Clusters and High Availability Features

One of the most important factors in both high availability and performance is the use of **clusters**. A cluster is a set of processes running on single or multiple computers that share the same workload. Using a clustered configuration promotes scalability, high availability, and performance.

High availability refers to the ability of users to access a system without loss of service. Deploying a high availability system minimizes the time when the system is down, or unavailable and maximizes the time when it is running, or available. See

Details about using clusters and other high availability features can be located in the application-specific guides listed in [Table 30–1](#):

Table 30–1 Clusters and High Availability Information in Oracle Fusion Middleware Documentation

Component	Location of Information
Oracle Fusion Middleware	<i>Oracle Fusion Middleware Administrator's Guide</i>
Oracle WebLogic Server	<i>Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server</i> <i>Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server</i>
Oracle SOA Suite	<i>The Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite</i> <i>The Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite</i>
Oracle WebCenter	<i>The Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal</i> <i>The Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Portal</i>
Oracle ADF	<i>The Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework</i> <i>The Oracle Fusion Middleware Web User Interface Developer's Guide for Oracle Application Development Framework</i>
Oracle Fusion Middleware Backup and Recovery	<i>The Oracle Fusion Middleware Administrator's Guide</i>
Oracle Web Cache	<i>The Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache</i>
Oracle Identity Management	<i>The Oracle Fusion Middleware Installation Guide for Oracle Identity Management</i> <i>The Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management</i>
Oracle Virtual Directory	<i>The Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory</i>
Oracle HTTP Server	<i>The Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server</i>
Oracle Internet Directory	<i>The Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory</i>
Oracle Repository Creation Utility (RCU)	<i>The Oracle Fusion Middleware Repository Creation Utility User's Guide</i>
Oracle Portal	<i>The Oracle Fusion Middleware Administrator's Guide for Oracle Portal</i>

30.2 Using Clusters with Oracle Fusion Middleware

For production environments that require increased application performance, throughput, or high availability, you can configure two or more Managed Servers to operate as a cluster. A cluster is a collection of multiple Oracle WebLogic Server server instances running simultaneously and working together to provide increased scalability and reliability.

For more information on using clusters with Oracle Fusion Middleware, see the following:

- "Understanding Managed Servers and Managed Server Clusters" in *Oracle Fusion Middleware Administrator's Guide*
- *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*
- *Oracle Real Application Clusters Administration and Deployment Guide*

30.3 Using High Availability Features with Oracle Fusion Middleware

In addition to using a clustered architecture within your Fusion Middleware components, there are a number of high availability features built-in to ensure your applications are continuously accessible by the users. The following list provides a few options for setting up a comprehensive high availability system. The options that you integrate depend on your overall performance goals as well as your system architecture. This list is meant to provide examples only.

- **Process death detection and automatic restart**

Processes may die unexpectedly due to configuration or software problems. A proper process monitoring and restart system should constantly check the health of the applications and restart them when problems appear.

A system process should also maintain the number of restarts within a specified time interval. This is also important since continually restarting within short time periods may lead to additional faults or failures. Therefore a maximum number of restarts or retries within a specified time interval should also be designed as well.

- **State replication and routing**

For stateful applications, client state can be replicated to enable stateful failover of requests in the event that processes servicing these requests fail.

- **Failover**

With a load-balancing mechanism in place, the instances are redundant. If any of the instances fail, requests to the failed instance can be sent to the surviving instances.

- **Server load balancing**

When multiple instances of identical server components are available, client requests to these components can be load balanced to ensure that the instances have roughly the same workload.

- **Disaster Recovery**

Disaster recovery solutions typically set up two homogeneous sites, one active and one passive. Each site is a self-contained system. The active site is generally called the production site, and the passive site is called the standby site. During normal operation, the production site services requests; in the event of a site failover or switchover, the standby site takes over the production role and all requests are routed to that site. To maintain the standby site for failover, not only must the standby site contain homogeneous installations and applications, data and configurations must also be synchronized constantly from the production site to the standby site.

For more information see the *Oracle Fusion Middleware High Availability Guide*.

