

Oracle® Fusion Middleware Tuning Performance Guide



14c (14.1.2.0.0)

F85527-01

December 2024

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Oracle Fusion Middleware Tuning Performance Guide, 14c (14.1.2.0.0)

F85527-01

Copyright © 2015, 2024, Oracle and/or its affiliates.

Primary Author: Oracle Corporation

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, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. 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®, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside 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, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Audience	xv
Documentation Accessibility	xv
Diversity and Inclusion	xv
Related Documents	xv
Conventions	xvi

Part I Introduction

1 Top Performance Areas

Identifying Top Performance Areas	1-1
Securing Sufficient Hardware Resources	1-2
Tuning the Operating System	1-3
Tuning Java Virtual Machines (JVMs)	1-4
Tuning the WebLogic Server	1-4
Tuning Database Parameters	1-4
Tuning Database Parameters	1-4
Tuning Database Files	1-7
Configuring REDO Logs	1-8
Configuring UNDO Tablespace	1-8
Configuring TEMP Tablespace	1-8
Creating Additional Tablespaces	1-8
Tuning Automatic Segment-Space Management (ASSM)	1-9
Reusing Database Connections	1-9
Enabling Data Source Statement Caching	1-9
Controlling Concurrency	1-10
Setting Server Connection Limits	1-10
Setting MaxRequestWorkers / ThreadsPerChild	1-11
Setting KeepAlive	1-12
Tuning HTTP Server Modules	1-12
Configuring Connection Pools	1-12
Tuning the WebLogic Server Thread Pool	1-13

Setting Logging Levels	1-14
------------------------	------

2 Performance Planning

About Performance Planning	2-1
Performance Planning Methodology	2-1
Step 1: Defining Your Performance Objectives	2-1
Defining Operational Requirements	2-2
Identifying Performance Goals	2-2
Understanding User Expectations	2-2
Conducting Performance Evaluations	2-3
Step 2: Designing Applications for Performance and Scalability	2-3
Step 3: Monitoring and Measuring Your Performance Metrics	2-4

3 Monitoring

About Oracle Fusion Middleware Management Tools	3-1
Measuring Your Performance Metrics	3-2
Oracle Enterprise Manager Fusion Middleware Control	3-2
Oracle WebLogic Remote Console	3-2
WebLogic Diagnostics Framework (WLDF)	3-3
WebLogic Scripting Tool (WLST)	3-3
DMS Spy Servlet	3-3
Viewing Performance Metrics Using the Spy Servlet	3-4
Using the DMS Spy Servlet	3-4
Native Operating System Performance Commands	3-5
Network Performance Monitoring Tools	3-5

4 Using the Oracle Dynamic Monitoring Service

About Dynamic Monitoring Service (DMS)	4-1
Understanding Common DMS Terms and Concepts	4-1
DMS Sensors	4-2
DMS Nouns	4-4
DMS Tracing and Events	4-6
About DMS Availability	4-7
About DMS Architecture	4-7
Viewing DMS Metrics	4-8
Viewing Metrics By Using the Spy Servlet	4-8
Viewing Metrics with WLDF (WebLogic Diagnostic Framework)	4-9
Viewing Metrics with WLST (Oracle WebLogic Server)	4-9
Viewing Metrics with JConsole	4-10

Viewing Metrics with Oracle Enterprise Manager	4-11
About DMS Execution Context	4-11
DMS Execution Requests and Subtasks	4-11
DMS Execution Context Usage	4-12
DMS Execution Context Communication	4-12
DMS Tracing and Events	4-13
Configuring the DMS Event System	4-14
Adding and Editing Filters	4-15
Adding and Editing Destinations	4-16
Adding and Editing Event Routes	4-16
Compound Operations	4-16
Configuring Destinations	4-17
LoggerDestination	4-17
MBean Creator Destination	4-19
Request Tracker Destination	4-20
Java Flight Recorder Destination	4-21
Understanding the Format of DMS Events in Log Messages	4-24
Understanding DMS Event Actions	4-28
DMS Best Practices	4-28

Part II Core Components

5 Tuning Oracle HTTP Server

About Oracle HTTP Server	5-1
Monitoring Oracle HTTP Server Performance	5-1
Basic Tuning Considerations	5-2
Tuning Oracle HTTP Server Directives	5-2
Reducing Process Availability with Persistent Connections	5-8
Logging Options for Oracle HTTP Server	5-9
Access Logging	5-9
Configuring the HostNameLookups Directive	5-9
Error logging	5-10
Advanced Tuning Considerations	5-10
Tuning Oracle HTTP Server	5-10
Analyzing Static Versus Dynamic Requests	5-10
Limiting the Number of Enabled Modules	5-11
Tuning the File Descriptor Limit	5-11
Tuning Oracle HTTP Server Security	5-11
Tuning Oracle HTTP Server Secure Sockets Layer (SSL)	5-11

6 Tuning Oracle Metadata Service

About Oracle Metadata Services (MDS)	6-1
Monitoring Oracle Metadata Service Performance	6-2
Basic Tuning Considerations	6-2
Tuning Database Repository	6-2
Collecting Schema Statistics	6-2
Increasing Redo Log Size	6-3
Reclaiming Disk Space	6-3
Monitoring the Database Performance	6-3
Tuning Cache Configuration	6-3
Enabling Document Cache	6-4
Purging Document Version History	6-5
Using Auto Purge	6-5
Purging Manually	6-6
Using Database Polling Interval for Change Detection	6-6
Advanced Tuning Considerations	6-6
Analyzing Performance Impact from Customization	6-7

7 Tuning Oracle Fusion Middleware Security

About Security Services	7-1
Basic Tuning Considerations	7-2
Tuning Oracle Platform Security Services	7-2
JVM Tuning Parameters	7-2
JDK Tuning Parameters	7-3
Authentication Tuning Parameters	7-3
Authorization Tuning Properties	7-3
OPSS PDP Service Tuning Parameters	7-6
Oracle Web Services Security Tuning	7-9
Choosing the Right Policy	7-9
Policy Manager	7-10
Configuring the Log Assertion to Record SOAP Messages	7-10
Configuring Connection Pooling	7-10
Monitoring the Performance of Web Services	7-11

Part III Oracle Fusion Middleware Server Components

8 Tuning Oracle Application Development Framework (ADF)

About Oracle ADF	8-1
Basic Tuning Considerations	8-2
Oracle ADF Faces Configuration and Profiling	8-2
Performance Considerations for ADF Faces	8-3
Tuning ADF Faces Component Attributes	8-14
Performance Considerations for Table and Tree Components	8-17
Performance Considerations for autoSuggest	8-18
Data Delivery - Lazy versus Immediate	8-18
Performance Considerations for DVT Components	8-19
Advanced Tuning Considerations	8-20
ADF Server Performance	8-20
Tuning Session Timeout	8-21
Tuning View Objects	8-21
Enabling Batch Processing	8-25
Tuning RangeSize	8-26
Configuring Application Module Pooling	8-26
Using ADFc Regions	8-34
Deferring Task Flow Execution	8-34
Deferring Task Flow Creation in Popups	8-34
Configuring the Task Flow Inside Switcher	8-35
Reusing Static Data	8-35
Conditional Validations	8-35

9 Tuning Oracle TopLink

About Oracle TopLink and EclipseLink	9-1
Basic Tuning Considerations	9-2
SQL Statement and Query Tuning Parameters	9-2
Entity Relationships Query Tuning Parameters	9-4
Cache Configuration Tuning Parameters	9-7
About Cache Refreshing	9-12
Locking Mode Policy Options	9-13
About Mapping and Descriptor Configurations	9-14
About Data Partitioning	9-14
Advanced Tuning Considerations	9-14
Integrating with Oracle Coherence	9-14
Analyzing EclipseLink JPA Entity Performance	9-15

Part IV Oracle Identity and Access Management

10 Oracle Internet Directory Performance Tuning

About Oracle Internet Directory	10-1
Monitoring Oracle Internet Directory Performance	10-1
Monitoring Performance on UNIX and Windows Systems	10-2
Updating Database Statistics by Using oidstats.sql	10-3
Setting Performance-Related Replication Configuration Attributes	10-3
Managing System Configuration Attributes	10-4
Setting Garbage Collection Configuration Attributes	10-4
Modifying Changelog Purging Attributes by Using ldapmodify	10-4
Modifying Changelog Purging in Oracle Directory Services Manager	10-5
Basic Tuning Considerations	10-5
Database Parameters	10-6
LDAP Server Attributes	10-6
Database Statistics	10-8
Low-Priority Tuning Considerations	10-8
Number of Entries to be Returned by a Search	10-8
Enabling the Group Cache	10-8
Timeout for Write Operations	10-8
Advanced Tuning Considerations	10-9
Replication or Oracle Directory Integration Platform	10-9
Replication Server Configuration	10-10
Garbage Collection Configuration	10-11
Oracle Internet Directory with Oracle RAC Database	10-11
Password Policies and Verifier Profiles	10-11
Server Entry Cache	10-12
Benefits of Using the Entry Cache	10-12
Values for Configuring the Entry Cache	10-12
Result Set Cache	10-14
When to Use Result Set Cache	10-14
Benefits of Using Result Set Cache	10-15
Configuring Result Set Cache	10-15
Values for Configuring Result Set Cache	10-15
Tuning Security Event Tracking	10-15
Optimizing Searches	10-16
Optimizing Searches for Large Group Entries	10-16
Optimizing Searches for Skewed Attributes	10-17
Optimizing Performance of Complex Search Filters	10-17
Specific Use Cases That Require Additional Tuning	10-20
Bulk Load Operations	10-20
Bulk Delete Operations	10-20

11 Oracle Access Management Performance Tuning

About Oracle Access Management	11-1
Performance Considerations for Oracle Access Management Services	11-2
Understanding Your Current Environment	11-2
Controlling Network Latency	11-3
Enabling DMS Performance Instrumentation	11-4
Tuning Oracle Access Management Access Manager	11-5
Basic Tuning Considerations for Access Manager	11-5
Tuning the Web Tier	11-5
Managing Policy Components	11-7
Tuning Common Settings	11-7
Advanced Tuning Considerations for Access Manager	11-8
Tuning Oracle Coherence	11-8
Setting the Java Message Bean Pool Size	11-9
Tuning the Server Cache	11-9
Tuning Webgate Caches	11-10
Changing Request Cache Type	11-15
Tuning Authentication Plug-Ins	11-15
Specific Use Cases That Require Additional Tuning for Access Manager	11-15
Managing Access Manager Sessions	11-15
Audit Settings	11-15
Managing Monitor Account	11-16
Kerberos Latency Issues	11-16
Oracle Access Protocol over REST Connectivity Issues	11-16
Tuning Oracle Access Management Identity Federation	11-16
Basic Tuning Considerations for Identity Federation	11-17
Tuning the Load Balancer and HTTP Server	11-17
Tuning SOAP Connections	11-17
Tuning the Data Tier Connections	11-17
Advanced Tuning Considerations for Identity Federation	11-19
Tuning Oracle Coherence	11-19
Tuning Identity Store	11-19
Tuning Protocol Binding	11-19
Tuning the Browser POST and Artifact Single Sign-On Profiles	11-20
Specific Use Cases That Require Additional Tuning for Identity Federation	11-21
Message Signing versus Token Signing	11-21
Tuning	11-21
Basic Tuning Considerations for Security Token Service	11-21
Tuning the Load Balancer and HTTP Server	11-21

Tuning Outbound SOAP Connections	11-22
Tuning the Data Tier Connections	11-22
Advanced Tuning Considerations for Security Token Service	11-22
Tuning the WS-Security Policy	11-22
Tuning Oracle Access Management Mobile and Social	11-23
Basic Tuning Considerations for Mobile and Social	11-23
Tuning the Access Management Authentication Service Provider	11-23
Tuning the User Profile Service Provider	11-24
Database Tuning for Oracle Access Management	11-24
Automatic Optimizer Statistics Collection	11-24
Partitioning AM_SESSION table using Config Utility Command	11-24
Purging Inactive Sessions as a Recovery Mechanism from Peak Load	11-25

12 Oracle Identity Governance Performance Tuning

About Oracle Identity Governance	12-1
Monitoring Oracle Identity Governance Performance	12-1
Basic Tuning Considerations	12-3
Tuning and Managing Application Cache	12-3
Tuning Oracle Identity Governance Cache	12-3
Purging the Cache	12-6
Tuning the Application Server for Oracle Identity Governance	12-7
Tuning JVM Memory Settings for Oracle Identity Governance	12-7
Tuning the JDBC Connection Pool for Oracle Identity Governance	12-8
Tuning OIG-specific Work Manager Properties	12-8
Disabling the Reloading of Adapters and Plug-in Configuration	12-10
Changing the Number of Open File Descriptors for UNIX (Optional)	12-10
Tuning the JVM Garbage Collection for Solaris Sparc T3 or T4	12-10
Tuning Database Parameters for Oracle Identity Governance	12-11
Sample Instance Configuration Parameters	12-11
Physical Data Placement	12-12
Resolving enq: HW - contention	12-15
Tuning Oracle Internet Directory	12-16
Tuning Application Module (AM) for User Interface	12-16
JMS Tuning	12-16
Advanced Tuning Considerations	12-17
Reconciliation Tuning	12-17
Target System And Connector Tuning	12-17
Database Indexes For Recon Matching Rules	12-19
Oracle Identity Governance Post-processing for Reconciliation	12-22
Tuning LDAP Synchronization	12-22
Increasing the Max Connection Pool for Oracle Identity Governance	12-22

Part V SOA Suite Components

13 Tuning the SOA Infrastructure

About the SOA Infrastructure	13-1
Tuning SOA Work Managers	13-1
Configuring Database Connections with the SOADatasource Property	13-2
Configuring Work Managers with the SOAMaxThreadsConfig Attribute	13-2
Tuning SOA Infrastructure Parameters	13-4
Using Advanced Tuning Options	13-5
Using Composite Lazy Loading	13-5
Configuring Composite Lazy Loading for the Domain Level	13-6
Configuring Composite Lazy Loading at the Component Level	13-6
Changing Modularity Profiles	13-7
Tuning Your Database for SOA Processes	13-8
Collecting Optimizer Statistics	13-8
Tuning Temporary Tablespaces for SOA	13-9
Minimizing SOA Database Contention	13-9
Purging	13-13
Reclaiming Space	13-13
Tuning Event Delivery Network Parameters	13-14
Adding JMS Topics with Mapping	13-18
Tuning the WebLogic Server	13-19
Advanced Tuning for Work Managers	13-21
Configuring Fair Share Request Class for SOA Work Managers	13-22
Creating a New Work Manager Constraint	13-22

14 Tuning Oracle BPEL Process Manager

About BPEL Process Manager	14-1
Tuning BPEL Parameters	14-1
Tuning BPEL Engine	14-1
Tuning BPEL Engine Parameters	14-2
Tuning BPEL in a Composite	14-4
Using Other Tuning Strategies	14-5
Identifying Tables Impacted By Instance Data Growth	14-5

15	Tuning Oracle Mediator	
	About Oracle Mediator	15-1
	Tuning Mediator Parameters	15-1
	Using Resequencer for Messages	15-2
16	Tuning Oracle Managed File Transfer	
	About Managed File Transfer	16-1
	Tuning MFT Parameters	16-1
	Tuning Remote FTP / SFTP/ FILE Type Sources	16-3
	Minimizing MDS label	16-4
	Adjusting the Materialized Views Refresh Interval	16-4
17	Tuning Oracle Business Rules	
	About Oracle Business Rules	17-1
	Tuning Oracle Business Rules	17-1
	Exerting assertXPath Support	17-2
18	Tuning Oracle Business Process Management	
	About Oracle Business Process Management	18-1
	Tuning Business Process Management Parameters	18-1
	Using Other Tuning Strategies	18-2
	Tuning Oracle Workspace Applications	18-2
	Tuning Process Measurement	18-4
19	Tuning Oracle Human Workflow	
	About Oracle Human Workflow	19-1
	Tuning Human Workflow	19-1
	Using Other Tuning Strategies	19-3
	Improving Server Performance	19-3
	Completing Workflows Faster	19-4
	Tuning the Identity Provider	19-5
	Tuning the Database	19-5
20	Tuning Oracle Business Activity Monitoring	
	About Oracle Business Activity Monitoring	20-1
	Tuning BAM Server Parameters	20-1
	Other Tuning Strategies	20-3

Creating an Index Column	20-3
Tuning Loggers	20-3
Tuning Continuous Query Service	20-3

21 Tuning Oracle Service Bus

About Oracle Service Bus	21-1
Tuning OSB Parameters	21-1
Tuning Oracle Service Bus with Work Managers	21-2
Tuning OSB Operation Settings	21-2
Using Other Tuning Strategies	21-4
Tuning Resequencer in OSB	21-4
Considering Design Time for Proxy Applications	21-5
Tuning XQuery	21-7
Tuning Poller-based Transports	21-8
Setting the Polling Interval	21-9
Setting Read Limit	21-9

22 Tuning Oracle Enterprise Scheduler Service

About Enterprise Scheduler Service	22-1
Tuning Enterprise Scheduler Service Parameters	22-1

23 Tuning Oracle Business Intelligence Performance

About Oracle Business Intelligence	23-1
Tuning Oracle BI Server Query Performance	23-1
Tuning Oracle BI Server Query Cache Performance	23-2
Tuning Oracle BI Web Client Performance	23-2

Part VI Oracle WebCenter Components

24 Tuning Oracle WebCenter Portal

About Oracle WebCenter Portal	24-1
Basic Tuning Considerations	24-1
Setting System Limit	24-2
Setting JDBC Data Source	24-2
Using Content Compression to Reduce Downloads	24-3
Tuning Configuration for WebCenter Portal	24-4
Setting a Session Timeout for WebCenter Portal	24-4
Setting MDS Cache Size and Purge Rate	24-4

Configuring Concurrency Management	24-5
Tuning Tools and Services Configuration	24-7
Tuning Performance of Mail	24-7
Tuning Performance of RSS News Feeds	24-7
Tuning Policy Store Parameters	24-8
Tuning Identity Store Configuration	24-8
Tuning the Identity Store when Using SSL	24-8
Tuning Performance when Using OVD	24-9
Tuning Performance when Using Active Directory	24-9
Tuning Portlet Configuration	24-10
Tuning Performance of the Portlet Client	24-10
Configuring Supported Locales	24-10
Configuring Portlet Cache Size	24-11
Configuring Portlet Timeout	24-11
Customizing the Container Runtime Environment Options	24-12
Suppressing Optimistic Rendering for WSRP Portlets	24-12
Setting Portlet Container Runtime Options	24-12
Excluding Request Attributes for Portlets	24-13
Tuning Performance of Oracle PDK-Java Producers	24-13
Setting WSRP Attribute for Portlet-served Resources	24-13
Setting WSRP Attribute for Resources Not Served by the Portlet	24-14

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.

- [Audience](#)
- [Documentation Accessibility](#)
- [Diversity and Inclusion](#)
- [Related Documents](#)
- [Conventions](#)

Audience

Oracle Fusion Middleware Tuning Performance 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 customer access to and use of Oracle support services will be pursuant to the terms and conditions specified in their Oracle order for the applicable services.

Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers, and partners, we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

Related Documents

For more information, see the following documents in the Oracle Fusion Middleware documentation set:

- *Understanding Oracle Fusion Middleware*
- *Securing Applications with Oracle Platform Security Services*
- *High Availability Guide*
- *Understanding Oracle WebLogic Server*
- *Tuning Performance*
- *Administering Oracle SOA Suite and Oracle Business Process Management Suite*
- *Administering Oracle HTTP Server*
- *Administering Web Services*

Conventions

The following text conventions are used in this document:

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

Part I

Introduction

Performance tuning is essential for improving system performance. Therefore, it is important to understand the basic performance concepts and how to design applications for performance and scalability.

This part contains the following topics:

- [Top Performance Areas](#)
By identifying top performance areas, you can tune Oracle Fusion Middleware for optimal performance.
- [Performance Planning](#)
A clearly defined plan for achieving your performance objectives is essential for deciding what to trade for higher performance.
- [Monitoring](#)
Oracle Fusion Middleware provides a variety of technologies and tools that monitor server and application performance.
- [Using the Oracle Dynamic Monitoring Service](#)
The Oracle Dynamic Monitoring Service (DMS) publishes component performance data.

1

Top Performance Areas

By identifying top performance areas, you can tune Oracle Fusion Middleware for optimal performance.

- [Identifying Top Performance Areas](#)
One of the most challenging aspects of performance tuning is knowing where to begin. Therefore, it is important to identify the top performance areas for Oracle Fusion Middleware.
- [Securing Sufficient Hardware Resources](#)
Manage the performance of Oracle Fusion Middleware applications to ensure there is sufficient CPU, memory, and network resources to support the user and application requirements for installation.
- [Tuning the Operating System](#)
Each operating system has native tools and utilities that can be useful for monitoring and tuning purposes.
- [Tuning Java Virtual Machines \(JVMs\)](#)
How you tune your Java virtual machine (JVM) greatly affects the performance of Oracle Fusion Middleware and your applications.
- [Tuning the WebLogic Server](#)
Tune the WebLogic Server to match your application needs.
- [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.
- [Reusing Database Connections](#)
It is important to tune the connection pool attributes in the JDBC data sources in your WebLogic Server domain correctly to improve application and system performance.
- [Enabling Data Source Statement Caching](#)
Statement caching improves performance by caching executable statements that are used repeatedly.
- [Controlling Concurrency](#)
Limiting concurrency at multiple layers of the system to match specific usage needs can greatly improve performance.
- [Setting Logging Levels](#)
The amount of information that is logged can have a significant impact on the performance.

Identifying Top Performance Areas

One of the most challenging aspects of performance tuning is knowing where to begin. Therefore, it is important to identify the top performance areas for Oracle Fusion Middleware.

[Table 1-1](#) 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 a 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 [Monitoring](#).

Table 1-1 Top Performance Areas for Oracle Fusion Middleware

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

Securing Sufficient Hardware Resources

Manage the performance of Oracle Fusion Middleware applications to ensure there is sufficient CPU, memory, and network resources to support the user and application requirements for installation.

No matter how well you tune your applications, if the appropriate hardware resources are not used, your applications cannot reach their 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 [Verifying Certification](#),

System Requirements, and Interoperability in *Planning an Installation of Oracle Fusion Middleware*.

Sufficient hardware resources must 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 details on monitoring, see [Monitoring](#).

 **Tip:**

Your target CPU usage should never reach 100% utilization. Target the 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, maintaining a fast response time is important. High CPU usage (approaching 100% utilization) can increase response time 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 might 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 any traffic can be rescheduled. To reduce the peak load, you must schedule the 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, this mechanism is a temporary solution. See [Controlling Concurrency](#).

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

Tuning Java Virtual Machines (JVMs)

How you tune your Java virtual machine (JVM) greatly affects the performance of Oracle Fusion Middleware and your applications.

For more information on tuning your JVM, see [Tuning Java Virtual Machines \(JVM\) in *Tuning Performance of Oracle WebLogic Server*](#).

Tuning the WebLogic Server

Tune the WebLogic Server to match your application needs.

If your Oracle Fusion Middleware applications use the WebLogic Server, see [Tuning WebLogic Server in *Tuning Performance of Oracle WebLogic Server*](#).

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.

Note:

The information in these topics is a subset of database tuning information for Fusion Middleware. Make sure that you have also reviewed the *Database Performance Tuning Guide*.

Always review the tuning guidelines in your database-specific vendor documentation.

- [Tuning Database Parameters](#)
- [Tuning Database Files](#)
- [Tuning Automatic Segment-Space Management \(ASSM\)](#)

Tuning Database Parameters

The following tables provide common `init.ora` parameters and their descriptions. Follow these guidelines to set the database parameters. Ultimately, however, the database administrator must monitor the database health and tune parameters based on the need.

The database that is used for SOA is configured with the suggested values. Tuning the database involves adjusting the sizing parameters based on the available resource and load on the database.

The `sga_target`, `pga_aggregate_target`, and `processes` parameters from [Table 1-2](#) are examples of such parameters that must be tuned based on the System Global Area (SGA) and Parent Global Area (PGA) advisories and looking into the number of open processes during peak load.

Table 1-2 Important Oracle 12c Database Tuning Parameters

Parameter	Description	Tuning Recommendation
audit_trail Default: DB	Enables or disables database auditing.	Set to NONE if there is NO policy to audit database activity. Enabling auditing can impact performance.
plsql_code_type Default: INTERPRETED	Compilation mode for PL/SQL library units. Possible modes are as follows: <ul style="list-style-type: none"> INTERPRETED: PL/SQL library units are compiled to PL/SQL byte code format and executed by the PL/SQL interpreter engine. NATIVE: PL/SQL library units are compiled to native (machine) code. Such modules are executed natively without incurring any interpreter impacts. 	Set to NATIVE.
nls_sort Default: Derived from NLS_LANGUAGE	Collating sequence for ORDER BY queries. <ul style="list-style-type: none"> If the value is a named linguistic sort, the collating sequence is based on the order of the defined linguistic sort. Most languages supported by the NLS_LANGUAGE parameter also support a linguistic sort with the same name. If the value is set to BINARY, then the collating sequence is based on the numeric value of characters. This requires fewer system resources. 	Set to BINARY.
open_cursors Default: 50	Maximum number of open cursors that a session can have at once. Open cursors are handles to private SQL areas. The value of OPEN_CURSORS must be high enough to prevent your application from running out of open cursors.	Increase to 500.
session_cached_cursors Default: 50	Number of session cursors to cache. Repeated parse calls of the same SQL statement cause the session cursor for that statement to be moved into the session cursor cache. Subsequent parse calls locate the cursor in the cache. However, they do not reopen the cursor. Oracle uses a least recently used algorithm to remove entries in the session cursor cache to make room for new entries when needed. This parameter also constrains the size of the PL/SQL cursor cache, which PL/SQL uses to avoid having to reparse as statements are reexecuted.	Increase to 500.
_b_tree_bitmap_plans Default: TRUE	Enables or disables the use of bitmap access paths for b-tree indexes.	Set to FALSE.

Table 1-2 (Cont.) Important Oracle 12c Database Tuning Parameters

Parameter	Description	Tuning Recommendation
processes Default: 100	Maximum number of operating system processes that can be connected to the Oracle database concurrently. The value of this parameter must account for Oracle the background processes. The <code>SESSIONS</code> parameter is deduced from this value.	For most systems, increasing to 1500 must suffice. For a large-scale system, such as databases with a large number of users, the recommended value is 5000.
Memory_target	Oracle system-wide usable memory. The database tunes memory to the <code>MEMORY_TARGET</code> value, reducing or enlarging the SGA and PGA as needed.	Consider setting to <code>NONE</code> . Then set the SGA and PGA targets separately as setting <code>MEMORY_TARGET</code> does not allocate sufficient memory to SGA and PGA as needed.
sga_target Default: 0	A non-zero value enables Automatic Shared Memory Management. This can simplify configuration and improve performance.	For small systems, use a minimum of 2 GB. For large systems, set it to 18 GB.
pga_aggregate_target Default: 0	Target aggregate PGA memory available to all server processes attached to the instance.	For small systems, use a minimum of 1 GB. For large systems, set it to 8 GB.
Disk_asynch_io Default: TRUE	Controls whether I/O to data files, control files, and log files is asynchronous. It decides what parallel server processes can overlap I/O requests with CPU processing during table scans.	Set to <code>FALSE</code> only if your platform does not support asynchronous I/O.
Filesystemio_options Default: None	I/O operations for file system files.	Set to <code>SETALL</code> .
Secure_Files Default: PERMITTED	How to store LOB objects from tables.	Set to <code>ALWAYS</code> .
parallel_max_servers Default: <code>PARALLEL_THREADS_PER_CPU*CPU_COUNT*concurrent_parallel_users*5</code>	Maximum number of parallel execution processes and parallel recovery processes for an instance. As the demand increases, the Oracle database increases the number of processes from the number created at instance startup to this value.	Set to 12.
job_queue_processes Default: 1000	Maximum number of job slaves per instance that can be created for the execution of <code>DBMS_JOB</code> jobs and Oracle Scheduler (<code>DBMS_SCHEDULER</code>) jobs.	Set to 12.
shared_servers Default: 0 (or) 1	Number of server processes that you want to create when an instance is started.	Set to 0.

The table below describes the important `inti.ora` Database Tuning Parameters.

Table 1-3 Important inti.ora Oracle 12c Database Tuning Parameters

Database Parameter	Description
AUDIT_TRAIL	If there is no policy to audit database activity, consider setting this parameter to NONE. Enabling auditing can impact performance.
MEMORY_MAX_TARGET	Maximum value to which a database administrator can set the MEMORY_TARGET initialization parameter.
MEMORY_TARGET	Consider setting to NONE. Set SGA and PGA separately as setting the 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 monitor the production database daily and adjust SGA and PGA accordingly. If the database server has more memory, consider setting the PGA_AGGREGATE_TARGET to a value higher than 1G, based on usage needs.
SGA_MAX_SIZE	Consider setting the MEMORY_TARGET instead of setting SGA and the PGA separately.
SGA_TARGET	Consider using a value of 2G initially and then monitor the production database daily and adjust SGA and PGA accordingly. If the database server has more memory, consider setting the SGA_TARGET to a value higher than 2G, based on usage needs.

In addition, set a minimum value for SHARED_POOL_SIZE and DB_CACHE_SIZE to minimize frequent resizing.

Tuning Database Files

In addition to tuning the database parameters, the database administrator must configure the REDO logs, UNDO table space, and TEMP table spaces to meet the demands of the database workload. The recommendations here are intended to provide initial guidance in these areas.

The location of the database files must be optimized for I/O performance and growth. Segment Advisor must be leveraged to optimize the use of segment space and ensure that performance degradation does not occur. The advisor can provide historical growth trends of segments, which can be used to proactively plan for growth. See Using the Segment Advisor in *Oracle Database Administrator's Guide*.

- [Configuring REDO Logs](#)
 - [Configuring UNDO Tablespace](#)
 - [Configuring TEMP Tablespace](#)
 - [Creating Additional Tablespaces](#)
- Oracle recommends you to create additional tablespaces based on the requirement of the workload.

Configuring REDO Logs

Under demanding workloads, the size of the REDO log files can influence performance. Generally, larger REDO log files provide better performance. Undersized log files increase checkpoint activity and log file switches, which reduces performance. You can obtain sizing advice on the REDO Log Groups page of the Enterprise Manager.

Depending on your storage configuration and performance characteristics, redistribute the REDO logs to optimize I/O performance. The REDO log files must be placed on a disk separately from the data files to improve the I/O performance.

See *Managing the REDO Log* *Oracle Database Administrator's Guide*

Configuring UNDO Tablespace

The suggested minimum size for the UNDO tablespace is 6 GB with auto-extend enabled. Oracle recommends that the default mode of automatic undo management is leveraged to maximize performance and efficiency.

The Oracle Enterprise Manager Automatic Undo Management Advisor must be leveraged to set configuration details for UNDO tablespace and retention settings. This advisor also provides access to the Undo Advisor that assesses the effect and provides advice of a new undo retention setting. For more information about using advisors, see *The Undo Advisor PL/SQL Interface* *Oracle Database Administrator's Guide*.

Configuring TEMP Tablespace

Oracle recommends the use of locally managed temporary tablespaces with the allocation type set to UNIFORM extents and the default size of 1 MB.

For tuning TEMP tablespaces for SOA, see [Tuning Temporary Tablespaces for SOA](#).

Creating Additional Tablespaces

Oracle recommends you to create additional tablespaces based on the requirement of the workload.

You can increase the size of a tablespace by either of the following options:

- **Changing Data File Size:** You can alter the size of a data file. For example, you can increase the size of one or more data files when more space is needed in the database. For more information, see *Changing Data File Size*.
- **Creating Data Files and Adding Data Files to a Tablespace:** You can create data files and associate them with a tablespace using several different SQL statements. For more information, see *Creating Data Files and Adding Data Files to a Tablespace*.
- **Enabling and Disabling Automatic Extension for a Data File:** You can create data files or alter existing data files so that they automatically increase in size when more space is needed in the database. The file size increases in specified increments up to a specified maximum. For more information, see *Enabling and Disabling Automatic Extension for a Data File*.

Sample Script to create additional tablespaces:

```
CREATE TABLESPACE apps_tbs LOGGING
  DATAFILE '/u01/app/oracle/oradata/mynewdb/apps01.dbf'
```

```
SIZE 500M REUSE AUTOEXTEND ON NEXT 1280K MAXSIZE UNLIMITED
EXTENT MANAGEMENT LOCAL;
-- create a tablespace for indexes, separate from user tablespace (optional)
CREATE TABLESPACE indx_tbs LOGGING
DATAFILE '/u01/app/oracle/oradata/mynewdb/indx01.dbf'
SIZE 100M REUSE AUTOEXTEND ON NEXT 1280K MAXSIZE UNLIMITED
EXTENT MANAGEMENT LOCAL;
```

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`.

Oracle recommends to specify the allocation type to `SYSTEM`.

See *Free Space Management and Specifying Segment Space Management in Locally Managed Tablespaces* in *Oracle Database Administrator's Guide*.

Reusing Database Connections

It is important to tune the connection pool attributes in the JDBC data sources in your WebLogic Server domain correctly to improve application and system performance.

Creating a database connection is a resource-intensive process in any environment. Typically, a connection pool starts with a few connections. As client demands for more connections grow, there will not be enough in the pool to fulfill the requests. WebLogic Server creates more 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. This is 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 preproduction performance testing.

When the load is uneven, and has high number of connections at peak load than at typical load, set the initial number of connections equal to your typical load. In addition, set 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.

See *Tuning Data Source Connection Pool Options* in *Administering JDBC Data Sources for Oracle WebLogic Server*.

Enabling Data Source Statement Caching

Statement caching improves performance by caching executable statements that are used repeatedly.

When a prepared statement or callable statement is used in an application or EJB, it impacts the performance associated with the processing of the communication between the application server and 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.

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.

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 as per the data source. That is, the statement cache for each connection in a data source uses the statement cache options specified for the data source. 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. Analyze your database statement parse metrics to size the statement cache sufficiently for the number of statements you have in your application.

You can use the Remote Console to set statement cache options for a data source.

For details on using statement caching, see *Increasing Performance with the Statement Cache in Administering JDBC Data Sources for Oracle WebLogic Server*.

Controlling Concurrency

Limiting concurrency at multiple layers of the system to match specific usage needs can greatly improve performance.

When system capacity is reached, and a web server or an application server continues to accept requests, application performance and stability can deteriorate. Within the Oracle Fusion Middleware, 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](#)

Setting Server Connection Limits

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

For details on modifying the `httpd.conf` file, see *Configuring Oracle HTTP Server in Administering Oracle HTTP Server*.

Use the `MaxRequestWorkers` and `ThreadsPerChild` directives to limit incoming requests to WebLogic instances from the Oracle HTTP Server based on your expected client load and

system resources. There are several Oracle HTTP Server tuning parameters related to connection limits that must be tuned based on the expected client load. See [Tuning Oracle HTTP Server](#) for details on setting server connection limits and a complete list of tunable parameters.

- [Setting MaxRequestWorkers / ThreadsPerChild](#)
- [Setting KeepAlive](#)
- [Tuning HTTP Server Modules](#)

Setting MaxRequestWorkers / ThreadsPerChild

Note:

The `MaxRequestWorkers` parameter is applicable only to UNIX platforms. The same is achieved through the `ThreadsPerChild` and `ThreadLimit` properties on Microsoft Windows (`mpm_winnt`).

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

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

If the server handles persistent connections, then it requires sufficient concurrent `httpd` server processes to handle both active and idle connections. When you specify `MaxRequestWorkers` to act as a throttle for system concurrency, 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 few 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 `MaxRequestWorkers` setting. All these properties must be set appropriately to achieve the number of connections as specified by `MaxRequestWorkers`. See [Table 5-1](#) for a description of all the HTTP configuration properties.

Setting KeepAlive

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

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

 **Note:**

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

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 the `mod_wl_ohs` module are available in the `mod_wl_ohs.conf` file in the `config` directory.

See *Understanding Oracle HTTP Server Modules* in *Administering Oracle HTTP Server*.

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 pooling needs to be customized. For example; control the size of the pools and types of connections to be pooled.

You configure connection pooling by using several system properties at program startup time. These are system properties, not environment properties and they affect all connection pooling requests.

For applications that use a database, performance can improve when the connection pool that is associated with a data source limits the number of connections. 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. Thus, 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 the `MaxCapacity` attribute is set to 15. For

best performance, specify a value for the `MaxCapacity` attribute 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. Configure the database to allow at least 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.

For connection pool options, see *Configuring Services* in the *Oracle WebLogic Remote Console Online Help* and *Tuning Data Source Connection Pool Options* in *Administering JDBC Data Sources for Oracle WebLogic Server*.

Tuning the WebLogic Server Thread Pool

By default, WebLogic Server uses a single thread pool. All types of work are executed in this thread pool. WebLogic Server uses work managers to prioritize work based on rules that you can define, and runtime 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.

The WebLogic Server thread pool is sized automatically and hence in most situations you do not need to tune it. 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 executed 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 is given a higher priority over another.
- A response time goal is required.
- A minimum thread constraint is 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 are 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.

 **Note:**

For details 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 *Tuning Performance of Oracle WebLogic Server*
- Thread Management in *Tuning Performance of Oracle WebLogic Server*
- MDB Thread Management in *Tuning Performance of Oracle WebLogic Server*
- Using Work Managers to Optimize Scheduled Work in *Administering Server Environments for Oracle WebLogic Server*
- Avoiding and Managing Overload in *Administering Server Environments for Oracle WebLogic Server*

Use Oracle WebLogic Remote 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 the scope of the application and the 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.

The work manager and thread pool metrics can also be viewed from the Oracle Fusion Middleware Control.

Setting Logging Levels

The amount of information that is logged can have a significant impact on the performance.

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.

After you set the application and server logging levels, ensure that the debugging properties or other application level debugging flags are set correctly 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 has specific recommendations for logging levels.

2

Performance Planning

A clearly defined plan for achieving your performance objectives is essential for deciding what to trade for higher performance.

- [About Performance Planning](#)
To maximize performance, you must monitor, analyze, and tune all the components that are used by your applications.
- [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.

About Performance Planning

To maximize performance, you must monitor, analyze, and tune all the components that are used by your applications.

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

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.

- [Step 1: Defining Your Performance Objectives](#)
- [Step 2: Designing Applications for Performance and Scalability](#)
- [Step 3: Monitoring and Measuring Your Performance Metrics](#)

Step 1: Defining 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.

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 to determine 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 must focus on maximum performance requirements.

- 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 Introduction and Roadmap in *High Availability Guide*.
- 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 time, throughput, and load on specific hardware.

- [Defining Operational Requirements](#)
- [Identifying Performance Goals](#)
- [Understanding User Expectations](#)
- [Conducting Performance Evaluations](#)

Defining 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

Identifying 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 application components are expected to meet, such as:

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


Understanding 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 complicated operation, response time is slower than while performing a simple operation.

For example, ensure that 90% of the response time is not greater than 5 seconds and the maximum response time for all is 20 seconds. Usually, it's not that simple. Application may include various operations with differing characteristics and acceptable response time. Set measurable goals for each of these operations.

Determine how variances in the load can affect the response time. For example, users might access the system heavily between 9:00 am and 10:00 am and then again between 1:00 pm and 2:00 pm, as illustrated by the graph in [Figure 2-1](#). If the peak load occurs regularly, for example, daily or weekly, it is advised to configure and tune systems to meet the peak load requirements. Accessing application in off-time gives better response time than accessing it during peak-time. If your peak load is infrequent, higher response times at peak loads must be expected for the cost savings of smaller hardware configurations.

Figure 2-1 Adjusting Capacity and Functional Demand



Conducting 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 that 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 the actual resource consumption for a range of loads, you can conduct objective scalability studies and from these reports predict the resource requirements for anticipated load volumes. For details on Monitoring, see, [Monitoring](#).

Step 2: Designing 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.

Step 3: Monitoring and Measuring 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 the server activity, watch trends, diagnose system bottlenecks, debug applications with performance problems and gather data that can assist you in tuning the system.

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

3

Monitoring

Oracle Fusion Middleware provides a variety of technologies and tools that monitor server and application performance.

- [About Oracle Fusion Middleware Management Tools](#)
Monitoring enables you to evaluate server activity, watch trends, diagnose system bottlenecks, debug applications with performance problems, and gather data that can assist in tuning the system.
- [Oracle Enterprise Manager Fusion Middleware Control](#)
Fusion Middleware Control is a web browser-based, graphical user interface that you can use to monitor and administer your domain.
- [Oracle WebLogic Remote Console](#)
Oracle WebLogic Remote Console is a web browser-based, graphical user interface that you use to manage an Oracle WebLogic Server domain.
- [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.
- [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.
- [DMS Spy Servlet](#)
The DMS Spy Servlet provides access to DMS metric data from a web browser.
- [Native Operating System Performance Commands](#)
Each operating system has native tools and utilities that can be useful for monitoring purposes.
- [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.

About Oracle Fusion Middleware Management Tools

Monitoring enables you to evaluate server activity, watch trends, diagnose system bottlenecks, debug applications with performance problems, and gather data that can assist in tuning the system.

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

Each tool is described in Overview of Oracle Fusion Middleware Administration Tools in *Administering Oracle Fusion Middleware*.

 **Note:**

The Oracle Process Manager and Notification Server (OPMN) is no longer used in Oracle Fusion Middleware. Instead, system components are managed by the WebLogic Management Framework, which includes WLST, Node Manager and the pack and unpack commands. See *What Is the WebLogic Management Framework in Understanding Oracle Fusion Middleware*.

- [Measuring Your Performance Metrics](#)

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 that you are monitoring.

If you are new to Oracle Fusion Middleware or if you need additional information about monitoring your environment by using the Performance Summary pages, see *Viewing the Performance of Oracle Fusion Middleware in Administering Oracle Fusion Middleware*. 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.

Oracle Enterprise Manager Fusion Middleware Control

Fusion Middleware Control is a web browser-based, graphical user interface that you can use to monitor and administer your domain.

It can manage an Oracle WebLogic Server domain with its Administration Server, one or more Managed Servers, clusters, the Oracle Fusion Middleware components that are installed, configured, and running in the domain, and the applications that you deploy.

See *Getting Started Using Oracle Enterprise Manager Fusion Middleware Control in Administering Oracle Fusion Middleware*.

Oracle WebLogic Remote Console

Oracle WebLogic Remote 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 Getting Started Using Oracle WebLogic Server Administration Console in *Administering Oracle Fusion Middleware*.

Additional WebLogic Server Console Resources:

For details on the content contained in each summary table, see Monitor Servers in the *Oracle WebLogic Remote Console Online Help*.

For detailed information on using the WebLogic Server to monitor your domain, see *Tuning Performance of Oracle WebLogic Server*.

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.

See Understanding the Diagnostic Framework in *Administering Oracle Fusion Middleware*.

Note:

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

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.

See Getting Started Using the Oracle WebLogic Scripting Tool (WLST) in *Administering Oracle Fusion Middleware*.

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.

- [Viewing Performance Metrics Using the Spy Servlet](#)
- [Using the DMS Spy Servlet](#)

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`: `/modules/oracle.dms_12.1.2/dms.war`.

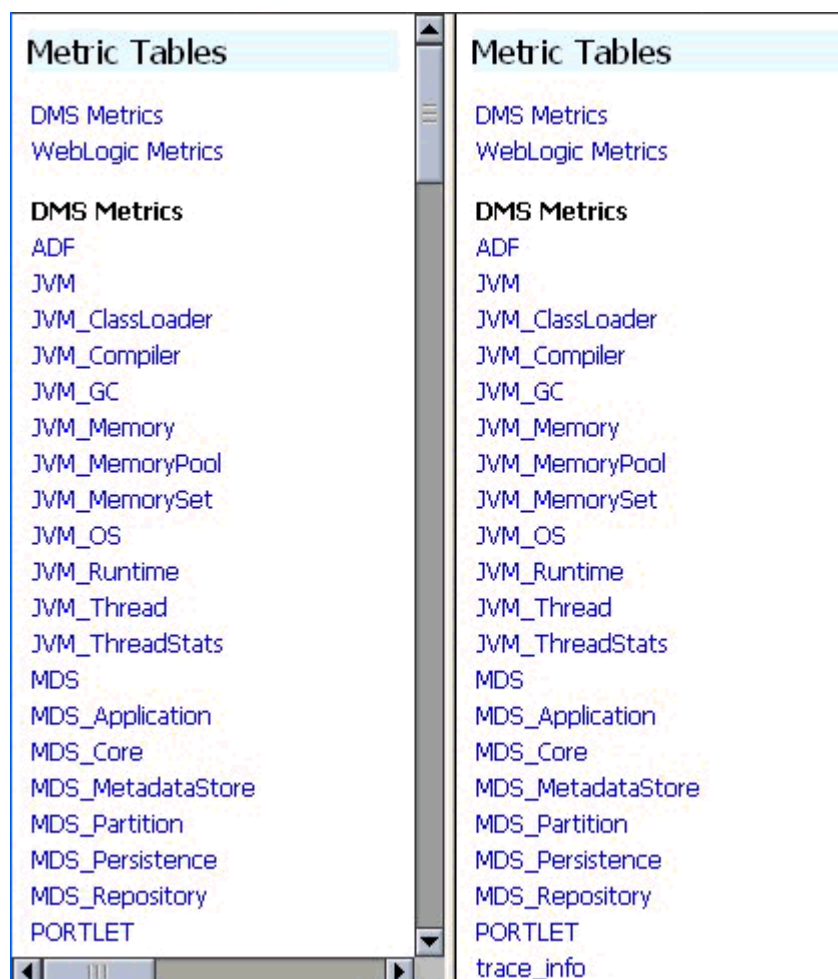
The DMS web application is deployed by default as part of a JRF-enabled server instance. The URL is: `http://host:port/dms/Spy`.

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

Using the DMS Spy Servlet

Figure 3-1 shows the initial page of the Spy servlet: both sides show the same list of metric tables.

Figure 3-1 Spy Servlet Page - Metrics Tables



The Spy servlet can display metric tables for WebLogic Server and also for non-Jakarta 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 3-2 shows the MDS_Partition table.

Figure 3-2 MDS Partition Table

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 0.106 254 1 5 0 27	0 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 0 0 0 0 0 0	0 10.66 100 1 69 4 1066	wsm-pm	oracle WLS_Spaces
webcenter		WLS_	active,	0 active,	0	webcenter(11.	WLS_

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

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 system activity information. For example CPU utilization, paging activity, swapping, and so on.

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

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.

4

Using the Oracle Dynamic Monitoring Service

The Oracle Dynamic Monitoring Service (DMS) publishes component performance data.

- [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.
- [About DMS Availability](#)
DMS functionality is available on all certified Jakarta EE servers.
- [About DMS Architecture](#)
It is important to understand the components of DMS and how they interact with other Oracle Fusion Middleware components.
- [Viewing DMS Metrics](#)
Oracle Fusion Middleware components are instrumented with DMS metrics to collect information that developers, system administrators, and support analysts can use to analyze system performance or monitor system status.
- [About DMS Execution Context](#)
The DMS execution context is the mechanism by which requests (such as RMI requests) can be uniquely identified and thus tracked as they flow through the system.
- [DMS Tracing and Events](#)
The DMS tracing feature can be used to diagnose issues or collect specific data at a specific time for a specific set of criteria.
- [DMS Best Practices](#)
Implement the following best practices when you use DMS metrics.

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. DMS measures and reports metrics, traces events and system performance, and provides a context correlation service for these components.

- [Understanding Common DMS Terms and Concepts](#)

Understanding Common DMS Terms and Concepts

There are common DMS terms and concepts related to DMS Sensors, DMS Nouns, and DMS Tracing and Events.

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

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 others are optional.

- [DMS PhaseEvent Sensors](#)
- [DMS Event Sensors](#)
- [DMS State Sensors](#)
- [Sensor Naming Conventions](#)

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 that are associated with a PhaseEvent, including the average, maximum, and minimum time that is spent in the PhaseEvent sensor.

[Table 4-1](#) lists the metrics that are available with PhaseEvent sensors.

Table 4-1 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 changes 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.

DMS Event Sensors

A DMS *event sensor* counts system events. Track system events through a DMS event sensor that has a short duration, or where the occurrence of the event is of interest.

[Table 4-2](#) describes the metric that is associated with an event sensor.

Table 4-2 DMS Event Sensor Metrics

Metric	Description
<code>sensor_name.count</code>	Specifies the number of times the event has occurred since the process started. <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.

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 the 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 4-3](#) 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 4-3 DMS State Sensor Metrics

Metric	Description
<code>sensor_name.value</code>	Specifies the metric value for <code>sensor_name</code> , by 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 for this <code>sensor_name</code> since startup. Optional metric.

Sensor Naming Conventions

The following list describes the DMS sensor naming conventions:

- Sensor names must be descriptive, but not redundant. Sensor names should not contain any part of the noun name hierarchy, or type, as it is redundant.
- Sensor names must avoid containing the value for the individual metrics.

- Where multiple words are required to describe a sensor, the first word must start with a lowercase letter, and the following words must start with uppercase letters. For 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 a `/` is used in a noun or sensor name, then when you use the sensor in a string with DMS methods, use an alternative delimiter, such as `,` or `_`, which does not appear anywhere in the path; it 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 by using the string:

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

where the delimiter is the `,` character.

- The Event sensor and PhaseEvent sensor names should have the form *verbnoun*. For example, `activateInstance` and `runMethod`. When a PhaseEvent monitors a function, method, or code block, it must be named to reflect the task performed as clearly as possible.
- The name of a state sensor must be a noun, possibly preceded by an adjective, which describes the semantics of the value that is tracked with this state sensor. For example, `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 4-1](#), [Table 4-2](#), and [Table 4-3](#).

DMS Nouns

DMS **nouns** organize performance data. Sensors, with their associated metrics, are organized in 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 the attribute that identifies the noun's type. Nouns that represent similar types of entities typically have the same noun type and usually record a common set of measurements for each of those entities.

- [General DMS Naming](#)
- [General DMS Naming Conventions and Character Sets](#)
- [Noun and Noun Type Naming Conventions](#)

General DMS Naming

A **noun name** is a string, which does not include 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. For example, `/dmsDemo/BasicBinomial/{http://mynamespace/}JAXWSHelloService` is a noun full name.

A **sensor name** is a string, which does not include 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. For example, `/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.

General DMS Naming Conventions and Character Sets

DMS names must 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 4-4 shows the DMS replacement for special characters in names.

Table 4-4 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.

Noun and Noun Type Naming Conventions

The following conventions are used when naming noun and noun types:

- A noun name must be unique.
- A noun name must identify a specific entity of interest.

- Noun types should have names that clearly reflect the set of metrics that are 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 must start with a capital letter to distinguish them from other DMS names. All nouns of a given type must contain the same set of sensors.
- The noun naming scheme uses a *l* as the root of the hierarchy, with each noun acting as a container under the root or under its parent noun.

DMS Tracing and Events

Conceptually, DMS generates a stream of events; each event is in response to one of the event-producing actions that are being performed on the DMS API by the components that integrate with DMS (such as a sensor being updated). That stream of events can be ignored or routed (and optionally filtered) to destinations that can respond in some way to events.

[Table 4-5](#) provides a list of DMS tracing and event terminology.

Table 4-5 DMS Tracing and Event Terminology

DMS Term	Definition
Condition	<p>A condition is the logic behind a condition filter. It determines which events might pass through a filter, based on the rules defined in the condition. Every condition filter has zero or one root condition, but conditions might 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 that is associated with a sensor or noun event. • Context Condition: operates on the values currently set within the current Execution Context. <p>See 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 logs events to a file, sends transformed copies of events to the Java Flight Recorder, renders information gathered from incoming events as data in an MBean.</p>
Event Route	<p>An event route connects a filter to a destination. Event routes can be enabled or disabled.</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 the events that are passed and the events that 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 <code>role-admin</code> • Only pass sensor updates from nouns of type <code>JDBC_Statement</code> <p>See DMS Tracing and Events.</p>
Listener	<p>A DMS listener is also known as the destination. See Configuring Destinations.</p>

About DMS Availability

DMS functionality is available on all certified Jakarta EE servers.

This includes both the runtime features and supporting commands. Also, several features of DMS operates in JSE applications and standalone C applications.

For details on which servers are certified, see the Oracle Fusion Middleware Certification Matrix.

About DMS Architecture

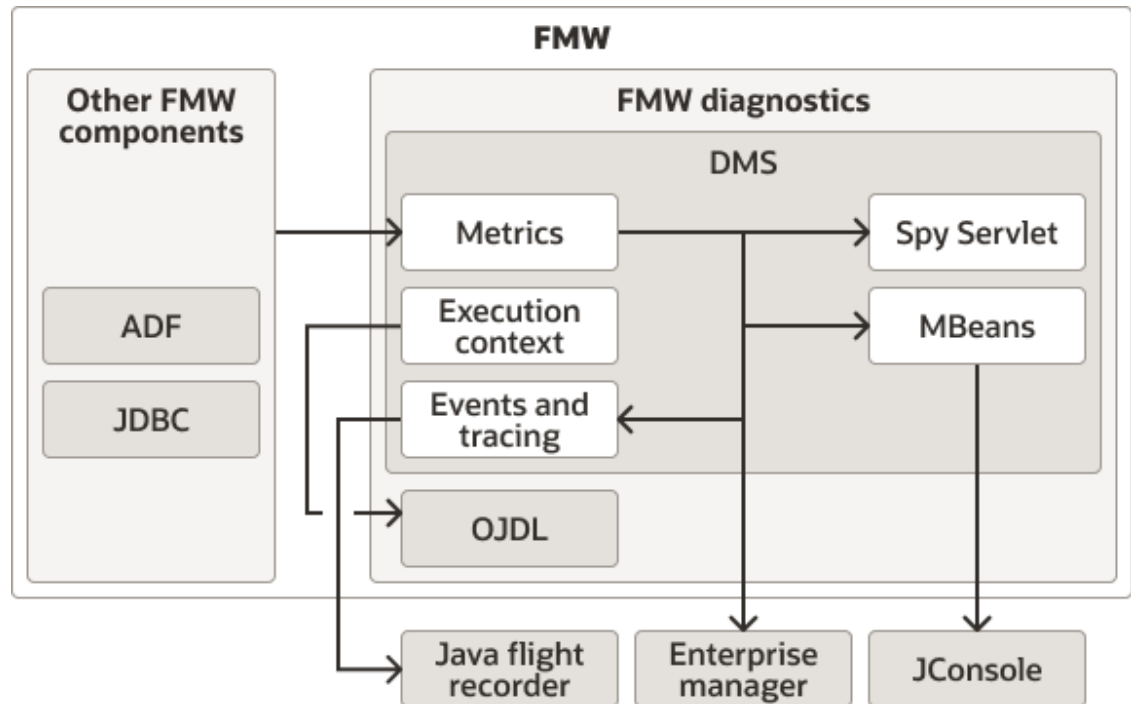
It is important to understand the components of DMS and how they interact with other Oracle Fusion Middleware components.

DMS consists of the following features:

- **DMS Metrics:** The DMS metrics feature provides Java and C APIs, which the Oracle Fusion Middleware components use for instrumenting code with performance measurements and other useful state metrics.
- **Execution Context:** Execution Context supports the maintenance and propagation of a specific context structure throughout the Oracle stack. By exploiting the propagated context structure, Oracle Fusion Middleware components can record diagnostic information (such as log records) that can be correlated between different components and products running on the same or different servers and hosts. See [About DMS Execution Context](#).
- **Events and Tracing:** Event Tracing enables you to configure live tracing with no restarts. DMS metrics that are updated by using Oracle Fusion Middleware products must be traced by using the DMS Event Tracing feature. The system has been designed to facilitate not only tracing, but also to support the other functionality that is driven from DMS activity.

[Figure 4-1](#) shows the components of DMS and how they interact with other Oracle Fusion Middleware components. The arrows show the direction in which information flows from one component to the next.

Figure 4-1 DMS Interactions with Oracle Fusion Middleware Components



Viewing DMS Metrics

Oracle Fusion Middleware components are instrumented with DMS metrics 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 *Viewing the Performance of Oracle Fusion Middleware* in *Administering Oracle Fusion Middleware* 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-Jakarta EE servers, such as Oracle servers.

You can use various tools to view the DMS metrics.

- [Viewing Metrics By 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 By 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 1521.

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_12.1.2/dms.war`.

See [DMS Spy Servlet](#).



Note:

The Spy Servlet is secured by using standard Jakarta EE declarative security in the web-application's `web.xml` file, and access is granted only to members of the Administrator's group.

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. See [Configuring the Harvester for Metric Collection in *Configuring and Using the Diagnostics Framework for Oracle WebLogic Server*](#).

Viewing Metrics with WLST (Oracle WebLogic Server)

DMS provides three commands to view metrics in WLST and they are detailed in the table below.

Table 4-6 DMS Commands

Use this command...	To do this...
<code>displayMetricTableNames()</code>	<p>List the names of the available metric tables.</p> <p>If there are many metric tables, consider using the <code>outputfile</code> parameter with <code>displayMetricTableNames()</code>. It is useful when the output is expected to be large. When <code>displayMetricTableNames()</code> has the <code>outputfile</code> parameter, it returns null to the script instead of the whole output. This prevents the command from running out of memory.</p> <p>NOTE: The command syntax for <code>displayMetricTableNames()</code> differs slightly for system components (such as OHS). After you connect WLST to Node Manager by using the <code>nmConnect()</code> command, you must specify both the server name and the server type explicitly.</p> <p>For example:</p> <pre>displayMetricTableNames (servertype="OHS", servers="ohs1")</pre>

Table 4-6 (Cont.) DMS Commands

Use this command...	To do this...
<code>displayMetricTables()</code>	<p>Show the content of the DMS metric tables.</p> <p>If you have many DMS metric tables, consider using the <code>outputfile</code> parameter with <code>displayMetricTables()</code>. This is useful when the output is expected to be large. When <code>displayMetricTables()</code> has the <code>outputfile</code> parameter, it returns null to the script instead of the whole output. This prevents the command from running out of memory.</p> <p>NOTE: The command syntax for <code>displayMetricTables()</code> differs slightly for system components (such as OHS). After you connect WLST to Node Manager by using the <code>nmConnect()</code> command, you must specify both the server name and the server type explicitly.</p> <p>For example:</p> <pre>displayMetricTables(servertype="OHS", servers="ohs1")</pre>
<code>dumpMetrics()</code>	<p>Display metrics in the internal format. Valid formats for the <code>dumpMetrics</code> command include raw, xml, and pdml.</p> <p>If you have many DMS metric tables, consider using the <code>outputfile</code> parameter with <code>dumpMetrics()</code>. This is useful when the output is expected to be large. When <code>dumpMetrics()</code> has the <code>outputfile</code> parameter, it returns null to the script instead of the whole output. This prevents the command from running out of memory.</p> <p>NOTE: The command syntax for <code>dumpMetrics()</code> differs slightly for system components (such as OHS). After you connect WLST to Node Manager by using the <code>nmConnect()</code> command, you must specify both the server name and the server type explicitly.</p> <p>For example:</p> <pre>dumpMetrics()(servertype="OHS", servers="ohs1")</pre>

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

For details on using these commands, see the following:

- Getting Started Using the Oracle WebLogic Scripting Tool (WLST) in *Administering Oracle Fusion Middleware*
- DMS Metric Commands in *WLST Command Reference for Infrastructure Components*

Viewing Metrics with JConsole

To provide a standards-based way to access metrics, DMS exposes them through MBeans. An MBean is created and registered for each type 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). 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 Jakarta 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-Jakarta EE component 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.

Viewing Metrics with Oracle Enterprise Manager

Oracle Fusion Middleware automatically and continuously measures data regarding the component's performance, state, and the on-going behavior. The metrics are automatically enabled; there is no need to set options or perform any extra configuration to collect them. See [Oracle Enterprise Manager Fusion Middleware Control](#).

About DMS Execution Context

The DMS execution context is the mechanism by which requests (such as RMI requests) can be uniquely identified and thus tracked as they flow through the system.

It also provides the means by which context information can be communicated between cooperating Fusion Middleware components involved in fulfilling requests.

- [DMS Execution Requests and Subtasks](#)
- [DMS Execution Context Usage](#)
- [DMS Execution Context Communication](#)

DMS Execution Requests and Subtasks

The DMS execution context has been developed with the understanding that a single request (or task) might create many subtasks that are coordinated to complete the request or root task. Consider the following examples of requests and their associated subtasks:

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

A DMS execution context is composed of the following:

- A unique identifier, the Execution Context ID (ECID).
The ECID is unique for each new root task and is shared across the tree of tasks that are associated with the root task.
- A relationship identifier, the Relationship ID (RID).
The 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 a request is sent from an Oracle 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 Server:
 - New ECID = B5C094FA...BE4AE8
 - Root RID = 0
2. Single subtask on Oracle WebLogic Server:
 - Same ECID = B5C094FA...BE4AE8
 - Sub-task RID = 0:1
3. Two subtasks, 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

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 that are associated with that task by querying the log files for messages that contain that ECID.

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

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

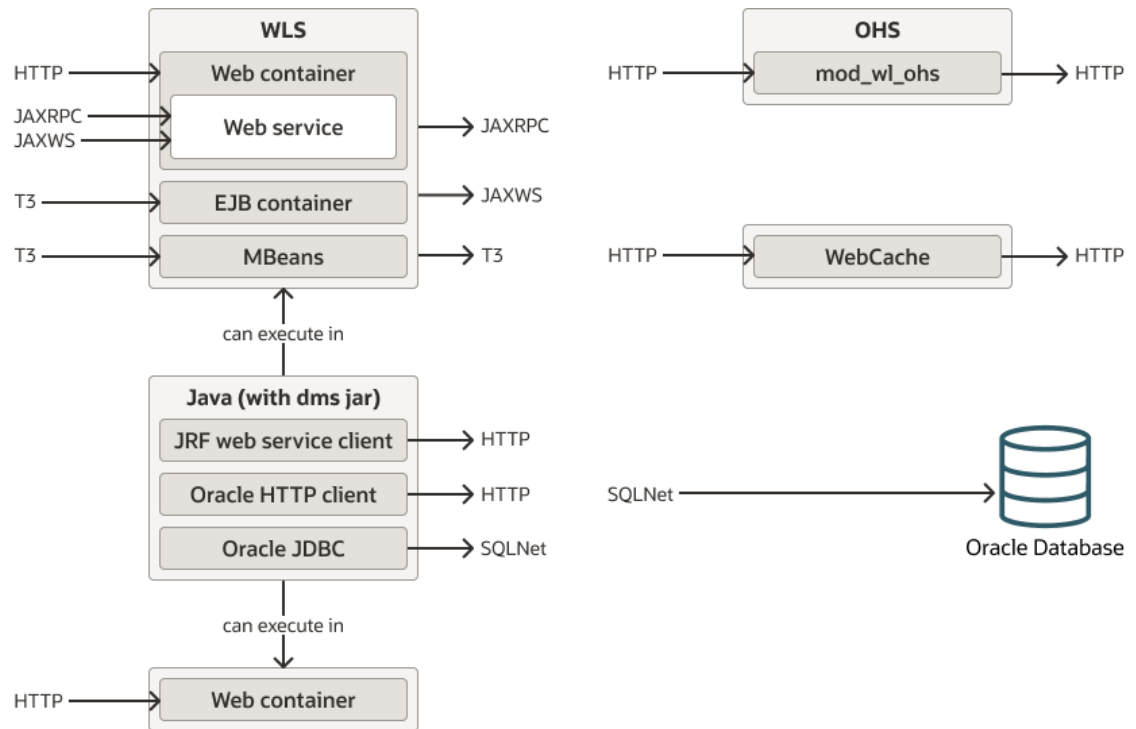
In this example, any log files with messages that contain the ECID B5C094FA...BE4AE8 is displayed.

DMS Execution Context Communication

Figure 4-2 shows the components that cooperate 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 4-2 DMS Execution Context Communication Protocols



DMS Tracing and Events

The DMS tracing feature can be used to diagnose issues or collect specific data at a specific time for a specific set of criteria.

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)
- 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 to define a filter rule once and have the resulting subset of all possible events processed on one or more destinations.

The configuration can be modified by 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.

See *Configuring Selective Tracing Using WLST* in *Administering Oracle Fusion Middleware*.

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 conditions mentioned
- [Configuring the DMS Event System](#)
- [Configuring Destinations](#)
- [Understanding the Format of DMS Events in Log Messages](#)
- [Understanding DMS Event Actions](#)

Configuring the DMS Event System

Configuration is recorded in each server `dms_config.xml` file. MBean updates can be made at runtime by using the 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=JMXEventConfig
```

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

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

The resulting output looks similar to:

```
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

- [Adding and Editing Filters](#)
- [Adding and Editing Destinations](#)
- [Adding and Editing Event Routes](#)
- [Compound Operations](#)

Adding and Editing Filters

Filters define the rules that select the events that 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 that are associated with JDBC operations are performed on nouns of types whose names begin with `JDBC_`.

If the rule must be modified, the filter must 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 by using either the shortened or longer name.

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 works, but is discouraged.

Table 4-7 DMS Operators

Noun Type Operators	Details
<code>equals, eq</code>	<code>notEquals, ne</code>
<code>contains</code>	<code>in</code>
<code>startsWith, sw</code>	-
Context Operators	Details
<code>equals, eq</code>	<code>notequals, ne</code>
<code>isnull</code>	<code>isnotnull</code>
<code>startswith, sw</code>	<code>contains</code>
<code>lt</code>	<code>gt</code>

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 details on the syntax used to describe a filter's rule (the condition property), refer to the *WebLogic Scripting Tool Command Reference* or the command help.

Adding and Editing Destinations

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

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

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

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

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

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

Adding and Editing Event Routes

The following example shows how to join the filter and create a destination.

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

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 that are 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')
```

Compound Operations

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

Note:

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 by 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"
```

Configuring Destinations

DMS offers several types of destinations.

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

LoggerDestination

Table 4-8 Logger Destination

Properties	Details
Description	The <i>LoggerDestination</i> writes each event to the associated logger.
Implementing Class	<i>oracle.dms.trace2.runtime.LoggerDestination</i>
Properties	
loggerName	The name of the ODL logger to which events are 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 [Static Loggers and Handlers](#)). 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 [Dynamic Loggers and Handlers](#)).

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

- [Static Loggers and Handlers](#)
- [Dynamic Loggers and Handlers](#)
- [Default Locations of the logging.xml File](#)
- [Using a CLI Command to Query the Trace Log File](#)

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`. It allows you to define the name of the log file, the maximum size, format, file rotation, and policies.

Oracle recommends using commands (like the example here) 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 details on logging configuration, see *Managing Log Files and Diagnostic Data in Administering Oracle Fusion Middleware*.

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

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

See [Understanding the Format of DMS Events in Log Messages](#).

Dynamic Loggers and Handlers

If the named logger has no associated handler defined in `logging.xml`, then the logger destination dynamically creates a handler object that writes 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 is the logger's name followed by `-event.log`. For instance, in the example in [Static Loggers and Handlers](#), DMS events would be written to `myTraceLogger-event.log`.

Default Locations of the logging.xml File

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

Table 4-9 Default locations of the logging.xml file

Platform	Server	Location
Oracle WebLogic Server	AdminServer	ORACLE_HOME/WLS_Home/ user_projects/domains/ base_domain/config/fmwconfig/ servers/AdminServer/logging.xml

Using a CLI Command to Query the Trace Log File

If the logger destination's logger and handler are defined in the `logging.xml` file then you can take advantage of the `displayLogs()` command to 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 that have an ECID of 0000HpmSpLWEkJQ6ub3FEH194kwB000004:

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

- To display only the log messages for myTraceLogger that 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 file name must be included:

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

MBean Creator Destination

Table 4-10 MBean Creator Destination Details

Properties	Details
Description	The MBean creator destination make nouns accessible as MBeans, exposing their metrics as attributes, for access through WLDF, JConsole, and so on.
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 creates 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 types that are of interest to the administrator.

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

WebLogic Diagnostic Framework (WLDF) can be used to harvest DMS metrics exposed by the MBean creator destination. See *Configuring and Using the Diagnostics Framework for Oracle WebLogic Server*.

- [Metric MBean Object Name](#)

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.
```

Request Tracker Destination

Table 4-11 Request Tracker Destination Details

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

Use in the default configuration: An instance of the request tracker destination is enabled by default. When a DFW incident is generated, the active request list is dumped automatically, allowing an administrator to correlate the failure with a specific request.

For each request the following information is dumped:

- Uniform Resource Identifier (URI)
- Start time of the request
- Execution Context ID (ECID)
- Query string
- Headers

When the request tracker is not enabled the Request Dump outputs the following:

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

- [Executing the Request Tracker Dump](#)

Executing the Request Tracker Dump

The information maintained by the request tracker can be accessed manually. When connected to a server, to execute the dump that reports the request information the WLST `executeDump` command can be used, as follows:

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

StartTime: 2009-12-14 02:24:41.870
ECID: 0000IMChyqEC8xT6uBf9EH1B9X9^000009,0
URI: /myApp/Welcome.jsp
QueryString:
Headers:
  Host: myHost.example.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
```

Java Flight Recorder Destination

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

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

1. Application level events and JVM level events can be reported as a single sequence. This avoids the need to combine such events from separate log files based only on the timestamp (which might not tick over fast enough to 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 so the JVM exits 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.
- [Dynamically Derived JFR Event Types – Names, Values, and Descriptions](#)

Dynamically Derived JFR Event Types – Names, Values, and Descriptions

A DMS noun type is associated with a JFR `InstantEvent` event type:

- The name of the JFR event type for a noun type is the noun type name with the suffix `state`.
- The path of the JFR event type for a noun type is `dms/` followed by the `producer-name`, followed by the event type name.
- Event sensors do not contribute any values to the JFR event type.
- The values of the JFR event for a noun type are described in [Table 4-12](#):

Table 4-12 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 4-12 (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 is populated with the full path of the noun. The field name assumes that the noun_type meaningfully categorizes all objects measured by the nouns of that type.
<state-sensor-name>	The value of the state sensor.	No	Each state sensor belonging to the noun contributes one of these values to the instant event. There might 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 counting 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 is associated with a JFR DurationEvent event type in the following ways:

- The name of the JFR event type for a phase sensor belonging to a noun of a particular noun type is the **noun type** name followed by the phase sensor name.
- The path of the JFR event for a noun type is `dms/` followed by the `producer-name`, followed by the event type name.
- The values of the duration event is as mentioned (except for the sensorName value). For example, the **stop** of a phase event results in a JFR duration event being reported to JFR that contains the state information of the phase event parent noun.

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

- Noun type description is used in creation of the JFR event type.
- State and event sensor descriptions are not applied—there is nowhere to apply them.
- Phase sensor descriptions are applied to their JFR event type.
- [Examples of Dynamically Derived Producers and Events](#)

Examples of Dynamically Derived Producers and Events

[Table 4-13](#) provides examples for the rules described in [Dynamically Derived JFR Event Types – Names, Values, and Descriptions](#):

Table 4-13 Examples of Dynamically Derived Producers and Events

DMS	Java 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 <i>noun type</i> State Event Type Path: dms/JDBC/JDBC_Connection_State <i>dms/leading component of noun path/noun type/_State</i></p> <p>Fields:</p> <ul style="list-style-type: none"> • ECID • RID • JDBC_Connection name Value is the full path of the noun • JDBC_Connection_Url Value of the state sensor of this name at the time of the event • JDBC_Connection_Username Value of the state sensor of this name at the time of the event • Event Name Value is 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: <i>dms/JDBC/JDBC_Connection_CreateStatement</i></p> <p>Fields:</p> <ul style="list-style-type: none"> • ECID • RID • JDBC_Connection name • JDBC_Connection_Url • JDBC_Connection_Username

Table 4-13 (Cont.) Examples of Dynamically Derived Producers and Events

DMS	Java Flight Recorder (JFR)
-	Producer Name: JDBC Event Type 3 Event Type Name: JDBC_Connection CreateNewStatement Event Type Path: dms/JDBC/JDBC_Connection_CreateNewStatement Fields: <ul style="list-style-type: none"> • ECID • RID • JDBC_Connection name • JDBC_Connection_Url • JDBC_Connection_Username
-	Producer Name: JDBC Event Type 4 Event Type Name: JDBC_Connection DBWaitTime Event Type Path: dms/JDBC/JDBC_Connection_DBWaitTime Fields: <ul style="list-style-type: none"> • ECID • RID • JDBC_Connection name • JDBC_Connection_Url • JDBC_Connection_Username

Understanding the Format of DMS Events in Log Messages

Table 4-14 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 4-14 Event Formatting Descriptions

Applicable Events	Field Number	Name	Description
All	1	Version number	The version number of the event format. For example: v1:1280737384058: _REQUE ST:STOP:/MyWebApp/emp
All	2	Event time	The time at which the event occurred. For example: v1:1280737384058: _REQUE ST:STOP:/MyWebApp/emp

Table 4-14 (Cont.) Event Formatting Descriptions

Applicable Events	Field Number	Name	Description
All	3	Source object type	<p>The type of object on which an action was performed to produce the event including:</p> <ul style="list-style-type: none"> • NOUN • EVENT_SENSOR • STATE_SENSOR • PHASE_SENSOR • EXECUTION_CONTEXT • _REQUEST <p>For example: v1:1280737384058: _REQUE ST: STOP:/MyWebApp/emp</p>
All	4	Action type	<p>The type of action that resulted in the generation of this event. A given source object type might not produce events for every action type:</p> <ul style="list-style-type: none"> • CREATE • UPDATE • DELETE • START • STOP • ABORT <p>For example: v1:1280737384058: _REQUE ST: STOP:/MyWebApp/emp</p>
Nouns	5	Noun type	<p>The name of the noun type.</p> <p>For example: v1:1281344803506:NOUN:C REATE: JDBC_Connection:/JDBC/JDBC Data Source-0/CONNECTION_1</p>
	6	Noun path	<p>The full path identifying the noun to which the sensor belongs</p> <p>For example: v1:1281344803506:NOUN:C REATE:JDBC_Connection:/JDBC/JDBC Data Source-0/CONNECTION_1</p>

Table 4-14 (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 4-14 (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>
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:_REQUEST:START:/myWebApp/showEmployees</pre> <pre>v1:1280737384058:_REQUEST:STOP:/myWebApp/showEmployees</pre>

Table 4-14 (Cont.) Event Formatting Descriptions

Applicable Events	Field Number	Name	Description
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 might contain : but do not interpret them as event field separators.</p> <p>For example:</p> <pre>v1:1280737384058:EXECUT ION_CONTEXT:STOP:bc4fd0 668f79d507:367c127f:12a 23f2013c:-8000-00000000 00000f73,0</pre>

Understanding DMS Event Actions

Table 4-15 shows the action types that can be performed on source object types.

Table 4-15 Actions Performed on Source Object Types

Object Type	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	-	-	-
Execution Context	-	-	-	Yes	Yes	-
Request	-	-	-	Yes	Yes	-

DMS Best Practices

Implement the following best practices when you use DMS metrics.

The use of DMS metrics can have an impact on application performance. When you add 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. Use a high resolution clock to time phase events accurately than 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 might not tick (change) once per millisecond. Instead, it might take up to 15 ms to tick as shown in the following example:

Table 4-16 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 4-16 shows a phase with a 12 ms 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 2 ms duration running from 12:00:00.014 to 12:00:00.016 would be reported in system time as having a duration of 15 ms.

 **Note:**

These behaviors are more evident on some operating systems than others. Use caution when you analyze individual periods of time that are shorter than the tick period of the system clock. Configuring DMS to use a higher resolution clock causes 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 by 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 you use the high resolution clock, then set the units property as follows:

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

Table 4-17 shows the supported values for the `oracle.dms.clock` property.

Table 4-17 The 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 PhaseEvents. The default value for the units for the default clock is MSECS.
HIGHRES	The Java Highres clock uses <code>System.nanoTime()</code> (no JNI required).

Table 4-18 shows the supported values for the `oracle.dms.clock.units` property.

Table 4-18 oracle.dms.clock.units Property Values

Value	Description
MSECS	Specifies that the time must 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 must be converted to microseconds and reported as usecs . A microsecond is 10^{-6} seconds.
NSECS	Specifies that the time must 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 you use 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`, and `HighRes`.
- DMS checks the property values at startup. When the clock property is set with a value that is not listed in Table 4-17, 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 4-18, DMS uses the default units for the specified clock.

Part II

Core Components

The core components in Oracle Fusion Middleware need to be tuned for optimal performance.

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



Note:

For information on performance tuning the Oracle WebLogic Server, see *Tuning Performance of Oracle WebLogic Server*.

- [Tuning Oracle HTTP Server](#)
You can tune Oracle HTTP Server (OHS) to optimize its performance as the web server component for Oracle Fusion Middleware.
- [Tuning Oracle Metadata Service](#)
You can tune Oracle Metadata Services (MDS) to optimize its performance as an application server and Oracle relational database.
- [Tuning Oracle Fusion Middleware Security](#)
You can tune Oracle Fusion Middleware security services to optimize the performance of security services through Oracle Platform Security Services (OPSS) and Oracle Web Services.

5

Tuning Oracle HTTP Server

You can tune Oracle HTTP Server (OHS) to optimize its performance as the web server component for Oracle Fusion Middleware.

Note:

The configuration examples and recommended settings are for illustrative purposes only. Consult your own use case scenarios to determine the configuration options that can provide performance improvements.

- [About Oracle HTTP Server](#)
Oracle HTTP Server (OHS) is the Web server component for Oracle Fusion Middleware.
- [Monitoring Oracle HTTP Server Performance](#)
Oracle Fusion Middleware automatically and continuously measures runtime performance for Oracle HTTP Server.
- [Basic Tuning Considerations](#)
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.
- [Advanced Tuning Considerations](#)
Advanced tuning recommendations may or may not apply to your environment. Review the following recommendations to determine if the changes would improve your Oracle HTTP Server performance.

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 2.4.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 on the Apache open-source software infrastructure, see the Apache Software Foundation at <http://www.apache.org/>.

Monitoring Oracle HTTP Server Performance

Oracle Fusion Middleware automatically and continuously measures runtime 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. See *Managing and Monitoring Server Processes* in *Administering Oracle HTTP Server*.

For monitoring, Oracle HTTP Server uses the 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 point 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.

See [Viewing Metrics with WLST \(Oracle WebLogic Server\)](#).

Basic Tuning Considerations

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.

- [Tuning Oracle HTTP Server Directives](#)
- [Reducing Process Availability with Persistent Connections](#)
- [Logging Options for Oracle HTTP Server](#)

Tuning Oracle HTTP Server Directives

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

See the Oracle HTTP Server, see *Understanding Oracle HTTP Server Management Tools* in *Administering Oracle HTTP Server*.

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

- **Worker:** It uses Multi-Process-Multi-Threads model and is the default MPM on all platforms other than Microsoft Windows platforms. Multithread support makes it more scalable by using fewer system resources and multiprocess 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. It is considered the least efficient MPM.
- **Event:** This MPM is designed to allow more requests to be served simultaneously by passing off some processing work to supporting threads, freeing up the main threads to

work on new requests. It is based on the worker MPM, which implements a hybrid multiprocess multithreaded server.

The directives for each MPM type are defined in the `DOMAIN_HOME/config/fmwconfig/components/OHS/<componentName>/httpd.conf`. The default MPM type is the event MPM. To use a different MPM (such as prefork MPM), edit the `ohs.plugins.nodemanager.properties` file found in the same directory.

 **Note:**

The information here is based on the use of worker and WinNT MPMs, which use threads. The directives listed might 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 <https://docs.oracle.com/en/middleware/webcenter/index.html>.

Table 5-1 Oracle HTTP Server Configuration Properties

Directive	Description
<code>ListenBackLog</code> This directive maps to the Maximum Queue Length field on the Performance Directives screen.	Specifies the maximum length of the queue of pending connections. Generally no tuning is needed. 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. Default Value: 511

Table 5-1 (Cont.) Oracle HTTP Server Configuration Properties

Directive	Description
<p><code>MaxRequestWorkers</code></p> <p>This directive maps to the Maximum Requests field on the Performance Directives screen.</p> <p>This parameter is not available in <code>mod_winnt</code> (Microsoft Windows). Winnt uses a single process, multithreaded model and is controlled by the <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>MaxRequestWorkers</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>MaxRequestWorkers</code> to increase server concurrency. Increase <code>MaxRequestWorkers</code> until your system becomes fully utilized (85% is a good threshold).</p> <p>Conversely, when system resources are saturated, increasing <code>MaxRequestWorkers</code> does not improve performance. In this case, the <code>MaxRequestWorkers</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 might require sufficient concurrent <code>httpd</code> or thread server processes to handle both active and idle connections. When you specify <code>MaxRequestWorkers</code> to act as a throttle for system concurrency, you must consider that persistent idle <code>httpd</code> connections also consume <code>httpd/thread</code> 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, during the connection, even if no requests are currently being processed for the connection.</p> <p>If you have sufficient capacity, <code>KeepAlive</code> must be enabled; using persistent connections improves performance and prevents wasting CPU resources reestablishing connections. Normally, you should not change <code>KeepAlive</code> parameters.</p> <p>The maximum allowed value for <code>MaxRequestWorkers</code> is 8192 (8K).</p> <p>Default Value: 150</p>

Table 5-1 (Cont.) Oracle HTTP Server Configuration Properties

Directive	Description
<p>StartServers</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 that are created on startup. If you expect a sudden load after restart, set this value based on the number of child servers that are required.</p> <p>The following parameters are inter-related and applicable only on UNIX platforms (worker_mpm):</p> <ul style="list-style-type: none"> • MaxRequestWorkers • MaxSpareThreads and MinSpareThreads • ServerLimit and StartServers <p>On the Windows platform (mpm_winnt), as well as UNIX platforms, the following parameters are important to tune:</p> <ul style="list-style-type: none"> • ThreadLimit • ThreadsPerChild <p>Each child process has a set of child threads that are defined for them and that can actually handle the requests. Use <code>ThreadsPerChild</code> with this directive.</p> <p>The values of <code>ThreadLimit</code>, <code>ServerLimit</code>, and <code>MaxRequestWorkers</code> can indirectly affect this value. Read the notes for these directives and use them with this directive.</p> <p>Default Value: 2</p>
<p>ServerLimit</p> <p>This parameter is not available in mod_winnt (Microsoft Windows). Winnt uses a single process, multithreaded 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. It is used to control the maximum number of server processes that can be created.</p> <p>Default Value: 16</p>
<p>ThreadLimit</p>	<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. It is used to control the maximum number of threads created per process to avoid conflicts or issues.</p> <p>Default Values:</p> <ul style="list-style-type: none"> • Windows Multi-Processing Module (mpm_winnt): 1920 • All others: 64

Table 5-1 (Cont.) Oracle HTTP Server Configuration Properties

Directive	Description
<p>ThreadsPerChild</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: 150 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 or worker), as specified in <code>ThreadsPerChild</code>, and a listener thread, which listens for requests and transfers the control to the worker or 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 <code>ServerLimit</code>, <code>ThreadLimit</code>, and <code>MaxRequestWorkers</code> directives affect the other directives as follows:</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>MaxRequestWorkers</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>MaxRequestWorkers</code>: Defines the upper limit on the number of server threads that can process requests simultaneously. This must be equal to the number of simultaneous connections that can be made. This value must be a multiple of <code>ThreadsPerChild</code>. If <code>MaxRequestWorkers</code> is greater than <code>ServerLimit</code> multiplied by <code>ThreadsPerChild</code>, it is automatically trimmed to the latter value.

Table 5-1 (Cont.) Oracle HTTP Server Configuration Properties

Directive	Description
<p>MaxConnectionsPerChild</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 that 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, it is not needed, but some UNIX systems have notable leaks in the libraries. For these platforms, set <code>MaxConnectionsPerChild</code> to 10000; a setting of 0 means unlimited requests.</p> <p>This value does not include <code>KeepAlive</code> 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: Windows systems <code>MaxConnectionsPerChild</code> must 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>These parameters are not available in <code>mod_winnt</code> (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 more server threads to handle transient load increases such as multiple simultaneous requests from a single browser.</p> <p>The server periodically checks how many server threads are waiting for a request. If there is fewer than <code>MinSpareThreads</code>, it creates a new spare. If there is more than <code>MaxSpareThreads</code>, some of the spares are removed.</p> <p>Default Values:</p> <p><code>MaxSpareThreads</code>: 75</p> <p><code>MinSpareThreads</code>: 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 to wait for an incoming request to be received before sending a 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 to allow persistent connections (more than one request per connection). Set to <code>Off</code> to deactivate.</p> <p>Default Value: <code>On</code></p>
<p>MaxKeepAliveRequests</p>	<p>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.</p> <p>Default Value: 100</p>
<p>KeepAliveTimeout</p> <p>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.</p>	<p>Number of seconds to wait for the next request from the same client on the same connection.</p> <p>Default Value: 5 seconds</p>

Table 5-1 (Cont.) Oracle HTTP Server Configuration Properties

Directive	Description
<code>limit</code> <code>ulimit</code>	<p>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.</p> <p>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 then adjusts the soft limit (<code>limit</code>) to match.</p> <p>Configuring file descriptor limits is platform-specific. Refer to your operating system documentation for details.</p>

Reducing Process Availability with Persistent Connections

If your browser supports persistent connections, you can support them on the server by 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. Consider the size and behavior of your own user population when you set 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.

The `KeepAlive` option should be used judiciously along with `MaxRequestWorkers` directive. The `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 the user population consistently higher than the specified `MaxRequestWorkers`.

Note:

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

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

- A high number of `MaxRequestWorkers` can overload the system resources and may lead to poor performance.

- For a high user population with fewer requests, consider increasing the `MaxRequestWorkers` to support the `KeepAlive` connections to avoid starvation. This can impact overall performance when the user concurrency increases. System performance is impacted by increased concurrency and can possibly cause the system to fail.

`MaxRequestWorkers` must 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 off the `KeepAlive` option or reduce it to a 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 you consider the nature of the workload and the system capacity.

Logging Options for Oracle HTTP Server

The logging options for Oracle HTTP Server include types of logging, log levels, and the performance implications for using logging.

- [Access Logging](#)
- [Configuring the `HostNameLookups` Directive](#)
- [Error logging](#)

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.

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 that is 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 offline, with the `logresolve` utility found in the `/Apache/Apache/bin/` directory.

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 `d.conf` file contains the error log configuration for OHS. The `OraLogMode` directive defines the logging mode. The default is `odl-text`, which produces the Oracle diagnostic logging format in a text file. Alternatively, change it 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 might also fill disk space. The log level control must be used based on need. For requests that use dynamic resources, like `mod_ossso` or `mod_plsql`, there is a performance cost associated.

Advanced Tuning Considerations

Advanced tuning recommendations may or may not apply to your environment. Review the following recommendations to determine if the changes would improve your Oracle HTTP Server performance.

- [Tuning Oracle HTTP Server](#)
- [Tuning Oracle HTTP Server Security](#)

Tuning Oracle HTTP Server

You can follow the topics to avoid or debug potential Oracle HTTP Server performance problems.

- [Analyzing Static Versus Dynamic Requests](#)
- [Limiting the Number of Enabled Modules](#)
- [Tuning the File Descriptor Limit](#)

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. When you configure 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.

Limiting the Number of Enabled Modules

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

The new architecture, Oracle HTTP Server invokes the service function of each module that is loaded (in the order of definition in the `d.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 `d.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. For 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. For 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.

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. You can tune the file descriptor limit by configuring the operating system's hard limit for the user who starts the OHS. Once configured, the OHS adjusts the soft limit to match the operating system limit.

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

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



Note:

This limit must be reconfigured after you apply a patch set.

Tuning Oracle HTTP Server Security

Tuning Oracle HTTP Server includes tuning the SSL and Port Tunneling.

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

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 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.

- [Caching SSL on Oracle HTTP Server](#)
- [Using SSL Application Level Data Encryption](#)
- [Tuning SSL Performance](#)

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 the `ssl.conf` file determines how long the server keeps a saved SSL session (the default is 300 seconds). The 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
```

Multiple Oracle HTTP Server processes can use a saved session cache file.

Saving the 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 the saved SSL session state has some performance costs. When the 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 you use persistent connections. Oracle HTTP Server uses persistent connections by default, assuming they are supported on the client-side. In over SSL as implemented by Oracle HTTP Server, the SSL session state is kept in memory while the associated connection is persisted, a process which essentially eliminates the performance impacts that are associated with SSL session reuse (conceptually, the SSL connection is kept open along with the connection). For more information, see [Reducing Process Availability with Persistent Connections](#).

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 the `ssl.conf` file. OHS 11g uses the 128 bit Encryption algorithm by default and no longer supports lower encryption.

 **Note:**

The previous release [10.1.3x] used 64 bit encryption for Windows. For UNIX, the 10.x releases used the 128 bit encryption by default.

Tuning SSL Performance

The following recommendations can assist you to determine performance requirements when you work with Oracle HTTP Server and SSL.

- The SSL handshake is a 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 the 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. If the application requires server authentication only, do not use client authentication. If additional hardware is required, the performance goals of an application cannot be met by this method.
- Design the application to use SSL efficiently. Group secure operations to take advantage of SSL session reuse and SSL connection reuse.
- Use persistent connections, if possible, to minimize the cost of SSL session reuse.
- Tune the session cache timeout value (the `SSLSessionCacheTimeout` directive in the `ssl.conf` file). 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, must be completed without incurring session creation more than once. The default value for the `SSLSessionCacheTimeout` attribute is 300 seconds. Test the application usability to help tune this setting.
- If large volumes of data are being protected through SSL, pay close attention to the cipher suite being used. The `SSLCipherSuite` directive specified in the `ssl.conf` file controls the cipher suite. If lower levels of security are acceptable, use a less-secure protocol by using a smaller key size (improves performance significantly). Finally, test the application by using each available cipher suite for the specified security level to find the optimal suite.
- 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.

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.

Except for 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; it 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. Ensure to check the number of open ports. Start processes equivalent to the number of open ports.

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 must be compared with a system without OracleAS Port Tunneling to determine whether your performance requirements can be met by using OracleAS Port Tunneling.

6

Tuning Oracle Metadata Service

You can tune Oracle Metadata Services (MDS) to optimize its performance as an application server and Oracle relational database.

- [About Oracle Metadata Services \(MDS\)](#)
Oracle Metadata Services (MDS) is an application server and Oracle relational database that keeps metadata in these areas: the `ClassPath`, the `ServletContext`, database repository and, sometimes, the file system.
- [Monitoring Oracle Metadata Service Performance](#)
MDS uses DMS sensors to provide tuning and diagnostic information, which can be viewed by using Enterprise Manager. This information is useful, for example, to see if the MDS caches are large enough.
- [Basic Tuning Considerations](#)
Tuning the MDS configuration is essential for improving performance.
- [Advanced Tuning Considerations](#)
After you have performed recommended modifications, you can make additional changes that are specific to your deployment. Consider carefully whether the advance tuning recommendations are appropriate for your environment.

About Oracle Metadata Services (MDS)

Oracle Metadata Services (MDS) is an application server and Oracle relational database that keeps metadata in these areas: the `ClassPath`, the `ServletContext`, database repository and, sometimes, the file system.

One of the primary uses of MDS is to store customizations and persisted personalization for Oracle applications. MDS is used by components such as 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 Services configuration parameters are immutable and cannot be changed at runtime unless otherwise specified.

Tuning MDS tablespace and cache size is important before you tune Oracle B2B and other Oracle products. If you are using the *Using Oracle B2B* to tune B2B, make sure you have completed the tuning described here first.

Monitoring Oracle Metadata Service Performance

MDS uses DMS sensors to provide tuning and diagnostic information, which can be viewed by 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.

Basic Tuning Considerations

Tuning the MDS configuration is essential for improving performance.

The default MDS configuration must be tuned in almost all deployments. It is important to review the requirements and recommendations carefully.

- [Tuning Database Repository](#)
- [Tuning Cache Configuration](#)
- [Purging Document Version History](#)
- [Using Database Polling Interval for Change Detection](#)

Tuning Database Repository

For optimal performance of MDS APIs, the database schema for the MDS repository must be monitored and tuned by the database administrator.

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

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

Collecting Schema Statistics

While MDS provides database indexes, they might 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 the 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 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 must be configured with automatic statistics recollection. For additional information on gathering statistics, see Automatic Performance Statistics in *Oracle Database Performance Tuning Guide*.

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.

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

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 that is consumed by MDS schema to grow. 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*.

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.

See:

- Generating Automatic Workload Repository Reports in *Oracle Database Performance Tuning Guide*.
- Monitoring Performance in *Oracle Database Performance Tuning Guide*.

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 might 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 might 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 *Administering Oracle Fusion Middleware*.

 **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 must be less than 1. If not, consider increasing the size of the MDS cache. For more information on viewing DMS metric information, see .

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. Manually update the cache-config in the `adf-config.xml` file by adding the following entry prior to packaging the Enterprise ARchive (EAR) file:

```
<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.

- [Enabling Document Cache](#)

Enabling Document Cache

In addition to the main MDS cache, MDS uses a document cache 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 small value, MDS uses a minimum size of 500 for document cache.

In general, the defaults must be sufficient usually. However, insufficient document cache size might impact performance. Set document cache size by adding this entry to the `adf-config.xml` file prior to packaging the Enterprise ARchive (EAR) file:

```
<metadata-store-usage id="dbl">
  <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 [Monitoring Oracle Metadata Service Performance](#)) must be less than 1. If not, consider increasing the document cache size.

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, you can purge version history automatically or manually.

 **Note:**

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

- [Using Auto Purge](#)
- [Purging Manually](#)

Using 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 the `adf-config.xml` file prior to packaging the EAR:

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

In the example above, the auto-purge is executed every T seconds and removes versions that are older than the specified time T (in seconds). For more information, see [Changing MDS Configuration Attributes for Deployed Applications](#) in *Administering Oracle Fusion Middleware*.

Purging Manually

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

See *Purging Metadata Version History in Administering Oracle Fusion Middleware*.

Using Database Polling Interval for Change Detection

MDS employs a polling thread, which queries the database to check if the data in the MDS in-memory cache is out of sync with data in the database. It happens 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. Configure the polling interval by adding this entry in the `adf-config.xml` file prior to packaging the Enterprise ARchive (EAR) file:

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

In the example mentioned, `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 is suitable for most purposes. See *Changing MDS Configuration Attributes for Deployed Applications in Administering Oracle Fusion Middleware*.

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 might stack up and polling can take longer to process.

Advanced Tuning Considerations

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

- [Analyzing Performance Impact from Customization](#)

Analyzing Performance Impact from Customization

MDS customization might 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 that 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:** 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):** are layers of customization corresponding to customization classes whose `getCacheHint` method returns `SINGLE_USER`, meaning the layer applies to one user. User customizations are cached on the user's session (Session) until the user logs out.

For details on customization concepts, writing customization classes, and configuring customization classes, see Customizing Applications with MDS in *Developing Fusion Web Applications with Oracle Application Development Framework*.

7

Tuning Oracle Fusion Middleware Security

You can tune Oracle Fusion Middleware security services to optimize the performance of security services through Oracle Platform Security Services (OPSS) and Oracle Web Services.

- [About Security Services](#)
Oracle Fusion Middleware provides security services through Oracle Platform Security Services (OPSS) and Oracle Web Services.
- [Basic Tuning Considerations](#)
Tuning considerations might improve the performance of the Oracle Fusion Middleware security services.
- [Tuning Oracle Platform Security Services](#)
Oracle Platform Security Services (OPSS) includes the following basic tuning configurations.
- [Oracle Web Services Security Tuning](#)
Oracle Web Services Security provides a framework of authorization and authentication for interacting with a web service by using XML-based messages. There are several factors that may affect performance of the web service.

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 Security 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 Jakarta 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:
 1. Configure and invoke a login module, as appropriate. You can use provided login modules, or you can use custom login modules.
 2. Authenticate the user attempting to log in, which is the role of the identity store service.
 3. Authorize the user by checking permissions for that role.
- **Oracle Web Services Security**
Oracle Web Services Security provides a framework of authorization and authentication for interacting with a web service by using XML-based messages.



Note:

The information here assumes that you have reviewed and understand the concepts and administration information for Oracle Fusion Middleware Security Services. See, *Administering Web Services* before you tune any security parameters.

Basic Tuning Considerations

Tuning considerations might improve the performance of the Oracle Fusion Middleware security services.

If you discover a performance bottleneck, you must 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, 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 might need to change how you manage your database or LDAP connections or how you secure the resources.

Tuning Oracle Platform Security Services

Oracle Platform Security Services (OPSS) includes the following basic tuning configurations.

- [JVM Tuning Parameters](#)
- [JDK Tuning Parameters](#)
- [Authentication Tuning Parameters](#)
- [Authorization Tuning Properties](#)
- [OPSS PDP Service Tuning Parameters](#)

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 runtime, all roles and permissions are stored in the in-memory cache. For more JVM tuning information, see [Tuning Java Virtual Machines \(JVMs\)](#).

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 you use 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.

See the JDK 7 release notes at <http://www.oracle.com/technetwork/java/javase/jdk7-relnotes-418459.html>

Authentication Tuning Parameters

For OPSS Authentication tuning, see "Improving the Performance of WebLogic and LDAP Authentication Providers" in *Oracle Fusion Middleware Securing Oracle WebLogic Server* guide on the Oracle Technology Network http://download.oracle.com/docs/cd/E12840_01/wls/docs103/secmanage/atn.html#wp1199087.

Authorization Tuning Properties

The following Java system properties can be used to optimize authorization:

Table 7-1 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 a JpsSubject, which contains the user/enterprise-role information, as well as the 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 the Platform subject and the JpsSubject. Two subjects could be confusing when 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 confusing 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 ready-to-use setting. • 5: Instead of using the whole subject as the key, this setting uses a subset of the principal set inside the subject as the key (actually use principals of the <code>WLSUserImpl</code> type). <p>This setting accelerates 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</p>

Table 7-1 (Cont.) Authorization Properties

Java System Properties	Default Value	Valid Values	Notes
-Djps.subject.cache.ttl	60000ms		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.
-Djps.subject.cache.ttl	60000ms		Cache's Time To Live (TTL) for case 5 (above). This system property controls how long the cache is valid. When the time has expired, the cached value is dumped. The setting is controlled by the -Djps.subject.cache.ttl=xxxxflag, where xxx is the duration in milliseconds. 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.
-Djps.combiner.optimize=true	True	True False	This system property is used to cache the protection domains for a given subject. Setting the flag -Djps.combiner.optimize=true can improve the Java authorization performance.
-Djps.combiner.optimize.lazyeval=true	True	True False	This system property is used to evaluate a subject's protection domain when a checkPermission occurs. Setting the flag -Djps.combiner.optimize.lazyeval=true can improve the Java authorization performance.

Table 7-1 (Cont.) Authorization Properties

Java System Properties	Default Value	Valid Values	Notes
- Djps.policystore.hybrid .mode=true	True	True False	<p>This hybrid mode property is used to facilitate transition from SUN java.security.Policy 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>. When starting the Weblogic server, the Hybrid mode can be disabled by setting the system property <code>jps.policystore.hybrid.mode</code> to <code>false</code>. Setting <code>-Djps.policystore.hybrid.mode=false</code> can reduce the runtime overhead.</p>
-Djps.authz=ACC	ACC	ACC SM	<p>Delegates the call to JDK API <code>AccessController.checkPermission</code>, which can reduce the performance impact at runtime or while debugging.</p> <p>ACC: Delegate the call to <code>AccessController.checkPermission</code>.</p> <p>SM: If <code>SecurityManager</code> is set, delegate the call to <code>SecurityManager</code>.</p>

OPSS PDP Service Tuning Parameters

[Table 7-2](#) describes OPSS tuning parameters for policy store:

Table 7-2 OPSS PDP Service Tuning Parameters

Parameter	Default Value	Valid Values	Notes
oracle.security.jps.pol icystore.rolemember.cac he.type	STATIC	STATIC, SOFT, WEAK	<p>This parameter specifies the type of role member cache. Valid only in Jakarta 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 best performance.</p>
oracle.security.jps.pol icystore.rolemember.cac he.strategy	FIFO	FIFO NONE	<p>The type of strategy used in the role member cache. Valid only in Jakarta 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 small. <p>Consider maintaining the default value for best performance.</p>

Table 7-2 (Cont.) OPSS PDP Service Tuning Parameters

Parameter	Default Value	Valid Values	Notes
oracle.security.jps.pol icystore.rolemember.cac he.size	1000		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.
oracle.security.jps.pol icystore.policy.lazy.lo ad.enable	True	True False	Enables or disables the policy lazy loading. If this parameter is set to False , the server initial startup time takes longer - especially in a large policy store. For faster start-up time, the recommended value is True .
oracle.security.jps.pol icystore.policy.cache.s trategy	PERMISSION_FIFO	PERMISSION_FIFO NONE	The type of strategy used in the permission cache. Valid only in Jakarta 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 small. Consider using the default value for the best performance.
oracle.security.jps.pol icystore.policy.cache.s ize	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.pol icystore.cache.updatabl e	True	True False	This property is used to enable refresh. Consider maintaining the default value for the best performance.

Table 7-2 (Cont.) OPSS PDP Service Tuning Parameters

Parameter	Default Value	Valid Values	Notes
oracle.security.jps.pol icystore.refresh.enable	True	True False	This property is used to enable refresh. Consider maintaining the default value for performance.
oracle.security.jps.pol icystore.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.lda p.policystore.refresh.i nterval	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 Jakarta EE and J2SE applications.
oracle.security.jps.pol icystore.rolemember.cac he.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 are higher than the number of application roles set to True ; set to False otherwise, that is, when the number of application roles are high.

Oracle Web Services Security Tuning

Oracle Web Services Security provides a framework of authorization and authentication for interacting with a web service by using XML-based messages. There are several factors that may affect performance of the web service.

- [Choosing the Right Policy](#)
- [Policy Manager](#)
- [Configuring the Log Assertion to Record SOAP Messages](#)
- [Configuring Connection Pooling](#)
- [Monitoring the Performance of Web Services](#)

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 multistep transactions. Application-level security has more performance implications, but provides end-to-end security.

See Determining Which Predefined Policies to Use in *Securing Web Services and Managing Policies with Oracle Web Services Manager* to determine which security policies are required for a deployment.

Policy Manager

There is an inherent performance impact when you use 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.

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. It includes the SOAP header, the body, and all attachments, probably the 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.

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"
```

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 [Oracle Fusion Middleware Control](#).

For detailed information on using Oracle Fusion Middleware Control to monitor Oracle Web Services, see *Overview of Performance Monitoring, Auditing, and Tuning* in *Administering Web Services*.

Part III

Oracle Fusion Middleware Server Components

The Oracle Fusion Middleware server components need to be tuned for optimal performance.

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

- [Tuning Oracle Application Development Framework \(ADF\)](#)
You can tune Oracle Application Development Framework (ADF) to optimize its performance and scalability with design, configuration, and deployment considerations.
- [Tuning Oracle TopLink](#)
You can tune EclipseLink, an open-source persistence framework used with Oracle TopLink, to optimize its performance as the Java Persistence API (JPA) implementation.

8

Tuning Oracle Application Development Framework (ADF)

You can tune Oracle Application Development Framework (ADF) to optimize its performance and scalability with design, configuration, and deployment considerations.

Note:

- *Developing Fusion Web Applications with Oracle Application Development Framework*
- *Developing Web User Interfaces with Oracle ADF Faces*

- [About Oracle ADF](#)
Oracle Application Development Framework (Oracle ADF) is an end-to-end application framework that builds on Java Platform, Enterprise Edition (Jakarta EE) standards and open-source technologies to simplify and accelerate implementing service-oriented applications.
- [Basic Tuning Considerations](#)
To achieve optimal performance, you can follow tuning recommendations before you build, configure, and deploy ADF applications.
- [Advanced Tuning Considerations](#)
After you have performed the recommended tuning modifications, you can make additional changes that are specific to your ADF Server deployment. Consider carefully whether the advanced tuning recommendations are appropriate for your environment.

About Oracle ADF

Oracle Application Development Framework (Oracle ADF) is an end-to-end application framework that builds on Java Platform, Enterprise Edition (Jakarta 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 by using web, wireless, desktop, or web services interfaces. If you develop enterprise solutions that search, display, create, modify, and validate data by using web, wireless, desktop, or web services interfaces, Oracle ADF can simplify your job. Used in tandem, Oracle JDeveloper 11g and Oracle ADF give you an environment that covers the full development lifecycle from design to deployment, with drag-and-drop data binding, visual UI design, and team development features built-in.

For more information, see Introduction to Oracle ADF in *Developing Fusion Web Applications with Oracle Application Development Framework*.

Basic Tuning Considerations

To achieve optimal performance, you can follow tuning recommendations before you build, configure, and deploy ADF applications.

- [Oracle ADF Faces Configuration and Profiling](#)
- [Performance Considerations for ADF Faces](#)
- [Tuning ADF Faces Component Attributes](#)
- [Performance Considerations for Table and Tree Components](#)
- [Performance Considerations for autoSuggest](#)
- [Data Delivery - Lazy versus Immediate](#)
- [Performance Considerations for DVT Components](#)

Oracle ADF Faces Configuration and Profiling

Configuration options for Oracle ADF Faces are set in the `web.xml` file. Most of these options have default values that are tuned for performance. [Table 8-1](#) describes some of these configuration options.

Table 8-1 ADF Configuration Options

Parameter	Description
Compression View State <code>org.apache.myfaces.trinidad.COMPRESS_VIEW_STATE</code>	Controls whether the page state is compressed. If the size of the data is compressed, latency can be reduced. This parameter should be set to <code>True</code> .
Enhanced Debug <code>org.apache.myfaces.trinidad.resource.DEBUG</code>	Controls whether output should be enhanced for debugging. This parameter should be removed or set to <code>False</code> .
Check File Modification <code>oracle.adf.view.rich.CHECK_FILE_MODIFICATION</code>	Controls whether ADF faces check for modification date of JSP pages and discards 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 the type of saving (<code>all</code> or <code>token</code>) that should be used when client-side state saving is enabled. The default value is <code>token</code> .
Client-Side 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>False</code> .
Assertion Processing <code>oracle.adf.view.rich.ASSERT_ENABLED</code>	Specifies when to process assertions on the client-side. The default value is <code>OFF</code> . This parameter should be removed or set to <code>False</code> .

 **Note:**

When you are profiling or measuring client response by 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/>.

Performance Considerations for 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, is never cached in the 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 by using <code>af:resource</code> tag.</p> <p>TIP: Consider using <code>af:resource</code> rather than <code>trh:script</code> when possible.</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 <code>web.xml</code>:</p> <pre><servlet> <servlet-name> oraclejsp <init-param> <param-name> jsp_timeout </param-name> <param-value> x </param-value> </init-param> </servlet-name> </servlet></pre>

 **Note:**

Set parameter `x` based on your own use case scenarios.

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

Configuration Recommendation	Description
Create a single toolbar item with a drop-down popup.	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 input fields. Create a single toolbar item with a drop-down and put all the input fields inside it. This drop down should have deferred child creation and <code>contentDelivery="lazy"</code> .
Remove unknown rowCount.	A table that has an unknown <code>rowCount</code> can impact performance because getting the last set of rows takes excessive scrolling from the user and the application can appear to be very slow. Remove unknown <code>rowCount</code> by setting <code>DeferEstimatedRowCountProperty="false"</code> on the view object (VO).
Disable pop-ups that cannot be displayed by the user.	The fn:attachment component, when stamped in a table, can generate an excessive amount of DOM and client component. The amount of DOM + Client component is ~8K per cell, which impacts the performance of the entire page especially on slower browsers. Most cells have no attachments initially and only one popup can be displayed by the user. Therefore, pop-ups that cannot be displayed by the user should have <code>renderer="false"</code> . This cuts down the unnecessary DOM or client components sent to the browser. Similarly, the DOM has a <code>panelGroupLayout</code> with a number of cells that are empty. There is no need to send DOM for empty cells.
Do not use hover pop-ups on navigation links.	A hover popup on a navigation link causes the navigation to wait for the hover to be fetched first. Consider removing the hover popup on the compensate workforce table navigation link column and, instead, place it on a separate column or on an icon inside the cell.
Increase table scrolling timeout.	Tables send a fetch request to the server on a scroll after a timeout. The timeout, before the fetch is sent to the server, is typically only 20ms if the user scrolls a short distance, but can increase to 200ms if the user scrolls further. Therefore, performance can be impacted when the user scrolls to the bottom of a page and the table sends multiple requests to the server. To prevent the performance impact, consider increasing the timeout limit to 300ms.

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

Configuration Recommendation	Description
Use a timeout to call <code>_prepareForIncompleteImages</code> .	<p>During Partial Page Rendering (PPR) some images may not load completely. When this occurs, the parent component must be notified that the size of one of its descendants has changed. In the past this was done by using the <code>complete</code> attribute on the image tag. Now with Internet Explorer 8 the <code>complete</code> attribute is always false to alleviate performance issues with Internet Explorer 7 and 8. The attribute shows as false event for cached images immediately after the PPR content is fetched.</p> <p>For Internet Explorer 8 use a timeout (10ms) to call <code>_prepareForIncompleteImages</code> so that the image tag called right after the <code>.xml</code> request is processed. Note that this is not an issue for Mozilla Firefox or Google Chrome.</p>
Cache the <code>GetFirstVisibleRowKeyandRow</code> .	<p>Performance can be improved by locally caching the first visible Rowkey and row. This cached value can be deleted on a scroll or a resize.</p>
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 by 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 reinterprets and reexecutes Javascript libraries, and does not spend time for cleanup or 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 <code>metaContainerfacet</code> (the <code>HEAD</code> 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.

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

Configuration Recommendation	Description
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 on 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. For example, adding round corners to the template, can impact the performance for every page.• When building complex templates, sometimes it is easier to build them in multiple pieces and include them in the top-level template by using the <code><f:subview></code> tag. However, from a performance perspective, this is not typically recommended since it can impact memory usage on the server side. The <code><f:subview></code> tag 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 the <code><f:subview></code> tag unless it is required. If you can ensure that all IDs are unique, it is not necessary to use the <code><f:subview></code> around <code><jsp:include></code>. 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 the <code><f:subview></code> tag. If you are including content developed by someone else, use the <code><f:subview></code> tag if you do not know which IDs the developer used. In addition, you do not have to put the <code><f:subview></code> tag 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.

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

Configuration Recommendation	Description
Enable ADF rich client 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> <p>PanelAccordion PanelStretchLayout PanelTabbed BreadCrumbs NavigationPane PanelSplitter Toolbar Toolbox Table Train</p> <p>Notes:</p> <ul style="list-style-type: none"> • When you use geometry management, try to minimize 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.
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. The Toolbar component should be contained in a Toolbox to handle the overflow.</p> <p>While there were several optimizations done 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. Sometimes it is a good practice to set a big enough initial size of the overflow component such that overflow does not happen in most cases.</p>

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

Configuration Recommendation	Description
Use ADF Rich Client Partial Page Rendering (PPR).	<p>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>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>Two main Ajax patterns are implemented with partial page rendering (PPR):</p> <ul style="list-style-type: none"> • native component refresh • cross-component refresh <p>While the framework builds in native component refresh, cross-component refresh has to be done by developers in certain cases.</p> <p>Cross-component refresh is implemented declaratively or programmatically by the application developer defining the components that are to trigger a partial update and the other components that are to act as partial listeners, and so be updated. Using cross-component refresh and implementing it correctly is one of the best ways to improve client-side response time. While designing the UI page always think about what should happen when the user clicks a command button. Is it needed for the whole page to be refreshed or only the output text field? What should happen if the value in some field is updated? For more information, see <i>Developing Fusion Web Applications with Oracle Application Development Framework</i>.</p> <p>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, user 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 style="list-style-type: none"> • 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. • Define the components that 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

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

Configuration Recommendation	Description
Use ADF rich client navigation.	<p>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> <p>The steps above achieve PPR by using a command button to trigger the partial page refresh.</p> <p>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>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>
Cache resources.	<p>ADF Rich Client has an extensive support for navigation. One of the common use cases is tabbed navigation. This 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. It 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> <p>Developers are strongly encouraged to ensure that any resources that can be cached (images, CSS, and 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 that all the resources are present on the server.</p> <p>Consider using the <code>ResourceServlet</code> to configure the <code>web.xml</code> file to enable resource caching:</p> <pre data-bbox="818 1478 1360 1703"> <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>

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

Configuration Recommendation	Description
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 is forgotten for the least recently viewed pages, which can impact users that actively use the Back button or have multiple windows that are open simultaneously.</p> <p>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>

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

Configuration Recommendation	Description
Define custom styles at the top of the page.	<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, define styles at the top of the page and possibly wrap them inside the ADF group tag.</p> <p>An HTML page basically has two parts, the head and the body. When you add an <code>af:document</code> component to a page, this component creates both parts of the page for you. Any child component of the <code>af:document</code> is in the body part of the page. To get a component (or static CDATA content) to show up in the head, use the <code>metaContainer</code> facet.</p> <p>To get a component (or static CDATA content) to display in the head, use the <code>metaContainer</code> facet as follows:</p> <pre data-bbox="818 789 1321 1499"> <af:document title="#{attrs.documentTitle}" theme="dark"> <f:facet name="metaContainer"> <af:group><![CDATA[<style type="text/css"> .TabletNavigationGlobal { text-align: right; padding-left: 0px; padding-right: 10px; white-space: nowrap; } HTML[dir=rtl] .TabletNavigationGlobal { text-align: left; padding-left: 10px; padding-right: 0px; } </style>]]> <af:facetRef facetName="metaContainer"/> </af:group> </f:facet> <af:form ...> <af:facetRef facetName="body"/> </af:form> </af:document> </pre> <p>If you use page templates, consider including <code>af:document</code> and <code>af:form</code> in the template definition and expose anything that you may want to customize in those tags through the page template attributes and page template <code>af:facetRef</code>. Your templates are then able to utilize the <code>metaContainer</code> facet if they have template-specific styling as shown above. Also, your usage pages do not have to repeat the same document and form tags on every page.</p> <p>See <i>Developing Fusion Web Applications with 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 a 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 you customize the Javascript code the 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 <code>debug-output</code> 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 <code>debug-output</code> 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 pop-up 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 the <code>trinidad-config.xml</code> file.</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 to fetch the profiler data. Disable the profiler in the <code>web.xml</code> file to avoid this extra round-trip.</p>

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

Configuration Recommendation	Description
Disable resource debug mode.	<p>When resource debug mode is enabled, the response headers do not tell the browser 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 the jsp or the 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 the CSS file modification checks are made. To aid in performance, this configuration option defaults to <code>false</code>-does 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>
Enable content compression.	<p>By default, style classes that are rendered are compressed to reduce the page size. In production environments, make sure that 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>

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

Configuration Recommendation	Description
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 builds. 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> <p>There are two ways to check that the code is obfuscated by using Firefox with Firebug enabled:</p> <p>Check the download size:</p> <ol style="list-style-type: none"> 1. Ensure that <code>All</code> or <code>JS</code> is selected on the <code>Net</code> tab. 2. Locate the <code>all-11-version.js</code> entry. 3. Check the size of the column. It should be about 1.3 MB (as opposed to 2.8 MB). <p>Check the source:</p> <ol style="list-style-type: none"> 1. From the <code>Script</code> tab select <code>all-11-version.js</code> from the drop-down menu located above the tabs. 2. Examine the code. If there are comments and long variable names, the library is not obfuscated. <p>Note: Copyright comments are kept even in the obfuscated version of the JS files.</p>
Enable library partitioning.	<p>In the Oracle 11g release, library partitioning is ON by default. In the 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>

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>immediate</code> attribute.	<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 that are 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">• To avoid processing the data from the current screen while navigating to the new page set the <code>immediate</code> attribute to <code>TRUE</code> in the <code>commandNavigationItem</code> in the <code>navigationPane</code>.• 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 Using the JSF Lifecycle with ADF Faces in <i>Developing Web User Interfaces with Oracle ADF Faces</i>.</p> <p>There are some important issues that are associated with the <code>immediate</code> attribute. Refer to Using the Immediate Attribute in <i>Developing Web User Interfaces with Oracle ADF Faces</i> for more information.</p> <p>Note that this is an advanced feature. Most of the performance improvements can be achieved by using the <code>af:subform</code> component.</p>
Use the <code>visible</code> and <code>rendered</code> 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 whether the component is to be displayed on the page or 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 that 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>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 style="list-style-type: none"> 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. By default, the events generated on the client-side 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.
Use the <code>id</code> attribute.	<p>The <code>id</code> attribute should not be longer than 7 characters in length. This is particularly important for naming containers. A long <code>id</code> can impact performance as the amount of HTML that must be sent down to the client is impacted by the length of the <code>ids</code>.</p>
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 <i>What Happens When You Set clientComponent to true in Developing Web User Interfaces with Oracle ADF Faces</i>.</p>
Set 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> postpones creating any child components of the popup and you cannot refer to them until after the popup is shown.</p>

Performance Considerations for 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> .	When using <code>editingMode="editAll"</code> all content of the editable values holders and their client components is sent. This can significantly increase the payload and the Document Object Model (DOM) content on the client. Consider switching to <code>editingMode="clickToEdit"</code> to reduce the amount of transmitted data and potentially improve user interaction.
Reduce <code>fetchSize</code> when possible.	A larger <code>fetch size</code> 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.
Modify table fetch size.	Tables have a <code>fetch size</code> , 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. In addition, consider keeping the table <code>fetch size</code> and iterator <code>range size</code> in sync. By default, the table <code>fetch size</code> is set to the EL expression <code>#{bindings.<name>.rangeSize}</code> and should be equal to the iterator size. For more information, see <i>Using Tables, Trees, and Other Collection-Based Components</i> in <i>Developing Web User Interfaces with Oracle ADF Faces</i> .
Disable column stretching.	Columns in the <code>table</code> 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 you enable this feature with complex tables.
Consider using header rows and frozen columns only when necessary.	The <code>table</code> 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.

Table 8-4 (Cont.) Table and Tree Component Configurations

Configuration Recommendation	Description
Consider using <code>visitTree</code> instead of <code>invokeOnComponent</code> .	<p>A partial visit using <code>visitTree</code> is always at least as fast as <code>invokeOnComponent</code>. In addition, for components that control visiting, providing both <code>invokeOnComponent</code> and <code>visitTree</code> implementations is a source of errors. Consider deprecating <code>invokeOnComponent</code> and use <code>visitTree</code> instead.</p> <p>For more information, see Using Tables, Trees, and Other Collection-Based Components in <i>Developing Web User Interfaces with Oracle ADF Faces</i>.</p>

Performance Considerations for autoSuggest

The `autoSuggest` feature 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. 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.

Data Delivery - Lazy versus Immediate

Data for Table, Tree, and other stamped components can be delivered immediately or lazily. 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:

Delivery option	Description
Lazy Delivery (default)	<p>Lazy delivery should be used for tables, or other stamped components, which are known to have slow fetch time. For example, the 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 is shorter, and the user perceives better performance.</p> <p>Lazy delivery is implemented by using the 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>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 mechanism to deliver data in lazy mode, parallel execution of data controls is not. To enable parallel execution, open the page definition and change <code>RenderHint</code> on the iterator to <code>background</code>.</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 requests in parallel and possibly more database connections are 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 task flows). When there is a single data control frame, parallel execution may not improve performance, since parallel execution synchronizes on the data control frame level.</p>
Immediate Delivery	<p>Immediate delivery (<code>contentDelivery="immediate"</code>) should be used if the table data control is fast, or if it returns a small set of data. In these cases, the response time is 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>

Performance Considerations for DVT Components

DVT components are data visualization components 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.	The <code>RangeSize</code> attribute defines the number of rows to return simultaneously. A <code>RangeSize</code> value of <code>-1</code> causes the iterator to return all the rows. Using a lower value may improve performance, but it may be harder to stop the data and any data beyond <code>RangeSize</code> is not available in the view.
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. The problem is that vertical text is not native in 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>

Advanced Tuning Considerations

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

- [ADF Server Performance](#)

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.

Note:

When you use ADFm, consider using deferred execution and monitor the refresh conditions to maintain performance.

- [Tuning Session Timeout](#)
- [Tuning View Objects](#)
- [Enabling Batch Processing](#)

- [Tuning RangeSize](#)
- [Configuring Application Module Pooling](#)
- [Using ADFc Regions](#)
- [Deferring Task Flow Execution](#)
- [Deferring Task Flow Creation in Popups](#)
- [Configuring the Task Flow Inside Switcher](#)
- [Reusing Static Data](#)
- [Conditional Validations](#)

Tuning Session Timeout

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 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>
```

Tuning View Objects

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. *Developing Fusion Web Applications with Oracle Application Development Framework* provides detailed information on tuning View Objects. Provided here are some tips pertaining to View Object performance.

- [Creating View Objects](#)
- [Configuring View Object Data Fetching](#)
- [Setting Additional View Object Configurations](#)

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.

Table 8-6 Types of View Objects

View Object Type	Description
Read-only View Objects	<p>If the View Object does not have to insert or update data, consider using a read-only View Object. 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 reexecuting 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 that are 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 <code>updatable</code> 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 you use 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 <code>View Objects Overview</code> tab. This sets <code>MaxFetchSize</code> to 0 (zero) for the View Object definition.</p>
run time-created View Objects	<p>View Objects can be created at runtime by 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>

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

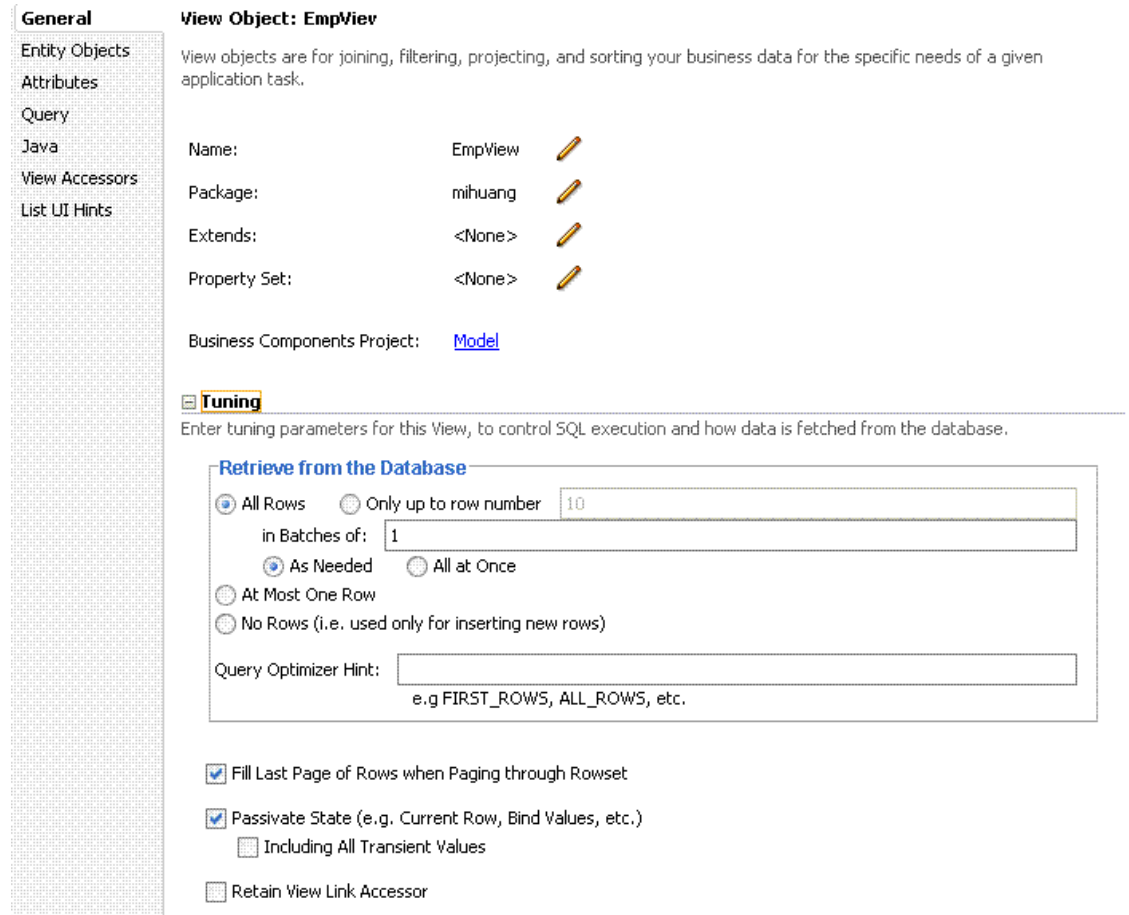


Table 8-7 View Object Configurations

Fetch Option	Description
Fetch Mode	The default fetch option is the All Rows option, which is retrieved as needed (<code>FetchMode="FETCH_AS_NEEDED"</code>) or all at once (<code>FetchMode="FETCH_ALL"</code>), depending on which option is appropriate. The As Needed option ensures that an <code>executeQuery()</code> 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.
Fetch Size	In conjunction with the fetch mode option, the Batches field controls the number of records fetched simultaneously from the database (<code>FetchSize</code> 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 to the $n+1$ value where n is the number of rows that are displayed in the user interface. Note that for DVT objects, Fetch Size should be $n+1$ where n is either <code>rangeSize</code> or the likely maximum rowset size if <code>rangeSize</code> is -1.

Table 8-7 (Cont.) View Object Configurations

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 <code>WHERE</code> clause expects to retrieve a single row, set the option <code>At Most One Row</code>. 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>

Setting Additional View Object Configurations

[Table 8-8](#) provides additional tuning considerations when you use the View Object:

Table 8-8 Additional View Object Configurations

Configuration Recommendation	Description
Optimize large data sets.	<p>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 you navigate 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).</p>

Table 8-8 (Cont.) Additional View Object Configurations

Configuration Recommendation	Description
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. When defining the database connection, the SQL style is set. By default, when you define an Oracle database connection, the SQL style is <code>Oracle</code> . To manually override the SQL style, pass the parameter - <code>Djbo.SQLBuilder="SQL92"</code> to the JVM at startup.
Use bind variables for view object queries.	If the query that is associated with the View Object contains values that may change from execution to execution, consider using bind variables. This may help to avoid reparsing the query on the database. Bind variables can be added to the View Object in the Query section of the View Object definition.
Use query optimizer hints for view object queries.	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.
Use dynamic SQL generation.	View Objects can be configured to dynamically generate SQL statements at runtime instead of defining the SQL at design time. A View Object instance, that is configured with generating SQL statements dynamically, can avoid requerying 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.

Enabling 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:**

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"`).

Tuning RangeSize

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.

Configuring 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 the `bc4j.xcfgfile`, 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 8-9 Application Module (AM) Pool Tuning

Configuration Recommendation	Description
Optimize the number of AM pools in the application.	<p>Parameters that are applied at the system level are applied per AM pool. If the application uses more than 1 AM pool, then the 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 four 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 the <code>bc4j.xcfgfile</code>.</p>

Table 8-9 (Cont.) Application Module (AM) Pool Tuning

Configuration Recommendation	Description
Optimize the number of database connections.	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.

 **Note:**

If you have an AM pool that needs to be used as the 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.

- [General AM Pool Configurations](#)
- [Configuring Application Module Pool Sizing](#)
- [Configuring Application Module Pool Resource Cleanup](#)
- [Designing an Application Module](#)

General AM Pool Configurations

Use the following guidelines as a general starting point when tuning AM and AM pool behavior. More specific tuning for memory or CPU usage can be found in [Configuring Application Module Pool Sizing](#).

Table 8-10 AM Pool Tuning Parameters

Parameter	Description
Initial Pool Size <code>jbo.ampool.initpoolsize</code>	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 that require an AM instance. Configure this value to 10% more than the anticipated number of concurrent AM instances that are required to service all users.
Maximum Pool Size <code>jbo.ampool.maxpoolsize</code>	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 the specified limit. A general guideline is to configure this to 20% more than the initial pool size to allow for some additional growth.

Table 8-10 (Cont.) AM Pool Tuning Parameters

Parameter	Description
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>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, set to 0 (zero).</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 always get cleaned up at this time. Then, additional available instances are removed, if necessary, to bring the number of available instances down to this size.</p> <p>The default maximum available size is 25 instances.</p> <p>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 during 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 8-10 (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 a 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 minimum available size.</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 never exceeds 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>

Table 8-10 (Cont.) AM Pool Tuning Parameters


Parameter	Description
Failover jbo.dofailover	<p>Specifies whether to disable or enable failover. By default, failover is disabled. To enable failover, set the parameter to <code>true</code>.</p> <div style="border: 1px solid #0070C0; padding: 10px; margin-top: 10px;"> <p> Note:</p> <p>When you enable application module state passivation, a failure can occur when Oracle WebLogic Server is configured to forcibly release the 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> </div> <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. 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>
Locking Mode jbo.locking.mode	<p>Specifies the locking mode (<code>optimistic</code> or <code>pessimistic</code>). The default is <code>pessimistic</code>, which means that a pending transaction state can be created on the database with row-level locks. With <code>pessimistic</code> locking mode, each time an AM is recycled, a rollback is issued in the JDBC connection. Web applications should set the locking mode to <code>optimistic</code> to avoid creating the row-level locks.</p>
Database Connection Pooling jbo.doconnectionpooling	<p>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 <code>false</code>, which means that an AM instance can retain its database connection when the AM instance is returned to the AM pool. When set to <code>true</code>, 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.</p>

Table 8-10 (Cont.) AM Pool Tuning Parameters

Parameter	Description
Transaction Disconnect Level <code>jbo.txn.disconnect_level</code>	When used in conjunction with <code>jbo.doconnectionpooling=true</code> , it specifies BC4J behavior for maintaining JDBC ResultSets. By default, <code>jbo.txn.disconnect_level</code> is 0 and you can use passivation 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 8-11](#).

Table 8-11 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 while 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 8-12](#).

Table 8-12 AM Pool Sizing Configurations - CPU Considerations

Parameter	Description
<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.
<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 saves the CPU processing cost of needing to switch an AM instance from one user session to another.

Configuring Application Module Pool Sizing

The Application Module pool sizing configuration is largely dependant on the number of concurrent users you expect to have. To prevent performance issues, you need to make sure AM pool size is sufficient to serve all concurrent users.

Caution:

The following example assumes at least 100 concurrent users. Always consult your own use case scenarios to determine the appropriate settings for your deployment.

To configure these parameters, open the `setDomainEnv.sh` file for the WebLogic Server instance and find these lines:

```
JAVA_OPTIONS="{JAVA_OPTIONS}"  
export JAVA_OPTIONS
```

Replace these lines with the following:

```
JAVA_OPTIONS="-Djbo.ampool.doampooling=true  
-Djbo.ampool.minavailablesize=1  
-Djbo.ampool.maxavailablesize=120  
-Djbo.recyclethreshold=60  
-Djbo.ampool.timetolive=-1  
-Djbo.load.components.lazily=true  
-Djbo.doconnectionpooling=true  
-Djbo.txn.disconnect_level=1  
-Djbo.connectfailover=false  
-Djbo.max.cursors=5  
-Doracle.jdbc.implicitStatementCacheSize=5  
-Doracle.jdbc.maxCachedBufferSize=19 {JAVA_OPTIONS}"
```

Note:

To limit performance implications, set the `ampool.maxavailablesize` to a value that is at least 20% more than the maximum number of concurrent users you expect in your own use case scenarios.

Configuring Application Module Pool Resource Cleanup

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. However, reducing the number of available AM instances and increasing the frequency of cleanups can result in higher CPU usage and longer response times. See [Table 8-13](#) for more information.

Table 8-13 AM Pool Resource Cleanup Configurations - Memory Considerations

Parameter	Description
<code>jbo.ampool.minavailablesize</code>	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 recreating the AM pool.
<code>jbo.ampool.maxavailablesize</code>	A lower value generally results in more AM instances being removed from the pool during cleanup.
<code>jbo.ampool.timetolive</code>	A lower value reduces the time an AM instance can exist before it must be removed at the next resource cleanup.
<code>jbo.ampool.maxinactiveage</code>	A low value results in more AM instances being marked as a candidate for removal at the next resource cleanup.
<code>jbo.ampool.monitorsleepinterval</code>	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.

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.

Table 8-14 AM Pool Resource Cleanup Configurations - CPU Considerations

Parameter	Description
<code>jbo.ampool.minavailablesize</code> and <code>jbo.ampool.maxavailablesize</code>	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.
<code>jbo.ampool.timetolive</code>	A higher value increases the time an AM instance can exist before it must be removed at the next resource cleanup.
<code>jbo.ampool.maxinactiveage</code>	A higher value results in fewer AM instances being marked as a candidate for removal at the next resource cleanup.
<code>jbo.ampool.monitorsleepinterval</code>	Configuring a higher interval results in less frequent resource cleanups.

Designing an Application Module

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 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
```

Using ADFc Regions

Adding regions to a page can be a powerful addition to the application. While there is no limit to the number of remote regions that you can render in a JSF page, use this capability with caution. For simple pages, where tabs are not used, regions may be combined in the page such that the maximum number of regions is determined by the design of the region and the view object queries it executes. Alternatively, for complex pages that use tabs, limit the use of regions to achieve best performance. For complex tabbed pages, ADF does not deactivate task flow transactions once a region is loaded. When switching tabs, the ongoing transaction must be stopped to achieve best performance.

Deferring 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.

Note:

For regions and task flows, the amount of time it takes to evaluate the current viewId and the time it takes to calculate input parameters to the flow can impact your overall performance. Consider this during your design phase.

Deferring Task Flow Creation in Popups

By default, the child components under a pop-up are created even when pop-up is not accessed. To avoid this overhead, consider the following:

- Set `childCreation` to `deferred`
Set `childCreation="deferred"` on the popup
Set `activation="deferred"` on the task flow

 **Note:**

This approach cannot be used if any of the following tags are present inside the pop-up:

- `f:attribute`
- `af:setPropertyListener`
- `af:clientListener`
- `af:serverListener`

It also cannot be used if you need to refer to any child components of the popup before the popup is displayed. Setting `childCreation="deferred"` postpones creating any child components of the popup and you cannot refer to them until after the popup is shown. In that case, use Conditional Activation.

- Use conditional activation

Add property listener on the popup in the **jsff** to set a condition

Set `activation="conditional"` on the task flow

Set `activate=condition` on the task flow

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.

Reusing Static Data

If the application contains static data that can be reused across the application, the cache data can be collected by using a shared application module. For more information on creating and using shared application modules, see *Sharing Application Module View Instances in Developing Fusion Web Applications with Oracle Application Development Framework*.

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. For more information on specifying preconditions for validation, see *How to Set Preconditions for Validation in Developing Fusion Web Applications with Oracle Application Development Framework*.

9

Tuning Oracle TopLink

You can tune EclipseLink, an open-source persistence framework used with Oracle TopLink, to optimize its performance as the Java Persistence API (JPA) implementation.

- [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.
- [Basic Tuning Considerations](#)
To achieve optimal performance, you can follow the tuning recommendations that apply to your own use case scenarios.
- [Advanced Tuning Considerations](#)
After you have performed the recommended modifications, you can make additional changes that are specific to your deployment. Consider carefully whether the advanced tuning recommendations are appropriate for your environment.

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 information here assumes that you are familiar with the basic functionality of EclipseLink. Before you begin tuning, consider reviewing the following introductory information:

- [Understanding Queries](http://www.eclipse.org/eclipselink/documentation/2.6/concepts/queries.htm#CHDGCCJB) at <http://www.eclipse.org/eclipselink/documentation/2.6/concepts/queries.htm#CHDGCCJB>
- [Understanding Caching](http://www.eclipse.org/eclipselink/documentation/2.6/concepts/general004.htm#CHDEEBFG) at <http://www.eclipse.org/eclipselink/documentation/2.6/concepts/general004.htm#CHDEEBFG>
- [Understanding Mappings](http://www.eclipse.org/eclipselink/documentation/2.6/concepts/mappingintro.htm#CHDFEJIJ) at <http://www.eclipse.org/eclipselink/documentation/2.6/concepts/mappingintro.htm#CHDFEJIJ>

For more information on Oracle TopLink, see the [TopLink](#) page on the Oracle Technology Network (OTN).

Note:

The information here serves as a Quick Start guide to performance tuning JPA in the context of a Jakarta EE environment. While this information provides common performance tuning considerations and related documentation resources, it is not meant to be a comprehensive list of areas to tune.

Basic Tuning Considerations

To achieve optimal performance, you can follow the tuning recommendations that apply to your own use case scenarios.

- [SQL Statement and Query Tuning Parameters](#)
- [Cache Configuration Tuning Parameters](#)
- [About Mapping and Descriptor Configurations](#)
- [About Data Partitioning](#)

SQL Statement and Query Tuning Parameters

[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>By 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. Use the persistence property <code>eclipselink.jdbc.bind-parameters</code> in the <code>persistence.xml</code> file to configure this.</p> <p>See also "Understanding Caching" at http://www.eclipse.org/eclipselink/documentation/2.6/concepts/cache.htm#CDEFHHEH and "Understanding Querying" at http://www.eclipse.org/eclipselink/documentation/2.6/concepts/queries.htm#CHDGGCJB</p> <p>Default Value: <code>PERSISTENCE_UNIT_DEFAULT</code>, which is true by default.</p>	Leave parameterized SQL binding enabled for selected databases and JDBC drivers that support these options.

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

Tuning Parameter	Description	Performance Notes
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 by using a database.</p> <p>Note: For Jakarta 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 <code>Statement Cached Type</code> and <code>Statement Cached Size</code> configuration options.</p> <p>See also <i>Increasing Performance with the Statement Cache</i> in <i>Administering 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>If your JDBC driver supports this option, you should always enable statement caching. The Oracle JDBC driver supports this option.</p>
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 <code>eclipselink.jdbc.fetch-size</code>.</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>

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

Tuning Parameter	Description	Performance Notes
Batch Writing	<p>Batch writing can improve the database performance by sending groups of INSERT, UPDATE, and DELETE statements to the database in a single transaction, rather than individually.</p> <p>Use the persistence property "eclipselink.jdbc.batch-writing"="JDBC" in the persistence.xml file to configure this.</p> <p>Default Value: Off.</p>	Enable for the persistence unit.
Change Tracking	<p>This is an optimization feature that lets you tune the way EclipseLink detects changes in an entity.</p> <p>Default Value: If using weaving (Jakarta EE default) AttributeLevel otherwise Deferred.</p>	Leave at default AttributeLevel for best performance.
Weaving	<p>Can disable through persistence.xml properties eclipselink.weaving</p> <p>Default Value: On.</p>	Leave On for best performance.
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 by using the @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 preallocation enables a batch of ids to be queried from the database simultaneously to avoid accessing the database for an id on every insert.</p> <p>Default Value: 50.</p>	Always use sequence number preallocation for best performance for inserts. SEQUENCE or TABLE sequencing should be used for optimal performance, not IDENTITY, which does not allow preallocation.

- [Entity Relationships Query Tuning Parameters](#)

Entity Relationships Query Tuning Parameters

[Table 9-2](#) shows the entity relationship between the query parameters for performance tuning.

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

Tuning Parameter	Description	Performance Notes
Batch Fetching	<p>The <code>eclipselink.batch</code> 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: <code>JOIN</code>, <code>EXISTS</code>, and <code>IN</code>. The type is set through the query hint <code>eclipselink.batch.type</code>.</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 <code>eclipselink.batch</code>.</p> <p>Batch fetching can also be set by using the <code>@BatchFetch</code> annotation.</p> <p>Default Value: <code>Off</code>.</p>	<p>Use it to query the tables with columns mapping to the table data you need. You should only use either batch fetching or joining if you know that you are going to access all the data; if you do not intend to access the relationships, then let the 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; if you are using lazy-loading, each relationship is fetched separately when accessed 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 as <code>multilevel</code> in the query hint <code>eclipselink.join-fetch</code>. It also can be set in the mapping annotation <code>@JoinFetch</code>.</p> <p>Joining is part of the JPA specification, whereas batch fetching is not. And, joining works on queries that do 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 or max results, whereas batch fetching does not.</p> <p>See also "Join Fetching" at http://www.eclipse.org/eclipselink/documentation/2.6/solutions/performance001.htm#CHDEGCHH</p> <p>Default Value: <code>Not Used</code>.</p>	<p>Use it to query the tables with columns mapping to the table data you need. You should only use either batch fetching or joining if you know that you are going to access all the data; if you do not intend to access the relationships, then let the 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 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 "Using Lazy Loading" at http://www.eclipse.org/eclipselink/documentation/2.6/concepts/mappingintro001.htm#CEGBCJAG.</p> <p>Default Value: On for collection mapping (ToMany mappings, @OneToMany and @ManyToMany)</p> <p>Default Value: Off for reference (ToOne mappings, @OneToOne and @ManyToOne)</p>	<p>Use lazy loading for all mappings. Using lazy loading and querying the referenced objects by using batch fetching or Join is more efficient than Eager loading.</p> <p>You may also consider using optimized loading with <code>LoadGroups</code>, which allows a query to force instantiation of relationships.</p>

 **Note:**

Setting lazy loading to On for @OneToOne and @ManyToOne requires weaving, which is set to On by default for Jakarta EE.

Cache Configuration Tuning Parameters

You can tune the default internal cache that is provided by EclipseLink. Oracle Toplink or EclipseLink can also be integrated with Oracle Coherence. For information on configuring and tuning an EclipseLink Entity Cache by using Oracle Coherence, see .

The default settings for EJB3.0/JPA that is 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 "Understanding Caching" at <http://www.eclipse.org/eclipselink/documentation/2.6/concepts/cache.htm#CDEFHHEH>.

 **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, or applications that are 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 that are 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 in the <code>persistence.xml</code> file 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>

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

Tuning Parameter	Description	Performance Notes
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 eclipselink.query-cache.size eclipselink.query-cache.invalidation</pre> <p>See also "Understanding Caching" at http://www.eclipse.org/eclipselink/documentation/2.6/concepts/cache.htm#CDEFHHEH and "JPA Query Customization Extensions" at http://www.eclipse.org/eclipselink/documentation/2.6/jpa/extensions/queryhints.htm#sthref498.</p> <p>Default Value: Not Used.</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.

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 the following persistence properties: <code>eclipselink.cache.size.entity</code> <code>eclipselink.cache.size.default</code> <code>eclipselink.cache.type.default</code></p> <p>See also "About the Persistence Unit" at http://www.eclipse.org/eclipselink/documentation/2.6/concepts/appdeployment002.htm#BABHCJDG and "Class PersistenceUnitProperties" at http://www.eclipse.org/eclipselink/api/2.6/org/eclipse/persistence/config/PersistenceUnitProperties.html.</p> <p>Default Value: Type <code>SoftWeak</code>, Size 100 (per Entity). The default value may be different if Toplink is running on Exalobic. See Enable the Exalobic Automated Tuner in the <i>Solutions Guide for Oracle TopLink</i> for more information about the Exalobic default.</p>	<p>Based on your tolerance for stale data, 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.</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 <code>@Version</code> annotation, <code>eclipselink.read-only</code></p> <p>See "Descriptors and Locking" at http://www.eclipse.org/eclipselink/documentation/2.6/concepts/descriptors002.htm#CHEEEIEA.</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 <code>eclipselink.cache-usage</code> 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 by 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, which are 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 <p><code>setTransactionIsolation</code></p> <p>See also Isolated Cache at http://www.eclipse.org/eclipselink/documentation/2.6/concepts/cache001.htm#CDEEGICF.</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>See About Cache Refreshing .</p> <p>See also Understanding Caching at http://www.eclipse.org/eclipselink/documentation/2.6/concepts/cache.htm#CDEFHHEH.</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

- [About Cache Refreshing](#)
- [Locking Mode Policy Options](#)

About Cache Refreshing

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 refreshes stale objects in the cache on locking errors.

Locking Mode Policy Options

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 "Descriptors and Locking" at <http://www.eclipse.org/eclipselink/documentation/2.6/concepts/descriptors002.htm#CHEEEIEA>.

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 that the data has not changed since the user read the data. See also "Using Optimistic Locking" at http://www.eclipse.org/eclipselink/documentation/2.6/concepts/mappingintro005.htm#CEGDIIIB .	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 (<code>Shared=False</code>) for best performance.
Read Only	Setting an EJB3.0 JPA entity to <code>Read Only</code> ensures that the entity cannot be modified and enables EclipseLink to optimize the unit of work performance. Set at the entity level by using <code>@ReadOnly</code> class annotation. Can also be set at the query level through the query hint <code>eclipselink.read-only</code> .	Defining an entity as <code>Read Only</code> can perform better than an entity that is not defined as <code>Read Only</code> , yet does no inserts, updates, or deletes, since it enables EclipseLink to optimize the unit of work performance. Always use <code>Read Only</code> for all read-only operations.

About Mapping and Descriptor Configurations

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 “Mapping and Descriptors” at <http://www.eclipse.org/eclipselink/documentation/2.6/solutions/performance002.htm#sthref153>.

About Data Partitioning

EclipseLink allows you to configure data partitioning by 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 “Partitioning Annotations” at http://www.eclipse.org/eclipselink/documentation/2.6/jpa/extensions/annotations_ref.htm#CACHIIB.

Advanced Tuning Considerations

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

- [Integrating with Oracle Coherence](#)
- [Analyzing EclipseLink JPA Entity Performance](#)

Integrating with Oracle 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 multithreaded and Jakarta EE server hosted applications that are 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 using EclipseLink JPA with a Coherence Cache, see Grid Cache Configuration in *Integrating Oracle Coherence*.

For details on Oracle Toplink integration with Oracle Coherence, see Integrating Toplink Grid with Oracle Coherence in *Integrating Oracle Coherence*.

Analyzing EclipseLink JPA Entity Performance

The following features in EclipseLink can help you analyze your JPA application performance:

- For form monitoring performance, see "Performance Monitoring" at <http://www.eclipse.org/eclipselink/documentation/2.6/concepts/monitoring003.htm#BABJABIH>. Note that this tool is intended to profile and monitor information in a multithreaded server environment.
- For profiling performance, see "Task 1: Measure EclipseLink Performance with the EclipseLink Profiler" at <http://www.eclipse.org/eclipselink/documentation/2.6/solutions/performance002.htm#CHDIAFJI>. 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, remember to restore the logging levels to the default levels when you are done profiling or debugging.

Part IV

Oracle Identity and Access Management

This part describes tuning the Oracle Identity and Access 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. The Access Management products enable you 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 sign into and access online resources.

It contains the following chapters:

- [Oracle Internet Directory Performance Tuning](#)
This chapter provides guidelines for tuning and sizing an Oracle Internet Directory installation.
- [Oracle Access Management Performance Tuning](#)
- [Oracle Identity Governance Performance Tuning](#)
This chapter provides guidelines for tuning and sizing specific to Oracle Identity Governance (OIG).

10

Oracle Internet Directory Performance Tuning

This chapter provides guidelines for tuning and sizing an Oracle Internet Directory installation.

It contains these topics:

- [About Oracle Internet Directory](#)
- [Monitoring Oracle Internet Directory Performance](#)
- [Basic Tuning Considerations](#)
- [Advanced Tuning Considerations](#)
- [Specific Use Cases That Require Additional Tuning](#)

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, multiprocess, 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 ready-to-use 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.

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 [Monitoring](#).

- [Monitoring Performance on UNIX and Windows Systems](#)
- [Updating Database Statistics by Using oidstats.sql](#)
- [Setting Performance-Related Replication Configuration Attributes](#)

- [Managing System Configuration Attributes](#)
- [Setting Garbage Collection Configuration Attributes](#)

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 <code>top</code> 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`
- The ANALYZE function in the DBMS_STATS package

See Also:

- *Database Reference* in the Oracle Database Documentation Library for information about `utlbstat.sql` and `utlestat.sql`
- *Database Performance Tuning Guide* for information about stats package
- *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 [Updating Database Statistics by Using oidstats.sql](#).

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 *Reference for Oracle Identity Management*

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 *Administering Oracle Internet Directory*
- "Configuring Attributes of the Replication Configuration Set by Using `ldapmodify`" in *Administering 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 *Administering Oracle Internet Directory*
- "Configuring Replication Agreement Attributes by Using Idapmodify" in *Administering 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 *Administering Oracle Internet Directory* or information on setting replication attributes by using the Replication Wizard.
- "Configuring Attributes of the Replication Configuration Set by Using Idapmodify" in *Administering Oracle Internet Directory*.

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 *Administering 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"

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.

- [Modifying Changelog Purging Attributes by Using Idapmodify](#)
- [Modifying Changelog Purging in Oracle Directory Services Manager](#)

Modifying Changelog Purging Attributes by Using Idapmodify

The following example is an LDIF file used to configure change log purging.

 **See Also:**

"Change Log Purging" in *Administering 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 *Administering Oracle Internet Directory*.

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**.

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.

- [Database Parameters](#)
- [LDAP Server Attributes](#)
- [Database Statistics](#)
- [Low-Priority Tuning Considerations](#)

Database Parameters

The suggested minimum values for Oracle Database instance parameters are described in [Table 10-1](#):

Table 10-1 Minimum Values for Oracle Database Instance Parameters

Parameter	Value	Notes
<code>sga_target</code> and <code>sga_max_size</code>	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.
<code>db_cache_size</code>	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.
<code>shared_pool_size</code>	300M	Applicable when SGA Auto Tuning using <code>sga_target</code> and <code>sga_maxsize</code> is not being used
<code>session_cached_cursors</code>	100	
<code>processes</code>	500	
<code>pga_aggregate_target</code>	300M	Before performing a large <code>bulkload</code> operation, set this to 1-4GB, if sufficient RAM is available. Set it back after the operation has completed
<code>job_queue_processes</code>	1 or more.	Tune this parameter only if you are using -based multimaster replication
<code>max_commit_propagation_delay</code>	99 or lower	Tune this parameter only in Oracle RAC Database deployments, RDBMS v10.1.

See the *Oracle Database Performance Tuning Guide* for information on setting Oracle Database instance parameters.

LDAP Server Attributes

The recommendations in this section are summarized in [Table 10-2](#).

- 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 settings for `orclmaxcc` and `orclserverprocs` in [Table 10-2](#).
- 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 *Administering 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 10-2 LDAP Server Attributes to Tune

Attribute	Default	Recommended Value	Notes
<code>orclmaxcc</code>	2	10	Server restart required.
<code>orclserverprocs</code>	1	Number of CPU cores on the system	
<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 *Administering Oracle Internet Directory*.

For information about configuring `orclskiprefinsql` or `orclmatchdnenabled` with Oracle Enterprise Manager Fusion Middleware Control, see "Configuring Shared Properties" in the *Administering 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 *Administering Oracle Internet Directory*.

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 [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 the 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.

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.

Low-Priority Tuning Considerations

This section describes attributes that can sometimes improve performance, but are considered low-priority.

- [Number of Entries to be Returned by a Search](#)
- [Enabling the Group Cache](#)
- [Timeout for Write Operations](#)

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.

See "Setting System Configuration Attributes by Using `ldapmodify`" in the *Administering Oracle Internet Directory*.

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 *Administering Oracle Internet Directory*.

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 *Administering Oracle Internet Directory*.

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 Oracle RAC Database](#)
- [Password Policies and Verifier Profiles](#)
- [Server Entry Cache](#)
- [Result Set Cache](#)
- [Tuning Security Event Tracking](#)
- [Optimizing Searches](#)

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 *Administering 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 *Administering 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 *Administering Oracle Internet Directory*. and "Setting System Configuration Attributes by Using `ldapmodify`" in the *Administering 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.

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 10-3](#).

- 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 10-3](#) summarizes these recommendations.

Table 10-3 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 [Setting Performance-Related Replication Configuration Attributes](#) for information on setting these replication attributes.

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 10-4](#).

Table 10-4 Garbage Collection Configuration Parameters

Parameter	Value	Notes
<code>orclpurgetargetage</code>	Less than 10days (240 hours)	Only if there is no requirement to retain change logs
<code>orclpurgeinterval</code>	6–12 hours	

You can modify these attributes by using `ldapmodify` or Oracle Directory Services Manager. See [Setting Garbage Collection Configuration Attributes](#).

Oracle Internet Directory with Oracle RAC Database

As described in [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.

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.
- 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

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 *Administering Oracle Internet Directory*.
- The "Managing Password Verifiers" chapter in *Administering Oracle Internet Directory*.

Server Entry Cache

The Oracle Internet Directory server entry cache enables LDAP entries to be cached on the Oracle Internet Directory server process heap for better performance. Configuring the entry cache provides benefits if, and only if, all or most entries can be cached.

Note:

The server entry cache is beneficial for small directory deployments only. Some of the tuning recommendations here contradict the tuning recommendations in the earlier sections. Review the applicability of entry cache to a given deployment and incorporate the tuning mentioned in this section only if all considerations enumerated here are met.

- [Benefits of Using the Entry Cache](#)
- [Values for Configuring the Entry Cache](#)

Benefits of Using the Entry Cache

One of the key benefits of using the entry cache is that the LDAP search operations with base scope are about five times as fast. This applies only when all or most entries can be cached. A cache miss is more expensive than disabling the entry cache.

Values for Configuring the Entry Cache

You can configure and optimize the server entry cache by setting the values shown in [Table 10-5](#).

Table 10-5 Server Entry Cache Configuration

Attribute	Default	Recommended Value	Notes
<code>orclmaxcc</code>	2	10	Restart the server after changing this attribute.
<code>orclserverprocs</code>	1	Total number of cores on the system.	
<code>orclecacheenabled</code>	2	2	

Table 10-5 (Cont.) Server Entry Cache Configuration

Attribute	Default	Recommended Value	Notes
orcldcache_maxsize	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.
orcldcache_maxentries	100000	Total number of entries in the DIT	
orcldcache_maxentrysize	1000000	Size, in bytes, of the largest entry in the DIT	The largest entry is usually a group entry or an entry with binary attribute values.

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 `orcldcache_enabled` to 1, `orcldcache_maxsize` to 1,500,000,000, and `orcldcache_maxentries` to 300,000. If the size of the largest group entry or entry with binary value is 10M, you would set `orcldcache_maxentrysize` to 10,000,000.

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                : 1000000512
Max Entries configured        : 1000000
Max Entries cached            : 100000
Num Entries in Cache          : 100000
Num Entries in GC             : 0
Page size                     : 976556
Entry cache Hit count         : 6172127
Entry cache Mis count         : 99999
Hash Area bytes used          : 24497696
Hash Area blocks used         : 37
ResultSet cache bytes used    : 6799604
Resultset cache blocks used   : 300000
Entry cache bytes used        : 404047820
Entry cache blocks used       : 5900293
Cache memory used             : 435345120

```

To configure the attributes, see "Attributes of the Instance-Specific Configuration Entry" in the *Administering Oracle Internet Directory* and "Setting System Configuration Attributes by Using ldapmodify" in the *Administering Oracle Internet Directory*.

Result Set Cache

Result set cache allows complete result sets to be stored in memory. If an 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 (for example, latches) that might normally be incurred. Configuring the result cache can improve performance since most LDAP applications typically look up user entries such as `mail=john.doe@example.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 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.

- [When to Use Result Set Cache](#)
- [Benefits of Using Result Set Cache](#)
- [Configuring Result Set Cache](#)
- [Values for Configuring Result Set Cache](#)

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)

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.

Note that the result set cache database parameter can be configured on the client side or server side. When the server side cache is enabled, the result set cache can consume a significant amount of database memory and OID performance may be impacted.

- Performance improved by 3 to 5 times when compared to performance when result set cache is not used.

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`, and `orclguid`. Typically these attributes are not modified for the life of the entry.

To enable result set cache, set `orclecacheenabled=2`. Result set cache can be turned off by setting `orclecacheenabled=1` or `orclecacheenabled=0`.

Any change to these configuration attributes requires a restart of OID server (all the instances).

Values for Configuring Result Set Cache

Note that any change to the following configuration attributes requires a restart of OID server (all the instances).

Table 10-6 Result Set Cache Attributes to Tune

Attribute	Default	Recommended Value	Notes
<code>OrclRSCacheAttr</code>	<code>cn, mail, uid, orclguid</code>		Multi valued attribute, Value contains the name of the Attribute. Typically these attributes are not modified for the life of the entry.
<code>ResultSetMaxEntries</code>	4		Maximum number of entries for a given search that can be cached.
<code>ResultSetMaxCacheSize</code>	10 MB		Maximum memory that can be allocated in the shared memory for the result set cache.
<code>ResultSetMaxTime</code>	8 hours		Time to live for the result set cache when the cache is full.

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 *Administering 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.

Optimizing Searches

This section contains these topics:

- [Optimizing Searches for Large Group Entries](#)
- [Optimizing Searches for Skewed Attributes](#)
- [Optimizing Performance of Complex Search Filters](#)

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.

- [Entry Cache Enabled Configuration](#)
- [Entry Cache Disabled Configuration.](#)

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. To do this, increase the value of the `orclEcacheMaxEntSize` attribute in the instance-specific configuration entry:

```
cn=componentname,cn=osdldapd,cn=subconfigsentry
```

This attribute controls the maximum size of a cache entry.



Note:

If you expect frequent updates to large groups, then do not use this tuning methodology. Use the Entry Cache Disabled Configuration.

Entry Cache Disabled Configuration.

No action is required. This configuration is enabled by default.

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 `or ldapmodify`. See "Attributes of the Instance-Specific Configuration Entry" in the *Administering Oracle Internet Directory* and "Setting System Configuration Attributes by Using `ldapmodify`" in the *Administering Oracle Internet Directory*.

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 *Administering Oracle Internet Directory* and "Setting System Configuration Attributes by Using `ldapmodify`" in the *Administering 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 `orclinmemfiltprocess` 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 `orclinmemfiltprocess` as follows:

```
(objectclass=inetorgperson) (employeeetype=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 `orclinmemfilter` will significantly reduce database processing time. See the following example:

```
orclinmemfiltprocess;dn: cn=users,dc=acme,dc=com
```

Specific Use Cases That Require Additional Tuning

This section describes some specific use cases that require additional tuning, in addition to [Basic Tuning Considerations](#).

- [Bulk Load Operations](#)
- [Bulk Delete Operations](#)
- [High LDAP Write Operations Load](#)

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 for the duration of the operation, if sufficient RAM is available.
- 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.

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 [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 100MB.
- 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, [High LDAP Write Operations Load](#).

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 [Server Entry Cache](#).
- If neither Oracle Internet Directory replication nor DIP is deployed, disable change log generation. See [Replication or Oracle Directory Integration Platform](#).

[Table 10-7](#) summarizes the redo log and undo tablespace recommendations provided in this section.

Table 10-7 Redo Log and Undo Tablespace Values

Attribute	Value	Notes
Redo Log	3 logs, 100MB each	Many <code>bulkdelete</code> operations.
Redo Log	Total size 1000-15000MB	Large number of write operations.
Undo Tablespace	At least 1GB total	Many <code>bulkdelete</code> operations.
Undo Tablespace	2-4 GB	Large number of write operations.

11

Oracle Access Management Performance Tuning

This chapter provides guidelines for tuning and sizing the services that make up an Oracle Access Management 14c Release 14.1.2.0.0 installation.

- [About Oracle Access Management](#)
- [Performance Considerations for Oracle Access Management Services](#)
- [Tuning Oracle Access Management Access Manager](#)
- [Tuning Oracle Access Management Identity Federation](#)
- [Tuning](#)
- [Tuning Oracle Access Management Mobile and Social](#)
- [Database Tuning for Oracle Access Management](#)
- [Purging Inactive Sessions as a Recovery Mechanism from Peak Load](#)

About Oracle Access Management

Oracle Access Management includes a full range of services that provide Web perimeter security functions and Web single sign-on; identity context, authentication and authorization; policy administration; testing; logging; auditing; and more.

Oracle Access Management is a Java Platform, Enterprise Edition (Jakarta EE)-based enterprise-level security application that provides restricted access to confidential information and centralized authentication and authorization services. Many existing access technologies in the Oracle Identity Management stack converge in Oracle Access Management.

Starting with release 11.1.2, Oracle Access Management includes the following "services":

- Oracle Access Management Access Manager (formerly the standalone product named Oracle Access Manager)
- (formerly the standalone product named Oracle Secure Token Service)
- Oracle Access Management Identity Federation (formerly the standalone product named Oracle Identity Federation)
- Oracle Access Management Mobile and Social (formerly the standalone product named Oracle Identity Connect)

For more information on administering these services, see the Oracle Fusion Middleware Administrator's Guide for Oracle Access Management.



Note:

Prior to the Oracle Fusion Middleware 11.1.2 release some of the services discussed in this chapter, such as Oracle Identity Federation and Oracle Secure Token Service, were standalone products and tuned individually.

Performance Considerations for Oracle Access Management Services

Identifying the areas of your Oracle Access Management environment that may impact performance is the first step in effective performance tuning. This section provides information on some of the common areas to review. Always consult your specific usecase scenarios and performance requirements to determine which configurations are applicable.

Before you begin tuning Oracle Access Management services, review the following sections as well as the recommendations discussed in [Top Performance Areas](#) :

- [Understanding Your Current Environment](#)
- [Controlling Network Latency](#)
- [Enabling DMS Performance Instrumentation](#)

Understanding Your Current Environment

Before tuning Access Management services consider the tuning recommendations described in [Table 11-1](#):

Table 11-1 Understanding Your Current Environment: Tuning Considerations

Tuning Consideration	Description
Number of Users	Understanding the overall user population size; group, membership and attribute counts; data types, and configuration parameters of the LDAP and database is essential. See Performance Planning for more information on using population data to improve performance.
Daily Activity Usage	Access Manager: It is important to know how many users are active during a 24-hour period and the expected traffic. Spikes in usage may require additional tuning to avoid performance issues. See Monitoring for more information on collecting performance data. Identity Federation: It is important to know how many Federated SSO requests are processed in a 24-hour period and the expected traffic. Spikes in usage may require additional tuning to avoid performance issues
Hardware Resources and Topology	Like any application deployed for interactive use in a demanding environment, proper server sizing and configuration is critical for acceptable performance. Ensuring that your hardware is sufficient to prevent bottlenecks is a key factor in performance tuning. See Securing Sufficient Hardware Resources for more information on optimizing hardware resources.
Partners and Protocols	When tuning Identity Federation, knowing which partners are configured, how those partners are modeled and the federation protocol used are important considerations. Specifically you should understand how many partners this instance has and what protection policies are assigned to them.

Table 11-1 (Cont.) Understanding Your Current Environment: Tuning Considerations

Tuning Consideration	Description
Protected Applications	Knowing which applications are being protecting and how that protection is modeled is an important consideration when tuning. Specifically you should understand how the applications are being protected: using Webgates (10g,11g, or 11gPS1); mod_osso; custom AccessGates; or a combination.
JVM and Garbage Collection	<p>Optimal performance of the Access Management services depends on correctly tuning JVM heap sizes and garbage collection.</p> <p>NOTE: When uploading large Plugins or CRLs (10MB+) through the OAM Console UI, you need to ensure that the OAM Server heap size is optimally tuned to overcome OutOfMemory issues.</p> <p>For example, increase the <code>-Xmx</code> and <code>XX:MaxPermSize</code> if the following error message is seen in the OAM logs:</p> <pre>javax.management.RuntimeErrorException: GC overhead limit exceeded</pre> <p>Use Parallel, Concurrent Mark and Sweep GC modes with the JVM running in the Server Mode. In addition, Oracle recommends to set the Heap size to a large value and use the same values for Minimum and Maximum (<code>-Xms=-Xmx</code>).</p>

Controlling Network Latency

The performance of the overall network is a major factor in the performance of the system. A reduction in network latency can improve network performance.

To control network latency, consider the following:

- Keep database repositories close to the OAM servers. Installing OAM servers on a remote server may cause significant latency. Latency between the application tier and the database tier should be 5ms or less to maintain optimal performance.
- Add an SSL accelerator or load balancer outside of the Oracle Access Manager system to improve the performance of your network.
- Deploying a load balancer in front of the Web servers or application servers is a best practice for increasing availability and performance of Web-based applications, including Oracle Access Manager. However, load balancers are not recommended between the Oracle Access Manager components themselves.
- Place the Access Manager Servers closer to client applications than to the directory.

During normal operations there can be a considerable amount of traffic between Webgates and Access Manager Servers. Locating these managed servers closer to the applications can reduce the latency between devices in high-traffic parts of the network.

Access Manager provides keep alive, failover, and fallback functionality to handle LDAP and network outages, replication, and related activities. The built-in features of Oracle Access Manager are often the same or better than similar features provided by a load balancer.

 **Note:**

In addition to ensure fast failover, tune the settings for fast failover. The defaults rely on the OS TCP/IP settings which must be tuned for the OS on which the Webgate is running.

You may use Load Balancers to manage the Access Manager server communication information for OAP (Oracle Access Protocol) by virtualizing it. The benefits of using a Load Balancer between Webgates and Servers should be measured against the following constraining requirements:

- OAP connections are persistent and need to be kept open for a configurable duration even while idle.
- The Webgates need to be configured to recycle their connections proactively prior to the Load Balancer terminating the connections, unless the Load Balancer is capable of sending TCP resets to both the Webgate and the server ensuring clean connection cleanup.
- The Load Balancer should distribute the OAP connection uniformly across the active Access Manager Servers for each WG (distributing the OAP connections according the source IP), otherwise a load imbalance may occur.

 **Caution:**

If the above constraining requirements are not met, you can negatively impact the performance of Access Manager resulting in outages.

Ensure that the LDAP timeout under load are negligible. This requires ensuring that LDAP Server is appropriately patched and load testing be performed to simulate OAM LDAP queries (bind, user/group lookup, search queries). LDAP timeouts under load increases OAM Server SSO latencies and increase the risk of an OAM server outage.

Temporary latency blips (for example, increase in LDAP query latency, server processing due to increased Coherence latency) results in increased Webgate response times. If the Web Tier does not have adequate capacity to handle the incoming user requests (through queuing or throttling) especially during peak load, you may run into a situation where the entire Web Tier is blocked and unable to accept new requests. This results in end users not being able to login to access business application.

Enabling DMS Performance Instrumentation

For performance tuning purposes, consider enabling Dynamic Monitoring Service (DMS) performance instrumentation which can tell you the latency and throughput of functional and operational metrics. DMS can identify components that are either processing a heavier load or taking longer than usual to service requests. See [Viewing DMS Metrics](#) for more information on determining the overall time to process calls to various components.



Note:

If you are using Enterprise Manager Grid Control, create Dashboard Reports based on the OAM Metrics of most interest, which can then be email'd on a regular schedule.

Tuning Oracle Access Management Access Manager

Oracle Access Management Access Manager (Access Manager) is an enterprise level solution that centralizes critical access control services to provide an integrated solution that delivers authentication, authorization, Web single sign-on, policy administration and enforcement, agent management, session control, systems monitoring, reporting, logging, and auditing.

For more information on using Access Manager, see "Introduction to Oracle Access Management Access Manager" in the Oracle Fusion Middleware Administrator's Guide for

Oracle Access Management.

- [Basic Tuning Considerations for Access Manager](#)
- [Advanced Tuning Considerations for Access Manager](#)
- [Specific Use Cases That Require Additional Tuning for Access Manager](#)

Basic Tuning Considerations for Access Manager

Depending on your Access Manager usage and performance issues, you may consider tuning the following basic parameters. See [Top Performance Areas](#) for additional tuning considerations.

- [Tuning the Web Tier](#)
- [Managing Policy Components](#)
- [Tuning Common Settings](#)

Tuning the Web Tier

Tuning your Web application's server is essential to maintaining optimal performance for Access Manager. This section describes tuning configurations for the following:

- [Tuning Oracle HTTP Server](#)
You can tune Oracle HTTP Server (OHS) to optimize its performance as the web server component for Oracle Fusion Middleware.
- [Tuning Access Manager Webgate](#)
- [Tuning OAM Agents](#)

Tuning Oracle HTTP Server

You can tune Oracle HTTP Server (OHS) to optimize its performance as the web server component for Oracle Fusion Middleware.



Note:

The configuration examples and recommended settings are for illustrative purposes only. Consult your own use case scenarios to determine the configuration options that can provide performance improvements.

Tuning Access Manager Webgate

Webgate is an out-of-the-box access client for Access Manager. This Web Server access client intercepts HTTP requests for Web resources and forwards them to the Access Manager Server. Webgates for various Web Servers are shipped with Access Manager.

Consider tuning the following parameters to increase the number of connections from the Webgate Server to the Access Manager servers. Adding more connections enables the servers to process more concurrent requests.

Parameter	Description
Max Connections	Maximum number of connections that this Access Manager Agent can establish with all the Access Manager Servers.
Maximum Number Of Connections	Maximum number of connections that the Access ManagerAgent can establish with a specified Access Manager Server.

For more information on setting these parameters, see "Registering Agents and Applications" in the *Administering Oracle Access Management*.

Tuning OAM Agents

Once you have registered an OAM Agent, you can tune the following parameters to the recommended values:

Parameter	Recommendation
Cache Pragma Header	Delete the default of <code>no-cache</code> so that this field is empty.
Cache Control Header:	Delete the default of <code>no-cache</code> so that this field is empty.
AAA Timeout Threshold	Default is -1, which means that there no timeout. You should change this to a hard number. Too low a number means that the socket connection can be closed before a reply comes from the OIM server. Too high a number means the connection may hang while waiting for a response.
Max Number of Connections for Each Server	For each server under Server Lists , change the Max Number of Connections from 1 to 10

To find these parameters, navigate the following menus: OAM11g Admin console > SystemConfig (tab) > Access Manager Settings > ssoAgents > OAM agents > (search and select the agent).

For more information on these parameters, see "Understanding Registered OAM Agent Configuration Parameters in the Console" in *Administering Oracle Access Management*.

Managing Policy Components

In order to limit the Access Manager processing overhead, all resources that do not require security should be modeled as excluded resources as opposed to unprotected resources. Modeling these resources as excluded resources can substantially help with ADF Applications. Excluded resources use a one-time interaction between the Webgate and the Access Manager Server as opposed to a per request interaction for unprotected resources.

For more information, see "Managing Shared Policy Components" in the *Administering Oracle Access Management*.

To design authentication policies for optimal performance, do the following:

- Get an inventory of all attributes you want for `authZ` and pre-fetch them at `AuthN` time.
- Combine attributes in the supplementary list to reduce `AuthN` time LDAP load.

Note the following:

1. Change all OAM policy responses for `userid` return from `$user.attr.uid` to `$user.userid`. This is because the latter is computed at login time as opposed to the former which is computed onDemand during authorization

`OAM_REMOTE_USER` is populated by default.

2. To design authorization policies for optimal performance, do the following:

- Use `$session` namespace. Attributes used for authorization must be retrieved and stored in the user's OAM session during login. This ensures that the `authZ` latency is constant to make OAM responsive thereby improving the user experience.

For example, modify `ismemberof`, `loa` and any other attribute related policy response to get value at authentication time instead of `authZ` time.

```
[Authentication Policies]ismemberof -> SESSION -> $user.attr.ismemberof
loa -> SESSION -> $user.attr.loa
uid ->SESSION -> $user.userid
```

```
[Authorization Policies]Responses:
uid: $user.userid
ismemberof: $session.attr.ismemberof
loa: $session.attr.cmsRoles
```

- For Authorization policies involving attributes, store and use attributes in the `$session` namespace instead of query them on-the-fly by using the `$user.attr` namespace.
- Use group based policies instead of explicitly listing users.

Tuning Common Settings

All OAM Servers and services in the domain share a set of common settings. You can tune them from the Launch Pad. See "Managing Common Settings" in the *Administering Oracle Access Management* for how to find these settings.

This section provides tuning values for the following Common Settings:

- [Global Session Settings](#)
- [Default and System Identity Stores](#)

Global Session Settings

The recommended values for the following parameters:

`Session Lifetime = 5m`

`Maximum Number of Sessions` should be set somewhere between 0 and 8. The default for this setting is 8. Usually the default is sufficient, but this number should be as low as possible. Note that setting this parameter to 0 means that a user can have unlimited sessions running concurrently, which is inadvisable.

For descriptions of these parameters, see "About Global Session Lifecycle Settings" in the *Administering Oracle Access Management*.

Default and System Identity Stores

LDAP stores are accessed by connection pools maintained by Access Manager. Identity store definitions contain the exposed pool parameters. Middleware Control and the DMS Spy Servlet can expose per-operation counts and latency which can be used to identify bottlenecks.

Consider specifying an explicit time-out value (default=unlimited) and ensure that the initial and maximum number of connections in the pool are appropriate for the deployment.

See [Tuning the Data Tier Connections](#) for more advanced tuning recommendations.

For more information on the how to find these settings, see "Defining the User Identity Store Registration Settings" in the *Administering Oracle Access Management* for more information.

Advanced Tuning Considerations for Access Manager

The following Access Manager 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.

- [Tuning Oracle Coherence](#)
- [Setting the Java Message Bean Pool Size](#)
- [Tuning the Server Cache](#)
- [Tuning Webgate Caches](#)
- [Changing Request Cache Type](#)
- [Tuning Authentication Plug-Ins](#)

Tuning Oracle Coherence

Oracle Access Manager uses Oracle Coherence to replicate session states within a distributed installation. Coherence is used to communicate state changes between the Oracle Access Manager Console and Access Manager Servers.

This section contains the following topic:

- [Updating Optimization Interval Time](#)

Updating Optimization Interval Time

In the `oam-config.xml` file, the value of the `OptimizedSessionUpdatesIntervalInMillis` element should be less than the value configured for `Idle Timeout` parameter, which is 15 minutes by default.

In the configuration file, the `OptimizedSessionUpdatesIntervalInMillis` element appears as follows:

```
<Setting Name="DBSMEConfig" Type="htf:map">
  <Setting Name="SessionCreationLockAcquirePercentage"
Type="xsd:integer">60</Setting>
  <Setting Name="SessionPurgeLockExpiryIntervalSeconds"
Type="xsd:long">3000</Setting>
  <Setting Name="SessionCreationLockExpiryIntervalSeconds"
Type="xsd:long">10</Setting>
  <Setting Name="SessionConcurrencyHardLimit" Type="xsd:long">5</
Setting>
  <Setting Name="OptimizedSessionUpdatesIntervalInMillis"
Type="xsd:long">180000</Setting>
</Setting>
```



Note:

This configuration should be under the XPath `/DeployedComponent/Server/NGAMServer/Profile/Sme`.

Based on the configured interval, the session updates during authorization will be optimised and it should be within the limit mentioned above. It is recommended to configure a minimum interval value to avoid any unexpected behaviour. In most of the cases, default value of 3 minutes (180000 ms) is sufficient to handle the load.

Setting the Java Message Bean Pool Size

By default, the Access Manager Proxy is set to handle 100 concurrent Webgate requests.

If necessary, consider adjusting the pool settings to reflect the maximum Webgate request load for the deployment. This is achieved by setting the `max-beans-in-free-pool` element to an appropriate value.

You can also calculate the appropriate value for the `max-beans-in-free-pool` based on the Web Tier settings discussed in [Tuning the Web Tier](#). This value should be greater than the `Max Number of connections (in Webgate)` multiplied by the `ServerLimit (in Oracle HTTP Server)` multiplied by the `Number of Webgates`.

Tuning the Server Cache

The following server caches can be tuned to improve Access Manager performance:

- [Tuning Identity Store Cache](#)

Tuning Identity Store Cache

Authorization policy administration allows authoring of grants to users or groups. Administrators can search within specific identity stores, selecting certain users or groups and granting or denying them access. Search results provide canonical identifiers for users and groups such that those values are stored as principals of the Identity Constraint component of Access Manager Authorization policy. The console displays the names and the Identity Store of origin.

To maximize performance, review configuration settings of the following Identity Store caches:

- **Group Membership Cache**

The Group Membership cache stores indirect membership data which is essentially a group's membership in another group. The number of entries and entry time-to-live are configurable parameters. The cache should be tuned if your deployment includes groups that will be checked against or exported as responses, such as groups that are set in identity constraints, for example.

CAUTION: The Group Membership Cache is populated by a recursive search of the entire LDAP tree of nested groups without any loop detection. Consider disabling this cache if you are experiencing degraded Access Manager Server performance.

- **User Attribute Cache**

User Attributes, once fetched, are always cached. Pre-fetching of attributes during authentication is controlled by specifying the attribute list in the SUPPLEMENTAL_RETURN_ATTRIBUTES parameter value of the Identity Store.

Supplemental attribute return values are useful when you do not require the user to make a list selection for the attributes, yet you want those attributes values, as determined by the current row, to participate in the update.

 **Note:**

All LDAP Attribute Condition used in Authz Policy must be retrieved during login and be cached. This improves authz latency and throughout while reducing the burden on the LDAP tier.

Tuning Webgate Caches

Webgate caches information on authentication and on whether or not a resource is protected. Webgate cache tuning sets the total number of unique URLs expected over the timeout interval. Default is 0 URLs, but this means that the cache is not automatically updated and is flushed only when the administrator manually updates the cache. While this is a good option for performance in some scenarios, it may not apply to your individual use cases.

For more information, see "Reviewing OAM Agent Metrics" in *Administering Oracle Access Management*.

This section provides the following topics:

- [Introducing Webgate Caches](#)
- [Reducing Network Traffic Between Components](#)
- [Changing the Webgate Polling Frequency](#)

Introducing Webgate Caches

Webgate caches various information related to resources, authentications and authorizations to improve performance. It uses the cached information to avoid trips to 11g Server for requesting same information. [Table 11-2](#) are the caches used by Webgate to maintain this information.

<<<used for AccessClients too?>>>

Table 11-2 Webgate Cache Types

Cache Type	Description
Resource to Authentication Scheme	This cache maintains information related to authentication schemes being used. Default: 100000 elements
Authentication Scheme	This cache maintains information related to authentication schemes being used. Default: 25 elements Typically Authentication Scheme cache elements require less than 2 Kb of memory per element. See Also: " Tuning Cache Timeout Values ".
Resource to Authorization Policy 11g Webgate only	This cache maintains information related to resources accessed and associated authorization policy. Default: 100000 elements See Also: " Tuning Maximum Cache Elements " and " Tuning Cache Timeout Values ".
Authorization Result 11g Webgate only	This cache maintains information related to authorizations associated with user sessions. Default: 1000 elements See Also: " Tuning Authorization Result Cache ".

About the 11g Webgate Diagnostics Page

This page displays useful information related to currently effective Cache configuration parameters. It also displays runtime information about the caches that include information on the number of cached elements, number of hits and misses so far, and current memory usage of individual caches. The page is found at the following URL:

`http://webserver:port/ohs/modules/webgate.cgi?progid=1`

After upgrading Oracle Webgate 10.1.4.3.0 to Bundle Patch 13 (BP13), the output of the Diagnostic page is a blank page. Starting with Bundle Patch 13, the Diagnostic Page is disabled by default.

To enable this page, per Webgate registration, add the parameter/value : `enableDiagnosticPage=true` in the list of user parameter of the webgate. With a Webgate instance already registered:

- Go to OAM Console > System Configuration > Access Manager > SSO Agents > OAM Agents : Search and Select your Webgate profile
- Add in the end of the list of the "User Defined Parameters" : `enableDiagnosticPage=true`.
- Click on Apply: a pop-up window mentions where the new artifacts are located.

- Copy the newly `ObAccessClient.xml` in the OHS configuration instance.
- Restart the OHS instance and check that the Diagnostic Page is displayed.



Note:

Changes to Webgate parameters are not reflected on Webgate until the next configuration refresh. For 11g Agents, the default configuration refresh interval is 10 minutes.

Tuning Maximum Cache Elements

By default, the Resource to Authentication Scheme and Resource to Authorization Policy caches are created to store 100000 elements. Typically, elements of these caches require less than 1 Kb of memory per element. Therefore, with 100000 elements in each of these caches, typical memory requirement for the caches will be 100000 Kb or 100 Mb each.

Considering memory requirements and your deployment, the Web Server being used and number of unique URLs in your application, you might want to increase or decrease the maximum number of elements to be cached.



Note:

Increase or decrease the Maximum Cache Elements parameter value as needed. If this is set to a value of -1, all Webgate caches are disabled.

For both 10g and 11g Webgates, you can tune the maximum number of elements to be cached property, by changing the Maximum Cache Elements parameter. Updates to this parameter require a Webgate restart.

How to tune the maximum number of elements to be cached

1. Locate and open the desired 10g or 11g Webgate registration page in the .
2. Set the Maximum Cache Elements parameter as desired.
3. Restart Webgate Web server.

Tuning Authorization Result Cache

By default, the Authorization Result cache is created to store 1000 elements. Authorization Result cache elements store the user session identifier, authorization policy identifier, and associated authorization result including any processed policy responses. Therefore, Authorization Result cache elements are bulky and generally require more than 2Kb of memory per element.

Considering memory requirements and the number of concurrent user sessions in your deployment, you might want to increase the number of elements to be cached.

How to tune the number of elements to be cached

1. Locate and open the desired 11g Webgate registration page in the .
2. In User Defined Parameters, add or update `maxAuthorizationResultCacheElems` as desired.

3. Restart Webgate Web server.

Tuning Cache Timeout Values

By default, the following caches are created with a timeout value of 1800 seconds or 30 minutes:

- Resource to Authentication Scheme
- Authentication Scheme
- Resource to Authorization Policy

Elements in these caches are stored with an expiry time that forces these caches to be flushed on expiry.

Considering the frequency of updates to Authentication Schemes, and Authentication and Authorization Policies in your deployment, you might want to increase or decrease the default timeout value.

How to tune the cache timeout

1. Locate and open the desired 10g or 11g Webgate registration page in the .
2. Set the Cache Timeout parameter as desired.
3. Restart Webgate Web server.

Tuning Authorization Result Cache Timeout

By default, the Authorization Result Cache timeout value is set at 15 seconds. Elements in the Authorization Result Cache is stored with an expiry time that forces it to be flushed on expiry. A low timeout value ensures that authorization results are cached for a small amount of time only.

Considering average length of user sessions and frequency with which user sessions are created and destroyed, you might want to change the default timeout value. Unlike other caches and parameters, updates to this parameter do not require Webgate restart. Instead, the updated value is dynamically picked up by 11g Webgate and enforced immediately.



Note:

If `authorizationResultCacheTimeout` is set to 0, Authorization Cache is disabled.

How to tune the authorization result cache timeout

1. Locate and open the desired 11g Webgate registration page in the .
2. In User Defined Parameters, add or update `authorizationResultCacheTimeout` as desired.
3. Restart Webgate Web server.

Reducing Network Traffic Between Components

The Webgate-to-OAM Server configuration polling reduces the traffic between both the Webgate and OAM Server and the OAM Server and the registered data stores for Oracle Access Manager.

Process overview: Webgate-to-OAM Server configuration polling

1. When the Webgate is inactive for 60 seconds, it reduces the frequency of polling for its configuration information.

The polling frequency is determined by the parameter `InactiveReconfigPeriod`, which is a user-defined parameter that is set in the Webgate configuration page. The value for `InactiveReconfigPeriod` is specified in minutes. Within ten seconds of resuming activity, the Webgate performs reconfiguration polling once a minute.

2. At startup, the Webgate checks the bootstrap configuration to see if any important parameters have changed.

This makes the re-initialization process unnecessary in most cases and reduces the transient OAM Server load.

3. Webgate and AccessClient configurations are cached in the OAM Server.

The default cache timeout is 59 seconds. This should cause no modifications to the system behavior on non-Apache access clients. The Apache Web server with Webgate avoids unnecessary hits to the directory server. The caching parameters can be set in the Webgate registration page.

- `Max Cache Elements` sets the maximum size of the cache (default 9999)
- `Cache Timeout` determines the maximum lifetime of any element in the cache (default 59 seconds)

There are two ways to reduce off-time network traffic between both the Webgate and OAM Server and the OAM Server and the database:

- Changing the default configuration cache timeout for Webgate and AccessClient configurations that are cached in the OAM Server, as described in Step 3.
- Changing Webgate polling frequency for configuration information, as described next.

Changing the Webgate Polling Frequency

One way to reduce off-time network traffic between both the Webgate and OAM Server and between the OAM Server and the database is to change the Webgate polling frequency using the `InactiveReconfigPeriod` parameter.

The default is 1 minute. When the Webgate is inactive for more than 60 seconds (for example, when no authentication requests are being processed), it reduces the frequency of polling for its configuration information. Within ten seconds of resuming activity, the Webgate resumes reconfiguration polling once every minute:

- If set to -2, Webgate never polls.
- If set to a value greater than 0 it polls at the specified interval.
- If set to -1 and Webgate is inactive and has been for 1 minute, then Webgate does not poll. Webgate resumes reconfiguration polling when it returns to an active state.

For example, the OAM Server reads the shared secret from the directory at an interval of 10 minutes and this cached value is returned to Webgate. In the idle state the Webgate reads the shared secret from the OAM Server using the `InactiveReconfigPeriod` value. If this value is not set, the Webgate polls the OAM Server for the shared secret value at an interval of 1 minute even though the updated shared secret value will be returned only after 10 minutes.

To change the configuration polling frequency

1. Locate the desired Webgate registration page using instructions in "Searching for a Webgate Registration" in *Administering Oracle Access Management*.

2. Add the `InactiveReconfigPeriod` parameter as a user-defined parameter on the Webgate registration page.
3. Specify the value for `InactiveReconfigPeriod` in minutes.
4. Apply your changes to the Webgate registration page.

Changing Request Cache Type

The default Request Cache type is set to `COOKIE`, which relies on the use of cookies to cache an unauthenticated request state.

Changing the type to `BASIC` can improve performance, but it is important to consider the following: If the server being used for an authentication flow goes down in the middle of that flow, the user's current state in the flow will be lost on their next request as the load balancer sends them to a different server.

Changing the type to `FORM` can improve performance when lengthy URLs are being accessed.

Tuning Authentication Plug-Ins

Authentication plug-ins can affect performance. When you develop customizations for Access Manager, consider the following to minimize performance impact:

- Evaluate the sequence in which actions are executed
- Minimize the plug-in footprint and external dependencies whenever possible

Specific Use Cases That Require Additional Tuning for Access Manager

This section describes some specific use cases that require specialized tuning, in addition to the [Basic Tuning Considerations for Access Manager](#).

- [Managing Access Manager Sessions](#)
- [Managing Access Manager Sessions](#)
- [Audit Settings](#)
- [Managing Monitor Account](#)
- [Kerberos Latency Issues](#)
- [Oracle Access Protocol over REST Connectivity Issues](#)

Managing Access Manager Sessions

By default, there can only be a maximum of 8 concurrent sessions for a given user ID. It is possible to raise this limit, but it is important to note that as the limit increases the security value of the feature is eroded, and ultimately disappears. Further, there is a performance cost associated with the feature, which increases with the limit. Therefore, if there is a need to have more than 20 concurrent user sessions, then consider disabling this feature by setting the limit to 0.

Audit Settings

OAM tends to generate a lot of audit information. During peak business hours, OAM generates audit information at a rate that is faster than the rate at which the OPSS AuditLoader can move the information to the Audit Database.

Given that SSO is a security service, it is recommended to set the Audit Filter to a value of `MEDIUM` or `ALL`. Also, ensure that the Audit BusStop directory has no max size limit (`maxDirSize=0`) to avoid zero data losses. In addition, monitor and confirm that the Audit data is constantly being moved to the Audit Database even if the AuditLoader falls behind during peak business hours.

Managing Monitor Account

Enterprises use automated monitors to measure end user latency and generate alerts when thresholds are exceeded.

Oracle recommends the following best practices:

1. Monitors should logout when their work is done. This ensures that sessions do not pile up in memory.
2. Monitors should not use the same user credential. This ensures that a single user does not create a very large number of sessions in a short amount of time.
3. Prune monitor sessions periodically. This can be done through the OAM Console or by writing a program using the ASDK. It also ensures that you do not have to set the maximum number of sessions to a very large value to accommodate monitors.
4. Refrain from running monitors very frequently when problems are seen.

Typically, monitors are set to run very frequently when an exception condition is noted (for example, when login latencies exceed the threshold). This has the effect of putting additional load on the system especially if this happens under peak load and this increases the risk of a catastrophic failure.

Kerberos Latency Issues

Kerberos authentication, by default, uses the UDP protocol. However, UDP does not perform well when the connection between the OAM Server and Kerberos Server has to span subnets or the packet loss increases during business hours. As a result, it is recommended that Kerberos be configured to use TCP instead of UDP.

This can be done by setting `udp_preference_limit=1` in the `/etc/krb5.conf` file.

Oracle Access Protocol over REST Connectivity Issues

Oracle Access Protocol (OAP) over REST enables the use of HTTP infrastructure to route and load balance requests. This is a new feature introduced in WebGate starting with release 12.2.1.4.0. Under load, you may see connection errors in the WebGate and/or HTTP Server logs.

Oracle recommends the following best practices to reduce the connection errors:

- Ensure that the HTTP Server has enough idle or spare Server threads to receive incoming requests.
- Increase the OAM WebLogic Server `wm/OAPOverRestWM` work manager capacity.
- Optionally, increase the RAM allocated to the HTTP Server and the WebLogic Server.

Tuning Oracle Access Management Identity Federation

Oracle Access Management Identity Federation (Identity Federation) 11gR2 is an identity federation server built into the Oracle Access Manager server. All configuration is performed in

Oracle Access Manager; unlike the standalone 11gR1 version. Identity Federation provides a self-contained and flexible multi-protocol federation server that can be rapidly deployed with existing identity and access management systems. It enables you to securely share identities across vendors, customers, and business partners without the increased costs of managing, maintaining, and administering additional identities and credentials.

For more information on administering Oracle Access Management Identity Federation, see "Introduction to Identity Federation in Oracle Access Management" in the *Administering Oracle Access Management*.

- [Basic Tuning Considerations for Identity Federation](#)
- [Advanced Tuning Considerations for Identity Federation](#)
- [Specific Use Cases That Require Additional Tuning for Identity Federation](#)

Basic Tuning Considerations for Identity Federation

The following sections describe basic tuning configurations that you should also consider while tuning Identity Federation:

- [Tuning the Load Balancer and HTTP Server](#)
- [Tuning SOAP Connections](#)
- [Tuning the Data Tier Connections](#)

Tuning the Load Balancer and HTTP Server

As of Oracle Fusion Middleware Release 11gR2, some of the features of Identity Federation are embedded in Access Manager. To optimize Identity Federation performance, follow the Load Balancer and HTTP Server tuning guidelines discussed in [Tuning the Web Tier](#) for Access Manager.

Tuning SOAP Connections

Identity Federation uses the Simple Object Access Protocol (SOAP) to send Security Assertion Markup Language (SAML) requests and to receive SAML responses. To optimize performance, configure the following SOAP connections:

- Total maximum number of SOAP connections that Identity Federation and Security Token Service can open at the same time
- Maximum number of SOAP connections that Identity Federation and Security Token Service can open at the same time to a given remote server

Tuning the Data Tier Connections

LDAP stores are accessed by connection pools. Identity store definitions contain the exposed pool parameters. As discussed in [Default and System Identity Stores](#), Middleware Control and the DMS Spy Servlet can expose per-operation counts and latency. Identity Federation uses an RDBMS to store session and runtime data. The server uses a caching mechanism to improve performance at runtime. This enables the server to keep a reference to recently used objects in memory to avoid read access to the database. The RDBMS also has an asynchronous write and delete mechanism.



Note:

The following parameters typically do not need to be changed. Review the descriptions, however, to determine if an adjustment could improve performance for your deployment.

To optimize RDBMS session caching and asynchronous writes, configure the parameters as described in [Table 11-3](#):

Table 11-3 Asynchronous Write Settings

Parameter	Description
<code>rdbsmasynchronousmanagerinterval</code>	Execution interval for the asynchronous thread manager
<code>rdbsmasynchronousmanagersleep</code>	Sleep interval for the asynchronous thread manager, to check if execution should occur
<code>rdbsmasynchronousqueuesize</code>	Size of the queue containing RDBMS operations of the same type (create session, create artifact...) NOTE: It is important to size the <code>rdbsmasynchronousqueuesize</code> correctly. If it is made too large, it can cause a lag in the asynchronous write to the database and may cause SSO operation to fail.
<code>rdbsmasynchronousqueuesleep</code>	Sleep time before the calling thread can retry to add an operation to a queue, in case the queue is full
<code>rdbsmasynchronousqueueretries</code>	Number of retries when trying to add an operation to the queue
<code>rdbsmasynchronousthreadcore</code>	Number of default threads in the RDBMS thread executor module for RDBMS asynchronous operations
<code>rdbsmasynchronousthreadkeepalive</code>	Maximum amount of time to keep the extra threads in the RDBMS thread executor module for RDBMS asynchronous operation
<code>rdbsmasynchronousthreadmax</code>	Maximum number of threads in the RDBMS thread executor module for RDBMS asynchronous operation <code>rdbsmasynchronousthreadmax</code> should be adjusted to handle the maximum system load based on the size of your system.
<code>rdbsmasynchronousthreadpolicy</code>	Thread policy of the RDBMS thread executor module for RDBMS asynchronous operation
<code>rdbsmasynchronousthreadqueuesize</code>	Size of the thread queue of the RDBMS thread executor module for RDBMS asynchronous operation

[Table 11-4](#) describes the RDBMS memory cache settings for artifact and transient cache:

Table 11-4 Cache Settings

Parameter	Description
RDBMS Artifact memory	RDBMS Artifact memory cache settings, used in conjunction of the RDBMS asynchronous module:
artifactrdbmscachetimeout	Time to live in the memory cache
artifactrdbmsretries	Maximum number of time to retry to locate an entry in RDBMS before returning a failure
artifactrdbmssleep	Sleeping time between retrying lookup operations
RDBMS Memory cache	RDBMS Memory cache settings (except for Artifact):
transientrdbmscachesize	Size of the cache
transientrdbmscachetimeout	Time to live for the objects in the cache, before being invalid and thus forcing an RDBMS lookup operation when an object is searched
Interval for the RDBMS cleanup thread	Indicates the interval of sleep of the thread removes expired entries from OIF DB tables

Advanced Tuning Considerations for Identity Federation

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 Identity Federation deployment.

- [Tuning Oracle Coherence](#)
- [Tuning Identity Store](#)
- [Tuning Protocol Binding](#)
- [Tuning the Browser POST and Artifact Single Sign-On Profiles](#)

Tuning Oracle Coherence

Identity Federation, as part of Access Manager 11gR2, uses Oracle Coherence to replicate session states within a distributed installation. See [Tuning Oracle Coherence](#) for more information.

Tuning Identity Store

Identity Federation, as part of Access Manager 11.1.2.0.0, will benefit from tuning the identity store as discussed in [Tuning the Server Cache](#).

Tuning Protocol Binding

This section describes the protocol binding options:

- XML Digital Signatures

Identity Federation relies 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 Identity Federation or Security Token Service 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
- XML Encryption

Federated Single Sign-On allows the use of token and element level encryption to provide confidentiality to the message exchange. Disabling use of encryption improves the latency and throughput of Identity Federation.

Tuning the Browser 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.

- 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.

Since the Artifact profile involves an additional round trip between the Service Provider and the Identity Provider, you may be able to improve performance by avoiding use of the Artifact profile.

- [Outbound SOAP Connections](#)

Outbound SOAP Connections

OAM Federation can communicate with remote SAML Servers using different bindings, among them the `SOAP` binding. When OAM needs to send a message to a remote server using the `SOAP` protocol, it will directly open a connection and send a `SOAP` message.

You can configure the following connection settings:

`soapmaxconnections` - The maximum number of concurrent connections that OAM Federation can open when sending `SOAP` messages.

`soapmaxconnectionsperhost` - The maximum number of concurrent connections that OAM Federation can open when sending SOAP messages to a specific provider.

`soapsockettimeout` - The default socket timeout (`SO_TIMEOUT`) in milliseconds which is the timeout for waiting for data. A timeout value of zero is interpreted as an infinite timeout.

`soapconnectiontimeout` - Sets the timeout until a connection is established. A value of zero means the timeout is not used.

You can use `WLST` cmd to set the above properties. For example, `putLongProperty ("/fedserverconfig/{PropertyName}", {Value})`.

Specific Use Cases That Require Additional Tuning for Identity Federation

This section describes some specific use cases that may benefit from additional tuning.

- [Message Signing versus Token Signing](#)

Message Signing versus Token Signing

Message exchange between the Service and Identity providers may be signed. Message signature provide additional security when the request/response transits numerous intermediaries. Disabling message signatures can improve performance but this should be done only when the security risk of doing so is mitigated by other security mechanisms

Tuning

(Security Token Service) provides a centralized mechanism to broker trust between applications and web services by enabling seamless propagation of identities and security context.

For more information on administering Security Token Service, see Introduction to Oracle Access Management Secure Token Service in *Administering Oracle Access Management*.

- [Basic Tuning Considerations for Security Token Service](#)
- [Advanced Tuning Considerations for Security Token Service](#)

Basic Tuning Considerations for Security Token Service

The following sections describe basic tuning configurations that you should also consider while tuning Security Token Service:

- [Tuning the Load Balancer and HTTP Server](#)
- [Tuning Outbound SOAP Connections](#)
- [Tuning the Data Tier Connections](#)

Tuning the Load Balancer and HTTP Server

To optimize Security Token Service performance, follow the Load Balancer and HTTP Server tuning guidelines discussed in [Tuning the Web Tier](#) for Access Manager.

Tuning Outbound SOAP Connections

Security Token Service uses the Simple Object Access Protocol (SOAP) to send Security Assertion Markup Language (SAML) requests and to receive SAML responses. To optimize performance, configure the following SOAP connections:

- Total maximum number of SOAP connections that can open at the same time
- Maximum number of SOAP connections that can open at the same time to a given remote server

Tuning the Data Tier Connections

Security Token Service uses an RDBMS to store runtime data. The server uses a caching mechanism to improve performance at run time. This enables the server to keep a reference to recently used objects in memory to avoid read access to the database. In addition there is an asynchronous write and delete mechanism to the RDBMS. See [Tuning the Data Tier Connections](#), and review the tuning parameters discussed in [Table 11-3](#) and [Table 11-4](#) as these parameters should also be set for Security Token Service.

In addition, because the LDAP connections are made from Security Token Service when LDAP credential validation is enabled in a validation template in Security Token Service, the connections to that LDAP instance should be tuned with the following parameters:

- Setting the LDAP Inactivity setting which tells Security Token Service 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.

- 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, Security Token Service 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. Security Token Service closes the connection, open a new one and re-issue the LDAP command.

- Setting the High Availability (HA) LDAP Flag.

When integrated with LDAP Servers that are deployed in HA mode, STS must be configured to indicate that the LDAP Servers are in HA mode.

Advanced Tuning Considerations for Security Token Service

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 Security Token Service deployment.

- [Tuning the WS-Security Policy](#)

Tuning the WS-Security Policy

To optimize Security Token Service performance, consider following the recommendations below when configuring your WS-Security Policy:

- Optimal use of Integrity, Confidentiality and RequiredElements assertion
- Optimal use of security binding properties

- Use TransportBinding over SymmetricBinding, which in turn should be considered before AsymmetricBinding
- Avoid encrypting the token for the WS Provider

Tuning Oracle Access Management Mobile and Social

Oracle Access Management Mobile and Social (Mobile and Social) is a new intermediary between a user seeking access to protected resources, and the back-end Identity and Access Management services that protect the resources. Mobile and Social provides simplified client libraries that allow developers to quickly add feature-rich authentication, authorization, and identity capabilities to registered applications. On the back-end, the Mobile and Social pluggable architecture lets system administrators add, modify, and remove Identity and Access Management services without having to update user installed software

- [Basic Tuning Considerations for Mobile and Social](#)

Basic Tuning Considerations for Mobile and Social

The following sections describe basic tuning configurations that you should also consider while tuning Mobile and Social:

- [Tuning the Access Management Authentication Service Provider](#)
- [Tuning the User Profile Service Provider](#)

Tuning the Access Management Authentication Service Provider

Mobile and Social has an out-of-the-box Authentication Service Provider which connects to server using Access Manager SDK components. To optimize Mobile and Social, consider tuning Access Manager as described in [Tuning Oracle Access Management Access Manager](#).

In addition to tuning the Access Manager configuration parameters, there is one configuration parameter that should be tuned in Mobile and Social:

Table 11-5 Mobile and Social Tuning Parameters

Parameter	Description
OAM_SERVER_x_MAX_CONN	<p>Use the following steps to configure the maximum number of connections provided for the Access Management server:</p> <ol style="list-style-type: none"> 1. From the Access Manager 11g R2 console, click System Configuration tab and select Mobile and Social on the left panel 2. Under "Authentication Service Provider", select the target Access Manager service provider 3. Change the value for OAM_SERVER_x_MAX_CONN properly for your performance requirements. 4. Save the change. <p>NOTE: This parameter should be set to be the same value as defined for the "Max Connection for webgate agent" in Access Manager. If different values are provided then the setting in Access Manager server will take precedence.</p>

Tuning the User Profile Service Provider

The User Profile Service in Mobile and Social depends on IDS/libOVD to connect to the user repository. There are two IDS/libOVD configuration parameters that can be tuned for the production deployment as described below. These parameters can be changed via Mobile and Social Remote Console.

Table 11-6 User Profile Service Provider Tuning Parameters

Parameter	Description
Connection Pool Initial Size	Category: LDAP Adapter Properties Default: 5 Recommendation: The default value can be used.
Connection Pool Maximum Size	Category: LDAP Adapter Properties Default: 10 Recommendation: 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.

Database Tuning for Oracle Access Management

This section describes the tuning process for the OAM Database.

- [Automatic Optimizer Statistics Collection](#)
- [Partitioning AM_SESSION table using Config Utility Command](#)

Automatic Optimizer Statistics Collection

Ensure that this is performed where AM_SESSION table has data. Normally this table will have data during the working hours. **Automatic Optimizer Statistics Collection Job** should be configured to run at a time when this table has data. Configuring to run at midnight or off peak hours may cause the wrong statistics to be collected and in turn cause performance degradation of the OAM servers.

Follow the procedure below to check the job details.

1. Connect as **dba** and run the query.
2. `select * from dba_autotask_client where client_name = 'auto optimizer stats collection'`

Partitioning AM_SESSION table using Config Utility Command

By default, the AM_SESSION table is not partitioned.

It is recommended to partition the AM_SESSION table for stability when high load is expected on the system. Also, the database statistics should be gathered at regular intervals to ensure that queries on the AM_SESSION table perform well.

Run the following Config utility command to partition or non-partition the AM_SESSION table:

```
java -cp $MW_HOME/idm/oam/server/tools/config-utility/config-utility.jar:$MW_HOME/oracle_common/modules/oracle.jdbc.ojdbc8.jar
```

```
oracle.security.am.migrate.main.ConfigCommand $DOMAIN_HOME createAMSessionTable /  
scratch/config.properties
```

The following line should be a part of the properties file (`/scratch/config.properties`) along with other default properties required for executing config utility commands:

```
oam.sessionTable.type=<value>
```

Where `<value>` should be one of the following:

- `PARTITIONED` - To partition `AM_SESSION` table
- `NON-PARTITIONED` - To use non-partitioned `AM_SESSION` table

Purging Inactive Sessions as a Recovery Mechanism from Peak Load

Following is the sample REST API:

```
Method: POST Path: https://oam-policy-admin-host:oam-policy-admin-port/oam/  
services/rest/access/api/v1/sme/purge?allInactiveSessions=true
```

Note:

This operation should be executed only during a maintenance or low load window (For example: midnight). You must ensure that following conditions are met:

- `AM_SESSION` table is partitioned.
- Heavy load was observed in current day.
- Heavy load is anticipated in upcoming days.

If the peak load is expected only for a very short duration in a day, you should not perform this operation. The performance will be optimal with right tunings in place.

12

Oracle Identity Governance Performance Tuning

This chapter provides guidelines for tuning and sizing specific to Oracle Identity Governance (OIG).

Note:

As with any enterprise class business application, there is no simple procedure for tuning that works for all systems. The tuning sections in this chapter provide (in some cases) sample configurations and outline the principles for tuning Oracle Identity Governance. Consider your own use case scenarios to determine which settings are appropriate.

- [About Oracle Identity Governance](#)
- [Monitoring Oracle Identity Governance Performance](#)
- [Basic Tuning Considerations](#)
- [Advanced Tuning Considerations](#)

About Oracle Identity Governance

Oracle Identity Governance (OIG) provides operational and business efficiency through centralized administration and complete automation of identity and user provisioning events across the enterprise, as well as extranet applications.

For more information on using Oracle Identity Governance, see the [Administering Oracle Identity Governance](#).

Monitoring Oracle Identity Governance Performance

To identify performance bottlenecks, you can monitor real-time performance metrics for Oracle Identity Governance. For more information on how to monitor your Oracle Fusion Middleware components, see [Monitoring](#).

For Oracle Identity Governance it is recommended that you perform the following at regular intervals:

- Monitor real-time performance by using a performance-monitoring tool such as Oracle Enterprise Manager console or Automatic Workload Repository (AWR) in Oracle Database 11g.

 **Note:**

You can use Oracle Enterprise Manager 11g Fusion Middleware Control to monitor Oracle Identity Governance. To do so:

1. Under Identity Management, select **Oracle Identity Governance** to go to the home page. On the Home page, you can monitor Oracle Identity Governance.
2. From the Oracle Identity Governance menu, select **Performance** to view performance metrics.

- Collect routine statistics and report by using Oracle Database Enterprise Manager (EM), which is available in Oracle Database as a standard offering.

- Routine Statistics Gathering

Routine statistics gathering can be taken care by the 'Automated Maintenance Tasks', which is available in the following navigation path in Oracle Database:

Oracle EM, the **Server** tab, **Query Optimizer**, **Manage Optimizer Statistics**, the **Automated Maintenance Tasks** link

- Reporting requirements of statistics through Oracle Database 11g EM

To report on the state of the currently gathered statistics, EM provides a reporting interface in the following navigation path:

Oracle EM, the **Server** tab, **Query Optimizer**, **Manage Optimizer Statistics**, the **Object Statistics** link

This interface can be used for the reporting purpose for All Objects (of the Schema or even the Object of choice), which have Stale, Missing, or Locked states or are already analyzed.

- Collect complete schema statistics upon implementation of Oracle Identity Governance.

Update OIG schema and its dependent schemas (*_MDS, *_SOAINFRA, *_OPSS and *_ORASDPM). You must consider complete schema or table statistics on mass data change events such as bulkload of users or accounts, import of a new connector, a huge reconciliation run from a new target system, or use of an archival utility. You should collect statistics regularly for OIG and also OIG dependent schemas *_MDS, *_SOAINFRA, *_OPSS and *_ORASDPM.

This helps the CBO determine an efficient query execution plan that is based on the current state of data. The following is a sample SQL command to collect database statistics on a regular basis:

 **See Also:**

Gathering routine statistics and reporting can be done by performing the automated maintenance tasks available in Oracle Database 11g. See *Oracle Database Performance Tuning Guide 11g Release 1 (11.1)* for details.

```
DBMS_STATS.GATHER_SCHEMA_STATS(OWNNAME=> schema_owner,
Exec dbms_stats.gather_schema_stats(OWNNAME=>
'OIG_OIG',ESTIMATE_PERCENT=>DBMS_STATS.AUTO_SAMPLE_SIZE,degree
```

```
=>DBMS_STATS.DEFAULT_DEGREE,options=>'GATHER AUTO', no_invalidate  
=>FALSE,cascade=>TRUE);
```

- Look for relevant recommendations provided in advisory sections in the Automatic Database Diagnostic Monitor (ADDM) or Automatic Workload Repository (AWR) report, and adjust the instance configuration parameters according to the recommended settings. This is specially required after importing a new connector and completing a round of reconciliation from a new target system so that you can identify the need of any new indexes according to your matching rules.

Basic Tuning Considerations

Depending on your Oracle Identity Governance usage and performance issues, you may consider tuning the following basic parameters. See [Top Performance Areas](#) for additional tuning considerations.

- [Tuning and Managing Application Cache](#)
- [Tuning the Application Server for Oracle Identity Governance](#)
- [Tuning Database Parameters for Oracle Identity Governance](#)
- [Tuning Oracle Internet Directory](#)
- [Tuning Application Module \(AM\) for User Interface](#)
- [JMS Tuning](#)

Tuning and Managing Application Cache

Oracle Identity Governance allows caching of metadata, which reduces DB activities. This results in reduced network load and improved performance.

By default, caching for most of the configurations are disabled (set to false) so that the configuration changes are reflected immediately without having to restart the application servers in the development environments.

The following sections provide some recommended cache values for tuning Oracle Identity Governance:

- [Tuning Oracle Identity Governance Cache](#)
- [Purging the Cache](#)

Tuning Oracle Identity Governance Cache

Caching is configured in the `/db/oim-config.xml` configuration file, which is located in MDS where Oracle Identity Governance stores the configuration. You can use Oracle Enterprise Manager (EM) to turn on caching, or export the `oim-config.xml` to make changes and then import it back to turn on caching.

Oracle recommends the following settings for the production environments for optimal and better performance. Using EM, go to System Mbean > Application Defined Mbeans > oracle.iam > server:OIM_server1 > Application: OIM > XMLConfig > Config > XMLConfig.CacheConfig > Cache > XMLConfig.CacheConfig.CacheCategoryConfig, and do the following:

- Set the caching to `true` for all the components *except* the following two sections:

```
threadLocalCacheEnabled="false"
```

- For non-clustered installation, set `clustered="false"`. For clustered installation, set `clustered="true"`.



Note:

Changing this value gets saved into the MDS database schema used by the Oracle Identity Governance servers. Therefore, change only once for multi-node/clustered installations.

Enabling Cache Categories `User_Org_Membership_And_Chain` and `ObjectDefinition`

It is recommended that you enable the cache categories described in [Table 12-1](#), based on your Oracle Identity Governance version. Note that you do not need to enable these, if your Oracle Identity Governance version is not same as given in "**Applicable Release**" column in the following table:

Table 12-1 Instructions to Enable Cache Category

Cache Category Name	Applicable Release	Instructions
User_Org_Membership_And_Chain	Oracle Identity Governance 12c Release (12.2.1.4.0)	<p>You can enable this cache category using Oracle Enterprise Manager (EM) or by editing the <code>oim-config.xml</code> configuration file. To do this, complete the following steps:</p> <p>Using EM</p> <ol style="list-style-type: none"> 1. Log in to EM. 2. Go to mbean XMLConfig.CacheConfig under oracle.iam, and set the value of attribute <code>Enabled</code> to <code>true</code>, if not already set to <code>true</code>. Mbean's Object name is "oracle.iam:name=Cache,type=XMLConfig.CacheConfig,XMLConfig=Config,Application=OIM,ApplicationVersion=12.2.1.4.0". 3. Create a new cache category using mbean's createCacheCategoryConfig operation with the following parameters: <pre>enabled=true expirationTime=3600 name=User_Org_Membership_And_Chain</pre> <p>Using oim-config.xml File</p> <ol style="list-style-type: none"> 1. Go to <code>\$(OIM_HOME)/bin</code>. 2. Set the environment variable <code>OIM_ORACLE_HOME</code> appropriately. 3. Open the <code>weblogic.properties</code> file, and set the following properties in order to export the metadata file: <pre>wls_servername=OIM_server1 application_name=OIMAppMetadata metadata_to_loc=<TMP_DIRECTORY> metadata_files=/db/oim-config.xml</pre> 4. Run the following command script to export the <code>/db/oim-config.xml</code> metadata file: <pre>./weblogicExportMetadata.sh</pre> <p>When prompted, enter the WebLogic credentials and the JNDI URL.</p> 5. Open the <code>\$(TMP_DIRECTORY)/db/oim-config.xml</code> file, and add the following in the <code>cacheCategoriesConfig</code> tag: <pre><cacheCategoryConfig enabled="true" expirationTime="14400" name="User_Org_Membership_And_Chain"/></pre> 6. Open the <code>weblogic.properties</code> file, and set the following properties in order to import the modified metadata file:

Table 12-1 (Cont.) Instructions to Enable Cache Category

Cache Category Name	Applicable Release	Instructions
		<pre>wls_servername=OIM_server1 application_name=OIMAppMetadata metadata_from_loc=<TMP_DIRECTORY></pre> <p>7. Run the following command to import the modified <code>/db/oim-config.xml</code> metadata file into MDS:</p> <pre>./weblogicImportMetadata.sh</pre> <p>When prompted, enter the WebLogic credentials and the JNDI URL.</p>
ObjectDefinition	Oracle Identity Governance 12c (12.2.1.4.0)	<p>You can enable this cache category using Oracle Enterprise Manager (EM). To do so, complete the following steps:</p> <ol style="list-style-type: none"> 1. Log in to EM. 2. Go to mbean XMLConfig.CacheConfig under oracle.iam, and set the value of attribute <code>Enabled</code> to <code>true</code> for the cache category ObjectDefinition.

 **Note:**

For more information on configuration change using Enterprise Manager, see *Using Enterprise Manager for Managing Oracle Identity Governance Configuration in Administering Oracle Identity Governance*.

Purging the Cache

If you want to purge the cache, use the `PurgeCache` utility in the `OIM_HOME/server/bin/` directory. This utility purges all elements in the cache.

 **Note:**

- Purging is required when caching is enabled and if you make any system configuration changes. It is not required if caching is disabled.
- Before running the `PurgeCache` utility, navigate to the `OIM_HOME/server/bin/` directory.

Before running the `PurgeCache` utility, you must run the `DOMAIN_HOME/bin/setDomainEnv.sh` script.

To use the `PurgeCache` utility, run `PurgeCache.bat CATEGORY_NAME` on Microsoft Windows or `PurgeCache.sh CATEGORY_NAME` on UNIX. The `CATEGORY_NAME` argument represents the name

of the category that must be purged. For example, the following commands purge all FormDefinition entries from a system and its clusters:

```
PurgeCache.bat FormDefinition
PurgeCache.sh FormDefinition
```

To purge all Oracle Identity Governance categories, pass a value of "All" to the PurgeCache utility. It is recommended to clear all the categories.

Tuning the Application Server for Oracle Identity Governance

This section describes how to tune Oracle WebLogic Server for Oracle Identity Governance to improve performance. For additional Oracle WebLogic Server performance tuning information, see *Tuning Performance of Oracle WebLogic Server*.

Note:

- All tuning parameter suggestions and values in this section are for reference purposes only. Values should be modified based on your requirement, application usage patterns, loads, and hardware specifications.
- Changing any of the settings may require you to restart the server.

- [Tuning JVM Memory Settings for Oracle Identity Governance](#)
- [Tuning the JDBC Connection Pool for Oracle Identity Governance](#)
- [Tuning OIG-specific Work Manager Properties](#)
- [Disabling the Reloading of Adapters and Plug-in Configuration](#)
- [Changing the Number of Open File Descriptors for UNIX \(Optional\)](#)
- [Tuning the JVM Garbage Collection for Solaris Sparc T3 or T4](#)

Tuning JVM Memory Settings for Oracle Identity Governance

These settings should be used in addition to those described in [Tuning Java Virtual Machines \(JVMs\)](#).

It is recommended to increase the heap and permgen memory for production environments as in [Table 12-2](#) and monitor the memory usage pattern. Based on the usage, you can choose to increase or decrease the memory settings.

Table 12-2 JVM Parameters to be set for Tuning JVM Memory Settings

JVM Parameter	HotSpot JVM
Min. Heap Size (Xms)	4GB
Max Heap Size (Xmx)	8GB
MetaspaceSize (-XX:MetaspaceSize)	500m
MaxMetaspaceSize (-XX:MaxMetaspaceSize)	1GB

To change the JVM memory setting:

1. Use `DOMAIN_HOME/bin/setStartupEnv.sh` (Unix) or `set OIMDomainEnv.cmd` (Windows). If not, continue to use `DOMAIN_HOME/bin/setStartupEnv.sh` (Unix) or `setStartupEnv.cmd` (Windows) to change the heap size settings.
2. Change the value of the memory argument `SERVER_MEM_ARGS_xxHotSpot` for `OIM-MGD_SVRS` in `DOMAIN_HOME/bin/setStartupEnv.sh` (Unix) or `setStartupEnv.cmd` (Windows).
3. Restart all servers.

If you are using the console to restart all the servers, then you must also restart the node manager.

**Note:**

For a clustered or multi-node installation, repeat the above steps on all the install locations.

Tuning the JDBC Connection Pool for Oracle Identity Governance

Oracle Identity Governance uses the `ApplicationDBDS`, `oimOperationsDB` and `oimJMSStoreDS` data sources deployed on the Oracle WebLogic Server. You may have to increase the connection pool size for each data source, based on your requirements

To increase the capacity of the JDBC connection pools:

1. Open the WebLogic Remote Console.
2. Click **Services > Data Sources > Data Source Name** and then click the **Connection Pool** tab.
3. Adjust the Initial Capacity and Maximum Capacity based on requirement.
4. Set the `Inactive Connection Timeout` parameter to 300.
5. Navigate to **Advanced** tab and set `Seconds to Trust an Idle Pool Connection` to 30.
6. Save and activate the changes.

**Note:**

Ensure that any increase in number of connections on the application server connection pools are compensated by database configuration changes. You might have to increase the `MAX SESSIONS` settings on Oracle Database.

Tuning OIG-specific Work Manager Properties

This section describes some tuning options for OIG-specific Work Managers. By default, Work Managers are not optimized for production. Tuning them can help performance by prioritizing processes into a configuration more tailored to your use case.

While Oracle can recommend a few `MaxThreadsConstraint` values, as shown in [Table 12-3](#), you can determine the optimal value for your system configurations using calculations also given in [Table 12-3](#).

To calculate the optimal `Maximum Threads Constraint` for each Work Manager in your particular installation, you should first consult your DBA and ascertain the following values:

- Number of database CPU available for the OIG database
- Number of nodes in your OIG cluster
- Number of threads used in OIG Access Policy Scheduled task "Evaluate User Policy."

Once you know these values, calculate the following values:

1. Multiply the number of database CPU available for the OIG database by 8. The resulting number is the total number of database connections.
2. Divide the number of database connections by the number of nodes in your OIG cluster.
3. For the following equations in [Table 12-3](#), replace the following variables with the values you have calculated:
 - `d` = the total number of database connections
 - `n` = the number of nodes in your OIG cluster
 - `t` = the number of threads used in OIG Access Policy Scheduled task "Evaluate User Policy"

Table 12-3 Recommended Max Thread Constraints for OIG Work Managers

Work Managers	Role	Recommended Value for Max Thread Constraint
<code>OIMMDBWorkManager</code>	This Work Manager applies to most OIG Message Driven Beans (MDB) and limits the number of concurrent threads/MDB-processing JMS messages for all offline activities except audit.	$\text{Round}(1/3[(d-t)/n]-10)$
<code>OIMAuditWorkManager</code>	This Work Manager applies to audit MDBs. It limits the number of concurrent threads/MDB processing audit-related JMS messages.	5
<code>OIMWorkManager</code>	This Work Manager applies to all OIG Jakarta Enterprise Beans (EJBs), which implement underlying APIs. It also limits the number of concurrent threads processing incoming API calls.	$\text{Round}(2/3[(d-t)/n]-10)$
<code>OIMUIWorkManager</code>	This Work Manager limits the number of threads serving requests to and from the user interface.	10 (based on UI Concurrency)
<code>OIMAccessPolicyWorkManager</code>	N/A	6
<code>OIMRoleGrantRevokeWorkManager</code>	N/A	6

For more information on how to tune Work Managers, see *Using Work Managers to Optimize Scheduled Work in Administering Server Environments for Oracle WebLogic Server*.

Disabling the Reloading of Adapters and Plug-in Configuration

By default, reloading of adapters and plug-in configuration are enabled for ease of development. These should be disabled in the production environment. To do so:

1. Export the `/db/oim-config.xml` file from MDS as described in Exporting and Importing Configuration Files in *Administering Oracle Identity Governance*.

2. In the `oim-config.xml` file, replace the following:

```
<ADPClassLoaderConfig adapterReloadingEnabled="true" loadingStyle="ParentFirst"
reloadInterval="15" reloadingEnabled="true">
```

With:

```
<ADPClassLoaderConfig adapterReloadingEnabled="false" loadingStyle="ParentFirst"
reloadInterval="15" reloadingEnabled="false">
```

3. Replace the following:

```
<storeConfig reloadingEnabled="true" reloadingInterval="20"/>
```

With:

```
<storeConfig reloadingEnabled="false" reloadingInterval="20"/>
```

4. Save the `oim-config.xml` file and import it back to MDS.

Changing the Number of Open File Descriptors for UNIX (Optional)

WebLogic limits the number of open file descriptors in the `WEBLOGIC_HOME/common/bin/commEnv.sh` script to 1024. In some cases, if there is a large number of concurrent users, WebLogic may throw the "TOO MANY OPEN FILES" exception. If you receive this error, then consider increasing the limit beyond 1024 in the script. Ensure that the operating system is able to handle the increase in the number of open files. To set the number of open file descriptors, see Setting the Open File Limit and Number of Processes Settings on UNIX Systems in *System Requirements and Specifications*.

Tuning the JVM Garbage Collection for Solaris Sparc T3 or T4

To tune the JVM garbage collection for Solaris Sparc T3 or T4:

1. In a text editor, open the `setSOADomainEnv.sh` or `setSOADomainEnv.cmd` file in the `DOMAIN_HOME/bin/` directory.
2. Set the value of `USER_MEM_ARGS` similar to the following:

Note:

The values shown for `USER_MEM_ARGS` are examples. You can change the values based on your requirement.

```
USER_MEM_ARGS="-Xms3048m -Xmx3048m -Xmn1648m -Xss256k -XX:PermSize=384m -
XX:MaxPermSize=384m"
```

3. Set the value of `JAVA_OPTIONS` similar to the following:

 **Note:**

The values shown for JAVA_OPTIONS are examples. You can change the values based on your requirement.

```
JAVA_OPTIONS="-Xnoclassgc -XX:SurvivorRatio=8 -XX:TargetSurvivorRatio=90
-XX:PermSize=350m -XX:MaxPermSize=350m -XX:+AggressiveOpts
-XX:+UseParallelOldGC -XX:ParallelGCThreads=8 -XX:+PrintGCDetails
-XX:+PrintGCTimeStamps -XX:+PrintGCDateStamps -XX:ReservedCodeCacheSize=64m
-XX:CICompilerCount=8 -XX:+AlwaysPreTouch -XX:+PrintReferenceGC
-XX:+ParallelRefProcEnabled -XX:-UseAdaptiveSizePolicy
-XX:+PrintAdaptiveSizePolicy -XX:+DisableExplicitGC"
```

4. Save and close the file.

Tuning Database Parameters for Oracle Identity Governance

This section describes one sample configuration and outlines the principles for tuning Oracle Database for Oracle Identity Governance. For general database tuning information, see [Tuning Database Parameters](#).

Oracle Identity Governance has many configuration options. The best way to identify bottlenecks and optimize performance is to monitor key database performance indicators in your production environment and adjust the configuration accordingly. Review the monitoring tasks described in [Monitoring Oracle Identity Governance Performance](#) and then use the guidelines in this section to help you choose the initial baseline database configuration.

 **Note:**

It is important that you maintain the baseline database tuning parameters when working with Oracle Identity Governance. See the *Oracle Database Performance Tuning Guide 11g Release 1 (11.1)* for information on setting Oracle Database instance parameters.

- [Sample Instance Configuration Parameters](#)
- [Physical Data Placement](#)
- [Resolving enq: HW - contention](#)
The High Water enqueue contention (enq: HW - contention) event occurs when competing processes are inserting into the same table and try to simultaneously increase the high water mark of a table.

Sample Instance Configuration Parameters

[Table 12-4](#) provides information on some important performance-related database initialization parameters.

SGA,PGA size are limited by the underlying operating system restrictions on the maximum available memory in some platforms. See Support Note: Oracle Database Server and the Operating System Memory Limitations [ID 269495.1].

 **Note:**

For the Database Instance Parameters listed in [Table 12-4](#), following memory management approach should be used based on the Oracle Database versions.

Using Automatic Shared Memory Management (ASMM) available in Oracle Database 10g onward: Here, the SGA components can be managed by specifying the SGA_TARGET and SGA_MAX_SIZE parameters. PGA is managed separately through PGA_AGGREGATE_TARGET.

You should set the processes parameter to accommodate the following connection pool requirements and few extra connections for external programs:

- Connection pool size of XA datasource configured in Application Server
- Connection pool size for non-XA datasource configured in Application Server
- Direct database connection pool size configured in xlconfig.xml

Table 12-4 Sample Configuration Parameters

Parameter	Recommended Initial Settings for Oracle Database
db_keep_cache_size	800M
cursor_sharing	FORCE
open_cursors	800
session_cached_cursors	800
query_rewrite_integrity	TRUSTED
query_rewrite_enabled	TRUE
processes	Based on connection pool settings
MAX_DISPATCHERS	0
MAX_SHARED_SERVERS	0
distributed_lock_timeout	1400
filesystemio_options	SETALL
sga_target	6GB
pga_aggregate_target	2GB

Physical Data Placement

The basic installation of Oracle Identity Governance uses three physical tablespaces to store the OIG database objects:

- Data Tablespace to store the data of tables, their indexes and other objects.
- LOB Tablespace to store OIG Orchestration LOB data.
- Archival Tablespace to store OOTB Archival Tables of the OIG Entities catering to the Real-time Purge feature.

 **Tip:**

To minimize disk space consumption, Oracle recommends the following:

During the initial startup phase of the deployment, Oracle Identity Governance tablespace is expected to grow at the rate 20G for every hundred thousand users reconciled into Oracle Identity Governance. LOB tablespace grows at around 30% of the size of main Oracle Identity Governance tablespace for the same users. Depending on the usage of orchestration in Oracle Identity Governance, which affects the LOB tablespace growth, the LOB tablespace can grow at a rate of 60% to 100% of the main tablespace in scenarios where orchestration is widely used.

Database administrators must monitor the exact growth rate in the real system for efficient disk space management.

For better performance, create multiple locally managed tablespaces and store each category of database object in a dedicated tablespace. This storage optimization helps efficient data access. The tables that are frequently accessed and have potential growth are highlighted in the following sections. Oracle recommends that you place these tables in their own dedicated tablespace(s).

Note that the tables highlighted in the following sections generally grow bigger and are accessed frequently in a typical Oracle Identity Governance deployment. In addition, you can use performance metrics to identify tables that are accessed frequently (hot tables). To reduce I/O contention, move hot tables to dedicated tablespaces.

 **Note:**

Oracle Identity Governance offers archival and purge solution in both Real-time online mode and Command Line mode to contain the data growth in most of these tables. See "Using the Archival Utilities" in *Using the Archival and Purge Utilities for Controlling Data Growth* for more information.

- [Tasks Tables](#)
- [Reconciliation Tables](#)
- [OIG Orchestration LOB Tables](#)
- [Audit Tables](#)
- [Redo-Log Files](#)
- [Keep Pool Changes](#)

Tasks Tables

Oracle Identity Governance stores provisioning and approval task details in the following tables. These tables have lot of potential to grow big overtime. It is recommended to group these in one or more dedicated tablespaces.

- OSI
- OSH
- SCH

Reconciliation Tables

The reconciliation schema of Oracle Identity Governance has both static and dynamic tables. The following is a list of static tables. The dynamic tables can be identified by querying the RECON_TABLE_NAME column in the RECON_TABLES table.

- RECON_ACCOUNT_OLDSTATE
- RECON_BATCHES
- RECON_CHILD_MATCH
- RECON_EVENTS
- RECON_EVENT_ASSIGNMENT
- RECON_EXCEPTIONS
- RECON_HISTORY
- RECON_JOBS
- RECON_TABLES
- RECON_UGP_OLDSTATE
- RECON_USER_OLDSTATE
- RECON_ACCOUNT_MATCH
- RECON_ORG_MATCH
- RECON_ROLE_HIERARCHY_MATCH
- RECON_ROLE_MATCH
- RECON_ROLE_MEMBER_MATCH
- RECON_USER_MATCH
- RA_LDAPUSER
- RA_MLS_LDAPUSER
- RA_LDAPROLE
- RA_MLS_LDAPROLE
- RA_LDAPROLEMEMBERSHIP
- RA_LDAPROLEHIERARCHY

If your environment generates a large amount of reconciliation data, then move these tables to one or more dedicated tablespace(s).

OIG Orchestration LOB Tables

You can use the Archival and Purge Utilities to control data growth in Orchestration tables. For more information, see *Using the Archival and Purge Utilities for Controlling Data Growth in Administering Oracle Identity Governance*.

Audit Tables

Oracle Identity Governance audits the transactions based on the audit level setting. Most of the audit levels are likely to increase data growth significantly. Oracle recommends storing audit tables in their own tablespace. Oracle Identity Governance audit tables are of two categories.

Following are the tables that store audit data in XML format. In this list, UPA table is especially expected to grow big and it is important to place it in a dedicated tablespace.

- UPA
- GPA

The user profile audit data is stored in the following flat structured tables. These tables are used by Oracle Identity Governance historical reports for compliance reporting. It is recommended to store these tables and their indexes in a dedicated tablespace.

- UPA_FIELDS
- UPA_GRP_MEMBERSHIP
- UPA_RESOURCE
- UPA_USR
- UPA_UD_FORMS
- UPA_UD_FORMFIELDS

You can use the Archival and Purge Utilities to control data growth in Audit (UPA) table. For more information, see [Using the Archival and Purge Utilities for Controlling Data Growth in Oracle Fusion Middleware System Administrator's Guide for Oracle Identity Governance](#).

Redo-Log Files

Depending on the reconciliation processes configured in Oracle Identity Governance, the volume of database transactions and commits during a reconciliation run can be high. Oracle recommends that you use multiple redo-log files. The total allocated redo-log space should be 1 GB to 2 GB.

Oracle recommends:

1. Use at least three redo log groups with redo log members.
2. Start with an initial size of 1GB for each redo log member and continue to monitor redo logs for contention or frequent log switches.
3. The multiplexing and the exact number of members and disk space for each member can be considered in accordance with the planning for failure.
4. Adjust the size or add more redo log files based on your findings.

Keep Pool Changes

By default, Oracle Identity Governance assigns frequently referenced small tables to be cached in the database by using a keep pool buffer. See `db_keep_cache_size` in [Table 12-4](#). If your installation contains more than 50,000 users, then Oracle recommends that you use the default database buffer for USR and PCQ tables instead of the keep pool buffer. You can use the following commands to put these tables in default buffer pool.

```
ALTER TABLE USR STORAGE(buffer_pool default);  
ALTER TABLE PCQ STORAGE(buffer_pool default);
```

Resolving enq: HW - contention

The High Water enqueue contention (enq: HW - contention) event occurs when competing processes are inserting into the same table and try to simultaneously increase the high water mark of a table.

In an OIG database, this issue is experienced by tables that have large object (LOB) columns. Under a heavy load, LOB segments in these tables experience contention, which is seen in an AWR report as the wait event enq: HW - contention.

The default storage for LOBs in an Oracle database is BasicFiles. Frequently allocating extents or reclaiming chunks may cause contention for the LOB segment high water marks. This contention can also occur for LOB segments that are ASSM-managed, since space allocation only acquires one block at a time.

This contention can be eliminated by switching LOB storage from BasicFiles to SecureFiles. SecureFiles is an LOB storage architecture that provides performance benefits over traditional BasicFiles. See *About LOB Storage in Database SecureFiles and Large Objects Developer's Guide* for more information on these two architectures.

If you encounter the enq: HW – contention event on your OIG database, you can resolve it by migrating LOB storage to SecureFiles by setting the following database event:

```
ALTER SYSTEM SET EVENT='44951 TRACE NAME CONTEXT FOREVER, LEVEL 1024'  
scope=spfile;
```

 **Note:**

This fix should only be applied when you see the enq: HW – contention event for SOA-related SQLs during certification. This is similar to [Migrating BasicFiles to SecureFiles \(enq:HW - contention\)](#), which provides additional details on resolving contention issues.

Tuning Oracle Internet Directory

To ensure that the Oracle Identity Governance is performing at the optimal level, it is important to tune the Oracle Internet Directory as described in [Oracle Internet Directory Performance Tuning](#).

Tuning Application Module (AM) for User Interface

For more information on AM Pool tunings, see section [Application Module Pooling](#) in the *Oracle Fusion Middleware Performance and Tuning Guide*.

 **Note:**

The recommended settings assume 100 concurrent users per node. If your number of concurrent users is different, use the following formula to change

Djbo.ampool.maxavailablesize:

Djbo.ampool.maxavailablesize = # of concurrent users + 20%

JMS Tuning

It is recommended to change the defaults (-1) of **Message Buffer Size** and **Messages Maximum** properties. Set the **Message Buffer Size** to 1 GB (1073741824 bytes) and **Messages Maximum** 1000000 respectively.

Go to WebLogic Remote Console to change these properties.

- **Message Buffer Size:**
Services > JMS Servers > JRFWSAsyncJmsServer_auto_*
- **Messages Maximum:**
Services > JMS Servers > JRFWSAsyncJmsServer_auto_* > Thresholds tab

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 Oracle Identity Governance performance.

- [Reconciliation Tuning](#)
- [Tuning LDAP Synchronization](#)
- [Tuning Order Audit Messages To Eliminate Slow SQL](#)

Reconciliation Tuning

Three distinct process stages or functional modules come into play during the end-to-end reconciliation flow. The following are the three functional modules or stages that need to be optimized separately, but in relation to each other, to achieve complete performance optimization:

- **The Target System And The Connector**
The Connector fetches data from the target system, and invokes reconciliation create event APIs to create events and event data in reconciliation staging tables in the OIG database schema.
- **OIG Reconciliation Engine**
The OIG reconciliation engine extracts data from the staging tables and reconciles into OIG. The process includes verification, matching of data, and taking actions based on the rules. The engine uses database's bulk collection mechanism to do all of the above processing in bulk.
- **Oracle Identity Governance Post-processing for Reconciliation**
Post-processing stage kicks in after reconciliation engine has completed processing of incoming data from the target. During this stage, OIG kernel orchestrations get triggered to execute event-handlers to do things like default password generation as per policy, role assignment, resource provisioning, audit processing and so on.

This section includes the following topics:

- [Target System And Connector Tuning](#)
- [Database Indexes For Recon Matching Rules](#)
- [Oracle Identity Governance Post-processing for Reconciliation](#)

Target System And Connector Tuning

This section describes the tuning that needs to be applied on your target systems as well as Oracle Identity Governance Connectors.

Oracle Internet Directory

During a reconciliation run, all changes in the target system records are reconciled into Oracle Identity Governance. Depending on the number of records to be reconciled, this process may require a large amount of time. In addition, if the connection breaks during reconciliation, then the process would take longer to complete. It is recommended that "paged reconciliation" is configured to optimize performance.

To configure paged reconciliation, you must specify a value for the `PageSize` attribute of the user reconciliation scheduled task. The default value of 100 for `PageSize` suits for most of the scenarios.

 **Note:**

OID LDAP Server (the target system in this case) v10.1.4 or later versions support the paged reconciliation related LDAP operations.

SAP

It is recommended that you use a reconciliation batch size of 100.

Active Directory (11.1.1.5.0 and 11.1.1.6.0 Connector)

- Performance improvement patch
 - If you are using Active Directory 11.1.1.5.0, make sure that you apply patch # 15916848. You can download the patch from [My Oracle Support](#). For patching instructions, refer to the Readme that is available with the patch.
 - If you are using Active Directory 11.1.1.6.0, download the patch # 15916848 from [My Oracle Support](#). Import only the `ReconAttributeMap.xml` that is provided as part of the patch, using the deployment manager. You can ignore `ActiveDirectory.Connector.dll` provided in the patch, as it is updated in the 11.1.1.6.0 version itself. For patching instructions, refer to the Readme that is available with the patch.

- Configuring the reconciliation engine to skip the ignore event API

The default behavior would be to first check to create a recon event or to ignore it for each of the user records returned by the connector. This process involves comparing the values of all the attributes of the user coming in from the connector against the values stored in the OIG database. To ignore this, open the lookup definition `Lookup.Configuration.ActiveDirectory` and add below entry.

- Code Key: Ignore Event Disabled
- Decode: true

 **Note:**

You must evaluate the pros and cons of disabling the ignore event API call before you make the above changes.

- Batching

If batching is used in the AD connector, then the result set needs to be sorted. Therefore, batching can be used when number of records to be reconciled is less than 10000. The recommended batch size is 500.

- **Paging**
 - When number of records to be reconciled is more than 10000, use the `Page Size Configuration` property present in `Lookup.Configuration.ActiveDirectory` and `Lookup.Configuration.ActiveDirectory.Trusted`.
 - If paging is configured to be used, then you must make sure that no value is specified for the scheduled task parameters - `Batch Size`, `Batch Start`, `Number of Batches`, `Sort By`, and `Sort Direction`.
 - Paging splits the entire result set of a query into smaller subsets called, appropriately enough, pages. In general, it is recommended to set this value to the maximum page size for simple searches. By setting the page size to the maximum value, you can minimize the network round trips necessary to retrieve each page, which tends to be more expensive operation for simple searches. If you specify a `PageSize` greater than the `MaxPageSize` of the target system, the Active Directory server ignores it and uses the `MaxPageSize` instead. No exception is generated in this case. In some cases, you might need to specify a smaller page size to avoid timeouts or overtaxing the server. Some queries are especially expensive. Therefore, limiting the number of results in a single page can help avoid this. For the Active Directory Connector, use the default value 1000 for the best performance.
- **Filters**

It is recommended to use `Filters` and provide the value for the `Search Base`, if a specific set of records is to be retrieved from the target. Filter provided in the scheduled task is converted into LDAP query. The filters help narrow down the search, making the searching and processing of the data quicker. For more information about the filters, refer to the Active Directory Connector Documentation.
- For the reconciliation in the forest topology, you can use connector for reconciling the data from the complete forest (via Global Catalog Server) or you can use the connector for reconciling the data from the specific domain or domain controller. It is recommended to use the second approach whenever the data from the specific data center is to be reconciled, instead of using first option with search base.

For example:

Assume that there are 10 data centers in the Active Directory forest namely DC1, DC2, ... , DC10. To reconcile data from an organization (tempOrg) which is present on DC2, you have use one of the following approaches:

1. Use Global Catalog and provide the DN of the organization in the Search Base.
2. Use DC2 and provide the DN of the organization in the Search Base.

It is recommended to use the second approach for better performance.

Database Indexes For Recon Matching Rules

Reconciliation uses matching algorithm to find if the user/account/role/organization for which the change is requested, already exists in OIG. The matching algorithm compares the data in set of columns in OIG with the data in target staging table columns. The columns that contain the matching rules are defined in the reconciliation profile and they are defined at run-time. To improve the performance of the matching operation, there must be correct indexes created on the matching rule columns.

To illustrate the recommended method of identifying the appropriate indexes, a sample Active Directory (AD) user profile present in the Meta Data Store (MDS) repository is taken as an example. This example covers the following:

- [Selecting Indexes For Trusted Source Reconciliation](#)
- [Selecting Indexes For Target Source Reconciliation](#)
- [Selecting Indexes For Target Source Reconciliation With Multi-Valued Data](#)

 **Note:**

Starting OIG 11g Release 2 (11.1.2.1.0), the indexes are automatically created in some cases where possible. It is still recommended to follow the below procedure and make sure that all of the indexes required for reconciliation matching rule are in place.

Selecting Indexes For Trusted Source Reconciliation

To select indexes based on the matching rule criteria in trusted source reconciliation, you must complete the following steps:

1. Open the Active Directory user profile file in a text editor. You can open Active Directory user profile using `Validate Recon Profile test` present in the diagnostic dashboard, or by using `Validate Recon Profile MBean` present in EM.
2. Search for `ownerMatchingRuleWhereClause` or `matchingRule` for all entities:

```
ownerMatchingRuleWhereClause =
(((UPPER(USR.USR_LOGIN)=UPPER(RA_ADUSER7.RECON_USERID5A729570)) OR
(UPPER(USR.USR_UDF_OBGUID)=UPPER(RA_ADUSER7.RECON_OBJECTGUID))))
```

3. After identifying the columns constituting the matching rule in the profile, create the indexes accordingly.

For example, following indexes are needed for matching rule in the above example.

Table 12-5 Table Names and Columns to be Indexed

Table Name	Column to be Indexed
USR	UPPER(USR_LOGIN)
USR	UPPER(USR.USR_UDF_OBGUID)
RA_ADUSER7	UPPER(RECON_USERID5A729570)
RA_ADUSER7	UPPER(RA_ADUSER7.RECON_OBJECTGUID)

 **Note:**

- It is important that the indexes are created along with functions like `UPPER`, `SUBSTR` in the matching rule. In [Table 12-5](#), `UPPER` is the function used on all columns.
- Some of the columns and functions might have been indexed already. In [Table 12-5](#), `USR` table should already have function-based index on `UPPER(USR_LOGIN)`.

Selecting Indexes For Target Source Reconciliation

To select indexes based on the matching rule criteria in target resource reconciliation, you must complete the following steps:

1. Open the Active Directory user profile file in a text editor. You can open Active Directory user profile using `Validate Recon Profile` test present in the diagnostic dashboard, or by using `Validate Recon profile` MBean present in EM.

2. Search for account search tag `<matchingruleWhereClause>`:

```
<matchingruleWhereClause>((UD_ADUSER.UD_ADUSER_OBJECTGUID=RA_ADUSER7.RECON_OBJECTGUID))</matchingruleWhereClause>
```

3. After identifying the columns constituting the matching rule in the profile, create the indexes accordingly.

For example, following indexes are needed for matching rule in the above example.

Table 12-6 Table Names and Columns to be Indexed

Table Name	Column to be Indexed
UD_ADUSER	UD_ADUSER_OBJECTGUID
RA_ADUSER7	RECON_OBJECTGUID

 **Note:**

- It is important that the indexes are created along with functions like `UPPER`, `SUBSTR` in the matching rule.
- Some of the columns and functions might have been indexed already.

Selecting Indexes For Target Source Reconciliation With Multi-Valued Data

To select indexes based on the matching rule criteria in target resource reconciliation with multi-valued data, you must complete the following steps:

1. Open the Active Directory user profile file in a text editor. You can open Active Directory user profile using `Validate Recon Profile` test present in the diagnostic dashboard, or by using `Validate Recon profile` MBean present in EM.

2. For entitlements, search for the `<matchingruleWhereClause>` tag under

```
<childreconevedata>
```

```
<matchingruleWhereClause>((UD_ADUSRC.UD_ADUSRC_GROUPNAME=RA_UD_ADUSRC.RECON_MEMBEROF))</matchingruleWhereClause>
```

3. After identifying the columns constituting the matching rule in the profile, create the indexes accordingly. For example, following indexes are needed for matching rule in the above example.

Table 12-7 Table Names and Columns to be Indexed

Table Name	Column to be Indexed
UD_ADUSRC	UD_ADUSRC_GROUPNAME
RA_UD_ADUSRC	RECON_MEMBEROF

 **Note:**

- It is important that the indexes are created along with functions like UPPER, SUBSTR in the matching rule.
- Some of the columns and functions might have been indexed already.

Oracle Identity Governance Post-processing for Reconciliation

Table 12-8 lists some of the important out-of-the-box event handlers that are invoked during post-processing of reconciliation.

Table 12-8 Event Handlers and Their Descriptions

Event Handler	Description
AccountReconAuditHandler	Responsible for Auditing account/target reconciliation changes
ReconScheduledTaskAccountHandler	Trigger workflows associated with account/target reconciliation
ReconScheduledTaskUserHandler	Trigger workflows associated with trusted reconciliation
ReconUserDisplayNameHandler	Generates custom display name for trusted reconciliation
ReconUserLoginHandler	Generates custom login during for reconciliation
ReconUserPasswordHandler	Generates custom passwords for trusted reconciliation
UserCreateLdapPostProcessHandler	Creates user in LDAP if LDAP synchronization is enabled
UserUpdateLdapPostProcessHandler	Updates user in LDAP if LDAP synchronization is enabled

You can find the rest of out-of-the-box and custom event handlers in DMS metric page of WebLogic Application Server. Use the following URL to go to the DMS metric page:

```
http://<servername>:<port>/dms
```

In this URL, `port` refers to the WebLogic Administration Server port. To log in, you must use the WebLogic admin credentials.

After you log into the DMS metric page, click on **OIG_EventHandler** to see the list of event handlers and their processing time metrics. You can use these metrics to identify event handlers that may need to be optimized.

Tuning LDAP Synchronization

Tuning performance in Oracle Identity Governance involves the following:

- [Increasing the Max Connection Pool for Oracle Identity Governance](#)

Increasing the Max Connection Pool for Oracle Identity Governance

To increase the max connection pool for Oracle Identity Governance:

1. Login to Oracle Identity System Administration.

2. On the left pane, under Configuration, click **IT Resource**. The Manage IT Resource page is displayed in a new window.
3. From the IT Resource Type list, select **Directory Server**, and then click **Search**.
4. For the Directory Server IT resource, click **Edit**. The Edit IT Resource Details and Parameters page is displayed.
5. Change the value of the following configuration parameters to 500:
 - Initial pool size: 500
 - Minimum pool size: 500
 - Maximum pool size: 500
6. Click **Update**.
 - [Increasing the LDAP Synchronization Batch Size](#)
 - [Setting Configuration Parameters in OVD](#)
 - [Setting Configuration Parameters in OID](#)
 - [Setting Configuration Parameters in Identity Virtualization Library \(libOVD\)](#)
 - [Setting Configuration Parameters in WebLogic Server and JDBC](#)

Increasing the LDAP Synchronization Batch Size

To increase the LDAP synchronization batch size, set the batch size of the following LDAP synchronization reconciliation scheduled jobs to 1000:

- LDAP User Create and Update Reconciliation
- LDAP Role Create and Update Reconciliation
- LDAP Role Hierarchy Reconciliation
- LDAP Role Membership Reconciliation

Setting Configuration Parameters in OVD

When LDAP synchronization with OVD configured for OID is enabled in Oracle Identity Governance, the configuration parameters in OVD, as listed in [Table 12-9](#), must be set:

Table 12-9 Configuration Parameters in OVD

Name	Parameter	Value
OVD general	Listeners - LDAP Endpoint	50
-	Listeners - LDAP SSL Endpoint	50
User Adapter	Max Pool Size	500
-	Operation Timeout	1500000
-	Max Pool Wait	1000
Changelog adapter	Max Pool Size	500
-	Operation Timeout	1500000

Setting Configuration Parameters in OID

When LDAP synchronization with OVD/OID is enabled in Oracle Identity Governance, the configuration parameters in OID, as listed in [Table 12-10](#), must be set:

Table 12-10 Configuration Parameters in OID

Name	Parameter	Value
Max Number of DB Connections	orclmaxcc	10
Number of Processes	orclserverprocs	2 - 4
Skip Referral Process	orclskiprefinsql	1
LDAP Connection Timeout	orclldapconntimeout	60
Enable MatchDN Processing	orclmatchdnenabled	0
Enable Entry Cache	orclcacheenabled	0

To modify the attributes in [Table 12-10](#), use the following syntax:

```
ldapmodify -h HOST_NAME -p PORT_NUMBER -D cn=orcladmin -w PASSWORD -v <<EOF
dn: cn=oid1,cn=osldapd,cn=subconfigsentry
```

Setting Configuration Parameters in Identity Virtualization Library (libOVD)

When LDAP synchronization with Identity Virtualization Library (libOVD) configured for OID is enabled in Oracle Identity Governance, the configuration parameters in Identity Virtualization Library (libOVD), as listed in [Table 12-11](#), must be set:



Note:

You can manage the Identity Virtualization Library (libOVD) tuning parameter configuration by using the WLST command.

Table 12-11 Configuration Parameters in Identity Virtualization Library (libOVD)

Name	Parameter	Value
User Adapter	Max Pool Size	500
User Adapter	Operation Timeout	1500000
User Adapter	Max Pool Wait	1000
Changelog adapter	Max Pool Size	500
Changelog adapter	Operation Timeout	1500000

 **See Also:**

Enabling Access Logging in Identity Virtualization Library (libOVD) in the *Oracle Fusion Middleware Integration Guide for Oracle Identity Management* for information about enabling access logging in Identity Virtualization Library (libOVD) to capture all requests and responses flowing through Identity Virtualization Library (libOVD), which can be very useful in triaging performance issues.

Setting Configuration Parameters in WebLogic Server and JDBC

For information about setting configuration parameters in Oracle WebLogic Server and JDBC, see [Tuning the Application Server for Oracle Identity Governance](#).

Tuning Order Audit Messages To Eliminate Slow SQL

While running extremely heavy load and the rate of processing by audit job is slower than the changes happening, it leads to accumulation. In order to increase the rate of processing and eliminate slow SQL during Issue Audit Message task, set **Order Audit Messages** to **NO**.

Sysadmin > Scheduler > Issue Audit Message > Order Audit Message = **No**

 **Note:**

It is recommended to always set the default value to **True**. However, during heavy loads it could lead to huge accumulation. In such cases, you can choose to turn off the ordering clause. When the ordering clause is turned off, there might be failures due to out of order processing and that will leave some audit entries unprocessed in **aud_jms** table. These failed entries can be processed again by the job updating failed column to 1 for all the failed **aud_jms** rows.

Part V

SOA Suite Components

The Oracle SOA Suite components need to be tuned for optimal performance.

This part covers how to tune Oracle SOA Suite components to improve performance.

Tuning information for B2B, Healthcare Integration, and adapters are documented in other documents. You can find how to tune for performance by using the links provided.

The SOA Suite components are documented in the following topics:

- [Tuning the SOA Infrastructure](#)
You can tune the SOA Infrastructure to optimize its performance in managing composites and their lifecycle, service engines, and binding components in Oracle WebLogic Server, by using Work Managers and other tuning parameters.
- [Tuning Oracle BPEL Process Manager](#)
You can tune Oracle Business Process Execution Language (BPEL) Process Manager properties to optimize its performance at the composite, fabric, application, and server levels.
- [Tuning Oracle Mediator](#)
You can tune Oracle Mediator to optimize its performance as the framework for mediation between various providers and consumers of services and events.
- [Tuning Oracle Managed File Transfer](#)
You can tune Managed File Transfer (MFT) to optimize its performance as the managed file gateway.
- [Tuning Oracle Business Rules](#)
You can tune Oracle Business Rules to optimize its performance in enabling automation of business rules and extraction of business rules from procedural logic, such as Java code or BPEL processes.
- [Tuning Oracle Business Process Management](#)
You can tune Oracle Business Process Management to optimize its performance in providing a seamless integration of all stages of the application development life cycle from design-time and implementation to runtime and application management.
- [Tuning Oracle Human Workflow](#)
You can tune Oracle Human Workflow to optimize its performance in handling various aspects of human interaction with a business process.
- [Tuning Oracle Business Activity Monitoring](#)
You can tune Oracle Business Activity Monitoring (BAM) to optimize its performance in monitoring business services and processes in the enterprise.
- [Tuning Oracle Service Bus](#)
You can tune Oracle Service Bus (OSB) to optimize its performance in providing connectivity, routing, mediation, management, and also some process orchestration capabilities between two or more applications.
- [Tuning Oracle Enterprise Scheduler Service](#)
You can tune Oracle Enterprise Scheduler Service (ESS) to optimize its performance in enabling scheduling and running jobs.

- [Tuning Oracle Business Intelligence Performance](#)
You can tune Oracle Business Intelligence to optimize its performance in collecting, presenting, and delivering data.

13

Tuning the SOA Infrastructure

You can tune the SOA Infrastructure to optimize its performance in managing composites and their lifecycle, service engines, and binding components in Oracle WebLogic Server, by using Work Managers and other tuning parameters.

- [About the SOA Infrastructure](#)
The SOA Infrastructure is a Jakarta EE-compliant application running on Oracle WebLogic Server.
- [Tuning SOA Work Managers](#)
You can perform a few simple checks and configurations to take advantage of Work Managers.
- [Tuning SOA Infrastructure Parameters](#)
Tuning SOA infrastructure parameters is important for optimal performance.
- [Using Advanced Tuning Options](#)
You can configure additional performance tuning settings for SOA for specific scenarios.
- [Advanced Tuning for Work Managers](#)
Work Managers are mapped to SOA projects and specific components, and you can use some advanced configuration options to fine-tune the Work Manager performance.

About the SOA Infrastructure

The SOA Infrastructure is a Jakarta EE-compliant application running on Oracle WebLogic Server.

The application manages composites and their lifecycle, service engines, and binding components. For more information, see Introduction to the SOA Infrastructure Application in *Administering Oracle SOA Suite and Oracle Business Process Management Suite*.

The information presented here does not cover any diagnostic tools or methodologies that are needed for a holistic approach, but addresses isolated tuning options for isolated symptoms. For information on monitoring the SOA Infrastructure performance to pinpoint problem areas, see Monitoring the SOA Infrastructure in *Administering Oracle SOA Suite and Oracle Business Process Management Suite*.

Tuning SOA Work Managers

You can perform a few simple checks and configurations to take advantage of Work Managers.

Beginning with Oracle SOA Suite 12c (12.2.1), Work Managers handle most SOA-related work threads. For more details on how Work Managers manage threads and self-tune, see Understanding Work Managers in *Administering Server Environments for Oracle WebLogic Server*.

Before you attempt to configure Work Managers, you should have a good understanding of your environment and be able to quantify the following:

- Volume of incoming requests that you need processed.

- Internal processing requirements, including any SLA expectations for transactions.
- An understanding of the processes you have that do not use Work Managers, such as the Event Delivery Network and most adapters.

Based on the information collected above, you can take advantage of the Work Managers' self-tuning feature.

- [Configuring Database Connections with the SOADatabase Property](#)
- [Configuring Work Managers with the SOAMaxThreadsConfig Attribute](#)

Configuring Database Connections with the `SOADatabase` Property

The `SOADatabase` property determines the total number of concurrent database connections that are available for your SOA processes. Because SOA processes use the database for most of their activities, this is a very important setting and can create a bottleneck if not appropriately configured.

To tune this setting, it is important to understand your database resources and consult your DBA.

To tune the `SOADatabase`, do the following:

1. Log in to the Oracle WebLogic Remote Console.
2. Select **Edit Tree > Services > Data Sources**.
3. On the **DataSource** configuration page, select **SOADatabase**.
4. Select the **Connection Pool** tab and scroll down to find the **Maximum Capacity** attribute.

The default for the **Maximum Capacity** attribute is 50. For most practical use cases, you should set this value to 300 to increase the size of the entire `SOADatabase` connection pool.

The `SOADatabase` setting is leveraged by the `SOAMaxThreadConfig` configuration that is explained in [Configuring Work Managers with the SOAMaxThreadsConfig Attribute](#). The `SOADatabase` attribute defines the total number of connections that are available to all Work Managers, while the `SOAMaxThreadConfig` attribute defines what percentage of those connections are available to certain categories of Work Managers.

Configuring Work Managers with the SOAMaxThreadsConfig Attribute

SOA composites are associated with a group of Work Managers that handles various components and functional areas. The `SOAMaxThreadsConfig` attribute determines the number of threads allowed for different groups of SOA Work Managers in a domain.

The number of threads allotted to handle incoming requests, internal processes, and other SOA processes are defined as percentages of the `SOADatabase` property that is explained in [Configuring Database Connections with the SOADatabase Property](#). The default percentage values and categories of the `SOAMaxThreadsConfig` attribute are listed in [Table 13-1](#).

Table 13-1 Thread distributions for Work Managers determined by SOAMaxThreadsConfig

Group	Description
incomingRequestsPercentage Default: 20%	This parameter determines the percentage of threads that your system allocates to Work Managers that process incoming client requests such as EDN. The parameter is used for requests like Facade invocation, WebService client, Direct/ADF/Rest, and BulkRecovery requests.
internalBufferPercentage Default: 30%	This parameter determines the percentage of threads distributed to other SOA functions, such as adapters. This parameter is also used for inbound adapters that are not part of workmanagers in 12.1.3.
internalProcessingPercentage Default: 50%	This parameter determines the percentage of threads that your system allocates to Work Managers for internal processes. This parameter is also used for handling all SOA backend processing services like message processing by BPELEngine, Mediator Error Handling and Parallel Processing, Resequencer and QuartzScheduler.

This attribute is defined at the domain level and applies to all the Work Managers under that domain. You can set this attribute by using the `SoaInfraConfig` MBean in the Fusion Middleware Control MBean Browser.

To access the attribute:

1. Log in to Fusion Middleware Control.
2. Select System MBean Browser from the WebLogic Domain menu.
3. In the System MBean Browser folder structure, navigate through the following folders: **Application Defined MBeans** --> **oracle.as.soainfra.config** --> **Server: AdminServerName** --> **SoaInfraConfig** --> **soa-infra**.
4. When you click on **soa-infra**, its attributes are listed in the main pane on the right. Look for the `SOAMaxThreadsConfig` attribute and click it. You should then see the parameters and values listed in [Table 13-1](#).

When you are ready to make your changes, click **Apply**.

Remember that the values you are adjusting on this screen are percentages, not the discrete number of threads. You should ascertain the total number of threads available to you by checking the value of the `SOADataSource` property, which is described in [Configuring Database Connections with the SOADataSource Property](#).

In a sample scenario, where the `SOADataSource` attribute is set to 50 connections and if you kept the default `SOAMaxThreadConfig` percentages that are listed in [Table 13-1](#), you would have the following thread allocations:

- 20% of 50 = 10 threads to process incoming request
- 30% of 50 = 15 threads for processes not using work managers
- 50% of 50 = 25 threads to process internal processes

Tuning SOA Infrastructure Parameters

Tuning SOA infrastructure parameters is important for optimal performance.

Table 13-2 describes the optimal settings for parameters with the greatest impact on SOA Infrastructure performance.

Table 13-2 Essential SOA Infrastructure Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
AuditLevel Default: Production	<ul style="list-style-type: none"> High database CPU Contentions causing increased processing times in applications 	<p>To prevent possible performance degradation, decrease the audit level to “off”. Set the default value of <code>Production</code> only for audit purposes.</p> <p>This parameter can be set in the Enterprise Manager. You can find the <code>Audit Level</code> parameter page on the SOA Infrastructure Common Properties page.</p> <p>To find this page:</p> <ol style="list-style-type: none"> 1. Toggle the SOA folder in your left-hand Target Navigation. 2. Right-click on the <code>soa-infra (soa_server)</code> you want to tune. 3. Select SOA Administration --> Common Properties <p>For more information about this parameter, see <i>Configuring Oracle SOA Suite and Oracle BPM Suite Profiles in Administering Oracle SOA Suite and Oracle Business Process Management Suite</i>.</p>	<p>Keeping the default audit level will generate audit data to be captured in the database and hence cause database growth. Users should use the audit information for debugging errors.</p>

Table 13-2 (Cont.) Essential SOA Infrastructure Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
Audit Purge Policy Default: Everyday Midnight and purges records older than 7 days	<ul style="list-style-type: none"> Exponential growth in database size If configured at peak hours, purging can take resources from other processes 	<ul style="list-style-type: none"> Ensure that auto purge is enabled. Perform purges more often. Set the auto purge to kick off at a time when there is less resource contention from other processes. <p>For information on finding the Auto Purge page in the Oracle Enterprise Manager Fusion Middleware Control, see Deleting Large Numbers of Instances with Oracle Enterprise Manager Fusion Middleware Control in <i>Administering Oracle SOA Suite and Oracle Business Process Management Suite</i>.</p>	Disabling this feature makes maintaining on-going database growth more time-consuming.

Using Advanced Tuning Options

You can configure additional performance tuning settings for SOA for specific scenarios.

These options are presented here in no specific order. Before you change any of these properties, you should have a holistic knowledge of your environment, SOA processes, and non-SOA processes.

It is important to understand that any advanced performance optimization should be a customized approach for individual scenarios, settings, environments, and expectations. A customized approach requires detailed capturing of diagnostic information to pinpoint and isolate bottlenecks and areas that need optimization.

For information on monitoring the SOA Infrastructure performance to pinpoint problem areas, see Monitoring the SOA Infrastructure in *Administering Oracle SOA Suite and Oracle Business Process Management Suite*.

- [Using Composite Lazy Loading](#)
- [Changing Modularity Profiles](#)
- [Tuning Your Database for SOA Processes](#)
- [Tuning Event Delivery Network Parameters](#)
- [Tuning the WebLogic Server](#)

Using Composite Lazy Loading

Composite lazy loading is a new feature in 12c. It improves server startup time when there is a large number of composites deployed.

At server startup, composites are loaded minimally, meaning that they only create in-memory java models and MBeans. Any initializing tasks, such as loading components and resources used by composite, namely WSLD and Schema file, are loaded later at first-request time when they are needed.

This greatly improves server startup times and staggers the composite startup times for when they receive requests, reducing overhead from rarely used or retired composites.

Composite lazy loading is helpful for:

- Scenarios that require speedy disaster recovery times during a server failure
- Customers with a huge number of composites that use large WSDLs or schema files

Composite lazy loading is enabled by default and can be configured at the domain level and at the composite levels.

- [Configuring Composite Lazy Loading for the Domain Level](#)
- [Configuring Composite Lazy Loading at the Component Level](#)

Configuring Composite Lazy Loading for the Domain Level

Composite lazy loading is enabled by default at the domain level. This setting can be disabled from the System MBean Browser in Enterprise Manager for Fusion Middleware Control. Changes to this setting takes affect when the server restarts.

To change the setting for lazy loading feature for the domain level:

1. After you log into Enterprise Manager, right-click the domain that you want to tune from the list of the WebLogic domains in the **Target Navigation** browser.
2. Select **System MBean Browser** from the drop-down menu.
3. In the System MBean Browser folder structure, navigate through the following folders: **Application Defined MBeans --> oracle.as.soainfra.config --> Server: AdminServerName --> SoainfraConfig --> soa-infra**.
4. When you click **soa-infra**, its attributes are listed in the main pane on the right. Look for the `CompositeLazyLoading` attribute and click it.
5. On the `CompositeLazyLoading` page, you can set the value to `true` to enable it or `false` to disable it. When you are ready to make your changes, click **Apply**.

Configuring Composite Lazy Loading at the Component Level

By default, composites inherit the lazy loading setting from the domain level. If there is a use case where you would like to control this behavior at specific composite level, then this can be configured in the `composite.xml` file, which is a file that is generated when you create a new SOA Suite composite application.

You can find the `composite.xml` file in the home folder of the application that you want to edit. You can also edit the `composite.xml` file by accessing it in JDeveloper. For more information on the `composite.xml` file, see *What Happens When You Create a SOA Application and Project in Developing SOA Applications with Oracle SOA Suite*.

At the beginning of the `composite.xml` file of the application that you want to edit, you need to add the new property `lazyLoading="false"` to override the default behavior at the domain level. Then redeploy the composite.

Below is a sample code snippet:

```
<?xml version="1.0" encoding="UTF-8" ?>
<!-- Generated by Oracle SOA Modeler version 12.2.1.0.0 at [8/7/13 4:14 PM]. -->
<composite name="ValidatePayment"
  revision="1.0"
  label="2013-08-07_16-14-11_843"
  mode="active"
  state="on"
  lazyLoading="false"
  xmlns="http://xmlns.oracle.com/sca/1.0"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  .....
  .....
</composite>
```

Changing Modularity Profiles

Modularity is another 12c feature that helps improve your memory footprint and server startup times. Some profile options are limited to only components and features that are used by your selected composites. The modularity profile you select determines what components are loaded in memory.

12c has ready-to-use profiles that can be changed after you complete installation. By default, new 12c customers have `SOA_FOUNDATION` as their install profile. Existing customers upgrading to 12c have `SOA_CLASSIC` as their install profile by default.

[Table 13-3](#) shows the modularity profiles in the increasing order of memory footprint size.

Table 13-3 Modularity Profiles

Profile	Components
BPEL-ONLY	BPEL Components + SOA Common Infrastructure + Partial Adapter set
ORCHESTRATION	BPEL-Only + HWF + Partial Adapter set
SOA FOUNDATION Default for new 12c customers	Orchestration + Mediator + Rules + Partial Adapter set
SOA FOUNDATION ENTERPRISE	SOA Foundation + Full Adapter Set
SOA FOUNDATION WITH B2B	SOA Foundation Enterprise + B2B
SOA FOUNDATION WITH HEALTHCARE	SOA Foundation with B2B + Healthcare UI
SOA CLASSIC Default for upgrade customers	SOA Foundation with B2B + BPM Modules

If you are using a limited set of components or features in the SOA suite, you can change your profile to optimize your memory usage and server startup times. This can free up resources for crucial processes and can improve disaster recovery.

You can change your modularity profile from the SOA dashboard in Enterprise Manager for Fusion Middleware Control.

See, *Configuring SOA Infrastructure Properties* in *Administering Oracle SOA Suite and Oracle Business Process Management Suite* to find the **SOA Infrastructure Common Properties** page.

Then, see *Configuring Oracle SOA Suite and Oracle BPM Suite Profiles* in *Administering Oracle SOA Suite and Oracle Business Process Management Suite* for more information on the profiles.

Tuning Your Database for SOA Processes

If needed, you can adopt advanced strategies for tuning your database for SOA processes. Make sure you have already read and followed the general database tuning suggestions covered in [Tuning Database Parameters](#) of this book before you progress.

- [Collecting Optimizer Statistics](#)
- [Tuning Temporary Tablespaces for SOA](#)
- [Minimizing SOA Database Contention](#)
- [Purging](#)
- [Reclaiming Space](#)

Collecting Optimizer Statistics

Optimizer statistics provide details about the database and the objects in the database. The query optimizer uses these statistics to choose the best execution plan for each SQL statement. See Introduction to the Query Optimizer in *Oracle Database SQL Tuning Guide* for more information.

- [Gathering Statistics Automatically](#)
- [Gathering Statistics Manually](#)
- [Optimizing the MDS Database Repository With Statistics](#)

Gathering Statistics Automatically

Because objects in a database can change constantly, you must update statistics regularly so that they accurately describe these objects.

All SOA databases should use the Automatic Statistics Collection, which is enabled by default. This job runs every night. See Controlling Automatic Optimizer Statistics Collection in *Oracle Database SQL Tuning Guide* for more information.

Gathering Statistics Manually

Automatic optimizer statistics collection is sufficient for most database objects, but in a database that is close to going live or for tables that are modified or purged significantly, manual statistic gathering is needed. See Gathering Optimizer Statistics Manually in *Oracle Database SQL Tuning Guide* for more information.

For SOA databases that implement purging of stale data on regular basis, you should collect stats manually right after purging has completed. In these cases, use the `DBMS_STATS.GATHER_TABLE_STATS` procedure. See `DBMS_STATS` in *Oracle Database PL/SQL Packages and Types Reference* for how to do this.

Optimizing the MDS Database Repository With Statistics

Ensure that automatic statistics collection is enabled. See Controlling Automatic Optimizer Statistics Collection in *Oracle Database SQL Tuning Guide* for more information.

In most cases, the first 32 characters of `PATH_FULLNAME` in the `MDS_PATHS` table are the same. You can prevent the database from putting them in the same section of the histogram by doing the following:

1. Drop the histogram for `PATH_FULLNAME` column by executing a command structured like the following as system:

```
execute dbms_stats.delete_column_stats(ownname=>'mdsSchemaOwner',  
tabname=>'MDS_PATHS', colname=>'PATH_FULLNAME', col_stat_type=> 'HISTOGRAM');
```

2. Set table preferences to exclude collecting histogram for the `PATH_FULLNAME` column with a command structured like the following:

```
execute dbms_stats.set_table_prefs(mdsSchemaOwner, 'MDS_PATHS', 'METHOD_OPT', 'FOR  
COLUMNS SIZE 1 PATH_FULLNAME');
```

Tuning Temporary Tablespaces for SOA

See [Tuning Database Files](#) for general guidelines on tuning `TEMP` tablespaces for Oracle Fusion Middleware before you progress to this topic.

Some SOA queries can generate a large amount of disk sorts that require high amounts of temporary space. Therefore, the use of multiple temporary tablespaces and tablespace groups is recommended to meet these requirements and assure optimal performance.

The suggested minimum size for the `TEMP` tablespace or tablespace group that is assigned to the SOA schema owner is 6 GB with auto-extend enabled. See *Changing Data File Size in Oracle Database Administrator's Guide* for more information on how to resize a tablespace and enable auto-extend.

Minimizing SOA Database Contention

Most SOA workloads generate heavy DML activity in the database and are likely to experience contention on database objects.

Wait event data in Automatic Workload Repository (AWR) reports reveal various symptoms that might impact performance. The most common wait events that could occur in SOA database are as follows:

- DB CPU
- Db file sequential read, db file scattered read
- log file sync
- enq: HW - contention
- enq: TX - index contention
- buffer busy waits
- gc buffer busy acquire, gc buffer busy release (RAC)
- enq: SQ - contention
- [Tuning the Redo Log Performance \(log file sync\)](#)
- [Migrating BasicFiles to SecureFiles \(enq:HW - contention\)](#)
- [Creating Hash Partitioned Indexes \(enq: TX - index contention\)](#)

Tuning the Redo Log Performance (log file sync)

In a SOA database, it is very common to see the foreground wait event `log file sync` with a high average wait time. This is caused by the redo log performance. The possible reasons for high log file sync waits are as follows:

- The database log writer (LGWR) is unable to complete writes fast enough for one of the following reasons:
 - Disk I/O performance to log files is not good enough.
 - LGWR is starving for CPU resources.
- LGWR is unable to post the processes fast enough due to excessive commits.
- LGWR is suffering from other database contentions, such as enqueue waits or latch contention.

Tuning the redo log performance can improve the performance for applications that run in an Oracle Fusion Middleware environment.

See [Tuning Database Files](#) for general guidelines on tuning redo logs for Oracle Fusion Middleware before using the strategies here to tune for SOA processes.

Finding LGWR wait events

The first step in identifying the root cause is to find and break down LGWR wait events. You can query for LGWR wait events by using its SID, as shown in the following example:

```
SQL> SELECT sid, event, time_waited, time_waited_micro
       FROM v$session_event
       WHERE sid IN
          (SELECT SID FROM v$session WHERE type!='USER' AND program LIKE '%LGWR%' )
       ORDER BY time_waited;
```

Sizing Online Redo Logs to Control the Frequency of Log Switches and Minimize System Waits

The suggested minimum setting for redo logs is to have at least 3 log groups of 2 GB each. Monitor the redo log performance periodically. Then adjust the number of redo log groups and size of each member as appropriate to control the frequency of log switches and minimize system waits.

Size the redo log files according to the amount of redos that the system generates. A rough guide is to switch logs at most once every 20 minutes.

For example, if your online redo logs switches once every 5 minutes during peak database activity, the logs would each need to be 4 times larger than their current size to achieve the 20 minute guideline. The calculation for this is $20\text{min} / 5\text{min} = 4x$.

Optimizing the Redo Log Disk to Prevent Bottlenecks

A SOA database is highly write-intensive, which generates massive amount of redo per second and per transaction. Sometimes no amount of disk tuning may relieve redo log bottlenecks, because Oracle must push all updates for all disks into a single redo location.

If I/O bandwidth is an issue, doing anything other than improving I/O bandwidth is not useful. One way to relieve redo bottlenecks is to use faster redo storage. It is recommended to use Solid State Disk (SSD) redo log files. SSD has greater bandwidth than platter disk.

Determining the Optimal Sizing of the log_buffer

SOA applications insert, modify, and delete large volumes of data. Most of these operations are committed in a row-by-row fashion rather than in batch mode. Frequent commits cause a significant overhead on the redo performance, so sizing the log_buffer optimally is important for performance.

The statistic `REDO_BUFFER_ALLOCATION_RETRIES` from your AWR reports and/or from `V$` views reflects the number of times a user process waits for space in the redo log buffer. You can obtain this statistic through the dynamic performance view `V$SYSSTAT` with the following query:

```
SELECT NAME, VALUE
FROM V$SYSSTAT
WHERE NAME = 'redo buffer allocation retries';
```

The value of redo buffer allocation retries should be near zero over an interval. If this value increments consistently, then processes have had to wait for space in the redo log buffer. The wait can be caused by the log buffer being too small or by check pointing. You can improve this wait by attempting the following:

- Increase the size of the redo log buffer, if necessary, by changing the value of the initialization parameter `LOG_BUFFER`. The value of this parameter is expressed in bytes. A good starting rule of thumb for a write intensive workload is to configure the log buffer to 100mb. Use caution while increasing `log_buffer` setting, because excessive redo size can also cause high `log file sync` waits.
- Improve the check pointing or archiving process.

You can also check to see if the log buffer space wait event is a significant factor in the wait time for the instance. If not, the log buffer size is most likely adequate.

Tuning the LGWR Process

For most SOA workloads, the commit rate is very high, and decreasing commits is not an option. If previous strategies to address high log file sync did not improve redo log performance, try increasing the priority of LGWR or increasing the priority class of LGWR to RT from the command line.

Using Smart Flash Logging for ExaData

If your database is on ExaData machine, it should have a minimum of Bundle Patch 11 (BP11) installed to take advantage of the Smart Flash Logging feature.

Exadata Smart Flash Logging is an additional feature that is implemented in Exadata Storage software 11.2.2.4.2 and database version 11.2.0.2 + BP11. With this feature, 512 MB of flash storage is reserved for redo writes and the LGRW process adopts a different pattern of behavior.

In a system which does not use this feature, LGWR writes in parallel to multiplexed copies of the redo logs and then waits for all writes to complete. This means that the time taken to perform these writes (indicated by the Oracle wait interface statistics `log file parallel write`) is the time taken for the slowest disk to complete the write.

With Exadata Smart Flash Logging, the redo log files remain on disk, but the additional reserved 512 MB of space is created on flash storage. When issuing a write call, LGWR writes to the redo logs on disk as usual but also makes a parallel write to the flash area. LGWR then waits for whichever of these writes completes first to post it, after which it continues without waiting for the other.

Migrating BasicFiles to SecureFiles (enq:HW - contention)

The High Water enqueue contention (enq: HW - contention) occurs when competing processes are inserting into the same table and are trying to increase the high water mark of a table simultaneously.

In a SOA database, this issue is experienced by tables that have large object (LOB) columns, such as `CUBE_SCOPE`, `XML_DOCUMENT`, `AUDIT_DETAILS`, and so on. Under a heavy load, LOB

segments in these tables experience contention, which is seen in an AWR report as the wait event enq: HW contention.

The default storage for LOBs in an Oracle database is BasicFiles. Frequently allocating extents or reclaiming chunks may cause contention for the LOB segment high water marks. This contention can also occur for LOB segments that are ASSM-managed, since space allocation only acquires one block at a time.

This contention can be eliminated by switching LOB storage from BasicFiles to SecureFiles. SecureFiles is an LOB storage architecture that provides performance benefits over traditional BasicFiles. See *About LOB Storage in Database SecureFiles and Large Objects Developer's Guide* for more information on these two architectures.

Migrating BasicFiles to SecureFiles can be done by using one of the following methods:

- Set the database parameter `SECURE_FILES = ALWAYS`.

This method is applicable for new installations prior to creating SOA tables by using RCU. Once this parameter is set at the instance level, any new LOB segments created uses SecureFiles automatically.

- Use the online redefinition method.

This method is applicable for installations that already have SOA tables created in them. In such cases, LOB segments from tables in an SOA database experiencing enq: HW contention can be migrated to SecureFiles.

Using the online redefinition method to migrate to SecureFiles can be done with very little downtime.

- Set the database event value to 44951 by using the following script:

```
ALTER SYSTEM SET EVENT='44951 TRACE NAME CONTEXT FOREVER, LEVEL 1024' scope=spfile;
```

This method helps SOA installations using an Oracle version older than 11g to avoid enq: HW contention on LOB segments.

You can use your AWR and Automatic Database Diagnostic Monitor (ADDM) reports to identify which LOB objects are suffering from enq:HW - contention. For most systems, however, it is highly recommended to move the LOB columns listed in the following table to SecureFiles.

Table 13-4 LOB Storage Attributes

Table Name	Column Name	Recommended LOB Storage Attributes
ATTACHMENT	ATTACHMENT	COMPRESS CACHE
AUDIT_DETAILS	BIN	COMPRESS CACHE
CUBE_SCOPE	SCOPE_BIN	COMPRESS CACHE

Creating Hash Partitioned Indexes (enq: TX - index contention)

In most SOA scenarios, multiple database sessions insert thousands of rows into SOA tables. In these situations, the number of index keys is constantly increasing, particularly the primary key indexes.

Though the number of primary key indexes increases over time, B-tree structure indexes only target a few database blocks for key insertions. These B-tree index insertions can become problematic in a Real Application Cluster (RAC). This issue is seen in an AWR report as high `buffer busy waits`.

B-tree indexes create other contentions for RAC environments that show in an AWR as `gc buffer busy acquire` and `gc buffer busy release` wait events. These occur when a transaction inserting a row in an index has to wait for the end of a different transaction's index block split, forcing the session to wait as well. When many concurrent inserts lead to excessive index block splits, performance decreases.

The solution for these contentions is to create global, hash partitioned indexes. This forces a random distribution of index keys across many database blocks to avoid these contentions or hot spots.

Hash partitioning has proven to be the best tuning method to address index contention. You should use your AWR and ADDM reports to identify indexes that need to be partitioned. Once you have identified hot indexes, consider hash partitioning them to reduce or avoid index contention.

Purging

The need for aggressive and continuous purging is a key aspect to improving performance and controlling disk space in SOA.

Managing the auto purge feature, enabled by default to help manage on-going database growth is described in [Table 13-2](#). SOA installations that accumulate a lot of data should also implement a purging strategy to clean up redundant data, to help the SQL query performance, and to save disk space.

To create a purging strategy, see *Developing a Purging and Partitioning Methodology in Administering Oracle SOA Suite and Oracle Business Process Management Suite*.

Reclaiming Space

SOA installations that implement frequent purging of unwanted data from SOA tables are more likely to experience disk space issues.

This problem occurs even with ASSM and locally managed tablespaces. When automatic purge scripts delete rows from database tables and indexes to release space within the data blocks for reuse, space is not released immediately after the rows are deleted. This causes fragmentation, with some space too small for reuse, particularly when the tables contain LOB columns.

To alleviate fragmentation and consolidate disk space, you should manually shrink tables and LOB columns to reclaim space on a routine basis.

Use the Segment Advisor to identify segments that would benefit from online segment shrink. Note that most SOA segments should be candidates for online segment shrink operations after constant purging. See *Using the Segment Advisor in Oracle Database Administrator's Guide* for more information on how to use the Segment Advisor.

Once you have identified the database tables and indexes that need shrinking, use the following commands to reclaim space manually:

```
ALTER TABLE CUBE_SCOPE ENABLE ROW MOVEMENT;  
ALTER TABLE CUBE_SCOPE SHRINK SPACE;  
ALTER TABLE CUBE_SCOPE MODIFY LOB (SCOPE_BIN) (SHRINK SPACE);  
ALTER TABLE CUBE_SCOPE DISABLE ROW MOVEMENT;
```

This shrink operation consolidates free space below the high water mark and compacts the segment. Then it moves the high water mark and deallocates space above the high water mark.

Tuning Event Delivery Network Parameters

The Event Delivery Network (EDN) delivers events published by Oracle Mediator, Oracle BPEL Process Manager, and external publishers such as Oracle Application Development Framework entity objects. See Introduction to the Event Delivery Network and JMS Provider Types in *Administering Oracle SOA Suite and Oracle Business Process Management Suite* for a more detailed description.

[Table 13-5](#) lists parameters that you can find in the Fusion Middleware MBean Browser and tune for improved event delivery.

Table 13-5 Event Delivery Network Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
numberOfPollerThreads Default: -1	<ul style="list-style-type: none"> Out-of-resource issues, for example, out of memory, system overload, transaction issue, and so on. Contention with other SOA threads 	<p>The default value of -1 means that the system uses <code>ThreadsPerSubscriber</code> to determine a poller thread count. This is optimal for most configurations.</p> <p>However, if you have a high number of subscribers, the default setting tries to assign a thread to each subscriber. This slows your system down. You should define a positive integer to limit the amount of poller threads created for this task.</p> <p>See <i>Updating the Local numberOfPollerThreads Value at the Service Component Level</i> in <i>Administering Oracle SOA Suite and Oracle Business Process Management Suite</i> for information on how to change this parameter's value in the Fusion Middleware MBean Browser.</p>	<p>If the value is too low for your system, then poller threads can cause event backlogs and long latencies between event publishing and composite instance creation.</p> <p>If the value is too high, then excess poller threads consume the server's resources needlessly.</p>

Table 13-5 (Cont.) Event Delivery Network Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
ThreadsPerSubscriber Default: 1 thread	<ul style="list-style-type: none"> Out-of-resource issues, for example, out of memory, system overload, transaction issue, and so on. Contention with other SOA threads 	<p>Typically, the default of 1 thread per subscriber is optimal.</p> <p>Note that <code>numberOfPollerThreads</code> should be adjusted first, since that parameter takes precedence over this value.</p> <p>See <i>Updating the ThreadsPerSubscriber Attribute in the System MBean Browser in Administering Oracle SOA Suite and Oracle Business Process Management Suite</i> for information on how to change this parameter's value in the Fusion Middleware MBean Browser.</p>	<p>If the value is too low for your system, then poller threads can cause event backlogs and long latencies between event publishing and composite instance creation.</p> <p>If the value is too high, then excess poller threads consume the server's resources needlessly.</p>

[Table 13-6](#) lists the parameters that you can modify for individual business events in JDeveloper. To modify these attributes, right-click the event that you want to edit to bring up the pop-up menu. From this menu, select **Edit Subscribed Events...** or **Edit Published Events...**, depending on the parameter that you are trying to edit.

For descriptions of the subscribed event parameters you can edit, see *How to Subscribe to a Business Event in Developing SOA Applications with Oracle SOA Suite*.

Table 13-6 Business Event Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
Consistency for a Subscribed Event Default: <code>oneAndOnlyOne</code>	<p>You are experiencing either one or both problems with business event delivery.</p> <ul style="list-style-type: none"> Unfulfilled delivery guarantee requirements to event subscribers Unnecessary system overhead from global transactions 	<p>Set the level for a selected business event to <code>guaranteed</code> in JDeveloper. A <code>guaranteed</code> delivery is performed in a local transaction with only one trip to the main queue.</p> <p>You can also edit this parameter on the Subscriptions page in the Oracle Enterprise Manager Fusion Middleware Control. See <i>Viewing Business Event Subscribers in Administering Oracle SOA Suite and Oracle Business Process Management Suite</i> for details.</p>	<p>The <code>oneAndOnlyOne</code> parameter guarantees delivery by taxing resources.</p> <p>If a <code>guaranteed</code> delivery fails, then there are no local retries and a system failure message is generated. Message duplication could occur in the event that the calling global transaction rolls back and retries since the message delivery is outside of that transaction.</p>

Table 13-6 (Cont.) Business Event Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
Durability for a Subscribed Event Default: Yes	You are experiencing either one or both problems with business event messages. <ul style="list-style-type: none"> Multiple dropped events Unnecessary retention of messages in the system 	Set the value under the <code>Durable</code> column to <code>No</code> to disable durability for a subscribed event by using <code>JDeveloper</code> . This frees the system from having to persist messages to storage.	If the subscriber is not running when events are published, setting the value to <code>No</code> causes the system to drop events. Setting the value to <code>Yes</code> retains events in the JMS server and incurs overhead.
Persistent Delivery for a Published Event Default: yes	<ul style="list-style-type: none"> Unreliable messaging High overhead 	Set this value to <code>No</code> to disable persistent delivery. This reduces overhead.	Setting the value to <code>No</code> causes less reliable messaging following an event publish since there is no persistence. Setting the value to <code>Yes</code> incurs overhead by guarding against a JMS server crash.

Table 13-6 (Cont.) Business Event Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
Time-to-live for a Published Event Default: 0 ms	<ul style="list-style-type: none"> Non-expired and unconsumed messages are occupying system resources and requiring manual cleanup. Messages are deleted before subscribers can read them. 	Specify a positive integer so that expired messages are automatically removed from the system and not consumed by the subscribers. The integer represents milliseconds. The best value depends on your system and can be determined by monitoring metrics.	If the message expiration duration value is too low, published messages can expire before an intended subscriber can read it. Once it is gone, it cannot be retrieved. If the value is too high, then lingering messages can occupy system resources.



N

o

t

e

:

T

h

e

d

e

f

a

u

l

t

v

a

l

u

e

o

f

0

m

e

a

n

s

t

h

a

t

m

e

s

s

a

Table 13-6 (Cont.) Business Event Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
			g e s n e v e r e x p i r e s .

- [Adding JMS Topics with Mapping](#)

Adding JMS Topics with Mapping

By default, all events are mapped to a single WLS topic.

If you have a large backlog of events or are experiencing latency or slowness in event processing due to single or limited JMS topics, you should create additional JMS topics and modify events to JMS mapping so that event types of different performance characteristics may be grouped or managed separately.

However, if you do this, the system will have additional JMS topics and JMS artifacts to manage, and you will have mapping changes to consider.

- [Choosing a JMS Topic Type](#)
- [Creating JMS Topics](#)
- [Mapping Events to JMS Topics](#)

Choosing a JMS Topic Type

You can create either a `WLSJMS` topic or an `AQJMS` topic.

`WLSJMS` is the default JMS topic type. It does not provide database indexing, LOB streaming, embedded rules engines, and lock management as well as `AQJMS`.

`AQJMS` is typically not faster than `WLSJMS`, but if your system has high concurrences, `AQJMS` works well because it is single-threaded. `AQJMS` can also get constrained by lower and storage nodes in Exalogic.

Creating JMS Topics

You can create a new `WLSJMS` topic under the `SOAJMSModule` in the WebLogic Remote Console if you are logged in as an Administrator. See *Create a JMS System Module* in the *Oracle*

WebLogic Remote Console Online Help for details on navigating to the **Create a New JMS System Module Resource** and creating a JMS topic.

You can create an AQJMS topic by using the Database Navigator in JDeveloper or SQL Developer as `soainfra` user by running the following script:

```
define edn_user=your_soainfra_schema_username
define topic=your_custom_aqjms_topic_name, e.g. 'EDN_AQJMS_TOPIC_2'
define topic_table=your_custom_aqjms_topic_table, e.g. 'EDN_AQJMS_TOPIC_TABLE_2'

begin
  DBMS_AQADM.stop_queue(queue_name => '&edn_user..&topic');
  DBMS_AQADM.drop_queue(queue_name => '&edn_user..&topic');
  DBMS_AQADM.drop_queue_table(queue_table => '&edn_user..&topic_table');
end;
/
begin
  dbms_aqadm.create_queue_table(queue_table => '&edn_user..&topic_table',
                                queue_payload_type => 'SYS.AQ$_JMS_MESSAGE',
                                multiple_consumers => true);
  dbms_aqadm.create_queue(queue_name => '&edn_user..&topic',
                          queue_table => '&edn_user..&topic_table',
                          max_retries => 256);
  dbms_aqadm.start_queue(queue_name => '&edn_user..&topic');
end;
/
commit;
```

You can reference *Create a JMS Queue or Topic* in *Administering JMS Resources for Oracle WebLogic Server* for information about AQ JMS topics.

Mapping Events to JMS Topics

When you have created new JMS topics, you can map business events to specific topics. Note that one event type can be mapped to only one JMS topic, whereas one JMS topic can store multiple event types.

For more information on using the Enterprise Manager for Fusion Middleware Control to map events, see *Mapping Business Events to JMS Topic Destinations* on the Business Events Page in *Administering Oracle SOA Suite and Oracle Business Process Management Suite*.

Tuning the WebLogic Server

The performance of the SOA Infrastructure depends on the WebLogic Server. Though tuning the WebLogic Server is a separate task not thoroughly addressed in this book, you can use [Table 13-7](#) to check the tuning knobs that affect the SOA Infrastructure.

Table 13-7 Essential WebLogic Server Tuning for SOA Infrastructure

Parameter	Tuning Recommendation	Resource
<p>ProductionModeEnabled</p> <p>Default: The mode you set during domain creation.</p>	<p>Production mode maximizes performance. You should enable this if you are not developing applications. You can enable the <code>ProductionModeEnabled</code> MBean in Oracle Fusion Middleware Control.</p>	<p>See Configure General Settings in <i>Administering Oracle WebLogic Server with Fusion Middleware Control</i>.</p> <p>Changing the domain mode also changes certain security and autodeployment settings. See Development vs. Production Mode Default Tuning Values in <i>Tuning Performance of Oracle WebLogic Server</i> for more information on domain modes.</p>
<p>WebLogic Server Logging Levels</p> <p>Default: Notification</p>	<p>To reduce the volume of logging requests, use the lowest acceptable logging level, such as <code>ERROR</code> or <code>WARNING</code> whenever possible. You can set log levels for handlers and loggers in a variety of ways.</p>	<p>See Using Log Severity Levels in <i>Configuring Log Files and Filtering Log Messages for Oracle WebLogic Server</i> for these methods.</p>
<p>HTTP Access Logging</p> <p>Default: Enabled</p>	<p>By default, the HTTP subsystem keeps a log of all HTTP transactions in a text file. Turn off HTTP access logging to improve performance. You can disable this property by using the Oracle WebLogic Server Remote Console.</p>	<p>See Configure HTTP in the <i>Oracle WebLogic Remote Console Online Help</i>.</p>
<p>JMS Persistence and Persistence Storage</p> <p>Default: Enabled</p>	<p>Ensure that the right persistence level is set for the JMS services destinations.</p> <ul style="list-style-type: none"> For persistent JMS scenarios, there are two choices: <code>File Store</code> and <code>JDBC Store</code>. 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 non-persistent JMS scenarios, turn off persistence at the JMS server level by un-checking the <code>Store Enabled</code> flag from the Advanced section of the General tab for the JMS server in the WebLogic Remote Console. You can also override the persistence mode at the JMS destination level. 	<p>See Using Custom File Stores and JDBC Stores in <i>Tuning Performance of Oracle WebLogic Server</i> for more information on creating and managing persistent JMS stores.</p>
<p>Connection Backlog Buffering</p>	<p>You can tune the <code>Accept Backlog</code> parameter when dealing with a large number of concurrent clients.</p> <p>The <code>Accept Backlog</code> parameter specifies how many TCP connections can be buffered in a wait queue. You can tune the number of connection requests that a WebLogic Server instance accepts before refusing additional requests.</p>	<p>For more information, see Tuning Connection Backlog Buffering in <i>Tuning Performance of Oracle WebLogic Server</i>.</p>

Advanced Tuning for Work Managers

Work Managers are mapped to SOA projects and specific components, and you can use some advanced configuration options to fine-tune the Work Manager performance.

When SOA Suite is installed, it creates a set of default Work Managers, global Work Managers, and application Work Managers to manage various areas of the SOA Infrastructure.

High priority composites can be associated with a Work Manager group that has been configured for higher priority. [Table 13-8](#) lists the set of Work Managers that are created when SOA is installed and describes the work area they manage.

Table 13-8 Work Manager Descriptions

Work Manager Name	Responsible Area
SOA_Request_WM	SOA synchronous request clients, such as the following: <ul style="list-style-type: none"> • Facade invocation • WebService client requests • Direct/ADF/Rest requests • B2B
SOA_Notification_WM	All SOA notification requests.
<i>WorkManagerName_dspSystem</i>	BPEL-specific system dispatcher messages.
<i>WorkManagerName_dspInvoke</i>	BPEL-specific engine process invocation dispatcher messages.
<i>WorkManagerName_dspEngine</i>	BPEL engine process dispatcher messages.
<i>WorkManagerName_dspNonBlocking</i>	BPEL engine process non-blocking invocation dispatcher messages.
<i>WorkManagerName_Analytics</i>	BPEL analytics.
<i>WorkManagerName_MediatorParallelRouting</i>	Mediator parallel routing.
<i>WorkManagerName_MediatorErrorHandling</i>	Mediator error handling.
<i>WorkManagerName_bpmnSystem</i>	BPM system dispatcher messages.
<i>WorkManagerName__bpmnInvoke</i>	BPM engine process invocation dispatcher messages.
<i>WorkManagerName__bpmnEngine</i>	BPM process engine dispatcher messages.
<i>WorkManagerName__bpmnNonBlocking</i>	BPM process non-blocking invocation dispatcher messages.
SOA_DataSourceBound_WM	All SOA backend processing services that access <code>SOADataSource</code> , including Workflow Jakarta Enterprise Beans (EJBs).
SOA_Default_WM	All SOA services that do not access the <code>SOADataSource</code> connection pool. It also handles Case Management.
SOA_EDN_WM	Event Delivery Network (EDN).
<i>WorkManagerName_Adapter</i>	Adapter framework.

The `SOAMaxThreadsConfig` property, discussed in [Configuring Work Managers with the SOAMaxThreadsConfig Attribute](#), determines the number of connections that are used by Work Managers to process incoming requests, internal processes, and other processes. This configuration determines the optimal usage for each of these processing categories when the system is functioning at its full potential.

Minimum and Maximum Constraints can also be set on Work Managers to control upper and lower limit of connections for Work Managers. A Fair Share Request class for a Work Manager can be created to determine the relative priority assigned to a Work Manager. The constraints and request class mentioned here are the ones most commonly configured for SOA Work Managers.

All SOA Work Managers are preconfigured with request classes and constraints that make most sense. It is strongly recommended to run with the default settings and make any essential changes after an evaluation period.

For information on all Work Manager constraints and request classes you can create and their default behaviors, refer to Managing Work Manager Groups in *Administering Oracle SOA Suite and Oracle Business Process Management Suite*.

- [Configuring Fair Share Request Class for SOA Work Managers](#)
- [Creating a New Work Manager Constraint](#)

Configuring Fair Share Request Class for SOA Work Managers

A Fair Share Request Class allows you to specify the relative priority of a given Work Manager. All SOA Work Managers managing internal process have been configured to one of the two Fair Share Classes that are created by default: `soa_fairShare_20` and `soa_fairShare_80`, with fair share values set to 20 and 80 respectively. A Fair Share value is a relative value from 1 to 1000.

If you want to further tune SOA Work Manager priorities, you need to create new Fair Share classes. For more information on how to do this, see Viewing and Creating Work Manager Groups in *Administering Oracle SOA Suite and Oracle Business Process Management Suite*.

Creating a New Work Manager Constraint

In addition to the default categories available in the `SOAMaxThreadConfig` property, you can create new categories to address specific scenarios.

Some processes in SOA do not require database connections. These processes do not depend on SOA Data Source allocation and hence do not have to wait for available connections.

The SOA Infrastructure automatically creates Work Managers that manage most of your processes and allocate resources accordingly. For most cases, performance can be improved by leveraging existing Work Managers and tuning their performance by using some of the knobs described above.

If you have special scenarios where you would like to handle uniquely, you can create a new Work Manager and configure it to meet special circumstances. You will be either creating a new application or a web application Work Manager. See Viewing and Creating Work Manager Groups in *Administering Oracle SOA Suite and Oracle Business Process Management Suite* for detailed procedures.

Tuning Oracle BPEL Process Manager

You can tune Oracle Business Process Execution Language (BPEL) Process Manager properties to optimize its performance at the composite, fabric, application, and server levels.

- [About BPEL Process Manager](#)
Oracle BPEL Process Manager offers a comprehensive and easy-to-use infrastructure for creating, deploying, and managing BPEL business processes.
- [Tuning BPEL Parameters](#)
You can tune BPEL parameters for optimal performance.
- [Using Other Tuning Strategies](#)
You can locate the Oracle BPEL Process Manager tables that are impacted by instance data growth and purge them for optimal performance.

About BPEL Process Manager

Oracle BPEL Process Manager offers a comprehensive and easy-to-use infrastructure for creating, deploying, and managing BPEL business processes.

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.

For an overview of Oracle BPEL Process Manager, see Oracle Business Process Execution Language (BPEL) Process Manager under Key Components in *Understanding Oracle SOA Suite*.

Tuning BPEL Parameters

You can tune BPEL parameters for optimal performance.

Tuning recommendations for BPEL parameters described here are *likely* or *highly likely* to improve performance. For descriptions of the other tuning parameters available for SOA Components, see the component-specific topics in this guide.

For detailed information on how to monitor, configure, and manage BPEL process service components and service engines, see *Administering BPEL Process Service Components and Engines in Administering Oracle SOA Suite and Oracle Business Process Management Suite*. Also see *Using the BPEL Process Service Component in Developing SOA Applications with Oracle SOA Suite* for how to use sensors to monitor select BPEL activities.

- [Tuning BPEL Engine](#)
- [Tuning BPEL in a Composite](#)

Tuning BPEL Engine

You can configure the performance tuning properties at the BPEL engine level by using the Enterprise Manager Fusion Middleware Control. For information on using Oracle Enterprise Manager Fusion Middleware Control to configure and monitor parameters, see *Getting Started with Administering Oracle SOA Suite and Oracle BPM Suite and Accessing the System MBean*

Browser from the Component Property Pages in *Administering Oracle SOA Suite and Oracle Business Process Management Suite*.

- [Tuning BPEL Engine Parameters](#)

Tuning BPEL Engine Parameters

[Table 14-1](#) lists the essential tuning parameter that you can adjust to improve performance for the BPEL engine.

Table 14-1 Essential BPEL Engine Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
auditLevel Default: Inherit	You are experiencing low performance because of frequent database inserts into the audit_trail table.	Use the Off value to stop storing audit information. Note that the auditLevel is set at the SOA Infrastructure level. See <i>Configuring BPEL Process Service Engine Properties</i> in <i>Administering Oracle SOA Suite and Oracle Business Process Management Suite</i> to see how to find and tune this parameter.	This property sets the audit trail logging level for both durable and transient processes. If you turn this off, both business flow and payload tracking is disabled. You cannot view the state of BPEL processes in the Oracle Enterprise Manager Console.

[Table 14-2](#) describes additional BPEL engine parameters that can be tuned for small performance improvements. Note that for most use cases, the default value is the recommended value.

Table 14-2 Other BPEL Engine Tuning Knobs

Parameter	Description
SyncMaxWaitTime Default: 45 seconds.	You can decrease this parameter's value to improve performance. The SyncMaxWaitTime property sets the maximum time the process result receiver waits for a result before returning. This property is required for synchronous interactions and is applicable to transient processes. See <i>How To Specify Transaction Timeout Values</i> in <i>Developing SOA Applications with Oracle SOA Suite</i> for instructions on how to find this property in the System MBean Browser of Oracle Enterprise Manager Fusion Middleware Control.

Table 14-2 (Cont.) Other BPEL Engine Tuning Knobs

Parameter	Description
<p>largedocumentthreshold</p> <p>Default: 10000 (100 kilobytes).</p>	<p>You can decrease this parameter's value to improve performance.</p> <p>This property sets 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. It is applicable to both durable and transient processes.</p> <p>Large XML documents can slow down the performance if they are constantly used while processing an instance.</p> <p>See <i>Configuring BPEL Process Service Engine Properties in Administering Oracle SOA Suite and Oracle Business Process Management Suite</i> to see how to find and tune this parameter in the Enterprise Manager Fusion Middleware Control.</p>
<p>validateXML</p> <p>Default: False.</p>	<p>You should set this parameter to the default value of <code>False</code> to improve performance.</p> <p>This property can make the Oracle BPEL Process Manager intercept nonschema-compliant payload data by validating incoming and outgoing XML documents. However, XML payload validation can slow performance.</p> <p>You can find this parameter in the System MBean Browser. See <i>Configuring BPEL Process Service Engine Properties in Administering Oracle SOA Suite and Oracle Business Process Management Suite</i> for information on how to find advanced BPEL properties by using the More BPEL Configuration Properties... button from the BPEL Service Engine Properties page in Enterprise Manager Fusion Middleware Control.</p>
<p>InstanceKeyBlockSize</p> <p>Default: 10000 keys.</p>	<p>You can increase the instance key block size to a value greater than the number of updates to the <code>ci_id_range</code> table to improve performance.</p> <p>The <code>InstanceKeyBlockSize</code> property controls the instance ID range size. Oracle BPEL Server creates instance keys (a range of process instance IDs) in batches by using the value specified. After creating this range of in-memory IDs, the next range is updated and saved in the <code>ci_id_range</code> table.</p> <p>See <i>Configuring BPEL Process Service Engine Properties in Administering Oracle SOA Suite and Oracle Business Process Management Suite</i> to see how to find and tune this parameter by using the System MBean Browser in Enterprise Manager Fusion Middleware Control.</p>
<p>Audit Level Threshold</p> <p>Default: 10000.</p>	<p>You can decrease this parameter's value to improve performance.</p> <p>This property sets the maximum size (in kilobytes) of an audit trail details string before it is stored separately from the audit trail. Strings larger than the threshold setting are stored in the <code>audit_details</code> table instead of the <code>audit_trail</code> table. In cases where the variable is very large, performance can be severely impacted by logging it to the audit trail.</p> <p>See <i>Configuring BPEL Process Service Engine Properties in Administering Oracle SOA Suite and Oracle Business Process Management Suite</i> to see how to find and tune this parameter in Enterprise Manager Fusion Middleware Control.</p>

Tuning BPEL in a Composite

You can tune BPEL properties for individual composites to improve performance. 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.

BPEL composite properties can be modified in the `composite.xml` file by using JDeveloper, or in the System MBean Browser of Oracle Enterprise Manager Fusion Middleware Control. For in-depth descriptions of each property's function, see Deployment Descriptor Properties in *Developing SOA Applications with Oracle SOA Suite*.

The BPEL tuning considerations listed in [Table 14-3](#) 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. See How to Define Deployment Descriptor Properties in the Property Inspector in *Developing SOA Applications with Oracle SOA Suite* for information on how to find and edit the parameters listed below.

Table 14-3 Essential BPEL in a Composite Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
<code>OneWayDeliveryPolicy</code> Default: <code>async.persist</code>	Slow performance because resources are being used to persist delivery messages.	Set value to <code>async.cache</code> . Incoming delivery messages for durable processes are kept only in the in-memory cache. By default, incoming requests are saved in the delivery service database table <code>dlv_message</code> .	This setting has a high risk of losing messages or overloading the system. It also changes the threading model for adapter.
<code>Audit Policy</code> Default: All activities	Slow performance because every activity is being audited.	Audit only key activities.	Lower level activities do not have an audit trail.
<code>inMemoryOptimization</code> Default: False	Slow performance because the <code>completionPersistPolicy</code> parameter has been activated at the BPEL component level, causing the BPEL server to dehydrate either all or some instances.	Set value to <code>False</code> to tell the Oracle BPEL Server that this process is a transient process and dehydration is not required.	No dehydration means that activities in the instance are lost if the system crashes.

[Table 14-4](#) describes additional BPEL parameters that can be tuned for small performance improvements, but in most cases, the default value is the recommended value. For in-depth descriptions of each property's function, see Properties for the partnerLinkBinding Deployment Descriptors in *Developing SOA Applications with Oracle SOA Suite*.

Table 14-4 Other BPEL in a Composite Tuning Knobs

Parameter	Description
idempotent Default: True	An idempotent activity is an activity that can be retried. Keeping this parameter's value as <code>True</code> allows idempotent activities by preventing the BPEL server from dehydrating immediately after a failed activity. This parameter is configured in a partner link at runtime in BPEL.
validateXML Default: False	<code>False</code> means that the system does not validate all XML messages during a receive activity. This parameter is configured in a partner link at runtime in BPEL.

Using Other Tuning Strategies

You can locate the Oracle BPEL Process Manager tables that are impacted by instance data growth and purge them for optimal performance.

- [Identifying Tables Impacted By Instance Data Growth](#)

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.

You can use [Table 14-5](#) to locate tables that may be affected by instance data growth. See *Monitoring Space Usage, Hardware Resources, and Database Performance in Administering Oracle SOA Suite and Oracle Business Process Management Suite* for advice on how to monitor performance for the following database tables:

Table 14-5 Oracle BPEL Process Manager Tables Impacted by Instance Data Growth

Table Name	Table Description
audit_trail	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.
audit_details	Stores audit details that can be logged through the API. Activities such as an assign activity logs the variables as audit details by default. Audit details are separated from the <code>audit_trail</code> 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 <code>audit_trail</code> 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).

Table 14-5 (Cont.) Oracle BPEL Process Manager Tables Impacted by Instance Data Growth

Table Name	Table Description
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 the 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.

When you have determined which tables are causing slow performance, you can purge them. See Understanding Growth Management Challenges and Testing Strategies in *Administering Oracle SOA Suite and Oracle Business Process Management Suite* for more information on managing database growth.

Tuning Oracle Mediator

You can tune Oracle Mediator to optimize its performance as the framework for mediation between various providers and consumers of services and events.

- [About Oracle Mediator](#)
Mediator is a component of the Oracle SOA Suite 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.
- [Tuning Mediator Parameters](#)
You can tune the Oracle Mediator properties to improve performance if necessary.
- [Using Resequencer for Messages](#)
A Resequencer is used to rearrange a stream of related but out-of-sequence messages back into order.

About Oracle Mediator

Mediator is a component of the Oracle SOA Suite 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 Jakarta EE application.



Note:

For details about the SOA Suite, see *Developing SOA Applications with Oracle SOA Suite*.

For details about Oracle Mediator, see Administering Oracle Mediator Service Components and Engines in *Administering Oracle SOA Suite and Oracle Business Process Management Suite*.

Tuning Mediator Parameters

You can tune the Oracle Mediator properties to improve performance if necessary.

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 topics under Administering Oracle Mediator Service Components and Engines in *Administering Oracle SOA Suite and Oracle Business Process Management Suite*.

Table 15-1 describes the parameter values that can be tuned for performance. Note that the need to tune Mediator to improve performance is unlikely.

Table 15-1 Essential Mediator Tuning Knobs

Parameter	Tuning Recommendation
DeferredMaxRowsRetrieved Default: 20 rows	Increase the default value to retrieve more deferred processing messages from the DB in one iteration. Note that in Mediator, this parameter is only used with parallel routing rules.
DeferredLockerThreadSleep Default: 2 seconds	If deferred messages constitute a small percentage of total messages, increase the default value to perform fewer trips to the DB to retrieve deferred messages. Some use case scenarios can benefit from an idle time of 3600 seconds (60 minutes).
metricsLevel Default: enabled	If you do not need to collect DMS metrics data, disabling this parameter can improve performance.

For more information about each parameter, see Configuring Oracle Mediator Service Components and Engines in *Administering Oracle SOA Suite and Oracle Business Process Management Suite*.

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 sends them to the target services in an orderly manner.

Table 15-2 lists the tunable parameters for Resequencer in Mediator. You can tune the following parameters by accessing the Mediator Service Engine Properties page or the System MBean Browser by using one of the methods described under Configuring Oracle Mediator Service Engine Properties in *Administering Oracle SOA Suite and Oracle Business Process Management Suite*.

Table 15-2 Essential Tuning Knobs for Resequencer in Mediator

Parameter	Tuning Recommendation
ResequencerMaxGroupsLocked Default: 4 rows	Increase the default value to lock more Resequencer groups from the database in one iteration.

Table 15-2 (Cont.) Essential Tuning Knobs for Resequencer in Mediator

Parameter	Tuning Recommendation
ResequencerLockerThreadSleep Default: 10 seconds	If resequencer groups constitute a small percentage of total groups and messages, increase the default value to perform fewer trips to the database to lock resequencer groups.
DeleteMessageAfterComplete Default: True	Set the value to <code>True</code> to delete message after successful execution. For a high load use case, this results in more database space. Changing the default value to <code>False</code> retains the resequenced messages in the resequencer database. This slows down the resequencer database queries, which in turn degrades the performance.

See Configuring Resequenced Messages in *Administering Oracle SOA Suite and Oracle Business Process Management Suite*.

16

Tuning Oracle Managed File Transfer

You can tune Managed File Transfer (MFT) to optimize its performance as the managed file gateway.

- [About Managed File Transfer](#)
Oracle Managed File Transfer (MFT) is a high performance, standards-based, end-to-end managed file gateway.
- [Tuning MFT Parameters](#)
You can tune MFT parameters to optimize performance.

About Managed File Transfer

Oracle Managed File Transfer (MFT) is a high performance, standards-based, end-to-end managed file gateway.

It features design, deployment, and monitoring of file transfers using a lightweight web-based design-time console that includes file encryption, scheduling, and embedded FTP and SFTP servers.

For more information about Managed File Transfer, see Understanding Oracle Managed File Transfer in *Using Oracle Managed File Transfer*.

Tuning MFT Parameters

You can tune MFT parameters to optimize performance.

[Table 16-1](#) lists and describes parameters that you likely need to tune to improve MFT performance. To diagnose problem areas in MFT, see Monitoring Oracle Managed File Transfer and Administering Oracle Managed File Transfer in *Using Oracle Managed File Transfer*.

Table 16-1 Essential MFT Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
Processor count Default: 2 for each type of processor	JMS messages are accumulating in message processing queues.	Increase the processor count for the queues where messages are accumulating. The optimal value depends on the meta data and incoming payload. You can calculate the optimal processor count by using DMS metrics. To enable DMS metrics, add the MBean property <code>enablePerformanceMetric</code> . To disable metrics later, set the value to <code>False</code> .	Having more processors requires more system resources for concurrent processing.

Table 16-1 (Cont.) Essential MFT Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
Maximum Concurrent Request and Max Logins settings for Embedded FTP/SFTP server Default: 10	<ul style="list-style-type: none"> Multiple connection requests in waiting status The message Too many users logged in, user will be disconnected occurs in the embedded server log file 	<p>Increase the maximum number of concurrent requests and maximum number of logins for embedded FTP/SFTP server.</p> <p>You can increase the count so long as performance continues to scale linearly.</p> <p>If the embedded server service (FTP/SFTP) is not being used, then disable this setting.</p>	Increased count requires more system resources for concurrent processing.
LDAP Max Pool Default: 10	Number of concurrent connections to the LDAP consistently reaches max limit.	<p>Increase count.</p> <p>Because LDAP is a shared resource for all deployed applications in WebLogic server, you should monitor LDAP connections and adjust this value accordingly.</p>	Increased count requires more system resources.
Max connections to MFTDataSource Default: 50	Number of concurrent connection to the data source consistently reaches max limit.	<p>Increase the connection count so long as performance continues to scale linearly.</p> <p>Optimal value can be determined based on the number of processors, listening source threading model, and max concurrent request settings of embedded servers.</p>	Increased count requires more system resources.
Generating checksum setting Default: Enabled	Overall MFT message processing is slow.	Disable this parameter if checksum validation for delivered payloads from MFT is not necessary.	Generating checksum is a time consuming operation.
Regular purge Default: Disabled	<ul style="list-style-type: none"> Disk space is approaching the maximum limit. Table space used by MFT tables reaches the max table space allotted. 	Run purge to free disk or table space.	Historical information or data is discarded.

Table 16-2 describes the tuning properties that do not regularly need to be tuned. Keeping their default values is recommended, so you can check these parameters see if their values have been changed.

Table 16-2 MFT Parameters with Low or Medium Importance

Parameter	Problem	Tuning Recommendation	Trade-offs
Processing function or callout usage recommendation for broadcasting use cases Default: <code>Target Level</code>	Associated processing function or callout is executed for each target that degrades the performance.	For broadcasting use cases, associate processing functions or callouts at source level instead of target level as much as possible.	None.
Sub-folder count MBean setting Default: 256	Degraded disk performance caused by MFT switching among a high number of sub-folders to store files.	Reduce the sub-folder count.	Reducing the number of sub-folders increases the number of files stored in each sub-folder. If the volume of incoming files is high, the number of the files inside a single sub-folder degrades performance.
Store Inline payload setting Default: <code>File System</code>	Slow performance because accessing inline payload for Web Service sources from the disk takes too much time.	Store inline payload in the database rather than the file system.	The table size used by MFT increases as inline payloads are stored in the database.
Always Save Modified Files setting at the target level Default: <code>False</code>	If you have changed this setting to <code>True</code> for auditing purposes, you will have increased the disk space usage.	The default value of <code>False</code> reduces disk space usage.	No audit information is available. Note that a target level resubmit does not work if there was any pre-processing associated with the target.
<code>minFileSizeForProgressMonitor</code> Default: 10 MB	Frequent updates about byte transfer.	Specify a minimum file size so that the transfer progress screen appears for larger files only.	For files smaller than the minimum specified, the file transfer progress is not displayed.
<code>progressMonitorTimeToCommit MBean</code> Default: 4 seconds	Frequent updates about byte transfer.	Specify a minimum file size so that the transfer progress screen appears for larger files only.	Database updates on bytes transferred for ongoing file transfers are slower.
<code>MaxMdsSessionCacheCount</code> Default: 100	Out-of-memory exceptions caused by MDS cache memory footprint.	Decrease this value.	Decreasing this will decrease the performance of the overall MFT message processing because accessing data from the cache is faster.

- [Tuning Remote FTP / SFTP/ FILE Type Sources](#)
- [Minimizing MDS label](#)
- [Adjusting the Materialized Views Refresh Interval](#)

Tuning Remote FTP / SFTP/ FILE Type Sources

If MFT is not able to pick up files even after the polling frequency is expired, you need to tune the remote FTP/SFTP/FILE type sources. MFT uses the JCA Adapters underneath for all these source types. Refer to the SOA adapter recommendations listed under Oracle JCA Adapter Framework Performance and Tuning in *Understanding Technology Adapters*.

[Table 16-3](#) lists the properties.

Table 16-3 Tuning Remote FTP/SFTP/FILE Type Source

Parameter	Problem	Tuning Recommendation	Trade-offs
ThreadCount Default: -1	A high priority endpoint is downloading files slowly because of insufficient threads in the global pool.	Specify a value greater than 0. This creates a dedicated thread pool for a given endpoint to download files.	A very high value may result in lots of threads assigned to one endpoint, which can lead to lower overall performance.
SingleThreaded Default: False	In rare cases, you may not want to use global threads or allocate a separate thread pool for a low-priority endpoint.	Set value to True.	If set to true, it can result in a delay in downloading files from the endpoint as now there is a single thread for polling as well as downloading new files.

Minimizing MDS label

Artifact deployment results in creation of new MDS labels. More MDS labels increases the memory footprint and time to retrieve the metadata.

In general, users should follow these best practices for deployments:

- Minimize frequent deployments and meta data creations.
- Use bulk deployment for WLST commands.
- Make all changes for metadata and deploy them at once.

Adjusting the Materialized Views Refresh Interval

Materialized views refresh every 1 minute. If there is a heavy load on the database server, you may want to increase the refresh frequency from 1 minute.

You can view data from materialized views on the MFT console. If a high load is observed on the database server, this refresh frequency can be adjusted by using the following command:

```
ALTER MATERIALIZED VIEW <<MV_NAME>> REFRESH NEXT <<REFRESH_INTERVAL>>;
```

The materialized views used by MFT are:

- MV_MFT_PAYLOAD_INFO
- MV_MFT_SOURCE_INFO
- MV_MFT_SOURCE_MESSAGE
- MV_MFT_TARGET_INFO
- MV_MFT_TRANSFER
- MV_MFT_TRANSFER_COUNT_INFO

Tuning Oracle Business Rules

You can tune Oracle Business Rules to optimize its performance in enabling automation of business rules and extraction of business rules from procedural logic, such as Java code or BPEL processes.

- [About Oracle Business Rules](#)
Oracle Business Rules provides an easy-to-use authoring environment as well as a very high-performance inference-capable rules engine.
- [Tuning Oracle Business Rules](#)
You can tune Oracle Business Rules to optimise performance.

About Oracle Business Rules

Oracle Business Rules provides an 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 is a core component of many Oracle products including both middleware and applications.

See *Designing Business Rules with Oracle Business Process Management* and *Getting Started with Oracle Business Rules in Developing SOA Applications with Oracle SOA Suite*.

Tuning Oracle Business Rules

You can tune Oracle Business Rules to optimise performance.

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.

Table 17-1 Essential Business Rules Tuning Strategies

Strategy	Description	Recommendation
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 the 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 are asserted as facts (and used in the rules).
Assert child facts instead of multiple dereferences	Expressions like <code>Account.Contact.Address</code> 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 the <code>assertXPath</code> method or other means to assert object hierarchy.

Table 17-1 (Cont.) Essential Business Rules Tuning Strategies

Strategy	Description	Recommendation
Avoid side effects in rule conditions	The tests in a rule condition may be evaluated a greater or lesser number of times than would occur in a procedural program.	Methods or functions, which have side effects such as changing a value or state should not be used in a rule condition. If a method or function has side effects, those side effects may be performed an unexpected number of times.
Avoid expensive operations in rule conditions	Expensive operations would include any operation that involves I/O (disk or network) or even intensive computations. These operations should be done externally to the rules engine.	Expensive operations should be avoided in rule conditions. In general, consider avoiding I/O or DBMS access from the rules engine directly. 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.
Consider pattern ordering	Reordering rule patterns can improve the performance of rule evaluation in time, memory use, or both. Finding the optimal order for your system requires some experimentation.	If a fact is not expected to change or does not change frequently during rule evaluation, order the fact clauses by the expected rate of change from least to greatest. 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, order the fact clauses from most restrictive (matches fewest facts) to least restrictive.
Consider the ordering of tests in rule conditions	Proper ordering can reduce the amount of computation required for facts that do not satisfy the rule condition.	The tests in a rule condition should be ordered so that a more restrictive test occurs before a less restrictive test. If the degree of restrictiveness is not known, or estimated to be equal for a collection of tests, then simpler tests should be placed before more expensive tests.

- [Exerting assertXPath Support](#)

Exerting assertXPath Support

The `assertXPath` method asserts the whole hierarchy in one call, but also asserts some XLink facts for children facts to link back to parent facts. Though very convenient, it may have a performance impact.

To improve the performance of the `assertXPath` method, select the `Enable improved assertXPath support for performanceCheck` box in the Dictionary Properties page in Rule Author. Taking advantage of this requires that the following conditions are met:

- The `assertXPath` method is only invoked with an XPath expression of `"/*"`. Any other XPath expression results in an `RLIllegalArgumentException`.

- XLink facts should not be used in rule conditions as the XLink facts are not asserted.

If XLink facts for children facts 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 the `assertXPath` method, 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());
}
```

Tuning Oracle Business Process Management

You can tune Oracle Business Process Management to optimize its performance in providing a seamless integration of all stages of the application development life cycle from design-time and implementation to runtime and application management.

- [About Oracle Business Process Management](#)
The Oracle Business Process Management Suite provides an integrated environment for developing, administering, and using business applications centered around business processes.
- [Tuning Business Process Management Parameters](#)
You can tune BPM performance parameters in the Enterprise manager, through the SOA Administration in BPMN properties.
- [Using Other Tuning Strategies](#)
You can consider using the following strategies to further improve performance.

About Oracle Business Process Management

The Oracle Business Process Management Suite provides an integrated environment for developing, administering, and using business applications centered around business processes.

Oracle Business Process Management 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.

See, *Oracle Fusion Middleware User's Guide for Oracle Business Process Management*.

For more details on tuning Oracle Business Process Management with your other Oracle Fusion Middleware components, see .

Tuning Business Process Management Parameters

You can tune BPM performance parameters in the Enterprise manager, through the SOA Administration in BPMN properties.

To tune the performance of the Oracle Business Process Management engine, you can reduce resource demands to reduce latency.

To reduce resource demands, you can tune the parameters listed in [Table 18-1](#):

Table 18-1 Essential Business Process Management Tuning to Reduce Resource Demands

Parameter	Problem	Tuning Recommendation	Trade-offs
largedocumentthreshold Default: 10000 (100 kilobytes)	Instances are being processed slowly because you are storing large BPMN Data objects.	Decrease the maximum size (in kilobytes) of this parameter to limit the size of BPMN Data Objects. If they surpass this limit, they are stored in a separate location from the rest of the instance scope data. This property is applicable to both durable and transient processes.	The overflow data is stored in an external append-only table. This adds to overall database size and can increase the overall workload when loading instances from the database.
auditLevel Default: Inherit from Infrastructure	You are seeing frequent database inserts into the audit_trail table. These are caused by audit events being logged by a process.	Reduce or disable audit. You can switch to any of the following settings: <ul style="list-style-type: none"> • Off to log no events or audit events • Minimal to log only events • Error to log only serious problems You can also consider expanding the size of the AuditKeyExtents.	You lose granular error reporting that you could use to diagnose problems later. Always choose the audit level according to your business requirements and use cases. For more information on how to use audit trails for monitoring, see <i>Monitoring BPMN Process Service Components and Engines in Administering Oracle SOA Suite and Oracle Business Process Management Suite</i> .

You can also try to purge completed instances as allowed by business requirements and add indexes for any flex fields.

Using Other Tuning Strategies

You can consider using the following strategies to further improve performance.

- [Tuning Oracle Workspace Applications](#)
- [Tuning Process Measurement](#)

Tuning Oracle Workspace Applications

Database performance and session state management are the primary drivers for performance. Effective database tuning and configuration of HTTP session timeout are important.

Application design is the next largest factor, especially if there are additional data controls used to render contextual data on task forms. In these cases, it is important to optimize data access from those data controls and when possible defer retrieving additional data unless it is needed. For more details on tuning ADF, see [Oracle ADF Faces Configuration and Profiling](#).

The following parameters can be changed in the `web.xml` descriptor in the `OracleBPMWorkspace` web application. Once they have been modified, you may have to redeploy.

Table 18-2 Workspace and Worklist Application Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
HTTP Session Timeout Default: 15 minutes	Memory is being allocated for users who may no longer be actively using the system.	<p>To better manage resource usage, decrease the session timeout value, in minutes, to the smallest value that preserves the expected user experience. This allows the system to reclaim any resources that are associated with unused sessions as soon as possible.</p> <p>This parameter is edited in the in the <code>web.xml</code> file. The following is a sample snippet of the <code>web.xml</code> file:</p> <pre><session-config> <session-timeout> 5 </session-timeout> </session-config></pre>	A short timeout value may mean users have to login more often if they let the time expire. They also may potentially lose session data.
ADF Client State Token Default: 15	The default value may consume too much memory.	<p>Decrease the value to 3 to minimize the memory footprint.</p> <p>Through this setting, you can control the number of pages users can navigate by using the browser Back button without losing information. To reduce CPU and memory usage, you can decrease the value in the <code>web.xml</code> file.</p> <p>The following is a sample snippet of the <code>web.xml</code> file:</p> <pre><context-param> <param-name> org.apache.myfaces.trinid ad.CLIENT_STATE_MAX_TOKEN S </param-name> <param-value> 3 </param-value> </context-param></pre>	<p>If the user clicks the Back button more than 3 times, there is no session data stored for that page.</p> <p>If the value is too small, users get an error when they click the Back button.</p>

Table 18-2 (Cont.) Workspace and Worklist Application Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
Compress_View_State Token Default: True	Slow performance on slower or higher latency networks.	Set this value to True to enable zipping. By default, this value is set to True . This setting controls whether the page state is compressed. Zipping greatly reduces the memory being taken up by page state in the session object. The following is a snippet of the <code>web.xml</code> file: <pre><param- name>org.apache.myfaces.t rinidad.COMPRESS_VIEW_STA TE</param-name> <param-value>true</ param-value></pre>	There is an additional CPU cost to zipping and unzipping the view state.
DISABLE_CONTENT_COMPRES SION Default: False	Slow initial load of pages.	In production environments, make sure you remove the DISABLE_CONTENT_COMPRES SION parameter from the web.xml file or set it to FALSE. By default, style classes that are rendered are compressed to reduce page size. The following is a snippet of the <code>web.xml</code> file: <pre><param- name>org.apache.myfaces.t rinidad.DISABLE_CONTENT_C OMPRESSSION</param-name> <param-value>false</ param-value></pre>	None.

Tuning Process Measurement

Process Analytics uses measurement events to sample the process and publish measurements to registered consumers. In 19c (19.1.0.0.0), these measurements can be enabled by setting the `DisableAnalytics` parameter to `False` in the BPM Enterprise Manager's Analytics Configuration MBean.

The two supported consumers for measurements in 12c are BAM 11g Monitor Express and BAM 12c Process Metrics. They can be enabled or disabled by using the `DisableProcessMetrics` and `DisableMonitorExpress` attributes of the `AnalyticsConfig` mbean.

 **Note:**

Only data that is useful should be published. The process design specifies what data (dimensions, measure, and counters) should be published and at what points. 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. To tune JMS for Measurements, consider changing the parameters listed in [Table 18-3](#), as needed, in a high volume environment:

Table 18-3 Essential JMS Resource Tuning for BPM

JMS Resource	Problem	Tuning Recommendation	Trade-offs
dist_MeasurementTopic_au to Default: Forwarding Policy Replicated	A distributed measurement topic in a cluster installation is configured by default with FORWARDING POLICY REPLICATED even though this is not the best performance option for BPM analytics.	Change the <i>Forwarding Policy</i> for this parameter to PARTITIONED. This parameter can be altered in the WebLogic console. You can find it from the front page with the following options: JMS Modules -> BPMJMSModule -> dist_MeasurementTopic_au to . You need to restart all SOA BPM cluster nodes for the changes to take effect.	A distributed topic with a Partitioned policy generally outperforms the FORWARDING POLICY REPLICATED. For more information on distributed topics versus other topic types, see Supported Topic Types in <i>Developing Message-Driven Beans for Oracle WebLogic Server</i> . For more information on partitioned and replicated forwarding policies, see Configuring Partitioned Distributed Topics in <i>Administering JMS Resources for Oracle WebLogic Server</i> .
MeasurementTopicConnect ionFactory Default: Send Timeout 200000	You have a high volume environment and you are receiving frequent resource allocation exceptions from message producers. For more information, see Defining a Send Timeout on Connection Factories in <i>Tuning Performance of Oracle WebLogic Server</i> .	Increase the Send Timeout for this parameter to 240000 in a high volume environment. The numerical value represents the maximum length of time in milliseconds. This parameter can be altered in the WebLogic console. You can find it from the front page with the following options: JMS Modules -> BPMJMSModule --> MeasurementTopicConnect ionFactory --> Default Delivery .	You may create a message backlog that consumes memory and resources.

Table 18-3 (Cont.) Essential JMS Resource Tuning for BPM

JMS Resource	Problem	Tuning Recommendation	Trade-offs
<p>MeasurementQuota</p> <p>Defaults: Message Maximum 1000000 and Bytes Maximum 800000000</p>	<p>Measurement messages cannot be published and fails with <code>javax.jms.ResourceAllocationException</code> thrown.</p>	<p>Set the Message Maximum and Bytes Maximum for this parameter equal to the amount of system memory available after you have accounted for the rest of your application load.</p> <p>The MeasurementQuota attributes can be altered in the WebLogic console. You can find it from the front page with the following options: JMS Modules -> BPMJMSModule -> MeasurementQuota.</p>	<p>Increasing this value consumes more memory. Message delivery may still fail if the aggregate size of messages pushed to the consumer is larger than the current protocol's maximum message size.</p> <p>For more information about measurement quotas, see <i>Tuning Performance of Oracle WebLogic Server</i>.</p>
<p>BPMJMSServer</p> <p>Default: MessageBuffer size 100000</p>	<p>The JMS server is frequently writing message bodies to disk.</p>	<p>Increase the Message Buffer Size for a given BPMJMSServer.</p> <p>Note that the BPMJMSServer uses Paging File and JMSFileStore.</p> <p>This parameter can be altered in the WebLogic console. You can find it from the front page with the following options: JMS Servers_auto_number.</p>	<p>The JMS server uses more memory.</p>

Tuning Oracle Human Workflow

You can tune Oracle Human Workflow to optimize its performance in handling various aspects of human interaction with a business process.

- [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.
- [Tuning Human Workflow](#)
You can tune Oracle Human Workflow to optimize its performance in handling the various aspects of human interaction with a business process.
- [Using Other Tuning Strategies](#)
You can consider using the following strategies to further improve performance.

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 any process. The Human Workflow service consists of a number of services that handle the various aspects of human interaction with a business process.

For more information, see *Using the Human Workflow Service Component* in *Developing SOA Applications with Oracle SOA Suite*.

See also the Oracle Human Workflow web site at <http://www.oracle.com/technetwork/middleware/human-workflow/overview/index.html>

Tuning Human Workflow

You can tune Oracle Human Workflow to optimize its performance in handling the various aspects of human interaction with a business process.

The suggestions presented here are all applicable to API usage.

Table 19-1 Essential Human Workflow Tuning Strategies

Name	Description	Recommendation
Minimize Client Response Time	<p>Since workflow client applications are interactive, it is important to have good response time at the client.</p> <p>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.</p>	Review your performance metrics to determine how response time can be improved.
Choose the Right Workflow Service Client	<p>Remote client 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 deserialization of data between the API methods.</p> <p>SOAP client is preferred for standardization (based on web services). There are additional performance considerations 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.</p>	<p>If the client application is based on Jakarta EE technology, then consider which client should be used based on your use case scenarios.</p> <p>Note that if the client application is based on .Net technologies, then only the SOAP workflow services can be used.</p>
Narrow Qualifying Tasks Using Precise Filters	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.	Use precise filters to improve response time.
Retrieve Subset of Qualifying Tasks (Paging)	The query API has paging parameters that control the number of qualifying rows returned to the user and the start row.	Decrease the <code>startRow</code> and <code>endRow</code> parameters to values that may limit the number of returned records. This decreases the query time, the application process time, and the amount of data returned to client.
Fetch Only the Information That Is Needed for a Qualifying Task	Typically only some of the payload fields are needed for displaying the task list.	<p>When you use the <code>queryTask</code> service, consider reducing the amount of optional information retrieved for each task returned in the list.</p> <p>In rare cases where the entire payload is needed, then the payload information can be requested.</p>
Reduce the Number of Return Query Columns	When you use the <code>queryTask</code> service, consider reducing the number of query columns to improve the SQL time.	Try to use the common columns as they are the most likely indexed columns. This allows the SQL to execute faster.

Table 19-1 (Cont.) Essential Human Workflow Tuning Strategies

Name	Description	Recommendation
Use the Aggregate API for Charting Task Statistics	Sometimes it is necessary to display charts or statistics to summarize task information.	Consider using the new aggregate APIs to compute the statistics at the database level rather than fetching all the tasks by using the query API and computing the statistics at the client layer.
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.	Call the <code>countTasks</code> API method, which returns only the number of matching tasks.
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.	Create indexes on these columns in certain cases where certain mapped flexfield columns are frequently used in query predicates.
Use the <code>doesTaskExist</code> Method	Sometimes it is necessary to check whether a task exists that matches a particular query criteria.	Consider using the <code>doesTaskExist</code> method instead of the default <code>countTasks</code> method. The <code>doesTaskExist</code> method performs an optimized query that checks if any rows exist that match the specified criteria. This method may achieve better results than calling the <code>countTasks</code> method.

Using Other Tuning Strategies

You can consider using the following strategies to further improve performance.

- [Improving Server Performance](#)
- [Completing Workflows Faster](#)
- [Tuning the Identity Provider](#)
- [Tuning the Database](#)

Improving Server Performance

Server performance essentially determines the scalability of the system under heavily loaded conditions. In [Tuning Human Workflow](#), strategy *Minimize Client Task 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 that is associated with querying. These techniques also reduce the database and service logic performance impacts on the server and can improve server performance. In addition, a few other configuration changes can be made to improve server performance:

Table 19-2 Essential server performance tuning strategies

Name	Description	Recommendation
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.	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.
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 task completion.	Ensure that there are sufficient resources to update the external system after the task is completed instead of after every workflow event. 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 that is associated with a BPEL callback since the callback method is executed in the same thread.
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.	Minimize the notification to alert a user only when a task is assigned instead of sending out notifications for each workflow event. Also 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.
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.

Completing Workflows Faster

The time it takes for a workflow to complete depends on the routing type that is specified for the workflow. The workflow functionality provides some options that can be used to decrease the amount of time it takes to complete workflows.

Table 19-3 Essential workflow completion tuning strategies

Name	Description	Recommendation
Use Workflow Reports to Monitor Progress	Several workflow reports (and corresponding views) are available that can make monitoring and proactive problem fixing easier.	By checking the 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 their tasks.

Table 19-3 (Cont.) Essential workflow completion tuning strategies

Name	Description	Recommendation
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. If the task must be escalated to some other user based on alternative routing logic, custom escalation rules can also be plugged in.	By specifying proper escalation rules, you can reduce workflow completion times.
Specify User and Group Rules for Automated Assignment	Rules can help significantly reduce workflow waiting time, which results in faster workflow completion.	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.
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 shift through the tasks or sort them to find out which one the user should work on next.	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.

Tuning the Identity Provider

The workflow service uses information from the identity provider in constructing the SQL query to determine the tasks that qualify for a user based on the role or 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 the user can 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 the nodes 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.
- If you use Oracle Internet Directory as the Identity Provider, ensure that you run the `oidstats.sql` to gather latest statistics on the database after the data shape has changed.

Tuning the Database

The Human Workflow schema is shipped with several indexes defined on the most important columns. 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 that 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 [Tuning Database Parameters](#).

Tuning Oracle Business Activity Monitoring

You can tune Oracle Business Activity Monitoring (BAM) to optimize its performance in monitoring business services and processes in the enterprise.

- [About Oracle Business Activity Monitoring](#)
Oracle Business Activity Monitoring (BAM) provides the tools for monitoring business services and processes in the enterprise.
- [Tuning BAM Server Parameters](#)
You can improve performance of the BAM server by following certain tuning recommendations.
- [Other Tuning Strategies](#)
If Oracle BAM is running more slowly than expected, you can try other tuning strategies.

About Oracle Business Activity Monitoring

Oracle Business Activity Monitoring (BAM) provides the tools for monitoring business services and processes in the enterprise.

It allows correlation of market indicators to the actual business process and to change business processes quickly or taking corrective actions if the business environment changes.

Oracle BAM also provides the necessary tools and runtime services for creating dashboards that display real-time data inflow and define rules to send alerts under specified conditions.

For information on how to monitor your BAM installation's performance, see *Monitoring Oracle BAM Performance* in *Monitoring Business Activity with Oracle BAM*.

Tuning BAM Server Parameters

You can improve performance of the BAM server by following certain tuning recommendations.

BAM performance largely depends on the performance of the following components:

- The Weblogic Server. See *Tuning Performance of Oracle WebLogic Server*.
- Metadata Service. See [Tuning Oracle Metadata Service](#).
- Coherence. See *Administering Oracle Coherence*.
- ADF. See [Tuning Oracle Application Development Framework \(ADF\)](#).
- Database Settings. See [Tuning Database Parameters](#).
- Java Virtual Machines (JVMs). See *Tuning Java Virtual Machines (JVM) in Tuning Performance of Oracle WebLogic Server*.
- Oracle Platform Security Service. See [Tuning Oracle Fusion Middleware Security](#).

BAM performance also depends on good data object design strategies at design time and on having good data object purging strategies at runtime.

While BAM 12c can support much larger transaction volumes (data arrival rates into BAM), BAM 12c is an operational analytics product, not a business intelligence product.

Hence, it is recommended that data that is of analytical value for operational decision-making be kept in BAM. For most customers, this means storing about 5-30 days of transactional data in BAM. Resting data sizes typically comparable to a data warehouse are not useful for operational decision-making, so such data volumes do not constitute a mainstream use case for BAM 12c.

The tuning suggestions listed and described in [Table 20-1](#) can be used to improve performance of the BAM Server:

Table 20-1 Essential BAM Server Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
Max connections to BAMDataSource Default: 50	The number of concurrent connection to the data source consistently reaches max limit.	Increase count as long as the performance continues to scale linearly. This is set at the WebLogic level. The value can be determined mainly based on the number of processors, listening source threading model and max concurrent request settings of embedded servers.	Increasing the count will most likely increase the system resources usage.
Viewset Expiry Time Default: 180 seconds	Viewsets are lingering after the DC connection is lost.	Decrease the expiry time value so that viewsets do not linger. See <i>Monitoring Viewsets in Monitoring Business Activity with Oracle BAM</i> for information on how to find and modify this parameter.	None.
DiagnosticLevel Default: Info	You need granular diagnostic logs to identify a problem. OR Your system is running fine and you do not need detailed logs.	Keeping the default of INFO will help performance. For more information on using the BAM Diagnostic Framework, see <i>Using the BAM Diagnostic Framework in Monitoring Business Activity with Oracle BAM</i> .	If your system slows down, you do not have detailed logs to identify a problem.
ASM (Automatic Server Migration) Default: WSM	You want to migrate a SOA Suite installation with BAM to High Availability. Because BAM is a real-time system, you should enable ASM.	ASM is used so High Availability can occur faster than WSM. Given that BAM is a real-time system, ASM is required for BAM HA.	None.
JVM heap size Default: -Xms768m -Xmx1536m	Oracle BAM is running slowly and an out-of-memory exceptions occur.	Increase the heap size to 2 GB. Use the following command with the -Xms2048m and -Xmx2048m arguments: setenv USER_MEM_ARGS "-Xms2048m -Xmx2048m -XX:PermSize=256m -XX:MaxPermSize=768m"	Increasing the JVM heap size for BAM could affect other SOA components. For more heap size tuning tips, see <i>Tuning Tips for Heap Sizes in Tuning Performance of Oracle WebLogic Server</i> .

Other Tuning Strategies

If Oracle BAM is running more slowly than expected, you can try other tuning strategies.

- [Creating an Index Column](#)
- [Tuning Loggers](#)
- [Tuning Continuous Query Service](#)

Creating an Index Column

If throughput of data into a data object from an Enterprise Message Source or other source is slow, create an index column for the primary key column. See *Adding Index Columns in Monitoring Business Activity with Oracle BAM* for more information.

Tuning Loggers

The default Oracle Diagnostic Logging Level for all loggers is `Notification`. For stress testing and production environments, consider using the lowest acceptable logging level, such as `ERROR` or `WARNING`.

The loggers in BEAM that can affect BAM performance are as follows:

```
oracle.beam.common.alertsengine
oracle.beam.server.service.alertsengine
oracle.beam.Common
oracle.beam.cqservice
oracle.beam.composer
com.oracle.beam
oracle.beam.datacontrol
oracle.beam.datacontrol.management
oracle.beam.server.service.ems
oracle.beam.messaging
oracle.beam.server.service.persistence
oracle.beam.server.service.reportcache
oracle.beam.security
oracle.beam.mbean
oracle.beam.shared
oracle.beam.server
oracle.beam.impexp.t2p
oracle.beam
```

For information about locating these loggers and changing their Oracle Diagnostic Logging Level, see *Configuring Log Files in Administering Oracle SOA Suite and Oracle Business Process Management Suite*.

Tuning Continuous Query Service

The Continuous Query Service (CQS) is a BAM-specific wrapper around the Continuous Query Language (CQL) engine within the Oracle Complex Event Processing Service Engine. The CQS is a pure push system: query results are delivered automatically. The CQS supports both stream (non-persistent) and archived relation (persistent) data objects.

When you create a query, the CQS sets up tables in the CQL engine, registers the query, and listens for data changes from the persistence engine. The query result is processed in the CQL engine, then pushed to the CQS and on to the report cache.

For information on how to monitor continuous queries for performance issues, see Monitoring Continuous Queries in *Monitoring Business Activity with Oracle BAM*. Once you understand how your current system is performing, you can try to improve performance by tuning the knobs described in Table 20-2. Note that for most of these parameters, tuning for performance means losing diagnostic information.

Table 20-2 Tuning the Continuous Query Service

Parameter	Problem	Tuning Recommendation
Data Object type Default: None	You have arbitrarily designated simple data objects as stream, archived stream, and archived relation, and are not sure what to do.	Categorize your data objects as stream if you do not care about historical data. See Data Object Types in <i>Monitoring Business Activity with Oracle BAM</i> for detailed descriptions of each data object type and relation.
Data Object purging Default: Disabled	By default, data object retention is not set. Many rows in the data object cause performance issues.	Customer can set Data Object retention in the Data Object Retention tab to specify how many days they want to keep the data in a Data Object. When the specified number of days has elapsed, the data rows are automatically purged. See Setting Data Retention in a Data Object in <i>Monitoring Business Activity with Oracle BAM</i> for information on how to find and change this setting.
Replay for Archived Stream Data Objects	Data parsing is slow for archived stream data objects.	Specify a smaller Replay Unit or a lower Replay Amount to reduce the amount of past data retained in memory. This reduces the time and memory to parse data retrieved from the database.
Time Window on Input Streams	You have chosen to turn an Active Data query into a continuous query and are receiving out-of-memory exceptions.	Decrease the time window size on the Active Data stream. This restricts the amount of memory the window uses to store elements. To get an idea of how much the window size affects memory usage, consider a scenario where the Window Size = 1 hour (RANGE 1 hour) and the event size = 100 bytes. If the event rate is 1000 events / second, then the window will contain 1000 * 3600 events when it is full. The memory consumed is 1000 * 3600 * 100 bytes = ~340 MB. See Enabling Active Data in a View in <i>Monitoring Business Activity with Oracle BAM</i> for information on how to configure the window size on an active data view.

Table 20-2 (Cont.) Tuning the Continuous Query Service

Parameter	Problem	Tuning Recommendation
Active Data Collapsing Interval Default: Unchecked	You have checked the box for Active Data Collapsing to make data aggregation active. You want more frequent snapshots or need to free up memory.	<p>Define a smaller Interval to make the view update more frequently and to reduce the amount of aggregated data stored in memory.</p> <p>You can maximize your memory usage by taking note of the evaluation interval, the event size, and the event rate. Given the following values:</p> <p style="padding-left: 40px;">Interval: Every 5 minutes Event Size: 100 bytes Event Rate: 1000 events/second</p> <p>The maximum size of the aggregated view is $5 * 60 * 1000 = 300,000$ events = ~28 MB.</p> <p>See Using Active Data in <i>Monitoring Business Activity with Oracle BAM</i> for information on finding the Active Data Collapsing setting.</p>
Slow Changing Dimension for Data Object Dimension Tables Default: Unchecked	Continuous queries on dimension tables are slow and consuming memory.	<p>Check this property to activate it. This indicates that the data in this dimension table changes infrequently.</p> <p>For information on specifying slow-changing dimensions for a data object, see Specifying Slow-Changing Dimensions for a Data Object in <i>Monitoring Business Activity with Oracle BAM</i>.</p>
Query Type Default: SQL	You are experiencing out-of-memory exceptions and most of your queries are continuous.	Use schedule query (SQL) where you do not expect frequent output. This saves memory because SQL involves JDBC resources while CQL stores data in memory.

21

Tuning Oracle Service Bus

You can tune Oracle Service Bus (OSB) to optimize its performance in providing connectivity, routing, mediation, management, and also some process orchestration capabilities between two or more applications.

- [About Oracle Service Bus](#)
Within a SOA framework, Oracle Service Bus (OSB) provides connectivity, routing, mediation, management, and also some process orchestration capabilities.
- [Tuning OSB Parameters](#)
Oracle Service Bus performance largely depends on the performance of the other components.
- [Using Other Tuning Strategies](#)
After you have performed the recommended modifications, you can make additional changes that are specific to your deployment.

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 a 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 service, 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 responding to the client.

It is necessary 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 *Oracle Fusion Middleware Administrator's Guide for Oracle Service Bus*.

Tuning OSB Parameters

Oracle Service Bus performance largely depends on the performance of the other components.

The following components affect OSB performance:

- WebLogic Server
- Coherence

- [Adapters](#)

You can begin tuning Oracle Service Bus if you believe the above components are tuned to your satisfaction.

- [Tuning Oracle Service Bus with Work Managers](#)
- [Tuning OSB Operation Settings](#)

Tuning Oracle Service Bus with Work Managers

Oracle Service Bus can be tuned by several Oracle WebLogic Server Work Managers.

For example, Split-Join tuning can be accomplished by using Work Managers. By default, applications do not specify a Work Manager for Split-Joins, but Split-Joins can be assigned a Work Manager if there are strict thread constraints that need to be met, such as scheduling parallel tasks.

For optimal performance, strike a balance between the following Work Manager constraints:

- `min-threads-constraint` so that Split-Join operations are not starved of threads.
- `max-threads-constraint` so that Split-Joins do not starve other resources

By default, there is no minimum or maximum thread constraint defined, which could either slow Split-Join operations down or slow down other operations sharing the same thread pool.

Work Managers take Split-Join operations into account when allotting threads to system-wide processes so that this balance is met automatically.

For more information on tuning OSB with Work Managers, see *Using Work Managers with Oracle Service Bus* in *Developing Services with Oracle Service Bus*.

Tuning OSB Operation Settings

[Table 21-1](#) lists and describes the knobs you will most likely need to tune to improve performance. For more information on monitoring Oracle Service Bus to diagnose trouble areas, see *Monitoring Oracle Service Bus* in *Administering Oracle Service Bus*.

Table 21-1 Essential OSB Operation Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
Monitoring and Alerting Default: Disabled	The Monitoring and Alerting framework is designed to have minimal impact on performance, but all of these processes have performance impacts. In general, the more monitoring rules and pipeline actions you have defined, the larger the performance impact.	Keep the default of <code>Disabled</code> at the OSB level. Most settings can be defined globally or per service. The settings for monitoring and alerting can be configured in the Enterprise Manager Administrator Console. Note that monitoring must be enabled for SLA alerts but not for Pipeline alerts.	Disabling these processes to improve performance means you are sacrificing certain metrics and alerts that could help you troubleshoot issues in the future. For more information on the OSB Monitoring Framework, see Introduction to the Oracle Service Bus Monitoring Framework in <i>Administering Oracle Service Bus</i> .

Table 21-1 (Cont.) Essential OSB Operation Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
Tracing Default: Disabled	If you have large message sizes and high throughput scenarios, tracing may be slowing your system down.	Leave tracing disabled to improve performance. For more information, see <i>How to Enable or Disable Tracing in Oracle Fusion Middleware Administrator's Guide for Oracle Service Bus</i> .	If disabled, you lose metrics. Tracing prints the entire message context, including headers and message body. 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.
<code>com.bea.wli.sb.pipeline.RouterRuntimeCache.size</code> Default: 100	You may have one of the following issues: Proxy services are accessed slowly. This means you want to store more proxy services in the static portion of the OSB cache for pipeline service runtime metadata. The proxy services stored here are never garbage-collected, meaning they are accessed faster. OR You are seeing a lot of cache misses in DMS dumps.	If you want to include more proxy services in the static cache, increase this value as long as there is sufficient memory for runtime data processing for large number of proxy services. If you are seeing cache misses in DMS dumps, increase this value. This system property caps the number of proxy services in the static portion of the OSB cache for pipeline service runtime metadata. These services never get garbage collected. You set the size of this value in the <code>setDomainEnv.sh</code> file as an extra java argument, as follows: - <code>Dcom.bea.wli.sb.pipeline.RouterRuntimeCache.size={size}</code> For example, if you want to set this value to 3000, you would write: <code>EXTRA_JAVA_PROPERTIES="_Dcom.bea.wli.sb.pipeline.RouterRuntimeCache.size=3000\${EXTRA_JAVA_PROPERTIES}"</code>	Increasing this value decreases the time it takes to make initial calls to the proxy server. It can also preload the cache when a configuration session is committed. However, while caching proxy services helps reduce compilation costs, it also increases memory consumption. Decreasing this value may mean you free up memory, but making initial calls to the proxy server may take longer.

Table 21-1 (Cont.) Essential OSB Operation Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
reorderJsonAsPerXmlSchema Default: False	JSON input to REST service may not be ordered as expected by the schema definition. When converting from JSON to XML, OSB runtime uses the order in which JSON name or value appear to construct the corresponding XML element. While well-formed, this format is not valid according to XML schema.	Set this parameter to <code>True</code> by running the REST wizard and checking the box on the first page. Checking this option makes the REST service reorder the input JSON so that the response from the external REST endpoint can be ordered as per the valid schema definition.	Using this option adds significant performance overhead.

Using Other Tuning Strategies

After you have performed the recommended modifications, you can make additional changes that are specific to your deployment.

Consider carefully whether the additional tuning recommendations are appropriate for your environment.

- [Tuning Resequencer in OSB](#)
- [Considering Design Time for Proxy Applications](#)
- [Tuning XQuery](#)
- [Tuning Poller-based Transports](#)

Tuning Resequencer in OSB

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 random order and then sends them to the target services in an orderly manner.

You can fine-tune the Resequencer by setting the properties listed in [Table 21-2](#) using the Global operational settings page in the OSB EM console:

Table 21-2 Essential Resequencer Tuning

Parameter	Problem	Tuning Recommendation	Trade-offs
ResequencerMaxGroupsLocked Default: 4 groups	<p>This parameter defines the maximum number of message groups that can be locked by resequencer locker threads for parallel processing. The locked groups can then use worker threads to process their respective messages.</p> <p>If message processing is being delayed, identify which of the following situations is true:</p> <ul style="list-style-type: none"> • Incoming messages belong to many groups. • There are many messages and they belong to fewer groups. 	<p>If you have many groups with a small number of messages each, increase this parameter's value. Resequencer will lock more groups in one iteration.</p> <p>If you have a few groups with many messages, decrease this value. Resequencer will lock less number of groups for processing.</p>	<p>Increasing the <code>MaxGroupsLocked</code> value may result in locking more groups than there are available worker threads. This could result in groups getting blocked while waiting for the availability of the worker threads for message processing.</p> <p>Decreasing the default value may result in under utilization of resources.</p>
ResequencerLockerThreadSleep Default: 10 seconds	<p>The resequencer locker thread queries the database to lock groups for parallel processing. When no groups are available, the locker thread <i>sleeps</i> for the configured amount of time specified by this parameter.</p> <p>If you have either of the following situations, this parameter needs tuning:</p> <ul style="list-style-type: none"> • You have a high number of messages and processing time between database queries is slow. • You have few messages but frequent database queries. 	<p>Decrease this parameter value if you have a high number of messages to reduce the lag time during processing.</p> <p>If Resequencer locker threads are making frequent database round trips even though you do not have many incoming messages, increase this value.</p>	<p>If the sleep time is too short, there may not be enough worker threads available to process incoming messages of the locked groups. Too many database queries will also cause slow performance.</p> <p>If the time interval between incoming messages is already long, configuring a higher value is not beneficial.</p>
DeleteMessageAfterComplete Default: True	<p>The resequencer database is low on space. If you changed this parameter's value to false, processed messages remain in the resequencer database and slow down database inquiries.</p>	<p>Keep the default value of <code>True</code> to delete message after successful execution. This frees up database space.</p>	<p>You do not have a detailed history of processed messages.</p>

Considering Design Time for Proxy Applications

Consider the design configurations described in [Table 21-3](#) for proxy applications based on your OSB usage and use case scenarios:

Table 21-3 Tuning Design Time for Proxy Application

Strategy	Description	Recommendations
Avoid creating many OSB context variables that are used once within another XQuery	Context variables created by using an Assign action are converted to XmlBeans and then reverted to the native XQuery format for the next XQuery. Multiple <i>Assign</i> actions can be collapsed into a single Assign action by using a <i>FLWOR</i> expression. Intermediate values can be created by using <i>let</i> statements.	Avoiding redundant context variable creation eliminates overheads that are associated with internal data format conversions. This benefit has to be balanced against visibility of the code and reuse of the variables.
Transform contents of a context variable such as <i>\$body</i> .	Transforming the contents of a context variable could be time-consuming.	Use a Replace action to complete the transformation in a single step. If the entire content of <i>\$body</i> is to be replaced, leave the XPath field blank and select <i>Replace node contents</i> . This is faster than pointing to the child node of <i>\$body</i> (for example, <i>\$body/Order</i>) and selecting <i>Replace entire node</i> . Leaving the XPath field blank eliminates an extra XQuery evaluation.
Specify a special XPath.	A general XPath like <i>\$body/Order</i> must be evaluated by the XQuery engine before the primary transformation resource is executed. OSB treats <i>\$body/*[1]</i> as a special XPath that can be evaluated without invoking the XQuery engine.	Use <i>\$body/*[1]</i> to represent the contents of <i>\$body</i> as an input to a Transformation (XQuery / XSLT) resource. This is faster than specifying an absolute path pointing to the child of <i>\$body</i> .
Enable streaming for pure content-based routing scenarios.	OSB leverages the partial parsing capabilities of the XQuery engine when streaming is used in conjunction with indexed XPaths. See Tuning XQuery for additional details.	Enabling streaming means that the payload is parsed and processed only to the field referred to in the XPath. Streaming also eliminates the overhead that is associated with parsing and serialization of XmlBeans. Trade-offs: If the payload is accessed a large number of times for reading multiple fields, the gains from streaming can be negated. If all fields read are located in a single subsection of the XML document, a hybrid approach provides the best performance. 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.	OSB can invoke a back end HTTP service asynchronously if the QOS is <i>Best-Effort</i> . 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.	Do not set XA or Exactly-Once unless the reliability level required is once and only once and it is possible to use the setting. If the client is a HTTP client it is not possible to use this setting. If OSB initiates a transaction, it is possible to replace XA with LLR to achieve the same level of reliability.

Table 21-3 (Cont.) Tuning Design Time for Proxy Application

Strategy	Description	Recommendations
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.	Disable or delete all log actions.

Tuning XQuery

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>
```

You can use the tuning strategies listed in [Table 21-4](#) to tune XQuery.

Table 21-4 XQuery Tuning Strategies

Strategy	Description	Recommendations
Avoid the use of double front slashes (<code>//</code>) in XPaths.	<code>//</code> 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 <code>//</code> .	Use <code>//</code> only if the exact location of a node is not known at design time.

Table 21-4 (Cont.) XQuery Tuning Strategies

Strategy	Description	Recommendations
Index XPath's when applicable.	Indexing helps your system process only what is needed. When indexing, only the top part of the document is processed by the XQuery engine.	<p>Index an XPath by adding [1] after each node of the path.</p> <p>For example, the XPath <code>\$body/Order/CtrlArea/CustName</code> implies returning all instances <code>Order</code> under <code>\$body</code> and all instances of <code>CtrlArea</code> under <code>Order</code>. The entire document has to be read to correctly process the above XPath.</p> <p>But if you know that there is a single instance of <code>Order</code> under <code>\$body</code> and a single instance of <code>CtrlArea</code> under <code>Order</code>, you can index the above XPath by rewriting it as <code>\$body/Order[1]/CtrlArea[1]/CustName[1]</code>. This only returns the first instances of the child nodes.</p> <p>Note: Do not index when you need a whole array of nodes returned. Indexing only returns the first item node of the array.</p>
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.	Using intermediate variables consumes more memory but reduces redundant XPath processing.
Use a hybrid approach for read-only scenarios with streaming.	If the payload is accessed a large number of times for reading multiple fields, The gains from streaming can be negated. If all fields read are located in a single subsection of the XML document, a hybrid approach provides the best performance.	<p>Enable 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.</p> <p>The fields <code>Total</code> and <code>Status</code> can be retrieved by using three <code>Assign</code> actions:</p> <pre>Assign "\$body/Order[1]/Summary[1]" to "foo" Assign "\$foo/Total" to "total" Assign "\$foo/Status" to "total"</pre>

 **Note:**

Pipelines enabled for content streaming should use *XQuery 1.0*. Using *XQuery 2004* does work, but incurs a significant performance overhead, as there are *on-the-fly* conversions that happen to and from XQuery 1.0 engine. There is a design-time warning to that effect.

Tuning Poller-based Transports

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.

- [Setting the Polling Interval](#)
- [Setting Read Limit](#)

Setting the 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 <i>Developing Services with Oracle Service Bus</i>
FTP Transports	60 seconds	FTP Transport Configuration Page in <i>Developing Services with Oracle Service Bus</i>
MQ Transport	1000 milliseconds	MQ Transport Configuration Page in <i>Developing Services with Oracle Service Bus</i>
SFTP Transport	60 seconds	SFTP Transport Configuration Page in <i>Developing Services with Oracle Service Bus</i>
JCA Transport	60 seconds	JCA Transport Configuration Page in <i>Developing Services with Oracle Service Bus</i>

Setting Read Limit

The read limit determines the number of files or messages that are read per polling sweep. You can tune it with the information in [Table 21-5](#).

For more information, see Using the File Transport in *Developing Services with Oracle Service Bus*.

Table 21-5 Essential Read Limit Tuning

Parameter	Symptoms if not properly tuned	Tuning Recommendation	Performance Trade-offs
Read Limit Default: 10 for File and FTP transports	Excessive memory use or high memory use due to a large number of files read into memory simultaneously.	Set this value to the desired concurrency. It can be set to 0 to specify no limit. The read limit determines the number of files or messages that are read per polling sweep.	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. If the message size is large, this can lead to an out-of-memory error .

Tuning Oracle Enterprise Scheduler Service

You can tune Oracle Enterprise Scheduler Service (ESS) to optimize its performance in enabling scheduling and running jobs.

- [About Enterprise Scheduler Service](#)
Oracle Enterprise Scheduler enables scheduling and running jobs within a particular time frame, or workshift, by using rules to create work assignments.
- [Tuning Enterprise Scheduler Service Parameters](#)
You can tune the enterprise scheduler service parameters for optimal performance.

About Enterprise Scheduler Service

Oracle Enterprise Scheduler enables scheduling and running jobs within a particular time frame, or workshift, by using rules to create work assignments.

Oracle Enterprise Manager Fusion Applications Control allows you to define, control and manage Oracle Enterprise Scheduler job metadata, including job definitions, job requests, job sets (a collection of job requests), incompatibilities (job definitions and job sets that cannot run at the same time for a given application) and schedules governing the execution of job requests.

For more information, see Introduction to Administering Oracle Enterprise Scheduler in *Administering Oracle Enterprise Scheduler*.

Tuning Enterprise Scheduler Service Parameters

You can tune the enterprise scheduler service parameters for optimal performance.

[Table 22-1](#) describes the enterprise scheduler service tuning parameters.

`Maximum Poll Interval` is a dispatcher parameter that applies to the Oracle Enterprise Scheduler request dispatcher. The request dispatcher manages requests that are awaiting their scheduled execution. The request processor handles the job requests once they have dispatched.

`Thread Count` is a processor tuning parameter that applies to the Oracle Enterprise Scheduler request processor. The request processor manages job requests whose scheduled execution time has arrived, and are ready to execute.

Table 22-1 Essential Enterprise Scheduler Service Tuning

Name	Symptoms	Recommendations	Trade-offs
Maximum Poll Interval Default: 15 seconds	A high number of requests whose execution time has been reached and remain in <code>WAIT</code> state for an extended time.	If there is an excess of waiting requests that are eligible to be dispatched and processor threads are available, decrease this value .	Lowering the value increases CPU usage and database activity. Increasing the value may delay the dispatching of requests that are ready for processing.
Thread Count Default: 25	A high number of requests in <code>READY</code> state that are otherwise available for processing.	If there is a build up of requests that are ready to be executed and the increase system resource usage is acceptable, increase this value. Lower the value to reduce the amount of system resources used for request processing.	Increasing this value increases CPU usage, memory usage, and database activity. Lowering this value may result in a build up and potentially delay processing of requests.

Tuning Oracle Business Intelligence Performance

You can tune Oracle Business Intelligence to optimize its performance in collecting, presenting, and delivering data.

- [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.
- [Tuning Oracle BI Server Query Performance](#)
You can improve query performance by tuning and indexing underlying databases, by using aggregate tables, query caching.
- [Tuning Oracle BI Server Query Cache Performance](#)
You can configure the Oracle BI Server to maintain a local, disk-based cache of query result sets (query cache).
- [Tuning Oracle BI Web Client Performance](#)
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.

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.

Tuning Oracle BI Server Query Performance

You can improve query performance by tuning and indexing underlying databases, by using aggregate tables, query caching.

For detailed information on BI performance tuning, see *Managing Performance Tuning and Query Caching in 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 *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 *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 *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 by using Fusion Middleware Control to improve system performance. See *Setting Performance Parameters in Fusion Middleware Control* in *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 *System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.

Tuning Oracle BI Server Query Cache Performance

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 *System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Tuning Oracle BI Web Client Performance

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 *System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Part VI

Oracle WebCenter Components

The Oracle WebCenter components need to be tuned for optimal performance.

This part describes configuring Oracle WebCenter components to improve performance. It contains the following topic:

- [Tuning Oracle WebCenter Portal](#)
You can tune Oracle WebCenter Portal to optimize its performance as a deployed application.

Tuning Oracle WebCenter Portal

You can tune Oracle WebCenter Portal to optimize its performance as a deployed application.

- [About Oracle WebCenter Portal](#)
Oracle WebCenter Portal helps companies to build enterprise-scale intranet and extranet portals that provide a foundation for the next-generation user experience (UX) with Oracle Fusion Middleware and Oracle Fusion Applications.
- [Basic Tuning Considerations](#)
Tuning considerations apply to most WebCenter Portal application deployment scenarios.
- [Tuning Configuration for WebCenter Portal](#)
You can tune configuration parameters to improve the performance of WebCenter Portal.
- [Tuning Tools and Services Configuration](#)
You can tune the performance of tools and services used by WebCenter Portal.
- [Tuning Identity Store Configuration](#)
Performance-related configurations may be required for specific environments.
- [Tuning Portlet Configuration](#)
You can tune the performance of portlets in WebCenter Portal.

About Oracle WebCenter Portal

Oracle WebCenter Portal helps companies to build enterprise-scale intranet and extranet portals that provide a foundation for the next-generation user experience (UX) with Oracle Fusion Middleware and Oracle Fusion Applications.

Portals built with Oracle WebCenter Portal commonly support thousands of users who create, update, and access content and data from multiple back-end sources. Oracle WebCenter Portal delivers intuitive user experiences by leveraging the best UX capabilities from a significant portfolio of leading portal products and related technologies. From the user's perspective, the integration is seamless.

Business users can easily assemble new portals or composite applications by using Portal Composer and a page editor that includes a library of prebuilt reusable components. They can enhance user experience by wiring components together on the page, configuring content personalization, enabling the use of integrated social tools, and creating data visualizations.

For more information about Oracle WebCenter Portal, see:

- *Using Portals in Oracle WebCenter Portal*
- *Building Portals with Oracle WebCenter Portal*
- *Administering Oracle WebCenter Portal*
- *Developing for Oracle WebCenter Portal*

Basic Tuning Considerations

Tuning considerations apply to most WebCenter Portal application deployment scenarios.

It is highly recommended that you review these configurations and implement those that meet your particular usage requirements.

- [Setting System Limit](#)
- [Setting JDBC Data Source](#)
- [Using Content Compression to Reduce Downloads](#)

Setting System Limit

To run WebCenter Portal 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.

Setting JDBC Data Source

To determine the correct setting for the JDBC data source, use the Oracle WebLogic Remote Console to monitor the running system database connection usage as described in [Configuring Services](#). 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 the 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 Administering JDBC Data Sources for Oracle WebLogic Server](#).

The following data source settings are WebCenter Portal defaults for data sources **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>
```

For information on how to edit MDS data source settings, see Tuning Data Source Connection Pools in *Administering JDBC Data Sources for Oracle WebLogic Server*.

Using Content Compression to Reduce Downloads

If clients connect to your server using relatively slow connections, that is, by using modems or VPN from remote locations, consider compressing the content before it downloads to the client. While content compression increases the load on the server, the client's download experience is much improved.

Note:

Beginning with release 11.1.1.8.0, WebCenter Portal is preconfigured with an ADF caching filter, which automatically sets up caching for static resources and do compression. This preconfigured ADF caching filter is available only for use with WebLogic Server.

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 the `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/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 the `mod_deflate` module, refer to: http://httpd.apache.org/docs/2.0/mod/mod_deflate.html.

Tuning Configuration for WebCenter Portal

You can tune configuration parameters to improve the performance of WebCenter Portal.

- [Setting a Session Timeout for WebCenter Portal](#)
- [Setting MDS Cache Size and Purge Rate](#)
- [Configuring Concurrency Management](#)

Setting a Session Timeout for WebCenter Portal

The default session timeout for the WebCenter Portal application is 45 minutes. Administrators can customize the session time to suit their installation. For details see [Specifying Session Timeout Settings in *Using Portals in Oracle WebCenter Portal*](#).

Setting MDS Cache Size and Purge Rate

If you encounter the any of the following conditions, then you can increase the MDS cache size in the `adf-config.xml` file. The default MDS cache size is 100 MB.

- Error message `JOC region full`
- Frequent MDS database access after the page is warmed up
- Retained memory by ADF application is close to the `max-size-kb`

Post deployment, modify these properties through the System MBeans Browser. For more information, see [Changing MDS Configuration Attributes for Deployed Applications in *Administering Oracle Fusion Middleware*](#).

The following is a sample snippet of the `adf-config.xml` file:

```
<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 the `adf-config.xml` file. Use the following sample snippet of the `adf-config.xml` file 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-namePortal>
        ...
      </metadata-namespace>
      <auto-purge seconds-to-live="3600"/>
    </persistence-config>
  </mds-config>
</adf-mds-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)
```

Configuring Concurrency Management

Concurrency management includes global settings that impact the entire WebCenter Portal and the 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 code snippet describes the format of the global, service, and resource entries in the `adf-config.xml` file:

```
<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 the `connections.xml` file or the `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 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, see *Accessing the System MBean Browser in Administering 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
2. To view the current concurrency settings, select **listResource**, and then click **Invoke**.
3. To change a setting, select **setResource**, enter the resource details, and then click **Invoke**.

Take care to enter the correct values for **service**, **resource**, **name**, and **value**.

 **Note:**

If the resource parameter that you are attempting to modify already has a **value** setting, you must remove the setting first by invoking the **removeResource** operation.

4. To save changes, navigate to **Application Defined MBeans: ADFConfig:ADFConfig** -> **save**, and click **Invoke**.

Tuning Tools and Services Configuration

You can tune the performance of tools and services used by WebCenter Portal.

For information about how to tune and improve the performance of back-end servers, for example, mail servers, BPEL servers, content servers, and so on, refer to the appropriate product documentation for each server.

- [Tuning Performance of Mail](#)
- [Tuning Performance of RSS News Feeds](#)
- [Tuning Policy Store Parameters](#)

Tuning Performance of Mail

To manage the overall resource usage for mail, 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 by using WLST. For details, see:

- [Modifying Mail Server Connection Details Using Fusion Middleware Control in *Administering Oracle WebCenter Portal*](#)
- [Modifying Mail Server Connection Details Using WLST in *Administering Oracle WebCenter Portal*](#)

The following is a sample code snippet of the `connections.xml` file 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>
```

Tuning Performance of RSS News Feeds

To manage the overall resource usage for RSS news feeds, 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 the `adf-config.xml` file:

```
<rssC:adf-rss-config>
  <rssC:RefreshSecs>3600</rssC:RefreshSecs>
  <rssC:TimeoutSecs>3</rssC:TimeoutSecs>
  <rssC:Configured>true</rssC:Configured>
</rssC:adf-rss-config>
```

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 [OPSS PDP Service Tuning Parameters](#). Depending on your use case scenarios, performance of WebCenter Portal 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 portals in your WebCenter Portal deployment.

Note:

Only modify this parameter if your WebCenter Portal deployment expects to have more than 3000 active portals.

- Set `oracle.security.jps.policystore.policy.cache.size` to 5 times the expected number of portals.

Note:

Always refer to your own use case scenarios before you modify the policy store parameters. For more information, see *Administering Web Services* before tuning any security parameters.

Tuning Identity Store Configuration

Performance-related configurations may be required for specific environments.

- [Tuning the Identity Store when Using SSL](#)
- [Tuning Performance when Using OVD](#)
- [Tuning Performance when Using Active Directory](#)

Tuning the Identity Store when Using SSL

When you configure an identity store for WebCenter Portal, 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>

```

- 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.

Tuning Performance when Using OVD

For Oracle Virtual Directory (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.stdldap.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>

```

Tuning Performance when Using Active Directory

When the Portal Server is connected to Active Directory, logging in to Portal is delayed. To avoid delays and to enable best performance while using Active Directory in a production environment, Oracle recommends you complete the following configuration:

- Log into the **Enterprise Manager** as `admin`.
- In the navigation pane, select **Weblogic Domain**.
- Navigate to **Security > Security Provider Configuration**.
- Expand **Identity Store Provider**.

5. Click **Configure**.

The **Identity Store Configuration** page appears.

**Note:**

Configure is similar to **Configure parameters for User and Role APIs to interact with identity store**.

6. Under **Custom Properties**, click **Add**.

7. Add the following new property:

```
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
```

8. Click **OK**.
9. Restart the Admin Server.
10. After the restart, log into the **Enterprise Manager** and check the values you entered in the previous step.

Tuning Portlet Configuration

You can tune the performance of portlets in WebCenter Portal.

- [Tuning Performance of the Portlet Client](#)
- [Customizing the Container Runtime Environment Options](#)
- [Tuning Performance of Oracle PDK-Java Producers](#)
- [Setting WSRP Attribute for Portlet-served Resources](#)
- [Setting WSRP Attribute for Resources Not Served by the Portlet](#)

Tuning Performance of the Portlet Client

Several tuning options are available for Portlet Client.

- [Configuring Supported Locales](#)
- [Configuring Portlet Cache Size](#)
- [Configuring Portlet Timeout](#)

Configuring Supported Locales

To manage the 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 and ready-to-use. You can remove the locales that are unnecessary for your application.

If you must modify these properties, post deployment, you must edit the `adf-config.xml` file manually. See [Editing adf-config.xml in Administering Oracle WebCenter Portal](#).

The following is a sample snippet of the `adf-config.xml` file:

```
<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>
  <maximumTimeout>300</maximumTimeout>
  <parallelPoolSize>10</parallelPoolSize>
  <parallelQueueSize>20</parallelQueueSize>
  <cacheSettings enabled="true">
    <maxSize>10000000</maxSize>
  </cacheSettings>
</portletC:adf-portlet-config>
```

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 the `adf-config.xml` file manually.

For more information, see [How to Edit Portlet Client Configuration](#) in *Developing for Oracle WebCenter Portal*

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 the `adf-config.xml` file manually. See *Editing adf-config.xml in Administering Oracle WebCenter Portal*.

The following is a sample snippet of the `adf-config.xml` file:

```
<adf-portlet-config>
  ....
  <defaultTimeout>5</defaultTimeout>
  <minimumTimeout>2</minimumTimeout>
  <maximumTimeout>300</maximumTimeout>
</adf-portlet-config>
```

Customizing the Container Runtime Environment Options

Customizing container runtime options can improve overall performance.

For more information, see *How to Customize the Runtime Environment for JSR 286 Portlets in Developing for Oracle WebCenter Portal*.

- [Suppressing Optimistic Rendering for WSRP Portlets](#)
- [Setting Portlet Container Runtime Options](#)
- [Excluding Request Attributes for Portlets](#)

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 in the `portlet.xml` file to `true`. For example:

```
com.oracle.portlet.suppressWsrpOptimisticRender=true
```

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 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.

For more information, see *How to Customize the Runtime Environment for JSR 286 Portlets in Developing for Oracle WebCenter Portal*.

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 that 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.

For more information, see *How to Customize the Runtime Environment for JSR 286 Portlets in Developing for Oracle WebCenter Portal*.

Excluding Request Attributes for Portlets

The `excludedActionScopeRequestAttributes` is a multivalued, Portlet container runtime property, where each value is a regular expression.

If you use the `javax.portlet.actionScopedRequestAttributes` container runtime option with a portlet, it is possible to optimize the request attributes that are stored between portlet lifecycles by using the `com.oracle.portlet.excludedActionScopeRequestAttributes` container runtime option. Any request attributes that are unnecessary to store between lifecycles can be indicated to increase performance.

For more information, see *How to Customize the Runtime Environment for JSR 286 Portlets in Developing for Oracle WebCenter Portal*.

Tuning Performance of Oracle PDK-Java Producers

To manage the 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 by using WLST. For details, see:

- *Editing WSRP Producer Registration Details Using Fusion Middleware Control in Administering Oracle WebCenter Portal*.
- *Editing Producer Registration Details Using WLST in Administering Oracle WebCenter Portal*.

The following is a sample snippet of the `connections.xml` file:

```
<webproducerconnection producerName="wc-WebClipping" urlConnection="wc-WebClipping-  
urlconn" timeout="10000" establishSession="true" mapUser="false"/>
```

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.

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.