

Oracle® Application Server

Performance Guide

10g Release 3 (10.1.3.1.0)

B28942-01

September 2006

Oracle Application Server Performance Guide 10g Release 3 (10.1.3.1.0)

B28942-01

Copyright © 2001, 2006, Oracle. All rights reserved.

Primary Author: Thomas Van Raalte

Contributors: Eric Belden, Alice Chan, Greg Cook, Bill Danell, Marcelo Goncalves, Helen Grembowicz, Bruce Irvin, Pushkar Kapasi, Paul Lane, Sharon Malek, Valarie Moore, Carol Orange, Julia Pond, Leela Rao, Ed Rybak, Joan Silverman, Cheryl Smith, Zhunquin Wang, Brian Wright

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software--Restricted Rights (June 1987). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Contents

Preface	ix
Intended Audience.....	ix
Documentation Accessibility	ix
Related Documentation.....	x
Conventions	x
1 Performance Overview	
Introduction to Oracle Application Server Performance	1-2
Performance Terms.....	1-2
What Is Performance Tuning?	1-2
Response Time.....	1-3
System Throughput.....	1-4
Wait Time.....	1-4
Critical Resources.....	1-4
Effects of Excessive Demand	1-5
Adjustments to Relieve Problems.....	1-5
Performance Targets	1-6
User Expectations.....	1-6
Performance Evaluation.....	1-6
Performance Methodology	1-6
Factors in Improving Performance.....	1-7
2 Monitoring Oracle Application Server	
Oracle Enterprise Manager 10g Application Server Control Console	2-2
Oracle Application Server Built-in Performance Metrics	2-2
Centralized Management of Oracle Application Server Instances	2-3
Native Operating System Performance Commands	2-3
Network Performance Monitoring Tools	2-3
3 Top Performance Areas	
Top Performance Areas	3-2
Ensure Sufficient Hardware Resources	3-3
Ensure Sufficient Java Heap for OC4J.....	3-3
Tune the JVM Garbage Collection Options.....	3-4
Reuse Database Connections.....	3-6

Specify Sufficient Oracle HTTP Server Connections	3-7
Enable Statement Caching for Data Sources	3-8
Verify Database Tuning	3-8
Verify Logging Levels	3-10
Reuse EJB Instances	3-11
Advanced Performance Areas	3-11
Managing Concurrency and Limiting Connections.....	3-11
Load Balancing	3-17
Using the -XX:AppendRatio Option (on Standalone OC4J)	3-19
4 Additional Performance Areas	
Improving TopLink Performance	4-2
Improving JTA Performance	4-2
Configuring Two-Phase Commit Logging for Performance	4-2
Configuring JTA Data Sources for Performance	4-4
Monitoring JTA Resources.....	4-5
Improving EJB Performance	4-5
Improving MDB Performance	4-5
Improving EJB CMP 2.1 Performance.....	4-9
5 Optimizing PL/SQL Performance	
6 Optimizing Oracle HTTP Server	
Configuring Oracle HTTP Server Directives	6-2
How Persistent Connections Can Reduce httpd Process Availability	6-3
Oracle HTTP Server Logging Options	6-3
Access Logging	6-3
Configuring the HostNameLookups Directive	6-3
Error logging.....	6-4
Oracle HTTP Server Security Performance Considerations	6-4
Oracle HTTP Server Secure Sockets Layer (SSL) Performance Issues	6-4
Oracle HTTP Server Port Tunneling Performance Issues	6-6
Oracle HTTP Server Performance Tips	6-6
Analyze Static Versus Dynamic Requests	6-7
Analyze Time Differences Between Oracle HTTP Server and OC4J Servers.....	6-7
Beware of a Single Data Point Yielding Misleading Results	6-7
7 Oracle BPEL Process Manager Performance Tuning	
Performance Tuning Overview	7-2
Domain and Process Configuration Property Settings.....	7-2
Durable and Transient Processes	7-2
One-Way and Two-Way Invocations.....	7-3
Idempotent Activities	7-3
In-Flight Database Storage.....	7-3
JTA Transactions for Two-way Invocations.....	7-4
BPEL Threading Model.....	7-4

Process Level Performance Settings	7-5
completionPersistLevel BPEL Property	7-5
completionPersistPolicy BPEL Property	7-6
idempotent BPEL Property	7-7
inMemoryOptimization BPEL Property	7-8
nonBlockingInvoke BPEL Property	7-9
Tables Impacted By Instance Data Growth	7-10
Domain Level Performance Tuning	7-11
Oracle BPEL Control Properties That Cannot Be Edited	7-11
auditDetailThreshold BPEL Property	7-11
auditLevel BPEL Property	7-12
bpelcClasspath BPEL Property	7-12
datasourceJndi BPEL Property	7-13
deliveryPersistPolicy BPEL Property	7-13
dspAgentDelay BPEL Property	7-14
dspInvokeAllocFactor BPEL Property	7-14
dspMaxRequestDepth BPEL Property	7-14
dspMaxThreads BPEL Property	7-14
dspMinThreads BPEL Property	7-15
expirationMaxRetry BPEL Property	7-15
idempotentThreshold BPEL Property	7-16
instanceKeyBlockSize BPEL Property	7-16
instCacheHighWatermark BPEL Property	7-16
instCacheLowWatermark BPEL Property	7-17
instCachePolicy BPEL Property	7-18
invokerQueueConnectionPoolMinSize BPEL Property	7-18
largeDocumentThreshold BPEL Property	7-18
minBPELWait BPEL Property	7-19
optCacheOn BPEL Property	7-19
optIdempotentRouting BPEL Property	7-20
optSoapShortcut BPEL Property	7-20
processCheckSecs BPEL Property	7-20
relaxBpelAssignRules BPEL Property	7-21
slowPerfThreshold BPEL Property	7-21
statsLastN BPEL Property	7-21
syncMaxWaitTime BPEL Property	7-22
txDatasourceJndi BPEL Property	7-22
uddiLocation BPEL Property	7-22
validateXML BPEL Property	7-22
workerQueueConnectionPoolMinSize BPEL Property	7-22
Tuning OC4J for Oracle BPEL	7-23
Tuning JTA Transaction Timeout for Oracle BPEL Process Manager	7-23
Oracle BPEL Server EJB Configuration	7-23
Configuring Data Sources for Oracle BPEL	7-24
Java Virtual Machine Performance Tuning for Oracle BPEL Server	7-25
Dehydration Store Database Performance Tuning	7-26
Summary	7-26

8 Oracle Business Activity Monitoring Performance

Managing the Redo Log Files	8-2
Avoiding Frequent Log Switches and Checkpoints	8-3
Tuning the System Global Area	8-3
Database Re-Organization in the Presence of Deletion Activity.....	8-5
Configuring Multiple Plan Monitor Services and Enterprise Links	8-6

A Monitoring Using Built-in Performance Tools

Summary of Oracle Application Server Built-in Performance Metrics.....	A-2
Viewing Performance Metrics Using AggreSpy with Basic Installation	A-2
Using the AggreSpy Display	A-2
Viewing Performance Metrics Using AggreSpy with Web Server	A-4
Using the AggreSpy Display with Web Server	A-5
AggreSpy URL With a Proxy Server with Web Server	A-6
AggreSpy URL and Access Control with Web Server.....	A-7
AggreSpy Limitation When Using Load Balancing With Multiple Instances	A-8
Viewing Performance Metrics Using dmstool.....	A-8
Access Control for dmstool.....	A-9
Using dmstool to List the Names of All Metrics	A-11
Using dmstool to Report Values for Specific Performance Metrics.....	A-11
Using dmstool With the Interval and Count Options	A-12
Using dmstool to Report All Metrics with Metric Values.....	A-12
Using dmstool to Report All Metrics with Metric Values in XML Format	A-12
Using dmstool to Reset Metric Values	A-13
Using dmstool to View Metrics on a Remote Oracle Application Server System.....	A-13
Viewing Performance Metrics Using AggreSpy (for Standalone OC4J).....	A-13
Using Built-in Performance Metrics on Windows Systems.....	A-14

B Instrumenting Applications With DMS

Introducing DMS Performance Metrics	B-2
Instrumenting Applications With DMS Metrics.....	B-2
Monitoring DMS Metrics	B-2
Understanding DMS Terminology (Nouns and Sensors)	B-3
DMS Naming Conventions.....	B-6
Adding DMS Instrumentation To Java Applications.....	B-8
Including DMS Imports	B-9
Organizing Performance Data.....	B-9
Defining and Using Metrics for Timing.....	B-10
Defining and Using Metrics for Counting.....	B-11
Defining and Using Metrics for Recording Status Information (State Sensors)	B-12
Validating and Testing Applications Using DMS Metrics.....	B-13
Validating DMS Metrics.....	B-13
Testing DMS Metrics For Efficiency	B-14
Understanding DMS Security Considerations.....	B-15
Conditional Instrumentation Using DMS Sensor Weight.....	B-15
Dumping DMS Metrics To Files	B-15

Resetting and Destroying Sensors	B-16
DMS Coding Recommendations	B-16
Isolating Expensive Intervals Using PhaseEvent Metrics	B-17
Using A High Resolution Clock To Increase DMS Precision	B-17
Configuring DMS Clocks for Reporting Time for OC4J (Java)	B-18
Configuring DMS Clocks for Reporting Time for Oracle HTTP Server	B-20
Rolling Up DMS Data for Descendent Nouns	B-21

C Performance Metrics

Oracle HTTP Server Metrics	C-2
Oracle HTTP Server Child Server Metrics.....	C-2
Oracle HTTP Server Responses Metrics	C-3
Oracle HTTP Server Virtual Host Metrics.....	C-3
Aggregate Module Metrics	C-3
HTTP Server Module Metrics.....	C-4
Oracle HTTP Server mod_oc4j Metrics.....	C-4
Oracle HTTP Server SSL Metrics	C-6
JVM Metrics	C-6
JVM Properties Metrics	C-6
JDBC Metrics	C-7
JDBC Driver Metrics	C-7
JDBC Data Source Metrics	C-7
JDBC Driver Specific Connection Metrics	C-8
JDBC Data Source Specific Connection Metrics	C-8
JDBC Connection Source Metrics	C-9
JDBC Driver Statement Metrics	C-9
JDBC Data Source Statement Metrics.....	C-10
JDBC Connection Pool Stats Metrics	C-11
mod_plsql Metrics	C-11
Oracle Process Manager and Notification Server - OPMN Metrics	C-15
OPMN_PM Metric Table.....	C-15
OPMN_OC4J_PROC Table.....	C-16
OPMN_HOST_STATISTICS Metric Table	C-16
OPMN_IAS_INSTANCE Metric Table	C-16
OPMN_IAS_COMPONENT Table.....	C-17
OPMN_ONS Metrics	C-18
OPMN_APPCTX Table	C-20
DMS Internal Metrics	C-20

D OC4J Performance Metrics

JTA Resource Metrics	D-2
JCA Metrics	D-4
OC4J J2EE Application Metrics	D-6
Web Module Metrics	D-6
Web Context Metrics	D-7
OC4J Servlet Metrics.....	D-7

OC4J JSP Metrics	D-7
OC4J EJB Metrics	D-8
OC4J OPMN Info Metrics	D-12
OC4J Work Management Pool Metrics	D-12
OC4J JMS Metrics	D-13
JMS Metric Tables	D-13
JMS Stats Metric Table.....	D-14
JMS Request Handler Stats	D-16
JMS Connection Stats.....	D-16
JMS Session Stats	D-16
JMS Message Producer Stats.....	D-17
JMS Message Browser Stats	D-17
JMS Message Consumer Stats	D-18
JMS Durable Subscription Stats	D-18
JMS Destination Stats.....	D-18
JMS Temporary Destination Stats.....	D-19
JMS Store Stats	D-19
JMS Persistence Stats	D-20
OC4J Task Manager Metrics.....	D-20
Java Object Cache JOC Metrics	D-20

Index

Preface

This guide describes how to monitor and optimize performance, use multiple components for optimal performance, and write highly performant applications in the Oracle Application Server environment.

This preface contains these topics:

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

Intended Audience

Oracle Application Server Performance Guide is intended for Internet application developers, Oracle Application Server administrators, database administrators, and Web masters.

Documentation Accessibility

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

<http://www.oracle.com/accessibility/>

Accessibility of Code Examples in Documentation

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

Accessibility of Links to External Web Sites in Documentation

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

TTY Access to Oracle Support Services

Oracle provides dedicated Text Telephone (TTY) access to Oracle Support Services within the United States of America 24 hours a day, seven days a week. For TTY support, call 800.446.2398.

Related Documentation

For more information, see these Oracle resources:

- *Oracle HTTP Server Administrator's Guide*
- *Oracle Application Server Containers for J2EE User's Guide*
- *Oracle Containers for J2EE Enterprise JavaBeans Developer's Guide*
- *Oracle Containers for J2EE Servlet Developer's Guide*
- *Oracle Containers for J2EE JSP Tag Libraries and Utilities Reference*
- *Oracle Database Performance Tuning Guide, 10g*
- *Oracle Application Server PL/SQL Web Toolkit Reference*

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.

Performance Overview

This chapter discusses Oracle Application Server performance and tuning concepts.

This chapter contains the following sections:

- [Introduction to Oracle Application Server Performance](#)
- [What Is Performance Tuning?](#)
- [Performance Targets](#)
- [Performance Methodology](#)

1.1 Introduction to Oracle Application Server Performance

To maximize Oracle Application Server performance, all components need to be monitored, analyzed, and tuned. The chapters of this guide describe the tools used to monitor performance and the techniques for optimizing the performance of Oracle Application Server components, such as Oracle HTTP Server and Oracle Containers for J2EE (OC4J).

1.1.1 Performance Terms

Following are performance terms used in this book:

concurrency The ability to handle multiple requests simultaneously. Threads and processes are examples of concurrency mechanisms.

contention Competition for resources.

hash A number generated from a string of text with an algorithm. The hash value is substantially smaller than the text itself. Hash numbers are used for security and for faster access to data.

latency The time that one system component spends waiting for another component in order to complete the entire task. Latency can be defined as wasted time. In networking contexts, latency is defined as the travel time of a packet from source to destination.

response time The time between the submission of a request and the receipt of the response.

scalability The ability of a system to provide **throughput** in proportion to, and limited only by, available hardware resources. A scalable system is one that can handle increasing numbers of requests without adversely affecting response time and **throughput**.

service time The time between the receipt of a request and the completion of the response to the request.

think time The time the user is not engaged in actual use of the processor.

throughput The number of requests processed per unit of time.

wait time The time between the submission of the request and initiation of the request.

1.2 What Is Performance Tuning?

Performance must be built in. You must anticipate performance requirements during application analysis and design, and balance the costs and benefits of optimal performance. This section introduces some fundamental concepts:

- [Response Time](#)
- [System Throughput](#)
- [Wait Time](#)
- [Critical Resources](#)

- [Effects of Excessive Demand](#)
- [Adjustments to Relieve Problems](#)

See Also: ["Performance Targets"](#) on page 1-6 for a discussion on performance requirements and determining what parts of the system to tune.


1.2.1 Response Time

Because **response time** equals **service time** plus **wait time**, you can increase performance in this area by:

- Reducing **wait time**
- Reducing **service time**

[Figure 1–1](#) illustrates ten independent sequential tasks competing for a single resource as time elapses.

Figure 1–1 *Sequential Processing of Independent Tasks*



In the example shown in [Figure 1–1](#), only task 1 runs without waiting. Task 2 must wait until task 1 has completed; task 3 must wait until tasks 1 and 2 have completed, and so on. Although the figure shows the independent tasks as the same size, the size of the tasks will vary.

In parallel processing with multiple resources, more resources are available to the tasks. Each independent task executes immediately using its own resource and no **wait time** is involved.

The Oracle HTTP Server processes requests in this fashion, allocating client requests to available `httpd` processes (or threads). The `MaxClients` directive specifies the maximum number of `httpd` processes simultaneously available to handle client requests. When the number of processes in use reaches the `MaxClients` value, the server refuses connections until requests are completed and processes are freed.

See Also: [Chapter 6, "Optimizing Oracle HTTP Server"](#)

1.2.2 System Throughput

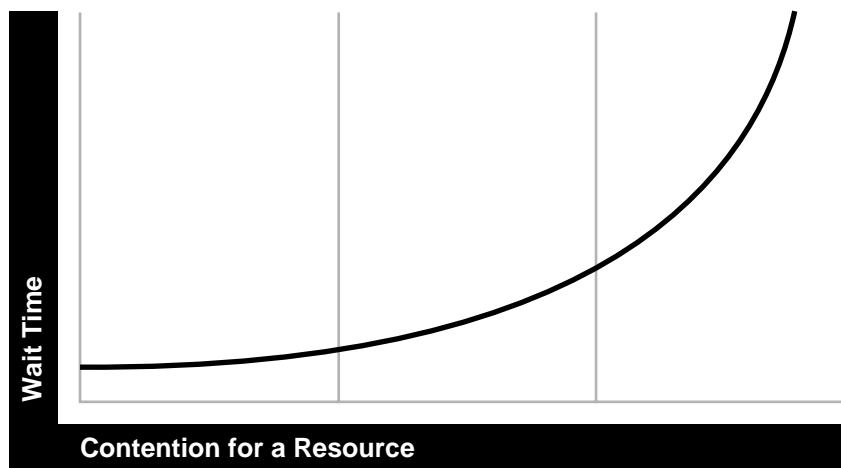
System **throughput** is the amount of work accomplished in a given amount of time. You can increase **throughput** by:

- Reducing **service time**
- Reducing overall **response time** by increasing the amount of scarce resources available. For example, if the system is CPU bound, then adding CPU resources should improve performance.

1.2.3 Wait Time

While the **service time** for a task may stay the same, **wait time** will lengthen with increased **contention**. If many users are waiting for a service that takes one second, the tenth user must wait 9 seconds. [Figure 1–2](#) shows the relationship between **wait time** and resource **contention**. In the figure, the graph illustrates that wait time increases exponentially as contention for a resource increases.

Figure 1–2 Wait Time Rising With Increased Contention for a Resource

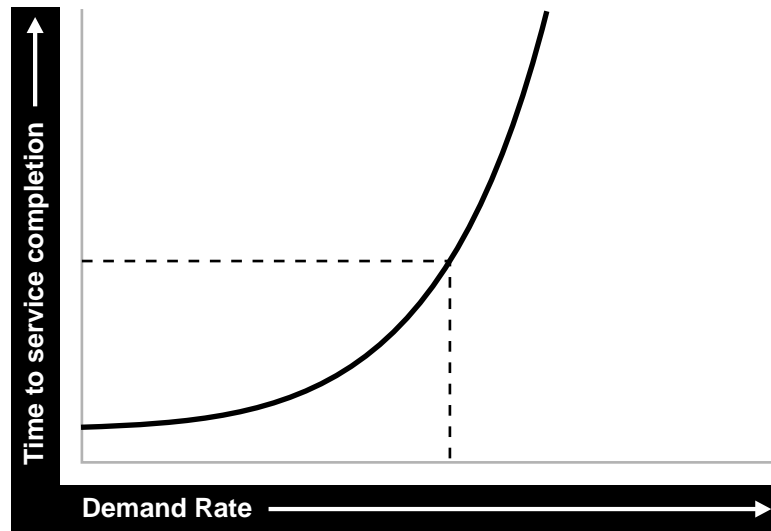


1.2.4 Critical Resources

Resources such as CPU, memory, I/O capacity, and network bandwidth are key to reducing **service time**. Adding resources increases **throughput** and reduces **response time**. Performance depends on these factors:

- How many resources are available?
- How many clients need the resource?
- How long must they wait for the resource?
- How long do they hold the resource?

[Figure 1–3](#) shows the relationship between time to service completion and demand rate. The graph in the figure illustrates that as the number of units requested rises, the time to service completion increases.


Figure 1–3 Time to Service Completion Versus Demand Rate

To manage this situation, you have two options:

- Limit demand rate to maintain acceptable response times
- Add resources

1.2.5 Effects of Excessive Demand

Excessive demand increases **response time** and reduces **throughput**, as illustrated by the graph in [Figure 1–4](#).

Figure 1–4 Increased Demand/Reduced Throughput

If the demand rate exceeds the achievable **throughput**, then determine through monitoring which resource is exhausted and increase the resource, if possible.

1.2.6 Adjustments to Relieve Problems

Performance problems can be relieved by making adjustments in the following:

- unit consumption
Reducing the resource (CPU, memory) consumption of each request can improve performance. This might be achieved by pooling and caching.
- functional demand
Rescheduling or redistributing the work will relieve some problems.
- capacity
Increasing or reallocating resources (such as CPUs) relieves some problems.

1.3 Performance Targets

Whether you are designing or maintaining a system, you should set specific performance goals so that you know how and what to optimize. If you alter parameters without a specific goal in mind, you can waste time tuning your system without significant gain.

An example of a specific performance goal is an order entry **response time** under three seconds. If the application does not meet that goal, identify the cause (for example, I/O **contention**), and take corrective action. During development, test the application to determine if it meets the designed performance goals.

Tuning usually involves a series of trade-offs. After you have determined the bottlenecks, you may have to modify performance in some other areas to achieve the desired results. For example, if I/O is a problem, you may need to purchase more memory or more disks. If a purchase is not possible, you may have to limit the **concurrency** of the system to achieve the desired performance. However, if you have clearly defined goals for performance, the decision on what to trade for higher performance is easier because you have identified the most important areas.

1.3.1 User Expectations

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

1.3.2 Performance Evaluation

With clearly defined performance goals, you can readily determine when performance tuning has been successful. Success depends on the functional objectives you have established with the user community, your ability to measure whether or not the criteria are being met, and your ability to take corrective action to overcome any exceptions.

Ongoing performance monitoring enables you to maintain a well tuned system. Keeping a history of the application's performance over time enables you to make useful comparisons. With data about actual resource consumption for a range of loads, you can conduct objective **scalability** studies and from these predict the resource requirements for anticipated load volumes.


1.4 Performance Methodology

Achieving optimal effectiveness in your system requires planning, monitoring, and periodic adjustment. The first step in performance tuning is to determine the goals you

need to achieve and to design effective usage of available technology into your applications. After implementing your system, it is necessary to periodically monitor and adjust your system. For example, you might want to ensure that 90% of the users experience **response times** no greater than 5 seconds and the maximum **response time** for all users is 20 seconds. Usually, it's not that simple. Your application may include a variety of operations with differing characteristics and acceptable response times. You need to set measurable goals for each of these.

You also need to determine variances in the load. For example, users might access the system heavily between 9:00am and 10:00am and then again between 1:00pm and 2:00pm, as illustrated by the graph in [Figure 1-5](#). If your peak load occurs on a regular basis, for example, daily or weekly, the conventional wisdom is to configure and tune systems to meet your peak load requirements. The lucky users who access the application in off-time will experience better **response times** than your peak-time users. If your peak load is infrequent, you may be willing to tolerate higher **response times** at peak loads for the cost savings of smaller hardware configurations.

Figure 1-5 Adjusting Capacity and Functional Demand



1.4.1 Factors in Improving Performance

Performance spans several areas:

- Sizing and configuration: Determining the type of hardware needed to support your performance goals.
- Parameter tuning: Setting configurable parameters to achieve the best performance for your application.
- Performance monitoring: Determining what hardware resources are being used by your application and what **response time** your users are experiencing.
- Troubleshooting: Diagnosing why an application is using excessive hardware resources, or why the **response time** exceeds the desired limit.

Monitoring Oracle Application Server

This chapter provides an overview and presents information on monitoring Oracle Application Server and its components. Monitoring Oracle Application Server and obtaining performance data can assist you in tuning the system and debugging applications with performance problems.

This chapter contains the following sections:

- [Oracle Enterprise Manager 10g Application Server Control Console](#)
- [Oracle Application Server Built-in Performance Metrics](#)
- [Centralized Management of Oracle Application Server Instances](#)
- [Native Operating System Performance Commands](#)
- [Network Performance Monitoring Tools](#)

2.1 Oracle Enterprise Manager 10g Application Server Control Console

Oracle Enterprise Manager 10g Application Server Control Console (Application Server Control Console) allows you to monitor Oracle Application Server and its components. Application Server Control Console shows performance metrics for Oracle Application Server components, including:

- Oracle Containers for J2EE (OC4J) and Applications running under OC4J

Using Application Server Control Console, you can also view performance metrics and other status information using Application Server Control Console.

See Also: *Oracle Application Server Administrator's Guide*

2.2 Oracle Application Server Built-in Performance Metrics

Oracle Application Server automatically measures runtime performance and collects metrics for Oracle HTTP Server and Oracle Containers for J2EE (OC4J) servers and components. The server performance metrics are measured automatically and continuously using performance instrumentation inserted into the implementations of Oracle Application Server components. The performance metrics are automatically enabled; you do not need to set options or perform any extra configuration to collect them (for performance reasons the JDBC metrics are enabled by setting options).

The Oracle HTTP Server performance metrics enable you to do the following:

- Monitor the duration of important phases of Oracle HTTP Server request processing.
- Collect status information on Oracle HTTP Server requests. For example, you can monitor the number of requests being handled at any given moment.

The OC4J performance metrics allow you to monitor the performance of J2EE containers and components and enable you to do the following:

- Monitor the number of active servlets, JSPs, EJBs, and EJB methods.
- Monitor the time spent processing an individual servlet, JSP, EJB, or EJB method.
- Monitor the sessions and JDBC connections associated with servlets, JSPs, EJBs, or EJB methods.
- Monitor OC4J JMS events and status.

You can use the performance metrics while troubleshooting Oracle Application Server components to help locate bottlenecks, identify resource availability issues, or help tune your components to improve throughput and response times.

Note: You can use the commands that access the built-in metrics in scripts or in combination with other monitoring tools to gather performance data or to check application performance.

See Also:

- [Appendix A, "Monitoring Using Built-in Performance Tools"](#)
- [Appendix C, "Performance Metrics"](#)

2.3 Centralized Management of Oracle Application Server Instances

While Application Server Control Console provides standalone management for an Application Server and its components, you can centrally manage all your Application Servers through one tool rather than through several Application Server Control Consoles by using the Oracle Enterprise Manager 10g Grid Control Console. For example, say you have 10 Application Servers deployed on five hosts. By deploying a Management Agent on each host, Enterprise Manager automatically discovers the Application Server on those hosts and automatically begins monitoring them using default monitoring levels, notification rules, and so on.

The Oracle Enterprise Manager 10g Grid Control Console contains an Application Server Home page which provides easy access to key information required by application server administrators, including the following:

- Links to Oracle Application Server component home pages
- Application server status, responsiveness, and performance data
- Alerts and diagnostic drill-downs so you can identify and resolve problems quickly
- Resource usage for the application server and its components
- A single view of all Java 2 Platform Enterprise Edition (J2EE) applications and web services
- Links to the Application Server Control Console for administration operations such as starting and stopping components, modifying configurations, and deploying applications.

See Also:

Oracle Enterprise Manager Concepts for more information on Oracle Enterprise Manager 10g Grid Control Console

Oracle Application Server Administrator's Guide

Oracle Enterprise Manager Grid Control Installation and Basic Configuration

2.4 Native Operating System Performance Commands

In order to solve performance problems or to monitor your system's activity, you can use the available native operating system commands. Native operating system commands allow you to gather and monitor CPU utilization, paging activity, swapping, and other system activity information.

See Also: Refer to the system level documentation for information on native operating system monitoring commands

2.5 Network Performance Monitoring Tools

You can use network monitoring tools to verify the status of requests that access your Oracle Application Server components. Tools are available that allow you to examine and save network traffic information. These tools can be helpful in analyzing and solving performance problems.

Top Performance Areas

This chapter provides a description of top tuning areas for Oracle Application Server and includes the following sections:

- [Top Performance Areas](#)
- [Advanced Performance Areas](#)

3.1 Top Performance Areas

This section covers critical Oracle Application Server performance issues and provides a quickstart for tuning J2EE applications that run on OC4J. [Table 3–1](#) lists a quick guide for performance considerations for Oracle Application Server.

Table 3–1 Top Performance Areas for Oracle Application Server Applications

Performance Area	Description and Reference
Ensure Sufficient Hardware Resources	<p>Oracle Application Server has minimum hardware requirements that enable you to use the product. You also need to run on hardware that supports the needs for your applications, including the database tier. For details on Oracle Application Server installation requirements, see the Requirements chapter in the <i>Installation Guide</i> for your platform.</p> <p>See "Ensure Sufficient Hardware Resources" on page 3-3 for more information on platform specific tools that can help you determine if your hardware resources are sufficient.</p>
Ensure Sufficient Java Heap Size	<p>To improve the performance of a J2EE application, you need to provide an adequate heap size for the JVM where the application runs. If the OC4J running your J2EE application does not have enough memory, performance can suffer due to the overhead required to manage limited memory.</p> <p>See "Ensure Sufficient Java Heap for OC4J" on page 3-3 for details on setting JVM heap size options.</p>
Tune the JVM for Garbage Collection	<p>To improve the performance of a J2EE application and ensure that your application performance is not being negatively impacted by JVM garbage collection, you need to manage the heap size and sometimes you also need to set JVM options that impact garbage collection frequency.</p> <p>See "Tune the JVM Garbage Collection Options" on page 3-4 for more information on garbage collection.</p>
Reuse Database Connections	<p>By default, OC4J data sources enable database connection pooling. Thus, OC4J manages database connections to avoid frequently reestablishing new connections. The Oracle Application Server data source facility provides options you can use to control the number of database connections maintained, and how long they are maintained.</p> <p>See "Reuse Database Connections" on page 3-6.</p>
Specify Sufficient HTTP Connections	<p>Tune the Oracle HTTP Server directives to set the level of concurrency by specifying the number of HTTP connections.</p> <p>See "Specify Sufficient Oracle HTTP Server Connections" on page 3-7.</p>
Enable JDBC Statement Caching Option	<p>By enabling statement caching to lower the overhead of repeated cursor creation and repeated statement parsing and creation, you can improve performance for applications using the database tier.</p> <p>See "Enable Statement Caching for Data Sources" on page 3-8.</p>
Ensure the Database is Properly Tuned	<p>For applications that access a database, ensure that your database is properly configured to support your application's requirements.</p> <p>See "Verify Database Tuning" on page 3-8.</p>
Verify Logging Levels	<p>You need to make sure 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.</p> <p>See "Verify Logging Levels" on page 3-10.</p>
Reuse EJB Instances	<p>Set the EJB options for creating and reusing instances to improve EJB performance.</p> <p>See "Reuse EJB Instances" on page 3-11.</p>

3.1.1 Ensure Sufficient Hardware Resources

A most crucial performance area is ensuring that there are sufficient CPU, memory and network resources to support the user population and application requirements for your Oracle Application Server installation. You need to monitor resource utilization over an extended period to determine if you have occasional peaks of usage or whether a resource is consistently saturated. You also need to define the acceptable response times and throughputs for applications running at your site, for both peak and extended periods. Also, check the system while running your application under normal load and monitor operating system statistics, including, CPU, memory, disk, and network performance to determine if any hardware resource is saturated.

To check the CPU, memory, and disk performance you can use the following commands:

On Linux systems use the `sar` or `mpstat` command.

On Windows systems use the `perfmon` command.

To check network performance, you can use the following commands:

On Linux and Windows systems:

```
% netstat
```

On Windows systems, you can also use the Windows Task Manager to check network performance.

If any of the hardware resources are saturated, this could be due to one or more of the following:

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

For a consistently saturated resource, the solutions are to reduce load or increase resources. For peak traffic periods, if the increased response time is not acceptable the alternatives are to again increase resources or to determine if there is traffic that can be rescheduled to reduce the peak load, such as scheduling batch or background operations during slower periods. Oracle Application Server provides a variety of mechanisms to help you control resource concurrency; this can limit the impact of bursts of traffic. However, for a consistently saturated system, these mechanisms should be viewed as temporary solutions.

See Also:

- ["Specify Sufficient Oracle HTTP Server Connections"](#) on page 3-7
- ["Managing Concurrency and Limiting Connections"](#) on page 3-11

3.1.2 Ensure Sufficient Java Heap for OC4J

If you have sufficient memory available on your system and your application is memory intensive, increase the JVM heap size from the default value. While the amount of heap required varies based on the application and on the available memory, for most OC4J server applications, if you have sufficient memory, then Oracle recommends using an initial heap size of 512 Megabytes or larger.

You can improve performance by setting the initial heap size equal to the maximum heap size.

Take the following steps to change the heap size values for an OC4J instance:

1. Navigate to the Home page for the OC4J instance.
2. Click Administration.
3. If necessary, expand the Properties section of the table by clicking the Expand icon. Then, click the Go to Task icon in the Server Properties row.
4. In the Command Line Options area, change the value in the **Maximum heap size** and **Initial heap size** fields.
5. Click Apply.
6. Navigate to the Cluster Topology page, select the OC4J instance that you modified, and click Restart. On the Confirmation page, click Yes.

This specifies the following JVM options and changes the size of the heap allocated to the OC4J process in an OC4J instance.

If your Oracle Application Server topology includes more than one JVM on the same system, then to maximize performance, set the maximum heap size to accommodate application requirements and make sure that the total memory consumed by all of the JVMs running on the system does not exceed the memory capacity of your system.

See Also:

- The JVM Metrics page in Application Server Control Console. This page is available from the OC4J home page, by clicking the **Performance** Secondary tab, and then, in the Related Links area, clicking **JVM Metrics**.
- You can find detailed information about JVM options and their impact on performance on the JVM vendor's Web sites, such as http://java.sun.com/performance/reference/whitepapers/5.0_performance.html

3.1.3 Tune the JVM Garbage Collection Options

JVM garbage collection is an expensive and can have an impact on application performance; inefficient garbage collection can severely degrade application performance. Therefore, it is important to understand how applications create and destroy objects.

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

In order to determine application memory requirements you can monitor JVM garbage collection and memory pool sizes using the following:

- The JVM command line options:

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

Look for "Full GC" to identify major collections.

- jstat tool
- visualgc tool
- The Application Server Control Console JVM Metrics page shows JVM memory pool and garbage collector information. This page is available from the OC4J home

page by clicking the **Performance** Secondary tab, and then, in the Related Links area, clicking **JVM Metrics**.

Set the `-XX:+AggressiveHeap` JVM option to tune internal VM parameters, and increase the total heap size, as described in "[Ensure Sufficient Java Heap for OC4J](#)" on page 3-3 to reduce the overhead associated with Full CG garbage collections. The `-XX:+AggressiveHeap` option tunes internal VM parameters to be optimal for long-running, memory-intensive workloads. This option should follow, on the command line, the heap sizing options `-Xms` and `-Xmx`. See "[Ensure Sufficient Java Heap for OC4J](#)" on page 3-3 for details on setting Java command line options in an Oracle Application Server managed environment (please disregard older versions of Sun documentation which advise against using `-XX:+AggressiveHeap`).

Note: The JVM provides a variety of parameters to allow you to more finely tune heap management and garbage collection behavior. See the following links for a detailed description of these topics.

To determine if an application uses explicit garbage collection, which can have a negative impact on performance, set the `-XX:+DisableExplicitGC` option. This debugging option disables explicit garbage collection. Applications should avoid the use of `system.gc()` calls. If you suspect an application may be explicitly triggering garbage collection, set this parameter and observe the differences in your garbage collection behavior. If you determine that the application is making explicit `system.gc()` calls, discuss with the application developer why this was done and the impact of disabling the calls. Application developers sometimes use `system.gc()` calls to trigger finalizers. This is not a recommended practice and can yield indeterminate behavior.

Take the following steps to change the JVM command line options:

1. Navigate to the Home page for the OC4J instance.
2. Click Administration.
3. If necessary, expand the Properties section of the table by clicking the Expand icon. Then, click the Go to Task icon in the Server Properties row.
4. In the Command Line Options area, modify or change the appropriate command line options in the Options table.
5. Click Apply.
6. Navigate to the Cluster Topology page, select the OC4J instance that you modified, and click Restart. On the Confirmation page, click Yes.

See Also:

- http://java.sun.com/j2se/1.5/pdf/jdk50_ts_guide.pdf
- <http://java.sun.com/docs/hotspot/gc5.0/ergo5.html>
- http://java.sun.com/docs/hotspot/gc5.0/gc_tuning_5.html

3.1.4 Reuse Database Connections

To obtain better performance in your application, by lowering the overhead of creating and recreating database connections, specify the connection pool `min-connections` attribute to set the minimum number of connections that the connection pool maintains.

By default, the value of `min-connections` is 0. For best performance, you should specify a value for `min-connections` other than 0. If `min-connections` is set to a value other than zero, the specified number of connections is maintained; OC4J maintains the connections when they are not in use and they do not time out when the specified `inactivity-timeout` is reached. The `min-connections` attribute does not specify that OC4J pre-create connections at startup. Specify the `initial-limit` attribute to set the number of connections in the connection pool when the pool is initially created or reinitialized. Oracle recommends that you set the `initial-limit` attribute to the same value as the `min-connections` attribute.

If the specified value for `min-connections` is less than `max-connections`, then you should set the `inactivity-timeout` to make sure that connections only time out after an appropriately long period of inactivity. The connection pool `inactivity-timeout` specifies the time, in seconds, to cache unused connections before closing the connection.

To improve performance, you can set the `inactivity-timeout` to a value that allows the connection pool to avoid dropping and then reacquiring connections while your J2EE application is running. The default value for the `inactivity-timeout` is 60 seconds, which is typically too low for frequently accessed applications where there may be some inactivity between requests. For most applications, to improve performance, Oracle recommends that you increase the `inactivity-timeout` to 120 seconds.

To determine if the default `inactivity-timeout` is too low, monitor your system. If you see that the number of database connections grows and then shrinks during an idle period, and grows again soon after that, you have two options: you can increase the `inactivity-timeout`, or you can increase the `min-connections`.

Notes for reusing database connections:

- Limiting the total number of open database connections to a number your database can handle is an important tuning consideration. You should check with your database administrator to make sure that the database is configured to support a number of connections that is greater than the following:

At least as large a number of connections as the sum of the values specified for all the connection pool `min-connections` that could be concurrently active, and as large as the maximum desired concurrency across all the datasources for the database.

- If the `min-connections` is set to a value other than zero, the specified number of connections is maintained; OC4J maintains the connections when they are not in use, and they do not time out when the specified `inactivity-timeout` is reached.

Once the specified connections are opened, you need to either stop OC4J or use the refresh operation to close the connections. Application Server Control shows the refresh operation in Connection Pool area on the JDBC Resources page. Click the icon in the **Refresh Connection Pool** field to initiate a refresh operation.

3.1.5 Specify Sufficient Oracle HTTP Server Connections

The Oracle HTTP Server `MaxClients` directive limits the number of clients that can simultaneously connect to your web server, and thus the number of `httpd` processes.

For Windows, the analogous parameter is `ThreadsPerChild`. The `Oc4jCacheSize` directive specifies the maximum number of idle connections that `mod_oc4j` maintains per OC4J JVM, relevant only on Windows.

You can use the `MaxClients`, `ThreadsPerChild`, and `Oc4jCacheSize` directives to limit incoming connections to the OC4J instances from the Oracle HTTP Server. This section covers the following topics:

- [Configuring the MaxClients Directive \(for UNIX\)](#)
- [Configuring the ThreadsPerChild Directive \(for Windows\)](#)
- [Configuring the Oc4jCacheSize Directive](#)

Note: The discussion in this section only applies for the default Oracle HTTP Server supplied with Oracle Application Server (based on Apache 1.3). This discussion does not apply for the Apache 2.0 based standalone version of Oracle HTTP Server.

3.1.5.1 Configuring the MaxClients Directive (for UNIX)

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

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

If the server handles persistent connections, then it may require sufficient concurrent `httpd` server processes to handle both active and idle connections. When you specify `MaxClients` to act as a throttle for system concurrency, you need to consider that persistent idle `httpd` connections also consume `httpd` processes. Specifically, the number of connections includes the currently active persistent and non-persistent connections and the idle persistent connections. A persistent, `KeepAlive`, HTTP connection consumes an `httpd` child process, or thread, for the duration of the connection, even if no requests are currently being processed for the connection.

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

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

When there are no `httpd` processes available, connection requests are queued in the TCP/IP system until a process becomes available, and eventually clients terminate connections.

3.1.5.2 Configuring the `ThreadsPerChild` Directive (for Windows)

You can configure the `ThreadsPerChild` directive in the `httpd.conf` file up to a maximum of 8K (the default value is 50). The `ThreadsPerChild` parameter on Windows systems works like the `MaxClients` parameter on UNIX systems.

See Also: ["Configuring the MaxClients Directive \(for UNIX\)"](#) on page 3-7

3.1.5.3 Configuring the `Oc4jCacheSize` Directive

The `Oc4jCacheSize` directive specifies the maximum number of idle connections that `mod_oc4j` maintains per OC4J JVM. On Windows only, it is sometimes useful to change the default value of this directive.

On UNIX systems where each Oracle HTTP Server process is single threaded, the only meaningful values are 1 which is the default value, and zero (0). A value of zero (0) specifies that Oracle HTTP Server should not maintain any connections and should open a new connection for every request. Since each process is single threaded, a process never needs more than one connection and hence a value of 1 or greater has the same effect on UNIX systems. For best performance, on UNIX systems, do not change the default value for `Oc4jCacheSize`.

On Windows systems, the default `Oc4jCacheSize` value is 75% of the value of `ThreadsPerChild`; the connection cache is shared among threads in the child process. If Oracle HTTP Server is serving a mixed load of static content along with OC4J requests, then the default should be adequate. If the user's load is all OC4J requests, that is, Oracle HTTP Server serves up little or no content and serves just as a front end for OC4J, then it is a good idea to set `Oc4jCacheSize` equal to `ThreadsPerChild`. This setting provides a dedicated connection per thread, if needed, and should give the best performance.

3.1.6 Enable Statement Caching for Data Sources

Enable statement caching to lower the overhead of repeated cursor creation and repeated statement parsing and creation by setting the `num-cached-statements` attribute to a value greater than 0 (the default value is 0, disabled). The number you set for `num-cached-statements` should be the number of SQL statements that you use in your application.

See Also: ["Statement Caching with Managed Data Sources"](#) in *Oracle Containers for J2EE Services Guide*

3.1.7 Verify Database Tuning

To achieve optimal performance in Oracle Application Server, for applications that use the database, the database tables you access need to be designed with performance in mind and you need to monitor and tune the database server to assure that the system is performant.

This section covers the following:

- [Tuning `init.ora` Database Parameters](#)
- [Tuning Redo Logs Location and Sizing](#)

- [Automatic Segment-Space Management \(ASSM\)](#)

See Also: *Oracle Database Performance Tuning Guide*

3.1.7.1 Tuning `init.ora` Database Parameters

[Table 3–2](#) shows tuning information for several of the `init.ora` database initialization parameters.

Table 3–2 Important `init.ora` Database Tuning Parameters

<code>init.ora</code> Parameter	Description
<code>DB_BLOCK_SIZE</code>	<p>The default block size of 8K is optimal for most systems. However, OLTP systems occasionally benefit from smaller block sizes, and DSS systems occasionally benefit from larger block sizes.</p> <p>See Also: table 8-3, "Block Size Advantages and Disadvantages" in the <i>Oracle Database Performance Tuning Guide</i>.</p>
<code>PGA_AGGREGATE_TARGET</code>	<p>Specifies the target aggregate PGA memory available to all server processes attached to the instance.</p> <p>See Also: "Memory Configuration and Use" in the <i>Oracle Database Performance Tuning Guide</i> for information on PGA memory management.</p>
<code>PROCESSES</code>	<p>Sets the maximum number of operating system processes that can be connected to Oracle concurrently. The value of this parameter must be 6 or greater (5 for the background processes plus 1 for each user process). For example, if you plan to have 50 concurrent users, set this parameter to at least 55. Many other initialization parameter values are deduced from this value.</p>
<code>SGA_MAX_SIZE</code>	<p>This parameter is the maximum size of the SGA for a running instance. Set this parameter to the amount of memory that you want dedicated for the SGA, which includes the following memory pools:</p> <ul style="list-style-type: none"> ■ Database buffer cache ■ Shared pool ■ Large pool ■ Java pool <p>It is a good practice to regularly monitor the buffer cache hit ratio and size the SGA so that the buffer cache has an adequate number of frames for the workload. The buffer cache hit ratio may be calculated from data in the view <code>V\$SYSSTAT</code>. Also the view <code>V\$DB_CACHE_ADVICE</code> provides data that can be used to tune the buffer cache.</p> <p>See Also: the chapter, "Memory Configuration and Use" in the <i>Oracle Database Performance Tuning Guide</i> for detailed information on how to set the <code>SGA_MAX_SIZE</code> parameter, including on how to use the <code>V\$SYSSTAT</code> and <code>V\$DB_CACHE_ADVICE</code> views to optimize the buffer cache hit ratio.</p>
<code>SGA_TARGET</code>	<p>Setting this parameter to a nonzero value enables Automatic Shared Memory Management. Oracle strongly recommends the use of automatic memory management, both to simplify configuration and to improve performance. Automatic Shared Memory Management was introduced with the Oracle Database 10g (10.1). For prior versions, you must manually configure individual SGA memory pools.</p> <p>See Also: the section, "Automatic Shared Memory Management" in the Chapter, "Memory Configuration and Use" in the <i>Oracle Database Performance Tuning Guide</i> for details on choosing a value for the <code>SGA_TARGET</code> parameter.</p>
<code>UNDO_TABLESPACE</code> <code>UNDO_MANAGEMENT</code>	<p>Oracle strongly recommends that you use automatic undo management (<code>UNDO_MANAGEMENT = AUTO</code>) and manage undo space using an <code>UNDO_TABLESPACE</code>. For backward compatibility reasons, the default value of <code>UNDO_MANAGEMENT</code> is <code>MANUAL</code>.</p> <p>See Also: <i>Oracle Database Performance Tuning Guide</i> for additional information on undo space management.</p>

3.1.7.2 Tuning Redo Logs Location and Sizing

Managing the database I/O load balancing is a non-trivial task. However, tuning the redo log options can provide performance improvement for applications running in an Oracle Application Server environment, and in some cases, you can significantly improve I/O throughput by moving the redo logs to a separate disk.

The size of the redo log files can also 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 by reducing checkpoint activity. It is not possible to provide a specific size recommendation for redo log files, but redo log files in the range of a hundred megabytes to a few gigabytes are considered reasonable. Size your online redo log files according to the amount of redo your system generates. A rough guide is to switch logs at most once every twenty minutes. Set the initialization parameter `LOG_CHECKPOINTS_TO_ALERT = true` to have checkpoint times written to the alert file.

The complete set of required redo log files can be created during database creation. After they are created, the size of a redo log size cannot be changed. However, new, larger files can be added later, and the original (smaller) ones can subsequently be dropped.

See Also: The chapters, "Configuring a Database for Performance" and "I/O Configuration and Design" in the *Oracle Database Performance Tuning Guide*

3.1.7.3 Automatic Segment-Space Management (ASSM)

For permanent tablespaces, Oracle recommends 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`.

See Also: *Oracle Database Concepts* for a discussion of free space management, and *Oracle Database Administrator's Guide* for more information on creating and using automatic segment-space management for tablespaces.

3.1.8 Verify Logging Levels

You need to assure that application and server logging levels are set appropriately, and that debugging properties or other application level debugging flags are set to appropriate levels or disabled. Set Oracle Application Server OC4J logger log levels to log messages at the `INFO` level (do not set log levels to levels that produce more diagnostic message, including the `FINE` or `TRACE` levels).

To configure OC4J component loggers through Application Server Control Console, do the following from the OC4J home page:

1. Click the Administration link
2. In the table, under **Properties**, click the task for **Logger Configuration**.
3. In the table, under **Loggers**, set the root Log Level to the desired value, or expand the tree to select individual Log Levels for specified loggers.
4. Click **Apply** to apply your changes to the OC4J runtime.

3.1.9 Reuse EJB Instances

This section describes EJB tuning options for creating and reusing instances to improve EJB performance; these options are specific to OC4J and apply for all types of EJBs (except Stateful Session EJBs). You can configure these options by setting attributes in `orion-ejb-jar.xml`.

The `min-instances` attribute specifies the minimum number of bean implementation instances to be kept instantiated or pooled. The default value is 0. For best performance, you should specify a value for `min-instances` other than 0. If `min-instances` is > 0 , OC4J maintains the `min-instances` number of instances in the pool when they are not in use. For instances above the `min-instances`, the instances are removed from the pool after the `pool-cache-timeout` specified timeout expires. The `pool-cache-timeout` cache expiration removes all bean instances that have not been accessed during the timeout window. For example, with the default value for `pool-cache-timeout`, all beans that have not been touched in 60 seconds would be removed from the pool and active or recently used beans are left in the pool.

The default value for the `pool-cache-timeout` is 60 seconds, which is typically too low for frequently accessed EJBs. If the `pool-cache-timeout` is 0 or negative, then the `pool-cache-timeout` is disabled and beans are not removed from the pool.

For performance tuning, try to reduce the frequency of the removal of beans from the pool by setting the `pool-cache-timeout` to a large value. You should set the `pool-cache-timeout` to a large enough value to allow OC4J to avoid destroying and then re-creating instances while your J2EE application is running.

3.2 Advanced Performance Areas

This section describes areas that can provide improved performance for some usage cases and environments.

This section covers the following topics:

- [Managing Concurrency and Limiting Connections](#)
- [Load Balancing](#)
- [Using the -XX:AppendRatio Option \(on Standalone OC4J\)](#)

3.2.1 Managing Concurrency and Limiting Connections

Oracle Application Server lets you limit concurrency at multiple layers of the system to match specific usage needs. In addition to controlling HTTP connections, you can control concurrency at additional levels of the product to meet specific usage requirements.

This section covers the following topics:

- [Using OC4J Thread Pools to Control Concurrency](#)
- [Setting the Maximum Number of Connections for Data Sources](#)
- [Controlling the Number of EJB Instances When Using EJBs](#)
- [Limiting Remote EJB Client Connections](#)

See Also: ["Specify Sufficient Oracle HTTP Server Connections"](#) on page 3-7

3.2.1.1 Using OC4J Thread Pools to Control Concurrency

By default, OC4J creates thread pools, as described in the following list. New threads are created and added to the pools on an as-needed basis.

- **http** thread pool: A thread pool serving HTTP and AJP requests and possibly RMI requests (if `rmi request` thread pool is not configured) and RMI connections (if `rmi connection` thread pool is not configured). The default behavior provides for a maximum of 1024 threads, and threads for this pool are created on demand.
- **jca** thread pool: contains the pool of threads that OC4J reserves specifically for use by deployed resource adapters. The default behavior provides for a maximum of 1024 threads, and threads for this pool are created on demand.
- **system** thread pool: contains internal OC4J threads. The default behavior provides for a maximum of 1024 threads and threads for this pool are created on demand.

Thread pools create and store threads for use and reuse by an OC4J process. The default thread pool management configuration should be sufficient for most common usage scenarios. Reusing existing threads from a thread pool improves performance and reduces the burden on the JVM and on the underlying operating system.

Note: When you use the thread pool management options, this is considered an expert-mode task. If you modify the default thread pool configuration, it is important to consider that the resulting concurrency is determined by the sum of the threads in all thread pools for the OC4J process.

There are cases where changing the default behavior, by specifying thread pool management options, can improve performance, including the following:

- A hardware resource is nearing maximum capacity on Oracle Application Server 10g. For example, consider the case where a site has a user population of 1000 users with the site often processing 20 concurrent requests; at this rate, the site is nearing full resource utilization. Occasionally, the site needs to handle peaks of 75 to 100 concurrent requests. Specifying a thread pool with a maximum number of threads value in the range of 20 to 25 is likely to yield better overall results (setting the value in the range of 75 to 100 is unlikely to improve peak performance, since at this level there are no additional resources available to service the threads).
- A hardware resource is nearing maximum capacity on the database or the database is a bottleneck due to other limitations, and you want to limit connections to the database using OC4J thread control (in addition, you can limit database connections using the `datasource max-connections` parameter).

Thread pool tuning does not improve performance in the following situations:

- There are sufficient hardware resources available on the system when the system is under maximum load and there are no other known software issues such as database locking problems. This includes both Oracle Application Server 10g tier and database systems. For example, consider the case where a site has a user population of 1000 users but typically only sees a few concurrent requests. In this case, using application server thread pool tuning to limit threads is not advantageous, since the number of threads is determined by the concurrent request rate, which for this case is very low.
- There already are sufficient concurrency limits specified elsewhere in Oracle Application Server 10g or in the database systems. For example when the Oracle HTTP Server `MaxClients` directive is set to control concurrency, at the HTTP

server level, or when the data source `max-connections` attribute is set to control concurrency for connections to the database.

Note: If the number of incoming requests is consistently higher than the request rate that the physical hardware can support, consider increasing the physical resources at the site. Likewise, if the response time at peak periods is unacceptable, then you may need to increase the hardware configuration to remedy this situation.

See Also:

- ["Specify Sufficient Oracle HTTP Server Connections"](#) on page 3-7
- ["Setting the Maximum Number of Connections for Data Sources"](#) on page 3-15

3.2.1.1.1 Controlling and Using Application Server Thread Pools This section describes how to manage and use the thread pool configuration options.

Note: This section describes the OC4J thread pool configuration options for Oracle Application Server 10g Release 3 (10.1.3.1.0). This release also supports the OC4J thread pool configuration options that were available in previous Oracle Application Server releases. The `<global-thread-pool>` and `<work-manager-thread-pool>` elements in `server.xml` configure thread pools in an older format. These elements are deprecated in OC4J 10g (10.1.3.1.0). For details, see *Oracle Containers for J2EE Configuration and Administration Guide*.

You can manage application server thread pools as follows:

- Adding the `<thread-pool>` element with appropriate attributes in `server.xml`. For example,


```
<thread-pool name="http" min="120" max="120" queue="240" \>
```
- Updating the attributes in the `ThreadPool` MBean, which is accessible through the system MBean Browser in Application Server Control Console.
- Using the Thread Pool Configuration page in Application Server Control Console which is accessible from the OC4J Home page by selecting the Administration secondary tab and then clicking the Thread Pool Configuration task (see [Figure 3-1](#)).

Figure 3–1 Application Server Control Console Thread Pool Configuration

ORACLE Enterprise Manager 10g
Application Server Control

Cluster Topology > Application Server: 10gSOA5.tvanraal-pc.us.oracle.com > OC4J: home >

Thread Pool Configuration

Page Refreshed Jul 21, 2006 8:33:07 AM PDT

Thread pools create and store threads for use and re-use by an OC4J process. Re-using existing threads rather than creating new threads on demand improves performance and reduces the burden on the JVM and underlying operating system.

Name Δ	Current Pool Size	Minimum Pool Size	Maximum Pool Size	Keep Alive		Queue Capacity	Current Queue Size	Debug Mode
				Duration	Unit			
http	64	0	1024	10	minutes	0	0	<input type="checkbox"/>
jca	14	1	1024	10	minutes	0	0	<input type="checkbox"/>
system	8	0	1024	10	minutes	0	0	<input type="checkbox"/>

Copyright © 1996, 2006, Oracle. All rights reserved.
Oracle, JD Edwards, PeopleSoft, and Retek are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.
[About Oracle Enterprise Manager 10g Application Server Control](#)

Notes for specifying thread pool options:

- Oracle recommends setting the `min` value (Minimum Pool Size) equal to the `max` value (Maximum Pool Size). When `min=max`, you should set the `queue` to a value equal to the maximum concurrent number of requests you expect. For example, if you are using Oracle HTTP Server with one OC4J instance and no direct RMI connections, then the value of `MaxClients` would represent your maximum concurrency. Note: if you set `MaxClients` to a very large number, a setting of 300 or less for the `queue` size is probably sufficient. See "[Validating and Monitoring Thread Pool Performance](#)" on page 3-15 for instructions on monitoring your thread pool and queues.

Oracle recommends starting with a smaller number of threads, for example, using Oracle Application Server 10g Release 3 (10.1.3.1.0) with the Basic install option, set `min` and `max` to a value in the range of 80 to 100 and then monitor your resulting performance. In a standalone OC4J configuration, set `min` and `max` to a value in the range of 15 to 40, and then monitor your performance. If the site has sufficient CPU and memory resources available, and the concurrent request rate is higher than the current thread count setting, try increasing the number of threads. If you observe that the system is resource saturated, then try reducing the number of threads. Specific thread settings depend on the characteristics of your application.

- A recommended starting value for setting thread pool options is to set the `min` value equal to the `max` value. However, if you set a larger `max` value to accommodate occasional peaks of traffic, OC4J will add a new thread on each request until the minimum number of threads are created. After the minimum number of threads are created, new threads will not be created until the queue is full. OC4J attempts to keep the number of threads at or near the minimum, unless the queue is full. If you set the `min` value less than the `max` value, it is generally advisable to keep the `queue` size small; the `queue` default value is 0, meaning requests are not queued and a new thread will be created if the number of threads is less than the specified `max` value. If you do not see your threads increasing beyond the minimum value, then you should decrease the `queue` size.

See Also: *Oracle Containers for J2EE Configuration and Administration Guide* for more details on using the application server thread pool

3.2.1.1.2 Using RMI Connection and RMI Request Thread Pools OC4J supports a thread pool for RMI connection threads. By default the RMI connection threads are allocated from the application server thread pool named `http`. Using the RMI connection thread pool provides a separate thread pool whose threads block-read on the RMI connections. Defining an RMI connection thread pool puts a separate limit on the RMI connections; they are not taken from the `http` thread pool. Specify the RMI connection thread pool, using the `rmi connection` name in the `<thread-pool>` element in `server.xml`, to separate potentially long-lived RMI connection threads from the threads used for application work. This configuration allows you to free other threads to do work, instead of being allocated to long lived RMI connections.

When an RMI connection thread pool is defined, the work from the RMI connections executes either from threads in the `http` thread pool, or if you also specify an RMI request thread pool, then the work from the RMI connections executes from threads in the RMI request thread pool (you specify the RMI request thread pool with the name `rmi request` specification in a `<thread-pool>` element in `server.xml`). Specify a RMI request thread pool if you want to control the work associated with the RMI connections, separately from other work (for example HTTP requests).

Note: Specifying the `rmi connection` thread pool and the `rmi request` thread pool is only needed if you anticipate many concurrent RMI connections, where the number varies greatly and you are trying to bound the active worker threads.

3.2.1.1.3 Using the Work Manager (JCA) Thread Pool Starting with 10g Release 3 (10.1.3), EJBs of type MDB use receiver threads with the JMS Resource Adapter. OC4J allocates these threads and threads for other JCA resource adapters from the work manager thread pool (`jca` thread pool). However, you need to consider receiver threads to control the overall concurrency on the system. Oracle recommends leaving the `jca` thread pool at the default setting (by default, JCA work manager threads are limited to a maximum value of 1024). When you want to control the concurrency for EJB MDBs, use the `JMSReceiverThreads` maximum values. The overall concurrency limit on your system includes the `max` threads specified for the `http` thread pool, plus the sum of all the MDB `receiverThreads` configured for your applications deployed to OC4J, or the `max jca` threads if smaller (plus the RMI connection `max` threads and the RMI request `max` threads, if they are configured).

3.2.1.1.4 Validating and Monitoring Thread Pool Performance You can use the Application Server Control Console **Current Pool Size** and **Current Queue Size** metrics to observe OC4J thread use on your system. You can access these metrics from OC4J home page by clicking the Administration secondary tab and then clicking the Thread Pool Configuration task (see [Figure 3-1](#)). The **Current Pool Size** shows the current number of threads in the pool. The **Current Queue Size** shows the current number of requests waiting in the queue for a thread to become available. Oracle recommends starting with the default settings and observing the behavior on your system. You should also observe these metrics after changing the thread pool or queue size limits for your instance.

3.2.1.2 Setting the Maximum Number of Connections for Data Sources

For applications that use a database, performance can improve when the connection pool associated with a data source limits the number of connections. You can use the

`max-connections` attribute to limit the database requests from Oracle Application Server so that incoming requests do not saturate the database, or to limit the database requests so that the database access does not overload the Oracle Application Server-tier resource.

The connection pool `max-connections` attribute specifies the maximum number of connections that a connection pool allows. By default, the value of `max-connections` is set to -1 (unlimited). For best performance, you should specify a value for `max-connections` that matches the number appropriate to your database performance characteristics.

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

3.2.1.3 Controlling the Number of EJB Instances When Using EJBs

You may want to limit the number of EJB instances to reduce memory usage or to control concurrency to reduce contention on resources that the EJBs use (for example a data source).

The `max-instances` parameter specifies the number of bean instances allowed in memory – either instantiated or pooled.

- For all types of EJBs except stateful session beans, when the `max-instances` value is reached, and a new EJB is requested, the container waits the number of milliseconds set in the `call-timeout` attribute to see if a bean instance becomes available in the pool. If no bean instance is available in the pool then a `TimeoutExpiredException` is thrown back to the client.
- For stateful session beans, when the `max-instances` value is reached, the container attempts to passivate the oldest bean instance from memory. If unsuccessful, the container waits the number of milliseconds set in the `call-timeout` attribute to see if a bean instance is removed from memory, either through passivation, using the `remove()` method, or by bean expiration before a `TimeoutExpiredException` is thrown back to the client.

To allow an unlimited number of bean instances, use `max-instances = 0` (the default value is 0).

Set `max-instances < 0`, for example to -1, to disable instance pooling. In this case OC4J creates a new bean instance or context when starting the EJB call, and releases the context and throws the instance away to Non-existence state at the end of the call.

The exception, `com.evermind.server.ejb.TimeoutExpiredException: timeout expired waiting for an instance`, occurs when there is no available EJB instance. To avoid this problem set the `max-instances` and `call-timeout` parameters appropriately.

3.2.1.4 Limiting Remote EJB Client Connections

To limit remote EJB client connections you can change the max value for the `http` thread pool to specify limits for all threads, or you can use the thread pool features that control the maximum number of threads that service incoming EJB clients. By default, the threads for remote EJB client connections come from the `http` thread pool. When you want use the RMI connection thread pool, configure the `rmi connection`

name in the `<thread-pool>` element in `server.xml`. By setting the max value of the RMI connection thread pool, you can limit remote EJB client connections.

See Also: *Oracle Containers for J2EE Configuration and Administration Guide* for more details on using the application server thread pool

3.2.2 Load Balancing

Oracle Application Server provides load-balancing features that spread the J2EE application load and incoming requests among multiple application server instances, which generally results in higher throughput and shorter response time. Using multiple application server instances with load-balancing allows you to improve performance by directing requests across the multiple application server instances. In addition, you can use multiple application server instances running on multiple hosts to handle high availability and failover needs.

This section covers the following topics:

- [Configuring Multiple Oracle Application Server Instances](#)
- [Web Application Load Balancing](#)
- [EJB Application Load Balancing](#)

Note: The Oracle Application Server features that provide replication for failover with Web sessions and for stateful session EJBs have a performance overhead; only use these features if you need replication for failover.

3.2.2.1 Configuring Multiple Oracle Application Server Instances

This section covers the following:

- [Determining the Number of OC4J Processes](#)
- [Partitioning Applications into Different OC4J Instances](#)

3.2.2.1.1 Determining the Number of OC4J Processes Determining the optimal ratio of OC4J processes to available CPUs is dependent on the characteristics of the applications you run, the OC4J configuration, the hardware configuration, and the type and number of expected incoming requests. In hardware configurations with a small number of CPUs, you may only need one OC4J instance.

Adding OC4J instances beyond the available resources of the system does not improve performance. For example, if one OC4J instance is sufficient to saturate the CPU resources of a system, adding additional OC4J instances is not likely to improve performance and may, in fact, degrade it. A good starting point is to initially configure one OC4J instance and measure the performance improvement from adding additional OC4J instances.

See Also:

- *Oracle Application Server High Availability Guide*
- *Oracle Containers for J2EE Configuration and Administration Guide*

3.2.2.1.2 Partitioning Applications into Different OC4J Instances Partitioning different applications to be run under different OC4J instances, each of which has different requirements, may help improve the performance of your applications. In this case, you may want to configure different OC4J instances to service the different

applications. After deploying the applications to different OC4J instances, you can monitor the performance to see if the overall throughput increases, or the response time decreases.

3.2.2.2 Web Application Load Balancing

In an Oracle Application Server environment, the Oracle HTTP Server uses `mod_oc4j` to load balance requests between the available OC4J instances. In this environment you can select `mod_oc4j` configuration options to choose the appropriate load balancing policies to improve performance. By default, the requests are routed using a roundrobin algorithm.

At many sites Oracle Application Server uses the Oracle HTTP Server module `mod_oc4j` to load balance incoming stateless HTTP requests. By selecting the appropriate load balancing policy for `mod_oc4j` you can improve performance on your site.

The `mod_oc4j` module supports several configurable load balancing policies, including the following:

- Round robin routing (this is the default `mod_oc4j` load balancing policy)
- Random routing
- Round robin or random with local affinity routing, using the `local` option
- Round robin or random with host-level weighted routing, using the `weighted` option

Note: For a session based request `mod_oc4j` always directs the request to the original OC4J which created the session, unless the original OC4J process is not available. In case of failure, `mod_oc4j` sends the request to another OC4J within the same group as the original request (either within same host if available, or on a different host).

Recommendations for Load Balancing with `mod_oc4j`:

1. On a Single Host for both Oracle HTTP Server and OC4J, the default load balancing policy, round robin load balancing is recommended. Random load balancing typically gives comparable performance.
2. With Oracle HTTP Server on a separate host from OC4J:
 - Using a Single OC4J host the default load balancing policy, round robin load balancing is recommended. Random load balancing typically gives comparable performance.
 - Using Multiple OC4J hosts, if all OC4J hosts provide comparable capacity, the default load balancing is recommended. Note that the number of OC4J processes started on a host will implicitly weight the number of requests sent to that host even with random or round robin load balancing. A host with 4 OC4J processes will receive 4 times as many requests as a host with 1 OC4J process.
 - With Multiple OC4J Hosts with varying hardware resources or capacity, you may wish to weight the number of requests sent to each host explicitly to match its capacity. In this case, use either the round robin with weighted option or the random with weighted option. If the hosts have comparable capacity, use simple random or round robin load balancing.

For example, to configure the `mod_oc4j` module in Oracle HTTP Server to specify round robin with a routing weight of 3 for `Host_A` and a routing weight of 1 for `Host_B`, add the following directives to `mod_oc4j.conf`:

```
Oc4jSelectMethod roundrobin:weighted
Oc4jRoutingWeight Host_A 3
```

In this example, you do not need to specify a routing weight for `Host_B`, since the default routing weight is 1.

3. With Multiple hosts with Oracle HTTP Server and OC4J on each host:
 - Multiple hosts with Oracle HTTP Server and OC4J on each host, and a hardware load balancer. Select the local affinity option to direct `mod_oc4j` to only select the local OC4J processes to service incoming requests. This will generally improve performance. When no local OC4J processes are available, `mod_oc4j` selects from the list of available remote OC4J processes.

For example, to select the round robin policy with local affinity, specify the following directive in `mod_oc4j.conf`:

```
Oc4jSelectMethod roundrobin:local
```

3.2.2.3 EJB Application Load Balancing

After an EJB application is deployed to multiple OC4J instances, an EJB client-side application can load balance its requests across the available OC4J instances. To use load balancing, the client-side application configures the JNDI properties to use load balancing. For good performance in some clients, you need to set `oracle.j2ee.rmi.loadBalance=context` to load balance for every `initialContext` call, rather than only once for the entire client.

See Also:

- "Configuring ORMI Request Load Balancing" in the *Oracle Containers for J2EE Services Guide*
- "Understanding OC4J EJB Application Clustering Services" in the *Oracle Containers for J2EE Enterprise JavaBeans Developer's Guide*

3.2.3 Using the -XX:AppendRatio Option (on Standalone OC4J)

With the Sun 5.0 JVM, under some circumstances under heavy load, synchronization in an application can result in thread starvation. This may cause some requests for an application to appear hung or to timeout after a long time.

In 10g Release 3 (10.1.3.1.0) the parameter: `-XX:AppendRatio=3` is specified by default for managed OC4J. For standalone OC4J, if you believe your installation has this problem, we recommend setting the JDK parameter: `-XX:AppendRatio=3` to avoid this problem.

See Also: See the SUN bug database for a description of this issue and the suggested workaround:

http://bugs.sun.com/bugdatabase/view_bug.do?bug_id=4985566

Additional Performance Areas

This chapter covers performance information for the following Oracle Application Server areas:

- [Improving TopLink Performance](#)
- [Improving JTA Performance](#)
- [Improving EJB Performance](#)

4.1 Improving TopLink Performance

Oracle TopLink (TopLink) provides features to optimize application performance, including the following important areas:

- **Cache Configuration:** In TopLink, to avoid stale data the cache must be properly configured, and configured in conjunction with locking and query refreshing. In addition, to provide applications with high performance and scalability you should understand how the cache works, its relationship to query and transaction processing, and the cache configuration options.
- **Efficient Querying:** In TopLink, it is important to understand the batch and join reading options on queries to ensure that minimal SQL executes, and to retrieve the graph of required objects. For searching use cases and the use of projections, you should use `ReportQueries` to show that only objects that may be modified or shared between requests need to be read in and cached as objects.
- **Efficient Transactions:** In TopLink, you should understand how to use the `UnitOfWork` to minimize transaction scope and thus commit cycles.

The TopLink documentation includes information on these important TopLink performance areas. See the appropriate chapters in the documentation for more information about tuning your application to optimize TopLink performance.

See Also:

- Chapter 11 "Optimization", in the *Oracle TopLink Developer's Guide*
- Chapter 90 "Understanding the Cache", in the *Oracle TopLink Developer's Guide*
- Chapter 96 "Understanding TopLink Queries", in the *Oracle TopLink Developer's Guide*

4.2 Improving JTA Performance

This section describes JTA performance options, including the following topics:

- [Configuring Two-Phase Commit Logging for Performance](#)
- [Configuring JTA Data Sources for Performance](#)
- [Monitoring JTA Resources](#)

4.2.1 Configuring Two-Phase Commit Logging for Performance

Using configuration options you can control the type and level of two-phase commit logging. To change the configuration options, you can modify the `transaction-manager.xml` file or use the JTA Resource MBean available from the Transaction Manager page in Application Server Control Console. When you configure two-phase commit logging, you need to be aware of the transactional ramifications of turning two-phase commit logging off.

Note: Two-phase commit logging is off by default. When you use the default logging level, JTA resources do *not* support recovery and full ACID properties.

Table 4–1 shows the two-phase commit logging configuration options that you can set in `transaction-manager.xml` or using the JTA Resource MBean.

Table 4–1 Two-phase Commit Logging Log Type Configuration Options

Log Type	Description	Performance Notes
none	Specifies no logging (or recovery). This is the default value.	When recovery and ACID properties are <i>not</i> required, use this option for best performance.
file	File logging specifies logging to the file system for transaction recovery.	File logging typically performs better than database logging, due to lower overhead.
database	Database logging specifies logging to an Oracle database for transaction recovery.	File logging typically performs better than database logging.

See Also: See the *Oracle Containers for J2EE Services Guide* for more information on setting two-phase commit logging configuration options, and for the details on the transactional and recovery ramifications of specifying different log type configuration options.

4.2.1.1 Setting JTA Store File Logging Options

[Table 4–2](#) describes the performance settings for file store logging that you can set in `transaction-manager.xml` or using the JTA Resource MBean. The default settings are adequate if the maximum concurrent number of two-phase commit transactions is less than 256.

To determine the maximum concurrent number of two-phase commit transactions, you can use the `TwoPhaseCommitCompletion.maxActive` metric from the JTAResource metric table.

See Also: [Table D–1](#) on page D-2 for details on JTA resource metrics

Table 4–2 JTA File Store Logging Parameters

Parameter	Description	Performance Notes
<code>maxOpenFiles</code>	Specifies the maximum number of file descriptors that can remain open or active; when this number is exceeded the oldest file descriptors are released until the <code>xid</code> is requested again. Avoid exceeding the <code>maxOpenFiles</code> if possible. Default Value: 256	The optimal value is large enough to cover the maximum number of concurrent requests that use two-phase commit transactions (plus a small additional number of files that may be required for recovery). The <code>maxOpenFiles</code> value is limited by the Operating System open file descriptor limit.
<code>minPoolSize</code>	Specifies the number of files that are pre-allocated to the pool during startup. Default Value: 40	The optimal value is large enough to handle the maximum number of concurrent two-phase commit requests. Note: if the maximum concurrency is large, then to avoid a high cost at startup, you can increase this value from the default, but set the value to something smaller than <code>maxOpenFiles</code> value.
<code>oldFileReleaseSize</code>	Specifies the number of the oldest file handles which are closed when <code>maxOpenFiles</code> is exceeded. Default Value: 20	If you expect that you will repeatedly exceed the <code>maxOpenFiles</code> value, which is not recommended, then increasing this value to release more file handles may help reduce the number of times that <code>maxOpenFiles</code> is exceeded.

Note: The number specified for `maxOpenFiles` does not limit the number of transactions; if `maxOpenFiles` is exceeded, old file handles are released, but new transactions can still be created (see the `oldFileReleaseSize` parameter).

4.2.2 Configuring JTA Data Sources for Performance

This section covers the following areas:

- [Specify the Data Source Type](#)
- [Use Last Resource Commit](#)
- [Use a Single Data Source Where Possible](#)

4.2.2.1 Specify the Data Source Type

Non-XA compliant data sources are generally faster than XA data sources. When a full XA-compliant two-phase commit is required, you must use an XA data source.

If transaction logging is set to none, any number of XA or non-XA compliant resources can be enlisted in a global transaction; however, in this case there are no ACID guarantees nor recovery.

If transaction logging is enabled, participants in a global transaction must be XA-compliant. The last resource commit feature allows for a single non-XA-compliant resource to participate in an XA transaction.

See Also: *Oracle Containers for J2EE Services Guide*

4.2.2.2 Use Last Resource Commit

In addition to allowing a single non-XA resource to participate in a global transaction, last resource commit can also be used as a performance optimization. By enlisting an XA-capable resource as a non-XA resource and using last resource commit, a gain in performance is achieved because the resource does not need to perform XA logging. Also, the resource would never be put in doubt, that is prepared, which would prevent resources from being locked. Although last resource commit can be used as a performance optimization, it is at the cost of guaranteed correctness.

See Also: *Oracle Containers for J2EE Services Guide*

4.2.2.3 Use a Single Data Source Where Possible

When your application uses multiple data sources to access a single resource, this can lead to unintended use of two-phase commit operations (using XA-transactions). In some cases you can improve performance by changing the configuration; this configuration change allows OC4J to eliminate two-phase commits and replace them with one-phase commits.

You can use the following metrics to check the number of one-phase commits, two-phase commits, and the count of global transactions that do not enlist any resources:

```
/oc4j/JTA/SinglePhaseCommitCompletion.completed  
/oc4j/JTA/TwoPhaseCommitCompletion.completed  
/oc4j/JTA/AverageCommitTime.completed
```

Note: The `/oc4j/JTA/AverageCommitTime.completed` metric shows all JTA involved transactions but does not show local transactions.

Whenever there are multiple data sources used within the same global transaction, use of XA two-phase commit transactions occurs, even if the data sources actually point to the same database. If you have deployed the data sources for your application in a

single database and schema, you may be able to reconfigure your application to use a single data source. This would improve performance by changing two-phase commits to single-phase commits.

Thus, transactional applications that use both traditional database resources, such as tables, and also use OJMS, where resources reside on the same database, can avoid some two-phase commits by specifying a single data source for each resource. This change should improve performance and requires that the OJMS data source configuration matches that specified for accessing the tables.

See Also: *Oracle Containers for J2EE Services Guide* for more information on Local and Global Transactions

4.2.3 Monitoring JTA Resources

When you monitor JTA resources, be aware that errors can cause performance problems. You can determine if there are JTA errors by looking for Rollback or Exception counts greater than 0 using the metrics in the JTAResource metric table. For example, look at the values of the following metrics: `RollbackExceptionCount`, `RolledbackCount`, or `SystemExceptionCount`.

Note that certain performance problems may also affect JTA errors. For example if performance is bad, timeout errors may occur. In this case, look at the value of the metric `RolledbackDueToTimedOutCount`.

See Also: [Table D-1](#) in "JTA Resource Metrics" on page D-2

4.3 Improving EJB Performance

This section includes the following topics:

- [Improving MDB Performance](#)
- [Improving EJB CMP 2.1 Performance](#)

4.3.1 Improving MDB Performance

This section covers some of the important performance related EJB configuration properties specified in the `orion-ejb-jar.xml` configuration file that apply for Message Driven Beans (MDBs), including the following:

- [Setting the JMS Connector Receiver Threads](#)
- [Using the `ejbCreate` Method for One Time Initialization](#)
- [Monitoring MDB Resources](#)

See Also: *Oracle Containers for J2EE Services Guide* and the chapter, Chapter, "Oracle Enterprise Messaging Service (OEMS)" for more information on using and configuring MDBs

4.3.1.1 Setting the JMS Connector Receiver Threads

When you set the number JMS Connector receiver threads for an MDB, this can improve performance either when there are many concurrent users sending messages to the queue of an MDB, or when significant processing occurs in the `onMessage` method. For example, if the `onMessage` method contains code to call another EJB and the EJB processing can occur concurrently while processing other messages, then setting the JMS Connector receiver threads to a value greater than one can improve performance. Depending on the underlying JMS Connector and the specific MDB,

some applications may see significant performance improvements when you increase the value of the JMS Connector `ReceiverThreads` configuration property.

For example, if a queue contains 100 messages, and the `ReceiverThreads` is set to the default value, 1, then only one MDB receiver thread processes the messages, in a serial fashion. When you set the `ReceiverThreads` to 5, this specifies that there can be a maximum of 5 MDB instances that take messages from the queue and process the messages in parallel. In this example, the total time required to complete the processing for the 100 messages may decrease, since in this case OC4J uses up to 5 MDB threads to dequeue and process the messages.

Note: The JMS Connector `ReceiverThreads` value specifies a maximum value for threads; not all the threads are necessarily used, depending on the load.

When you specify a JMS Connector `ReceiverThreads` value greater than 1, this enables multiple instances of the MDB to concurrently process messages from queues. However, in this case any performance improvement depends on the application and on the number of threads you specify. If you specify a value that is too large, this can cause performance to degrade due to resource contention.

Note: For JMS topics, always set the JMS Connector `ReceiverThreads` configuration property to the value 1 (only for queues are values over 1 meaningful).

4.3.1.1.1 Consider Message Processing Order Requirements for MDBs Use JMS Connector `ReceiverThreads` set to the value 1 if the messages must be processed in order. If you use `ReceiverThreads` with a value greater than 1, messages are still removed from a queue serially, but the order of processing the messages cannot be guaranteed since the MDB is processing the messages with multiple threads.

4.3.1.1.2 Coordinate Thread Pool and Bean Instance Settings OC4J allocates the threads used as JMS Connector receiver threads from the work manager thread pool (shown as the `jca` thread pool using Application Server Control Console and in `server.xml`) You can limit the number of JMS Connector receiver threads in the work manager thread pool using the `max` parameter for the `jca` thread pool. You can also use the `min` value to set the initial number of available work manager thread pool threads.

Performance Note: Oracle recommends leaving the `jca` work manager thread pool at the default setting. Thus, if you want to control concurrency for EJB MDBs, use the JMS Connector `ReceiverThreads` value.

You need to coordinate the following configuration options when you set the value for the JMS Connector `ReceiverThreads`:

- Overall Concurrency Limit: On the system this includes all the `thread-pool max threads` plus either:
 - a. The sum of all the MDB JMS Connector `ReceiverThreads` configured for your applications deployed to OC4J (if the `max` for the work manager thread pool, shown as the `jca` thread pool, is specified at the default value, plus the sum of all threads used by other JCA adapters).
 - b. The maximum allowed number of work manager thread pool threads, shown as `jca`, specified with the `max` parameter (if this is smaller than the maximum number specified for the sum of all of the JMS Connector `ReceiverThreads`).
- Set Minimum MDB Instances Appropriate for the Receiver Threads: There is a one-to-one correspondence between the JMS Connector `ReceiverThreads` created and the number of active MDB bean instances. When the initial processing time for an MDB may be significant, you should set the `min-instances` MDB setting to match the number of desired JMS Connector `ReceiverThreads`, so that these instances are initialized at startup.

Also, you should configure the MDB configuration `min-instances` value to be at most as large as the JMS Connector `ReceiverThreads` setting per MDB.

To maintain the desired number of instances, set the `pool-cache-timeout` to a value that is large enough so that the MDB instances are not removed when idle.

See Also: ["Using OC4J Thread Pools to Control Concurrency"](#) on page 3-12

4.3.1.1.3 Consider the Database Connections When Setting JMS Connector Receiver Threads

The number of JMS Connector `ReceiverThreads` also multiplies of the number of required database connections, if any, for the MDB. For example, if a particular MDB uses 5 database connections concurrently, and there are 5 active MDB instances, then the number of requested concurrent database connections would be 25. Thus, the number of JMS Connector `ReceiverThreads` must be included in the calculation of the data source `max-connections` count.

See Also: ["Reuse Database Connections"](#) on page 3-6

4.3.1.2 Using the `ejbCreate` Method for One Time Initialization

An MDB is stateless and contains no specific client state across invocations. However, for nonclient related state, an MDB instance can contain some state across client message handling. For example, state can be maintained for a lookup. In addition, other state information that you may want to cache across `onMessage` invocations, such as a reference to an EJB, can be initialized in the `ejbCreate` method and cached to optimize MDB performance.

Remember to destroy the state in the `ejbRemove` method in case idle MDB objects are removed from the pool and reallocated when needed.

4.3.1.3 Monitoring MDB Resources

When MDBs use OracleAS JMS as a message provider, DMS message related metrics are available from the Oracle Application Server performance monitoring tools.

For example, the OracleAS JMS `JMSStoreStats` metric table includes information for a destination corresponding to a queue that an MDB uses:

```
destination.value:      name
messageDequeued.count: x ops
messageEnqueued.count: x ops
messageCount.value:    n
```

These metrics show the destination name, the total messages enqueued, the total number of messages dequeued, and the total number currently in the queue.

You can also check the MDB `onMessage` metrics to check that the time in `onMessage` is as expected and use the `maxActive` metric to see the total number of concurrent receiver threads is as expected.

```
client.active: 1 threads
client.avg: 112 msecs
client.completed: 4 ops
client.maxActive: 1 threads
client.maxTime: 70 msecs
client.minTime: 130 msecs
client.time: 121 msecs
```

Note: When monitoring a JMS destination, other applications besides the MDB may access the destination. Thus, when you test the performance of an application, make sure that you know whether the application is responsible for the message activity that is reported in the metrics.

Application Server Control Console provides information for the performance of all MDBs and of individual MDBs.

To access the summary MDB information, do the following:

1. From the OC4J home page, select the Administration secondary tab.
2. In the table, under Services, select the **JMS Providers** task.
3. In the in the Performance area, Application Server Control Console displays the following summary information:

```
Active Connections
Messages Waiting for Read
Messages Waiting for Commit
Messages Enqueued per Second
Messages Dequeued per Second
Messages Paged In per Second
Messages Paged Out per Second
Messages Committed Since Startup
Messages Rolled Back Since Startup
Messages Expired Since Startup
```

To access the individual MDB information, use the Application Server Control Console performance area:

1. From the OC4J home page, select the Applications secondary tab.

2. Select the application that you want to monitor.
3. In the **Modules** table, select the appropriate EJB module.
4. In the **Message Driven Beans** area, select the MDB you want to monitor to see the following information:

Messages Dequeued
 Messages Rolled Back
 Average Message Processing Time (seconds)
 Number of Available Instances
 Number of Used Instances

See Also: ["OC4J JMS Metrics"](#) on page D-13

4.3.2 Improving EJB CMP 2.1 Performance

This section covers some of the available performance options for entity beans using CMP, and includes the following topics:

- [Using Efficient SQL Statements and Querying](#)
- [Cache Configuration Performance Tuning](#)
- [Monitoring CMP Resources](#)

Note: This section describes several EJBs options that you can use to improve performance of your applications. This section does not describe how to set these options. For details on configuring your EJBs to use various options, refer to the *Oracle Containers for J2EE Enterprise JavaBeans Developer's Guide* and the *Oracle TopLink Developer's Guide*.

See Also:

- "Optimization", in the *Oracle TopLink Developer's Guide*
- "Understanding the Cache" in the *Oracle TopLink Developer's Guide*
- "Understanding TopLink Queries" in the *Oracle TopLink Developer's Guide*
- "Using Advanced Unit of Work API" in the *Oracle TopLink Developer's Guide*
- "Database Transaction Isolation Levels" in the *Oracle TopLink Developer's Guide*

4.3.2.1 Using Efficient SQL Statements and Querying

This section covers using efficient SQL statements and SQL querying. [Table 4-3](#) and [Table 4-4](#) show tuning parameters and performance recommendations related to SQL statements and querying.

Table 4–3 CMP EJBs Using Efficient SQL Statements and Querying

Tuning Parameter	Description	Performance Notes
Parameterized SQL Binding	<p>Using parameterized SQL and prepared statement caching, you can improve performance by reducing the number of times the database SQL engine parses and prepares SQL for a frequently called query. TopLink and OC4J/CMP does not enable parameterized SQL and prepared statement caching by default, because not all databases and JDBC drivers support these options. Note that the Oracle JDBC driver bundled with OC4J does support these options</p> <p>Default Value: off</p> <p>See Also: <i>Oracle TopLink Developer's Guide</i> section, "Configuring Named Query Parameterized SQL and Statement Caching at the Project Level"</p>	<p>Turn SQL parameter binding on for all queries (at the project level). Oracle recommends that you enable parameterized SQL and prepared statement caching for selected databases and JDBC drivers that support these options.</p>
JDBC Statement Caching	<p>Statement caching is used to lower the overhead of repeated cursor creation and repeated statement parsing and creation; this can improve performance for applications using a database.</p> <p>Note: Use the data source statement caching (and do not use TopLink Statement Caching for CMP).</p> <p>Default Value: off</p> <p>See Also: <i>Oracle Containers for J2EE Services Guide</i> section, "Statement Caching with Managed Data Sources"</p>	<p>You should always enable statement caching if your JDBC driver supports this option. The Oracle JDBC driver supports this option. Set this option in <code>data-sources.xml</code> by setting <code>num-cached-statements</code>.</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. 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>Default Value: 0</p> <p>See Also: <i>Oracle TopLink Developer's Guide</i></p>	<p>The optimal fetch size is not always obvious. Usually, a fetch size of one half or one quarter of the total expected result size is optimal. Note that if you are unsure of the result set size, incorrectly setting a fetch size too large or too small can decrease performance.</p>
Batch Writing	<p>Batch writing can improve database performance by sending groups of INSERT, UPDATE, and DELETE statements to the database in a single transaction, rather than individually.</p> <p>Default Value: off</p> <p>See Also: <i>Oracle TopLink Developer's Guide</i> section, "Data Access Optimization"</p>	<p>Enable for all EJBs.</p>

4.3.2.1.1 Querying Container Managed Relationships Performance Tuning

Table 4–4 shows the CMR parameters for performance tuning.

Table 4–4 CMP EJBs and CMR Query Performance Options

Tuning Parameter	Description	Performance Notes
Batch Reading	<p>Batch reading propagates query selection criteria through an object's relationship attribute mappings. You can also nest batch read operations down through complex object graphs. This significantly reduces the number of required SQL select statements and improves database access efficiency.</p> <p>Default Value: off</p> <p>See Also: <i>Oracle TopLink Developer's Guide</i> section, "Using Batch Reading"</p>	<p>Use for queries of tables with columns mappings to table data you also need to retrieve.</p> <p>You should only use either batch-reading or joining if you know that you are going to access all of the data; if you do not intend to access the relationships, then just let indirection defer their loading.</p> <p>Batch reading is more efficient for 1-m and other relationships as it reads less data from the database, that is, n versus n*m. Batch reading is also more efficient on logical m-1 relationships as less data may be read. Batch reading also performs better with caching, because if the original object and its relationship were already cached then the batch query does not need to execute (where joining would have already read in all of the data). TopLink supports batch reading for most mappings (1-1, 1-m, m-m, dc, ac) at the query and mapping level.</p>
Join	<p>Join reading is a query optimization feature that allows a single query for a class to return the data to build the instances of that class and its related objects. Use this feature to improve query performance by reducing database access. By default, relationships are not join-read: each relationship is fetched separately when accessed if you are using indirection, or as a separate database query if you are not using indirection.</p> <p>Default Value: not used</p> <p>See Also: <i>Oracle TopLink Developer's Guide</i> section, "Using Join Reading"</p>	<p>Use for queries of tables with columns mappings to table data you also need to retrieve.</p> <p>You should only use either batch-reading or joining if you know that you are going to access all of the data; if you do not intend to access the relationships, then just let indirection defer their loading.</p> <p>TopLink supports joining for only 1-1 and 1-m at the query level, but only 1-1 (inner) at the mapping level. For performance, Join is recommended only for joining logical 1-1 relationships.</p> <p>Joining is not supported to related classes that use inheritance and have subclasses that span multiple tables.</p>
Indirection	<p>Without indirection on, when TopLink retrieves a persistent object, it retrieves all of the dependent objects to which it refers. When you configure indirection (also known as lazy reading, lazy loading, and just-in-time reading) for an attribute mapped with a relationship mapping, TopLink uses an indirection object as a place holder for the referenced object: TopLink 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>Default Value: Value Holder Indirection On for all CMRs</p> <p>See Also: <i>Oracle TopLink Developer's Guide</i> section, "Indirection"</p>	<p>Leave the indirection option at the default value. That is, use indirection for CMP for all situations. Querying the referenced object using Join or Batch Reading is more efficient.</p>

4.3.2.2 Cache Configuration Performance Tuning

Table 4–5 shows the cache configuration options.

Table 4–5 CMP EJBs and Cache Configuration Options

Tuning Parameter	Description	Performance Notes
Object Cache	<p>TopLink sessions provide an object cache. CMP 2.1 applications which use the TopLink persistence manager create TopLink sessions which 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 a TopLink application. Typically, a server session's object cache is shared by all client sessions acquired from it. Isolated sessions provide their own session cache isolated from the shared object cache.</p> <p>Default Value: enabled</p> <p>See Also: <i>Oracle TopLink Developer's Guide</i> section, "Object Cache"</p>	Disable (use isolated cache) for pessimistic queries.
Query Result Set Cache	<p>In addition to the object cache in TopLink, TopLink 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 allows for any query executed with the same parameters to obtain a query cache hit and return the same result set <p>Default Value: not used</p> <p>See Also: <i>Oracle TopLink Developer's Guide</i>, section, "Caching Query Results in the Session Cache"</p>	<p>Use for frequently executed non-primary key queries with infrequently changing result sets.</p> <p>Use with a cache validation timeout to refresh as needed.</p>
Cache Size	<p>Default Value: <code>softCacheWeakIdentityMap</code> size 100 (per EJB).</p> <p>See Also: <i>Oracle TopLink Developer's Guide</i> section, "Guidelines for Configuring the Cache and Identity Maps"</p>	Set the cache size to be as large as the maximum number of objects (of the same type) referenced within a transaction
Locking	<p>Oracle supports the locking policies shown in Table 4–6.</p> <p>Default Value: no locking</p> <p>See Also: <i>Oracle TopLink Developer's Guide</i> sections, "Configuring Locking Policy" and "Understanding Descriptors and Locking"</p>	.

Table 4–5 (Cont.) CMP EJBs and Cache Configuration Options

Tuning Parameter	Description	Performance Notes
Cache Usage	<p>Using Cache Usage conform does not check the session cache. For read-all it first checks the database, then conforms the result with the unit of work changes and new and deleted objects. For read-object is first checks the unit of work changes and new objects for a matching object, then checks the database and then conforms the results with changes and deleted objects.</p> <p>Conforming is more similar to the POJO default option <code>CheckCacheByPrimarykey</code> than <code>CheckCacheThenDatabase</code>, although somewhat similar to <code>CheckCacheThenDatabase</code> for read-object queries. Conforming adds additional overhead as it must conform the query results with unit of work changes and new and deleted objects.</p> <p>For CMP, the default is <code>ConformResultsInUnitofWork</code>. To avoid conforming, normally you can change this to <code>CheckCacheByPrimarykey</code>; for read-all queries this basically means the same as <code>DoNotCheckCache</code>.</p> <p>Default: <code>ConformResultsInUnitofWork</code></p>	<p>Turn Cache Usage Conform off if you do not need uncommitted data read in your transaction, especially for read-only operations, in each descriptor and at the query level if not needed.</p>
Isolation	<p>There is not a single tuning parameter that sets a particular database transaction isolation level in a CMP application that uses TopLink. In a typical CMP application, a variety of factors affect when database transaction isolation levels apply and to what extent a particular database transaction isolation can be achieved, including the following:</p> <ul style="list-style-type: none"> ■ Locking mode ■ Use of the Session Cache ■ External Applications ■ Database Login method <code>setTransactionIsolation</code> <p>See Also: <i>Oracle TopLink Developer's Guide</i> section, "Database Transaction Isolation Levels"</p>	

Table 4–5 (Cont.) CMP EJBs and Cache Configuration Options

Tuning Parameter	Description	Performance Notes
Cache Refreshing	<p>By default, TopLink 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, such as a read-all query, accesses the data source, if the objects corresponding to the records returned are in the cache, TopLink uses the cached objects. This default caching policy can lead to stale data in the application.</p> <p>Refreshing can be enabled at the descriptor level (<code>alwaysRefreshCache</code>) and the query level (<code>refreshIdentityMapResult</code>). If set at the descriptor level primary-read-object queries will still get cache hits unless you also set (<code>disableCacheHits</code>).</p> <p>Using an appropriate locking policy is the only way to ensure that stale or conflicting data does not get committed to the database. Thus, there are a few cases to consider for data refreshing in the cache, all with performance implications:</p> <ul style="list-style-type: none"> ▪ In the case where you never want cached data and always want fresh data, we recommend using an isolated cache. 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 <code>cache-invalidation</code> policy would be the recommended solution. In this case you should also use optimistic locking, and typically also need a method to refresh stale objects when a locking error occurs. In this case you need to define a finder such as <code>findAndRefreshByPrimaryKey</code> that had (<code>refreshIdentityMapResult</code>) set to force refreshing of the object. If using optimistic locking you could also enable <code>alwaysRefreshCache</code> and <code>onlyRefreshCacheIfNewerVersion</code> 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 finder operations when you know you want refreshed data, or even provide the option of refreshing something from the client that would call a refreshing finder. ▪ In the case when you are not concerned about stale data. In this case you should use optimistic locking and also write an error handler to commit operations to refresh stale objects on locking errors. <p>Default: No Cache Refreshing</p> <p>See Also: <i>Oracle TopLink Developer's Guide</i> section, "Configuring Cache Refreshing"</p>	<p>Try to avoid descriptor 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

Note: By default, TopLink assumes that your application has exclusive access to the data it is using (that is, there are no external, nonTopLink, applications modifying the data). If your application does not have exclusive access to the data, then you need to change some of the defaults from [Table 4–5](#).

The default settings for CMP2.1 used with the TopLink persistence manager and cache are no locking, no cache refresh, and cache-usage conform. To assure that your application doesn't read stale data from the cache when you don't have exclusive access, and gets the required data consistency level, you need to configure these and other isolation related settings appropriately.

The locking modes, as shown in [Table 4–6](#), along with TopLink cache-usage and query refreshing options, assures data consistency for EJB entity beans using CMP. The different combinations have both functional and performance implications, but often the functional requirements for up-to-date data and data consistency will lead to the settings for these options, even when it may be at the expense of performance.

Table 4–6 Locking 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 Bean is never updated concurrently or concurrent reads and updates to the same rows with read-committed semantics is sufficient. See Also: <i>Oracle TopLink Developer's Guide</i> sections, "Configuring Locking Policy" and "Understanding Descriptors and Locking"	In general, no locking will be faster, but may not meet your needs for data consistency
Optimistic	All users have read access to the data. When a user attempts to make a change, the application checks to ensure the data has not changed since the user read the data. See Also: <i>Oracle TopLink Developer's Guide</i> sections, "Configuring Locking Policy" and "Understanding Descriptors and Locking"	
Pessimistic	The first user who accesses the data with the purpose of updating it locks the data until completing the update. See Also: <i>Oracle TopLink Developer's Guide</i> sections, "Configuring Locking Policy" and "Understanding Descriptors and Locking"	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 bean level, it is recommended that you use it with an isolated cache for best performance
Read Only	Setting a CMP bean descriptor to read-only ensures that the entity bean cannot be modified and allows TopLink to optimize unit of work performance. See Also: <i>Oracle TopLink Developer's Guide</i> sections, "Configuring Locking Policy" and "Understanding Descriptors and Locking"	Defining a bean as read-only will perform better than a bean that is not defined as read-only, yet does no inserts, updates, or deletes, since it allows TopLink to optimize the unit of work performance.

4.3.2.3 Monitoring CMP Resources

Check the DMS metric table type named `oc4j_ejb_method` and the metrics `wrapper.avg` or `client.avg` to see if your EJB methods are taking significant time (check for unexpectedly large values, as some methods you may expect to take a long time to complete). For example, check the metric values for the methods `ejbCreate`, `create`, `findAll`, or your application specific EJB methods.

You can also check the metrics that Application Server Control Console provides to see how the response times and transactions per second are changing over time, as follows:

1. From the OC4J home page, select the Applications secondary tab.
2. In the table, under All Applications, select the application that you want to monitor.
3. In the **Modules** table, select the EJB module of that you want to monitor.

You can also view CMP TopLink metrics for additional information on CMP performance.

Note: You need to set the DMS configuration OC4J command line-property `-Doracle.dms.sensors` to the value `Heavy` or `All` to turn on collection of TopLink related DMS metrics.

See Also:

- [Table D-11](#) in "OC4J J2EE Application Metrics" on page D-6
- *Oracle TopLink Developer's Guide*, Table 11-1 TopLink DMS Metrics for details on the available metrics

Optimizing PL/SQL Performance

This chapter provides references to the information that describes improving PL/SQL performance for web applications. Most of this information is in the Oracle Application Server `mod_plsql` User's Guide.

See Also:

- *Oracle Application Server `mod_plsql` User's Guide* for information on optimizing PL/SQL performance
- [Appendix C, "Performance Metrics"](#) for information on `mod_plsql` metrics
- *Oracle HTTP Server Administrator's Guide* for details on DAD Parameters
- *Oracle Application Server PL/SQL Web Toolkit Reference* for information on the PL/SQL Web Toolkit that enables you to develop Web applications as PL/SQL procedures stored in an Oracle database server

Optimizing Oracle HTTP Server

This chapter discusses the techniques for optimizing Oracle HTTP Server performance in Oracle Application Server.

This chapter contains:

- [Configuring Oracle HTTP Server Directives](#)
- [Oracle HTTP Server Logging Options](#)
- [Oracle HTTP Server Security Performance Considerations](#)
- [Oracle HTTP Server Performance Tips](#)

6.1 Configuring Oracle HTTP Server Directives

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

Table 6–1 lists directives that may be significant for performance.

Table 6–1 Oracle HTTP Server Configuration Properties

Directive	Description
<code>ListenBackLog</code>	<p>Specifies the maximum length of the queue of pending connections. Generally no tuning is needed or desired. Note that some Operating Systems do not use exactly what is specified as the backlog, but use a number based on, but normally larger than, what is set.</p> <p>Default Value: 511</p>
<code>MaxClients</code>	<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 process becomes available).</p> <p>The maximum allowed value for <code>MaxClients</code> is 8192 (8K).</p> <p>Default Value: 150</p>
<code>MaxRequestsPerChild</code>	<p>The number of requests each child process is allowed to process before the child dies. The child will exit so as to avoid problems after prolonged use when Apache (and maybe the libraries it uses) leak memory or other resources. On most systems, this isn't really needed, but some UNIX systems have notable leaks in the libraries. For these platforms, set <code>MaxRequestsPerChild</code> to something like 10000 or so; a setting of 0 means unlimited.</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>Note: On Windows systems <code>MaxRequestsPerChild</code> should always be set to 0 (unlimited). On Windows there is only one server process, so it is not a good idea to limit this process.</p>
<code>MaxSpareServers</code> <code>MinSpareServers</code>	<p>Server-pool size regulation. Rather than making you guess how many server processes you need, Oracle HTTP Server dynamically adapts to the load it sees, that is, it tries to maintain enough server processes to handle the current load, plus a few spare servers to handle transient load spikes (for example, multiple simultaneous requests from a single Netscape browser).</p> <p>It does this by periodically checking how many servers are waiting for a request. If there are fewer than <code>MinSpareServers</code>, it creates a new spare. If there are more than <code>MaxSpareServers</code>, some of the spares die off.</p> <p>The default values are probably ok for most sites.</p> <p>Default Values:</p> <p><code>MaxSpareServers</code>: 10</p> <p><code>MinSpareServers</code>: 5</p>
<code>StartServers</code>	<p>Number of servers to start initially. If you expect a sudden load after restart, set this value based on the number child servers required.</p> <p>Default Value: 5</p>
<code>Timeout</code>	<p>The number of seconds before incoming receives and outgoing sends time out.</p> <p>Default Value: 300</p>

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

Directive	Description
<code>KeepAlive</code>	Whether or not to allow persistent connections (more than one request per connection). Set to <code>Off</code> to deactivate. Default Value: <code>On</code>
<code>MaxKeepAliveRequests</code>	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, you might want to increase this value. Default Value: 100
<code>KeepAliveTimeout</code>	Number of seconds to wait for the next request from the same client on the same connection. Default Value: 15 seconds

6.1.1 How Persistent Connections Can Reduce httpd Process Availability

The default settings for the `KeepAlive` directives are:

```
KeepAlive on
MaxKeepAliveRequests 100
KeepAliveTimeout 15
```

These settings allow enough requests per connection and time between requests to reap the benefits of the persistent connections, while minimizing the drawbacks. You should consider the size and behavior of your own user population in setting these values on your system. 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.

6.2 Oracle HTTP Server Logging Options

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

6.2.1 Access Logging

For static page requests, access logging of the default fields results in a 2-3% performance cost.

6.2.2 Configuring the `HostNameLookups` Directive

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

Performance degraded by about 3% (best case) in Oracle in-house tests with `HostNameLookups` set to `on`. Depending on the server load and the network connectivity to your DNS server, the performance cost of the DNS lookup could be high. Unless you really need to have host names in your logs in real time, it is best to log IP addresses.

On UNIX systems, you can resolve IP addresses to host names off-line, with the `logresolve` utility found in the `$ORACLE_HOME/Apache/Apache/bin/` directory.

6.2.3 Error logging

The server notes unusual activity in an error log. The `ErrorLog` and `LogLevel` directives identify the log file and the level of detail of the messages recorded. The default level is `warn`. There was no difference in static page performance on a loaded system between the `warn`, `info`, and `debug` levels.

For requests that use dynamic resources, for example requests that use `mod_ossso`, `mod_plsql`, or `mod_oc4j`, there is a performance cost associated with setting higher debugging levels, such as the `debug` level.

6.3 Oracle HTTP Server Security Performance Considerations

This section covers the following topics:

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

6.3.1 Oracle HTTP Server Secure Sockets Layer (SSL) Performance Issues

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

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

This section covers the following SSL Performance related information:

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

See Also: *Oracle Application Server Security Guide*

6.3.1.1 Oracle HTTP Server SSL Caching

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 Secure Sockets Layer (SSL) session information by default. With session caching, only the first connection to the server incurs high latency.

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

`SSLSessionCache` directive specifies the location for saved SSL session information, the default location on UNIX is the `$ORACLE_HOME/Apache/Apache/logs/` directory or on Windows systems, `%ORACLE_HOME%\Apache\Apache\logs\`. Multiple Oracle HTTP Server processes can use a saved session cache file.

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

The reuse of saved SSL session state has some performance costs. When SSL session state is stored to disk, reuse of the saved state normally requires locating and retrieving the relevant state from disk. This cost can be reduced when using HTTP persistent connections. Oracle HTTP Server uses persistent HTTP connections by default, assuming they are supported on the client side. In HTTP over SSL as implemented by Oracle HTTP Server, SSL session state is kept in memory while the associated HTTP connection is persisted, a process which essentially eliminates the overhead of SSL session reuse (conceptually, the SSL connection is kept open along with the HTTP connection).

6.3.1.2 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. For example, on a modern processor, RSA estimates its RC4 cipher to take in the vicinity of 8-16 machine operations per output byte. Standard DES encryption will incur roughly 8 times the overhead of RC4, and triple DES will take about 25 times the overhead of DES. However, when using triple DES, the encryption costs will not be noticeable in most applications. Oracle HTTP Server supports these three cipher suites, and other cipher suites as well.

Oracle HTTP Server negotiates a cipher suite with a client based on the `SSLCipherSuite` attribute specified in `httpd.conf`.

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

6.3.1.3 SSL Performance Recommendations

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

1. The SSL handshake is an inherently expensive process in terms of both CPU usage and response time. Thus, use SSL only where needed. Determine the parts of the application that require the security, and the level of security required, and protect only those parts at the requisite security level. Attempt to minimize the need for the SSL handshake by using SSL sparingly, and by reusing session state as much as possible. For example, if a page contains a small amount of sensitive data and a number of non-sensitive graphic images, use SSL to transfer the sensitive data only, use normal HTTP to transfer the images. If the application requires server authentication only, do not use client authentication. If the performance goals of an application cannot be met by this method alone, additional hardware may be required.

2. Design the application to use SSL efficiently. Group secure operations together to take advantage of SSL session reuse and SSL connection reuse.
3. Use persistent connections, if possible, to minimize cost of SSL session reuse.
4. Tune the session cache timeout value (the `SSLSessionCacheTimeout` attribute in `httpd.conf`). A trade-off exists between the cost of maintaining an SSL session cache and the cost of establishing a new SSL session. As a rule, any secured business process, or conceptual grouping of SSL exchanges, should be completed without incurring session creation more than once. The default value for the `SSLSessionCacheTimeout` attribute is 300 seconds. It is a good idea to test an application's usability to help tune this setting.
5. If large volumes of data are being protected through SSL, pay close attention to the cipher suite being used. The `SSLCipherSuite` directive specified in `httpd.conf` controls the cipher suite. If lower levels of security are acceptable, use a less-secure protocol using a smaller key size (this may improve performance significantly). Finally, test the application using each available cipher suite for the desired security level to find the most performant suite.
6. Having taken the preceding considerations into account, if SSL remains a bottleneck to the performance and scalability of your application, consider deploying multiple Oracle HTTP Server instances over a hardware cluster or consider the use of SSL accelerator cards.

6.3.2 Oracle HTTP Server Port Tunneling Performance Issues

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

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

The number of OracleAS Port Tunneling processes is based on the degree of availability required, and the number of anticipated connections. This number cannot be automatically determined because for each additional process a new port must be opened through the firewall between the DMZ and the intranet. You cannot start more processes than you have open ports, and you do not want less processes than open ports, since in this case ports would not have any process bound to them.

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

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

6.4 Oracle HTTP Server Performance Tips

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

- [Analyze Static Versus Dynamic Requests](#)
- [Analyze Time Differences Between Oracle HTTP Server and OC4J Servers](#)
- [Beware of a Single Data Point Yielding Misleading Results](#)

6.4.1 Analyze Static Versus Dynamic Requests

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

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

6.4.2 Analyze Time Differences Between Oracle HTTP Server and OC4J Servers

In some cases, you may notice a high discrepancy between the average time to process a request in Oracle Containers for J2EE (OC4J) and the average response time experienced by the user. If the time is not being spent actually doing the work in OC4J, then it is probably being spent in transport.

If you notice a large discrepancy between the request processing time in OC4J and the average response time, consider tuning the Oracle HTTP Server directives shown in the section, "[Configuring Oracle HTTP Server Directives](#)" on page 6-2.

6.4.3 Beware of a Single Data Point Yielding Misleading Results

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

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

Note that `handle.maxTime` is much higher than `handle.avg` for this module. This is probably because when the first request is received, a database connection must be opened. Later requests can make use of the established connection. In this case, to obtain a better estimate of the average service time for a PL/SQL module, that does not include the database connection open time which causes the `handle.maxTime` to be very large, recalculate the average as in the following:

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

For example, in this case this would be:

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

Oracle BPEL Process Manager Performance Tuning

Oracle BPEL Process Manager provides a number of property settings that can be configured to optimize performance at the process, domain, application server, Java Virtual Machine (JVM), and dehydration store database levels. This chapter describes these property settings and provides recommendations on how to use them.

This chapter contains the following sections:

- [Performance Tuning Overview](#)
- [Process Level Performance Settings](#)
- [Tables Impacted By Instance Data Growth](#)
- [Domain Level Performance Tuning](#)
- [Tuning OC4J for Oracle BPEL](#)
- [Java Virtual Machine Performance Tuning for Oracle BPEL Server](#)
- [Dehydration Store Database Performance Tuning](#)
- [Summary](#)

7.1 Performance Tuning Overview

This section provides an overview of key Oracle BPEL Process Manager tuning concepts. Review this section before attempting to configure any property settings.

This section contains the following topics:

- [Domain and Process Configuration Property Settings](#)
- [Durable and Transient Processes](#)
- [One-Way and Two-Way Invocations](#)
- [Idempotent Activities](#)
- [In-Flight Database Storage](#)
- [JTA Transactions for Two-way Invocations](#)
- [BPEL Threading Model](#)

7.1.1 Domain and Process Configuration Property Settings

Domain and process configuration properties can be set at two different levels in Oracle BPEL Process Manager:

- Domain level: enables you to configure *all* processes deployed in a specific domain.
- Process level: enables you to specify which processes to configure, and which not to configure, in a specific domain. If a setting at the domain level conflicts with the same setting at the process level, the process level setting take priority.

7.1.2 Durable and Transient Processes

Oracle BPEL Process Manager uses the dehydration store database to maintain long-running asynchronous processes and their current state information in a database while they wait for asynchronous callbacks. Storing the process in a database preserves the process and prevents any loss of state or reliability if a system shuts down or a network problem occurs. There are two types of processes in Oracle BPEL Process Manager. These processes impact the dehydration store database in different ways.

- Transient processes: this process type does not incur any intermediate dehydration points during process execution. If there are unhandled faults or there is system downtime during process execution, the instances of a transient process do not leave a trace in the system. Instances of transient processes cannot be saved in-flight (whether they complete normally or abnormally). Transient processes are typically short-lived, request-response style processes. The synchronous process you design in Oracle JDeveloper is an example of a transient process.
- Durable processes: this process type incurs one or more dehydration points in the database during execution because of the following activities:
 - Receive activity
 - OnMessage branch in a pick activity
 - OnAlarm branch in a pick activity
 - Wait activity

Instances of durable processes can be saved in-flight (whether they complete normally or abnormally). These processes are typically long-living and initiated

through a one-way invocation. Because of out-of-memory and system downtime issues, durable processes cannot be memory-optimized.

7.1.3 One-Way and Two-Way Invocations

There are two types of invocations into BPEL process instances:

- A one-way invocation: a request-only operation and has only an inbound message.
- A two-way invocation: a request-and-response operation. The caller thread is blocked until a response is ready.

Table 7–1 describes the use of one-way and two-way invocations.

Table 7–1 One-Way and Two-Way Invocations

Use	One-Way Invocation	Two-Way Invocation
WSDL file definition	<code><operation name="oneway"> <input message="in"/> </operation></code>	<code><operation name="twoway"> <input message="in"/> <output message="out"/> </operation></code>
Variable declarations in BPEL activities	<code><receive operation="oneway" variable="in"/></code>	<code><receive operation="twoway" variable="in"/> ... <reply operation="twoway" variable="out"/></code>
Through-delivery service	The request is saved in the delivery service. The caller thread does not block until the message is delivered to the targeted instance.	The request is delivered into Oracle BPEL Server and the targeted BPEL instance. The caller thread is blocked until the response is ready.

7.1.4 Idempotent Activities

An idempotent activity is an activity that can be retried (for example, an assign activity or an invoke activity). Oracle BPEL Server saves the instance after a nonidempotent activity.

See Also: ["idempotent BPEL Property"](#) on page 7-7 for additional details

7.1.5 In-Flight Database Storage

Over its life cycle, a BPEL instance in its current state of execution can be saved multiple times in the dehydration store database. There are two cases in which this occurs:

- When the instance is waiting for an event. It can be either an alarm or an invocation message. This happens when one of the following BPEL activities is being executed:
 - Wait activity
 - OnAlarm branch of a pick activity
 - Receive activity
 - OnMessage branch of a pick activity

When a BPEL instance is saved to the dehydration store database, the instance is known as being dehydrated. When the event later occurs (the alarm expires or the message comes in), the instance is read from the database and resumes execution.

- After a nonidempotent activity. Instance storage is necessary here if you want to retry the steps. The retry occurs from the steps *after* the nonidempotent activity.

7.1.6 JTA Transactions for Two-way Invocations

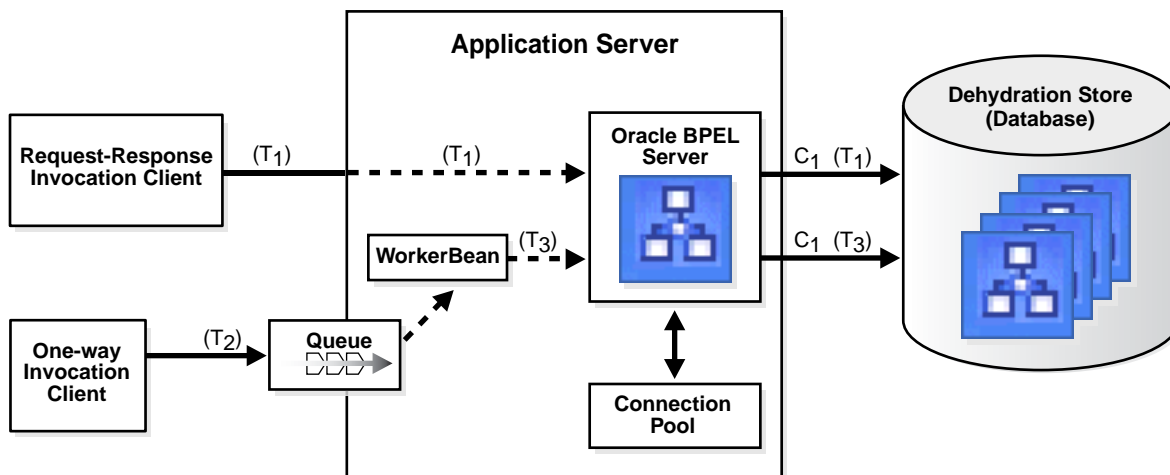
For two-way invocations, if the process being called is a transient process, Oracle BPEL Server honors the caller's Java Transaction API (JTA) transaction. If the process being called is a durable process, meaning an in-flight database save can be occurring, Oracle BPEL Server creates a new transaction.

7.1.7 BPEL Threading Model

Review the BPEL threading model details in this section before attempting to configure any property settings.

Figure 7-1 shows thread usage during a request-response and one-way process instance invocation.

Figure 7-1 Thread Usage



7.1.7.1 Request-Response Invocation

In Figure 7-1, the client is running in thread T_1 . When the caller initiates a process instance, the same thread is used during processing. Eventually, when database operations must be performed, the thread obtains a database connection (C_1 in Figure 7-1) from the connection pool.

7.1.7.2 One-Way Invocation

In Figure 7-1, the one-way invocation client is running in thread T_2 . When the client initiates a process instance, the invocation request is placed in a queue. At this point, thread T_2 is released by Oracle BPEL Server and the caller can continue its own processing. Inside Oracle BPEL Server, a message-driven bean (MDB), WorkerBean, monitors the queue for invocation requests. When a message is dequeued, Oracle BPEL Server allocates a separate thread (T_3) to process the message. This thread is used by Oracle BPEL Process Manager to process the instance. When database

operations must be performed, the thread obtains a database connection from the connection pool.

See Also: ["Oracle BPEL Server EJB Configuration"](#) on page 7-23 for additional details about WorkerBean

7.1.7.3 Threading and Connection Pool Relationships

From [Figure 7-1](#) and the previous sections, some important relationships can be derived for properly setting the threading and connection pooling parameters.

The number of concurrent instances being processed is determined by the number of request-response client requests and the number of WorkerBean threads allocated. The following relationship can be stated.

$$\text{Maximum DB Connections} \geq (\text{WorkerBean listener threads}) + (\text{Maximum concurrent request-response invocations})$$

The `dspMaxThreads` property allocates WorkerBean threads to various domains. This leads to the following relationships:

$$\sum_{\text{domains}} \text{dspMaxThreads} = (\text{WorkerBean listener threads})$$

$$\text{Maximum DB Connections} \geq (\sum_{\text{domains}} \text{dspMaxThreads}) + (\text{Maximum concurrent request-response invocations})$$

If only one domain exists, these formulas can be simplified further:

$$\text{dspMaxThreads} = (\text{WorkerBean listener threads})$$

$$\text{Maximum DB Connections} \geq (\text{dspMaxThreads}) + (\text{Maximum concurrent request-response invocations})$$

7.2 Process Level Performance Settings

This section describes process level performance tuning properties.

Process level performance properties are set in the `bpel.xml` file for a specific BPEL process. This file is in the same directory as the process's `.bpel` file. After modifying the settings in the `JDev_Oracle_Home\jdev\mywork\workspace_name\process_name\bpel\bpel.xml` file, the process must be redeployed for the new settings to take effect.

Note: You can also set these properties in the Deployment Descriptor Properties window of Oracle JDeveloper.

See Also: ["Domain and Process Configuration Property Settings"](#) on page 7-2

7.2.1 completionPersistLevel BPEL Property

This property controls the type (and amount) of data to save after instance completion.

When process instances complete, Oracle BPEL Server by default saves the final state (for example, the variable values) of the process. If you do not need to save these values after completion, you can set this property to save only instance metadata (completion state, start and end dates, and so on). This property is applicable to transient BPEL processes.

This property is used only when the `inMemoryOptimization` performance property is set to `true`. Use the `completionPersistLevel` property in conjunction with the `completionPersistPolicy` property.

This property can greatly impact database growth (in particular, the `cube_instance`, `cube_scope`, and `work_item` tables). It can also impact throughput (due to reduced I/O).

See Also:

- ["completionPersistPolicy BPEL Property"](#) on page 7-6
- ["inMemoryOptimization BPEL Property"](#) on page 7-8
- [Table 7-2](#) on page 7-10 for additional details about the `cube_instance`, `cube_scope`, and `work_item` tables

Values

This property has the following values:

- `all` (default): Oracle BPEL Server saves the complete instance, including the final variable values, work item data, and audit data. This setting causes the database to grow in size.
- `instanceHeader`: The Oracle BPEL Process Manager saves only the instance metadata.

Example

In the following example, only faulted instances are persisted (`completionPersistPolicy=faulted`). For the faulted instances, all variable values associated with the instance are saved (`completionPersistLevel=All`).

```
<BPELSuitcase>
  <BPELProcess src="HelloWorld.bpel" id="HelloWorld">
    ...
    <configurations>
      <property name="inMemoryOptimization">true</property>
      <property name="completionPersistPolicy">faulted</property>
      <property name="completionPersistLevel">All</property>
    </configurations>
  </BPELProcess>
</BPELSuitcase>
```

7.2.2 completionPersistPolicy BPEL Property

This property controls if and when to persist instances. If an instance is not saved, it does not appear in Oracle BPEL Control. This property is applicable to transient BPEL processes.

This property is only used when `inMemoryOptimization` is set to `true`. If you set `completionPersistPolicy` to a value other than `off`, you can then set `completionPersistLevel` to more finely tune the persistence data to save.

This parameter strongly impacts the amount of data stored in the database (in particular, the `cube_instance`, `cube_scope`, and `work_item` tables). It can also impact throughput.

See Also:

- ["completionPersistLevel BPEL Property"](#) on page 7-5
- ["inMemoryOptimization BPEL Property"](#) on page 7-8

Values

This property has the following values:

- `on`: (default): Completed instances are saved normally.
- `deferred`: Completed instances are saved with a different thread and in another transaction. If a server fails, some instances may not be saved.
- `faulted`: Only faulted instances are saved.
- `off`: No instances (and their data) are saved.

Example

In the following example, `completionPersistPolicy` is set to `deferred`:

```
<BPEL Suitcase>
  <BPELProcess src="HelloWorld.bpel" id="HelloWorld">
    . . .
    <configurations>
      <partnerLinkBinding name="PartnerService">
        <property name="inMemoryOptimization">true</property>
        <property name="completionPersistPolicy">deferred</property>
      </partnerLinkBinding>
    </configurations>
  </BPELProcess>
</BPEL Suitcase>
```

7.2.3 idempotent BPEL Property

A BPEL invoke activity is by default an idempotent activity, meaning that the BPEL process does not dehydrate instances immediately after invoke activities. Therefore, if `idempotent` is set to `true` and Oracle BPEL Server fails right after an invoke activity executes, Oracle BPEL Server performs the invoke again after restarting. This is because no record exists that the invoke activity has executed. This property is applicable to both durable and transient processes.

If `idempotent` is set to `false`, the invoke activity is dehydrated immediately after execution and recorded in the dehydration store. If Oracle BPEL Server then fails and is restarted, the invoke activity is not repeated, because Oracle BPEL Process Manager sees that the invoke already executed.

When `idempotent` is set to `false`, it provides better failover protection, but at the cost of some performance, since the BPEL process accesses the dehydration store much more frequently. This setting can be configured for each partner link in the `bpel.xml` file.

Setting this parameter to `true` can significantly improve throughput. However, as mentioned previously, you must ensure that the partner's service can be safely retried in the case of a server failure. Some examples of where this property can be set to `true` are read-only services (for example, `CreditRatingService`) or local EJB/WSIF invocations that share the instance's transaction.

Values

This property has the following values:

- `false`: activity is dehydrated immediately after execution and recorded in the dehydration store
- `true` (default): If Oracle BPEL Server fails, it performs the activity again after restarting. This is because the server does not dehydrate immediately after the invoke and no record exists that the activity executed.

Example

The following `bpel.xml` file example shows the `idempotent` property. This example shows a one-way invocation message being saved to the dehydration store database. This property can be set for each partner link.

```
<BPELSuitcase>
  <BPELProcess src="Invoke.bpel" id="Invoke">
    <partnerLinkBindings>
      . . .
      <partnerLinkBinding name="PartnerService">
        <property name="wsdlLocation">
          partner-wsdl
        </property>
        <property name="idempotent">false</property>
      </partnerLinkBinding>
    </partnerLinkBindings>
  </BPELProcess>
</BPELSuitcase>
```

7.2.4 inMemoryOptimization BPEL Property

This property indicates to Oracle BPEL Server that this process is a transient process and dehydration of the instance is not required. When set to `true`, Oracle BPEL Server keeps the instances of this process in memory only during the course of execution. This property can only be set to `true` for transient processes (that is, those that do not contain any middle process receive, pick, or wait activities).

The default for this property is `false`, which means that instances are persisted completely and recorded in the dehydration store database for a synchronous BPEL process.

When `inMemoryOptimization` is set to `true`, dehydration is deactivated, and Oracle BPEL Process Manager keeps instances in memory only. The settings for the `completionPersistPolicy` and `completionPersistLevel` properties are also examined to determine persistence behavior. The `inMemoryOptimization` property can improve throughput when set to `true` and, in conjunction with these two other properties, can minimize database growth.

See Also:

- ["completionPersistLevel BPEL Property"](#) on page 7-5
- ["completionPersistPolicy BPEL Property"](#) on page 7-6

Values

This property has the following values:

- `false` (default): instances are persisted completely and recorded in the dehydration store database for a synchronous BPEL process.
- `true`: Oracle BPEL Process Manager keeps instances in memory only.

Example

The following `bpel.xml` file example shows the `inMemoryOptimization` property for the synchronous Hello World BPEL process:

```
<BPEL Suitcase>
  <BPELProcess src="HelloWorld.bpel" id="HelloWorld">
    . . .
    <configurations>
      <property name="inMemoryOptimization">true</property>
    </configurations>
  </BPELProcess>
</BPEL Suitcase>
```

7.2.5 nonBlockingInvoke BPEL Property

This property can improve performance when executing multiple branches of a flow or flowN activity. By default, Oracle BPEL Process Manager executes in a single thread, executing the branches sequentially instead of in parallel. When this property is set to `true`, the process manager creates a new thread to perform each branch's invoke activity in parallel. This setting can be configured for each partner link in the `bpel.xml` file. This property is applicable to both durable and transient processes.

Consider setting this property to `true` if you have invoke activities in multiple flow or flowN branches. This is especially effective if the parallel invoke activities are two-way, but some benefits can be realized for parallel one-way invokes as well.

Values

This property has the following values:

- `true`: Oracle BPEL Server spawns a new thread to execute the invocation. This thread is essentially the `InvokerBean` message driven bean thread. If the process has additional nonblocking invoke activities, increase the `InvokerBean` thread value. You may also need to increase the connection pool maximum size:

```
connection pool size >= (InvokerBean listener threads + WorkerBean listener
threads + maximum concurrent request-response invocations)
```

- `false` (default): Oracle BPEL Server executes the invoke activity in the single process thread.

See Also: "[InvokerBean](#)" on page 7-24 for instructions on configuring the `InvokerBean`

Example

The following `bpel.xml` file example enables the `nonBlockingInvoke` property:

```
<BPELSuitcase>
  <BPELProcess src="Invoke.bpel" id="Invoke">
    <partnerLinkBindings>
      . . .
      <partnerLinkBinding name="PartnerService">
        <property name="wsdlLocation">
          partner-wsdl
        </property>
        <property name="nonBlockingInvoke">true</property>
      </partnerLinkBinding>
    </partnerLinkBindings>
  </BPELProcess>
</BPELSuitcase>
```

7.3 Tables Impacted By Instance Data Growth

Instance data occupies space in Oracle BPEL Process Manager schema tables. [Table 7–2](#) describes the tables that are impacted by instance data growth. A brief description is provided of each table. The values to which you can set some domain level performance properties described in "[Domain Level Performance Tuning](#)" on page 7-11 impact the growth of these tables.

Table 7–2 Oracle BPEL Process Manager Tables Impacted By Instance Data Growth

Table Name	Table Description
audit_details	<p>Stores audit details that can be logged through the API. Activities such as an assign activity log the variables as audit details by default. You can set this behavior through the auditLevel property in Oracle BPEL Control under Manage BPEL Domain > Configuration.</p> <p>Audit details are separated from the <code>audit_trail</code> table due to their large size. To view a detail, click a link on the Audit tab for a specific instance in Oracle BPEL Control and load the detail separately. The auditDetailThreshold property in Oracle BPEL Control under Manage BPEL Domain > Configuration is used by this table. 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.</p> <p>See Also: "auditDetailThreshold BPEL Property" on page 7-11 and "auditLevel BPEL Property" on page 7-12</p>
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.
cube_instance	Stores process instance metadata (for example, the instance creation date, current state, title, and process identifier)
cube_scope	Stores the scope data for an instance (for example, all variables declared in the BPEL flow and some internal objects that help route logic throughout the flow).
dlv_message	Stores callback messages upon receipt. This table only stores the metadata for a message (for example, current state, process identifier, and receive date).
dlv_subscription	Stores delivery subscriptions for an instance. Whenever an instance expects a message from a partner (for example, the receive or onMessage activity) a subscription is written out for that specific receive activity.
document_ci_ref	Stores cube instance references to data stored in the <code>xml_document</code> table.
document_dlv_msg_ref	Stores references to <code>invoke_message</code> and <code>dlv_message</code> documents stored in the <code>xml_document</code> table.
invoke_message	Stores incoming (invocation) messages (messages that result in the creation of an instance). This table only stores the metadata for a message (for example, current state, process identifier, and receive date).
schema_md	Stores metadata about columns defined in the Oracle BPEL Process Manager schema (<code>orabpel</code>).
task	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 <code>work_item</code> 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, <code>invoke_message</code> documents, <code>dlv_message</code> documents, and so on). 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.

7.4 Domain Level Performance Tuning

This section describes domain level performance tuning properties.

Oracle recommends that you modify these settings in Oracle BPEL Control under **Manage BPEL Domain > Configuration**. Oracle BPEL Control checks the existing settings and any new settings entered, and validates them without requiring a restart. Domain level performance settings are located in the `SOA_Oracle_Home\bpel\domains\domain_name\config\domain.xml` file. If you directly edit the `domain.xml` file, you must restart Oracle BPEL Server for the new settings to take effect.

See Also: ["Domain and Process Configuration Property Settings"](#) on page 7-2

7.4.1 Oracle BPEL Control Properties That Cannot Be Edited

The following properties display in Oracle BPEL Control under **Manage BPEL Domain > Configuration**. These properties have empty **Name** and **Comment** columns in the Configuration tab. Do *not* modify these properties; this has no impact on system performance tuning.

- `cbCacheHighWatermark`
- `cbCacheLowWatermark`
- `cbCachePolicy`
- `cbCacheUnits`
- `instCacheUnits`
- `invCacheHighWatermark`
- `invCacheLowWatermark`
- `invCachePolicy`
- `invCacheUnits`
- `subCacheHighWatermark`
- `subCacheLowWatermark`
- `subCachePolicy`
- `subCacheUnits`

Instead, see the following subsections for details about properties that can be set to optimize performance.

7.4.2 `auditDetailThreshold` BPEL Property

This property sets the maximum size (in kilobytes) of an audit trail details string before it is stored separately from the audit trail. If an audit trail details string is larger than the threshold setting, it is not immediately loaded when the audit trail is initially retrieved; a link is displayed with the size of the details string. Strings larger than the threshold setting are stored in the `audit_details` table, instead of the `audit_trail` table.

This property is applicable to durable processes.

The details string typically contains the contents of a BPEL variable. In cases where the variable is very large, performance can be severely impacted by logging it to the audit trail.

Values

The default value is 50 kilobytes.

See Also: [Table 7-2](#) on page 7-10 for additional information about the `audit_trail` and `audit_details` tables

7.4.3 auditLevel BPEL Property

This property sets the audit trail logging level. This process is applicable to both durable and transient processes.

This property controls the amount of audit events logged by a process. This setting greatly impacts performance because more audit events means more database inserts into the `audit_trail` table. This audit information is used only for viewing the state of the process from Oracle BPEL Control.

Use this property if you do not want to store all audit information. Choose the level according to your business requirement. Auditing information has a significant impact on database growth and throughput. For optimal performance, set this property to the lowest acceptable level.

See Also: [Table 7-2](#) on page 7-10 for additional information about audit level details and the `audit_trail` table

Values

This property has the following values:

- `off`: No audit events (activity execution information) are persisted and no logging is performed; this can result in a slight performance boost for processing instances.
- `minimal`: all events are logged; however, no audit details (variable content) are logged. This setting is recommended for larger payload processes.
- `production`: all events are logged. The audit details for assign activities are not logged; the details for all other activities are logged. This setting is recommended for smaller payload processes.
- `development` (default): all events are logged; all audit details for all activities are logged.

7.4.4 bpelcClasspath BPEL Property

This property sets the BPEL process compiler classpath.

This is the server-side BPEL process compiler classpath. Any user-specific classes and libraries used by a BPEL Java exec activity (that have not been packaged in the BPEL archive) must be specified in this classpath. This enables the server-side BPEL process compiler to successfully compile the BPEL process.

This process is applicable to both durable and transient processes.

Values

The default value is:

`Oracle_Home\bpel\system\classes;`

`Oracle_Home\bpel\lib\j2ee_1.3.01.jar`

7.4.5 datasourceJndi BPEL Property

This property sets the domain data source JNDI name.

This data source can refer to any data source (JTA is not required).

This process is applicable to both durable and transient processes.

Values

The default value is `jdbc/BPELServerDataSourceWorkflow`.

7.4.6 deliveryPersistPolicy BPEL Property

WARNING: Oracle recommends that this property remain set to the default value of `on`. If you set this property to `off` and your system fails, you lose messages. Exercise extreme care if changing this property setting from the default value.

This property enables and disables database persistence of messages entering Oracle BPEL Server. By default, incoming requests are saved in the following delivery service database tables:

- `dlv_message`
- `invoke_message`

These requests are later acquired by Oracle BPEL Server worker threads and delivered to the targeted BPEL process. In the case where performance is preferred over reliability, persisting the incoming messages in the database can be skipped. This property persists delivery messages and is applicable to durable processes.

One-way invocation messages are stored in the delivery cache until delivered. If the rate at which one-way messages arrive is much higher than the rate at which Oracle BPEL Server delivers them or if the server fails, some messages can get lost. In Oracle BPEL Control (under **Manage BPEL Domain >Threads**), you can monitor the size of the delivery cache by viewing the **New Instance Requests** and **Callback Requests** statistics in the **Pending Requests** section. The **Scheduled** column indicates the number of cached messages.

See Also: [Table 7-2](#) on page 7-10 for additional details about the delivery service database tables

Values

This property has the following values:

- `on` (default): delivery messages are persisted in the database
- `off`: incoming delivery messages are kept only in the in-memory cache. If more messages are delivered, the system can become overloaded (messages become backlogged in the scheduled queue) and you receive out-of-memory errors. Tune the number of WorkerBean threads to accommodate the number of incoming messages.
- `off.immediate`: directs Oracle BPEL Server to bypass the scheduling of messages in the invoke queue, and invokes the BPEL instance synchronously.

See Also: ["WorkerBean"](#) on page 7-24

7.4.7 dspAgentDelay BPEL Property

This property sets the number of seconds between triggers of the dispatcher agent. This agent cleans up any messages in the dispatcher layer that have not been processed due to a failure in the JMS layer.

This process is applicable to durable processes.

Values

The default value is 120 seconds.

7.4.8 dspInvokeAllocFactor BPEL Property

This property sets the percentage of active threads to be tasked to process incoming invocation messages. After a thread has finished processing a message, it can be tasked again to process an Oracle BPEL Server or invocation message, depending upon the current thread allocation situation.

This process is applicable to durable processes.

Values

The default value is 0.4 (40%).

7.4.9 dspMaxRequestDepth BPEL Property

This property sets the maximum number of in-memory activities to process within the same request. After processing an activity request, Oracle BPEL Process Manager attempts to process as many subsequent activities as possible without jeopardizing the transactionality of the request. Once the activity processing chain has reached this depth, the instance is dehydrated and the next activity is performed in a separate transaction.

If the request depth is too large, the total request time can exceed the application server transaction timeout limit.

This process is applicable to durable processes.

Values

The default value is 600 activities.

7.4.10 dspMaxThreads BPEL Property

This property sets the maximum number of active dispatcher threads that process messages during peak load times. This property is applicable to durable processes and is dependent on the application server configuration.

This is the simplest way to improve the performance and scalability of the domain. Oracle BPEL Server uses MDB threads to process Oracle BPEL Server messages. The maximum value for this property is dependent upon the Oracle BPEL Server's MDB J2EE listener threads setting. For Oracle Application Server, this count is configured in the `orion-ejb-jar.xml` deployment descriptor file.

For example, if the total number of MDB J2EE listener threads is 120, the value of **dspMaxThreads** can be set to 120 or less. If you have configured multiple domains,

the sum of the **dspMaxThreads** settings for all domains must not exceed the MDB J2EE listener threads setting.

If the CPU utilization of your application server and database hosts are well below capacity, try increasing this value and the MDB J2EE listener threads setting when necessary. If the CPUs are still not fully utilized, then consider running multiple Oracle BPEL Server instances.

Note: MDB J2EE listener threads configuration is specified in the following file:

- For the Oracle BPEL Process Manager for Developers installation type, this file is located at *SOA_Oracle_Home\j2ee\home\application-deployments\orabpel\ejb_ob_engine\orion-ejb-jar.xml* under WorkerBean.
 - For the Oracle BPEL Process Manager for OracleAS Middle Tier installation type, this file is located at *SOA_Oracle_Home\j2ee\home\application-deployments\orabpel\ejb_ob_engine\orion-ejb-jar.xml* under WorkerBean.
-
-

See Also:

- ["BPEL Threading Model"](#) on page 7-4 for complete details
- ["Oracle BPEL Server EJB Configuration"](#) on page 7-23 for MDB J2EE listener thread details

Values

The default value is 100 threads.

7.4.11 dspMinThreads BPEL Property

This property sets the minimum number of active dispatcher threads that process messages during peak load times.

If the current number of active threads is under this number, the load factor is not taken into consideration when determining whether or not to allocate a new thread.

This process is applicable to durable processes.

Values

The default value is 5 threads.

7.4.12 expirationMaxRetry BPEL Property

This property sets the maximum number of times a failed expiration call (in a wait activity or an onAlarm branch of a pick activity) is retried before failing.

If the activity or instance targeted by the expiration call cannot be found, the call is rescheduled again.

The retry count does not include the first (original) attempt by the expiration call.

This process is applicable to durable processes.

Values

The default value is 5.

7.4.13 idempotentThreshold BPEL Property

This property sets the maximum time (in seconds) in which an idempotent service must successfully complete an activity. If an idempotent service takes longer than this time to complete, the service is considered nonidempotent and the current transaction is committed to the database. This feature prevents lengthy services from having to redo work in case another service in the idempotent chain fails.

Values

The default value is 30 seconds.

See Also: ["Idempotent Activities"](#) on page 7-3

7.4.14 instanceKeyBlockSize BPEL Property

This property controls the instance ID range size. Oracle BPEL Server creates instance keys (a range of process instance IDs) in batches using this number. After creating this range of in-memory IDs, the next range is updated and saved in the `ci_id_range` table. For example, if `instanceKeyBlockSize` is set to 100, Oracle BPEL Server creates a range of instance keys in-memory (100 keys, which are later inserted into the `cube_instance` table as `cikey`). If the block size is smaller than the number of updates to the `ci_id_range` table, this may cause performance issues.

See Also: [Table 7-2](#) on page 7-10 for additional details about the `cube_instance` table

Values

The default value is 10000.

7.4.15 instCacheHighWatermark BPEL Property

Note: Oracle recommends that you do not change this parameter. Only change this parameter if you fully understand JVM issues.

This property sets the maximum number of in-flight instances that can be placed in the cache before pruning occurs. Once the high watermark is reached, the cache removes (prunes) enough older instances from cache to reach the low watermark value (set with the `instCacheLowWatermark` property). Pruned instances can be retrieved as needed from the dehydration store. This property is applicable to durable processes.

This value is only used when the `instCachePolicy` property is set to `lru` or `hybrid`. Consider the following factors when setting this property:

- The number of in-flight instances Oracle BPEL Process Manager is expected to handle at any point in time
- The amount of memory each process instance takes. The memory usage can be determined using a Java Profiler.

You can run a single instance through the system and measure the corresponding increase in memory utilization.

If this property is set too high, your system can encounter `OutOfMemoryException` error messages. The system can also actually slow down if this value is set too high because the garbage collector runs more frequently. To monitor the garbage collector,

use Sun's visual garbage collection (GC) tool
(<http://java.sun.com/performance/jvmstat>).

See Also:

- "[instCacheLowWatermark BPEL Property](#)" on page 7-17
- "[instCachePolicy BPEL Property](#)" on page 7-18
- "[optCacheOn BPEL Property](#)" on page 7-19

Values

The default value is 3000; zero implies no limit.

7.4.16 instCacheLowWatermark BPEL Property

Note: Oracle recommends that you do not change this parameter. Only change this parameter if you fully understand JVM issues.

This property sets the number of in-flight instances to which the cache is pruned when pruning occurs. This property is applicable to durable processes.

When the high watermark in the cache is reached, the cache removes enough instances to reach this level.

Cache pruning occurs when the cache size grows to the high watermark value (set with the **instCacheHighWatermark** property). This **instCacheLowWatermark** property controls how much pruning occurs. The default value is 75% of the high watermark setting. This indicates the cache is reduced to 75% of the high watermark value when pruning occurs. This value is only used when the **instCachePolicy** property is set to `lru` or `hybrid`.

Monitor instance cache statistics by going to **Manage BPEL Domain > Threads** in Oracle BPEL Control. In the **Server cache statistics** section at the bottom of this page is the **instance cache** entry. You can view the cache size and hit percentage. If the hit percentage is quite low, consider increasing your cache size or the low watermark value.

If this property is set too high, your system can encounter `OutOfMemoryException` errors. The system can also actually slow down if this value is set too high. This is because the garbage collector must run more frequently. To monitor the garbage collector, use Sun's visual GC tool
(<http://java.sun.com/performance/jvmstat>).

See Also:

- "[instCacheHighWatermark BPEL Property](#)" on page 7-16
- "[instCachePolicy BPEL Property](#)" on page 7-18
- "[optCacheOn BPEL Property](#)" on page 7-19

Values

The default value is 2250 (75%).

7.4.17 instCachePolicy BPEL Property

This property sets the eviction policy to use when removing in-flight instances from the cache. This property is applicable to durable processes.

This property takes effect only when the `optCacheOn` property is set to `true`.

If you want to fine tune cache management, use this property. If the number of process instances that must be kept in memory is well known, Oracle recommends the `lru` setting. When using the `lru` setting, the `instCacheHighWatermark` and `instCacheLowWatermark` properties must also be set.

Note: Some JVM implementations have been observed to display an `OutOfMemoryException` error message when cache values are set to `auto`. This happens because the auto caching setting relies on JVM soft references. If you encounter this error, set the caching value to `lru`.

See Also:

- ["instCacheHighWatermark BPEL Property"](#) on page 7-16
- ["instCacheLowWatermark BPEL Property"](#) on page 7-17
- ["optCacheOn BPEL Property"](#) on page 7-19

Values

This property has the following values:

- `lru`: least recently used; this setting first removes those instances that have not been accessed for the longest period of time. This setting is recommended.
- `auto` (default): delegates the removal decision to the JVM. Instances are removed when the garbage collector reaps soft references.
- `soft-lru`: combination of `lru` and `auto`.

7.4.18 invokerQueueConnectionPoolMinSize BPEL Property

This property sets the invoker queue connection pool minimum size. This value must match the number of invoker threads. If the invoker threads are set to 200, this value can be set to 200 to avoid JMS warm up.

This property is applicable to both durable and transient processes.

Values

The default value is 25.

See Also: ["InvokerBean"](#) on page 7-24 for details about invoker threads

7.4.19 largeDocumentThreshold BPEL Property

This property sets the large XML document persistence threshold. This is the maximum size (in kilobytes) of a BPEL variable before it is stored in a separate location from the rest of the instance scope data.

This property is applicable to both durable and transient processes.

Large XML documents impact the performance of the entire Oracle BPEL Server if they are constantly read in and written out whenever processing on an instance must be performed.

Values

The default value is 50 kilobytes.

7.4.20 minBPELWait BPEL Property

This property sets the minimum BPEL activity wait.

If the wait time for a wait activity or an onAlarm branch of a pick activity is less than the value defined here, the wait is ignored.

This property is applicable to durable processes.

Values

The default value is 2 seconds.

7.4.21 optCacheOn BPEL Property

This property sets the in-memory cache for in-flight instances. This property is applicable to durable processes.

If set to `true`, Oracle BPEL Process Manager attempts to load active instances from in-memory cache rather than looking them up from the database. To disable optimization, specify a value other than `true`.

Set this property to `false` if your process is long running and the subprocesses do not immediately call back. Consider this if you are dealing with shorter processes, which expect many callbacks.

Setting this property to `true` necessitates setting the following caching-related settings:

- `instCacheHighWatermark`
- `instCacheLowWatermark`
- `instCachePolicy`

If you can meet your performance goals without using the cache, Oracle recommends leaving this setting as `false` to simplify administration and tuning.

Note: Enabling the cache may adversely impact performance. This can happen if cache values are set too high, causing the JVM garbage collector to run at frequent intervals. Use Sun's visual GC tool (<http://java.sun.com/performance/jvmstat>) to monitor the garbage collector.

Values

This property has the following values:

- `true`: Oracle BPEL Server attempts to load active instances from in-memory cache rather than looking them up from the database.
- `false` (default): Oracle BPEL Server loads the instance from the database every time.

See Also:

- ["instCacheHighWatermark BPEL Property"](#) on page 7-16
- ["instCacheLowWatermark BPEL Property"](#) on page 7-17
- ["instCachePolicy BPEL Property"](#) on page 7-18

7.4.22 optIdempotentRouting BPEL Property

Note: Oracle recommends that you do not change this parameter.

This property sets a routing shortcut for idempotent services.

If set to `true`, Oracle BPEL Server attempts to process as many activities as possible within the same transaction if the activity services are idempotent.

This property is applicable to durable processes.

The default value is `true`. To disable optimization, specify a value other than `true`.

Values

The default value is `true`.

7.4.23 optSoapShortcut BPEL Property

Note: Oracle recommends that you do not change this parameter.

This property sets a short-cut for a local SOAP request.

Local SOAP calls are normally performed with an internal call instead of sending a message through the SOAP stack.

The default behavior for the Oracle BPEL Process Manager is to optimize all by bypassing the SOAP stack. To disable optimization, specify a value other than `true`.

This property is applicable to both durable and transient processes.

Values

- `true` (default): Local SOAP calls bypass the SOAP stack.
- `false`: Local SOAP calls go through the SOAP stack.

7.4.24 processCheckSecs BPEL Property

This property sets the number of seconds to wait since the last time Oracle BPEL Server checked the BPEL archive before checking it again. Checking means to check the last modified time stamp on the BPEL archive for a particular process. If the specified number of seconds has passed and the BPEL archive file has been modified since the last time checked, the process is refreshed from the new archive. If not enough time has passed since the last time the stale check was performed, the currently-loaded process classes are used.

This property is applicable to both durable and transient processes.

To disable process checking, use a value of -1. In this case, once a process has been loaded, Oracle BPEL Server never checks if a newer version of the same process has been deployed.

Values

The default value is 1 second.

7.4.25 relaxBpelAssignRules BPEL Property

Note: While this property does display in Oracle BPEL Control, Oracle recommends that you do *not* use this property. This property has been deprecated.

This property relaxes enforcement of the *Business Process Execution Language for Web Services Specification Version 1.1* assign rules. If set to `true`, Oracle BPEL Process Manager does not apply rules while assigning BPEL variables. For example, Oracle BPEL Process Manager does not display an error about null assignments (which are not allowed in the BPEL specifications).

This property is applicable to both durable and transient processes.

Values

This property has the following values:

- `false` (default): does not relax assignment rules.
- `true`: relaxes assignment rules.

7.4.26 slowPerfThreshold BPEL Property

This property sets the maximum time (in seconds) for a service to successfully complete an activity. If a service takes longer than this time to complete, the service is considered slow. Oracle BPEL Process Manager collects statistics on slow services.

This property is applicable to durable processes.

Values

The default value is 1 second.

7.4.27 statsLastN BPEL Property

This property sets the size of the most-recently processed request list. After each request is finished, statistics for the request are kept in a list. A value less than or equal to zero disables statistics gathering.

This property is applicable to both durable and transient processes.

You can view statistics from Oracle BPEL Control under **Manager BPEL Domain > Statistics**.

Values

The default value is 1000.

7.4.28 syncMaxWaitTime BPEL Property

This property sets the maximum time the process result receiver waits for a result before returning. Results from asynchronous BPEL processes are retrieved synchronously by a receiver that waits for a result from Oracle BPEL Server.

This property is applicable to transient processes.

Values

The default value is 45 seconds.

7.4.29 txDataSourceJndi BPEL Property

This property sets the domain transactional data source JNDI name. This data source must be configured for JTA support.

This property is applicable to both durable and transient processes.

Values

The default value is `jdbc/BPELServerDataSource`.

7.4.30 uddiLocation BPEL Property

This property specifies the inquiry URL of the Universal Description, Discovery, and Integration (UDDI) version 3-compliant registry.

If you use virtual locations, meaning that you only reference the abstract WSDL in your partner link and provide a property named `registryServiceKey` in the deployment descriptor on the partner link binding level, this property is used to connect to the UDDI registry to retrieve the information.

Values

There is no default value.

7.4.31 validateXML BPEL Property

This property validates incoming and outgoing XML documents.

If set to `true`, the Oracle BPEL Process Manager applies schema validation for incoming and outgoing XML documents.

This property is applicable to both durable and transient processes.

values

The default value is `false`.

7.4.32 workerQueueConnectionPoolMinSize BPEL Property

This property sets the worker queue connection pool minimum size. This value must match the number of worker threads. If the number of worker threads is 200, this value can be set to 200 to avoid JMS warm up.

This property is applicable to durable processes.

Values

The default value is 25.

See Also: ["WorkerBean"](#) on page 7-24

7.5 Tuning OC4J for Oracle BPEL

The parameters described in this section are set at the Oracle Application Server level. You must restart the OC4J instance for these parameters to take effect.

This section contains the following topics:

- [Tuning JTA Transaction Timeout for Oracle BPEL Process Manager](#)
- [Oracle BPEL Server EJB Configuration](#)
- [Configuring Data Sources for Oracle BPEL](#)

See Also: *Oracle Application Server Administrator's Guide* for instructions on starting and stopping Oracle Application Server

7.5.1 Tuning JTA Transaction Timeout for Oracle BPEL Process Manager

Oracle BPEL Server uses JTA to achieve atomicity. The transaction timeout value is set by default to 60000 milliseconds in the `transaction-manager.xml` file. The location of this file depends on the method by which you installed Oracle BPEL Process Manager:

- For Oracle Application Server SOA installations, the file is located in `SOA_Oracle_Home\j2ee\home\config`.
- For Oracle BPEL Process Manager installations, the file is located in `SOA_Oracle_Home\bpel\system\appserver\oc4j\j2ee\home\config`.

You can sometimes experience transaction rollback errors due to timeouts, especially when Oracle BPEL Server is under stress. The timeout can happen for many reasons:

- Insufficient resources (for example, not enough database connections in the connection pool, the server thread waits for 60 seconds and displays a timeout error, and so on).
- Large document manipulation (for example, database writes of very large documents can take longer than 60 seconds).

Change this value according to your process. The following example sets the timeout to 120 seconds:

```
<transaction-config timeout="120000" />
```

If your process invokes partners that take longer than the specified timeout threshold, call them using a one-way request or set the `nonBlockingInvoke` partner link property to `true` in the `bpel.xml` deployment descriptor file.

See Also: ["nonBlockingInvoke BPEL Property"](#) on page 7-9

7.5.2 Oracle BPEL Server EJB Configuration

To increase performance, Oracle recommends removing the `max-instances` attribute for all of Oracle BPEL Server's EJBs in the `orion-ejb-jar.xml` file. For the Oracle BPEL Process Manager for OracleAS Middle Tier installation type, this file is located in `SOA_Oracle_Home\j2ee\home\application-deployments\orabpel\ejb_ob_engine`.

This enables the application server to allocate more resources to heavily-used beans.

7.5.2.1 WorkerBean

Oracle BPEL Server uses an MDB called `WorkerBean` to perform processing. Therefore, it is important to allocate enough threads to this MDB. Otherwise, resource utilization is not optimal. The following code from the `orion-ejb-jar.xml` file shows an allocation of 70 threads.

```
<message-driven-deployment name="WorkerBean"
  destination-location="jms/collaxa/BPELWorkerQueue"
  connection-factory-location="jms/collaxa/BPELWorkerQueueFactory"
  listener-threads="70" min-instances="100">
  <ejb-ref-mapping name="ejb/local/DispatcherLocalBean" />
  ..
  ..
</message-driven-deployment>
```

7.5.2.2 InvokerBean

The invoker bean is used only for nonblocking invoke activities. If you set some invokes to be nonblocking, increase the number of threads allocated to the `InvokerBean`. The following `orion-ejb-jar.xml` code shows an allocation of 30 threads.

```
<message-driven-deployment name="InvokerBean"
  destination-location="jms/collaxa/BPELInvokerQueue"
  connection-factory-location="jms/collaxa/BPELInvokerQueueFactory"
  listener-threads="30" min-instances="100">
<ejb-ref-mapping name="ejb/local/ProcessManagerLocalBean" />
  ..
  ..
</message-driven-deployment>
```

Note: The sum of the `InvokerBean` and `WorkerBean` threads must be greater than or equal to the value specified for the `dspMaxThreads` domain property in Oracle BPEL Control under **Manage BPEL Domain > Configuration**. If you configured multiple domains, add the `dspMaxThreads` property for all your domains and compare that sum to the MDB total thread count.

See Also:

- ["nonBlockingInvoke BPEL Property"](#) on page 7-9
- ["dspMaxThreads BPEL Property"](#) on page 7-14

7.5.3 Configuring Data Sources for Oracle BPEL

Oracle BPEL Server obtains database connections using an application server JTA data source. Oracle BPEL Server by default is configured to use the Oracle Database Lite dehydration store. For stress testing and production, Oracle recommends that you use Oracle Database 10g. Oracle Database Lite packaged with the default installation to ease the initial developer experience.

Be aware of the following issues when configuring the Oracle BPEL Server data source entry. For the Oracle BPEL Process Manager for OracleAS Middle Tier installation type, the data source entry is located in the `SOA_Oracle_Home\j2ee\home\config\data-sources.xml` file.

- When configuring the data source, ensure that the connection pool has enough free connections to serve Oracle BPEL Server.
- The connection pool size must be greater than or equal to the sum of the **dspMaxThreads** property value in Oracle BPEL Control. If you have configured multiple domains, add all **dspMaxThreads** property values and compare that value with the data source's `max-connections` value. The default `max-connections` value is unlimited.
- For Oracle Database 10g, the data source must also use the thin driver. For the Oracle9i Database, Oracle Call Interface (OCI) performs slightly better.

When database persistence is enabled, Oracle BPEL Server generally performs better with JDBC statement caching enabled using the `num-cached-statements` attribute. Statement caching eliminates overhead due to repeated cursor creation and repeated statement parsing and creation. Statement caching also reduces the overhead of communication between the application server and the database server.

See Also:

- ["dspMaxThreads BPEL Property"](#) on page 7-14
- *Oracle BPEL Process Manager Installation Guide* for instructions on configuring an Oracle Database as the Oracle BPEL Process Manager dehydration store
- *Oracle Containers for J2EE Services Guide* for additional information on setting JDBC statement caching
- ["Enable Statement Caching for Data Sources"](#) on page 3-8

7.6 Java Virtual Machine Performance Tuning for Oracle BPEL Server

JVM parameters can have an impact Oracle BPEL Server performance. The major factors that impact performance relate to the heap size. The heap size controls the amount of memory the JVM uses. If your BPEL process instance runs on a dedicated host, set the heap size value as high as possible.

Another important heap configuration is the garbage collector's generational settings. The garbage collector optimizes collection by classifying objects by how long they live. Most Oracle BPEL Server objects are short-lived; thus they live in the Eden space. Oracle recommends sizing the Eden space to be 60 to 70 percent of the total heap size.

The JVM `-xmn` setting startup option sets an explicit value for the Eden space size. Do the following to set this option:

1. Calculate the 60-70 percent value, based on the specified maximum heap size.
2. Use the calculated value to set the JVM `-xmn` command line parameter.

Take the following steps to change JVM command line options:

1. Navigate to the Home page for the OC4J instance.
2. Click Administration.
3. If necessary, expand the Properties section of the table by clicking the Expand icon. Then, click the Go to Task icon in the Server Properties row.
4. In the Command Line Options area, modify or change the appropriate command line options in the Options table.
5. Click Apply.

6. Navigate to the Cluster Topology page, select the OC4J instance that you modified, and click Restart. On the Confirmation page, click Yes.

See Also:

- ["Ensure Sufficient Java Heap for OC4J"](#) on page 3-3
- ["Tune the JVM Garbage Collection Options"](#) on page 3-4

7.7 Dehydration Store Database Performance Tuning

Oracle BPEL Server performance is related to the dehydration store's capacity. Oracle recommends the following:

- Moving the redo logs into a separate RAID 1+0 disk
- Increasing the size of each redo log file to a large value (for example, 1 GB)
- Creating a separate database tablespace for Oracle BPEL Server

The database parameters shown in [Table 7-3](#) impact Oracle BPEL Process Manager performance. The specific values to use depend on your hardware configuration.

Table 7-3 Database Parameters Impacting Oracle BPEL Process Manager Performance

Parameter Name	Sample Value
LOG_BUFFER	1048576
SHARED_POOL_SIZE	400M
JOB_QUEUE_PROCESSES	1
DB_CACHE_SIZE	1000M
DB_FILE_MULTIBLOCK_READ_COUNT	8
UNDO_RETENTION	0
PROCESSES	300
SESSION_CACHED_CURSORS	100

See Also: *Oracle Database Tuning* and *Oracle Database Reference* for your Oracle Database release:

- For Oracle Database 10g Release 2 (10.2)
<http://www.oracle.com/technology/documentation/databases/e10gr2.html>
- For Oracle Database 10g Release 1 (10.1)
<http://www.oracle.com/technology/documentation/databases/e10g.html>
- For Oracle Database 9i Release 2 (9.2)
<http://www.oracle.com/technology/documentation/oracle9i.html>

7.8 Summary

This chapter describes how to configure Oracle BPEL Process Manager property settings to optimize performance at the process, domain, application server, Java

Virtual Machine (JVM), and dehydration store database levels. This chapter describes these property settings and provides recommendations on how to use them.

Oracle Business Activity Monitoring Performance

To achieve the highest performance for Oracle Business Activity Monitoring, you should maintain a database on its own hardware dedicated to the Oracle Business Activity Monitoring system. Additionally, following certain database administration practices maximizes the throughput of incoming Oracle Business Activity Monitoring transactions. These practices are based on specific testing and observations of the Oracle Business Activity Monitoring system. General database administration practices, as described in the *Database Performance Tuning Guide*, also apply to a database dedicated to Oracle Business Activity Monitoring.

This chapter includes the following sections:

- [Managing the Redo Log Files](#)
- [Avoiding Frequent Log Switches and Checkpoints](#)
- [Tuning the System Global Area](#)
- [Database Re-Organization in the Presence of Deletion Activity](#)
- [Configuring Multiple Plan Monitor Services and Enterprise Links](#)

8.1 Managing the Redo Log Files

When Oracle Business Activity Monitoring receives input data at a high rate, the Oracle Business Activity Monitoring Active Data Cache (ADC) sends the data it receives to the database with little or no loss of input bandwidth incurred in the ADC server. Thus, the limiting factor for input data throughput is the rate at which the database can capture data it receives from the ADC server. At high data rates, the database throughput is limited by the ability of the database to write redo log records to disk, also called a log file sync. When an application, in this case the ADC server, commits data at a rate that is faster than the rate at which redo log data can be written to disk, subsequent requests to commit or rollback a transaction must wait.

The overall goal of redo log file tuning is to reduce or eliminate log file sync waits since they limit insertion throughput. There are two main strategies to reduce the number of log file sync waits in an Oracle database:

- You can increase the I/O bandwidth available to redo log sync activity. To do this, move the redo log files onto their own physical disk so they do not have to compete with other I/O activity in the system.
- You can modify the database insert workload to increase the number of rows inserted into the database per transaction committed. This strategy involves modifying the application to issue fewer commits. By batching a larger number of operations in each transaction, the demand for commit processing is reduced. The downside of such an approach, however, is that when there is a system failure, all data input since the previous commit is lost. With frequent commits, a small amount of data is lost when there is a system failure. With less frequent commits, a larger amount of recently captured data is lost. This tradeoff between performance and reliability implies that selecting a larger number of input operations per transaction must be consistent with the reliability requirements of the application. In the case of Oracle BPEL Process Manager sending data to Oracle Business Activity Monitoring, batching is an option that can be used and even if higher numbers are used for batching, there is no data loss in the case of Oracle Business Activity Monitoring failure, because Oracle BPEL Process Manager can retry the operations.

Redo log sync bandwidth may also be increased by striping the redo log files across multiple disks, using either a RAID or operating system based striping mechanism. Because log sync activity involves sequential writing, a large stripe depth such as 32K or 64K provides the best performance.

Note: Avoid RAID-5 usage for redo logs because it is known that redo log sync activity performs poorly when the redo logs are stored on a RAID-5 storage system. If you must use RAID-5, consider removing one or more disks from the RAID-5 configuration on which you store the redo log groups.

See Also:

- Chapter 6, "Managing the Redo Log" in the *Database Administrator's Guide* for more information on managing redo logs.
- Chapter 8, "I/O Configuration and Design" in the *Database Performance Tuning Guide* for more information on high performance I/O architectures, including the management of redo logs for high performance.

8.2 Avoiding Frequent Log Switches and Checkpoints

An Oracle database must have at least two redo log files so that when one file is full, the database can switch to another and the space in the file can be reclaimed asynchronously. Because these log switches have a significant negative impact on database performance, the redo log files should be sized to reduce the frequency of the switches to an acceptably low rate. A good rule of thumb is to have no more than 1 log switch every 20 minutes.

You can determine the rate of log switches by setting the database initialization parameter `LOG_CHECKPOINTS_TO_ALERT` to `TRUE`. This writes log switch events to the database alert log with a timestamp, so that you can monitor the alert log for an excessive frequency of log switches.

Incremental checkpoints also degrade performance. In order to limit the amount of the redo log that must be processed during database crash recovery (and thus limit the time to startup a database after a system failure), the database checkpoints the redo log at regular intervals. These checkpoint events are written to the alert log when the `LOG_CHECKPOINTS_TO_ALERT` parameter is set to `TRUE`.

Normally, the frequency of incremental checkpoints is set implicitly. When you set the `FAST_START_MTTR_TARGET` initialization parameter, this specifies the target instance startup time for crash recovery after a crash. Setting a large value for `FAST_START_MTTR_TARGET` allows for infrequent incremental checkpoints, which benefits runtime performance (this is at the expense of crash recovery time during instance startup). Setting a small value for `FAST_START_MTTR_TARGET` causes frequent incremental checkpoints that can decrease the required crash recovery time during instance startup, at the expense of runtime performance. Application requirements should dictate the analysis of this tradeoff. When runtime performance is the dominant consideration, setting the initialization parameter, `FAST_START_MTTR_TARGET` to be large enough to completely eliminate incremental checkpoints is desirable. In this case, checkpoints are made at the time of a log switch.

See Also:

- ["Tuning Redo Logs Location and Sizing"](#) on page 3-10 for more information on tuning the redo log options.
- The chapters, "Configuring a Database for Performance" and "I/O Configuration and Design" in the *Oracle Database Performance Tuning Guide*

8.3 Tuning the System Global Area

Oracle Business Activity Monitoring database performance is sensitive to the size specified for the `db_cache` and `log_buffer` in the database System Global Area (SGA), as well as the overall SGA size (the size specified should be relative to the size of the objects managed within the SGA).

Note: The optimal SGA size is dependent on the complete database workload; thus, these guidelines for SGA tuning for an Oracle Business Activity Monitoring database only apply if the database is dedicated to Oracle Business Activity Monitoring data.

Configure a dedicated Oracle Business Activity Monitoring database with a minimum SGA size of 1024MB. Set the SGA size using the `SGA_MAX_SIZE` initialization parameter. Use either SGA autotuning to set the values of memory pools within the

SGA, or use fixed sized pools. To specify autotuning, set the `SGA_TARGET` initialization parameter to the same value as `SGA_MAX_SIZE`.

You may obtain better performance by manually setting the internal SGA memory pools. Autotuning of the SGA is disabled by setting the `SGA_TARGET` initialization parameter to 0. For an installation with an SGA maximum size set to 1024MB, the settings shown in [Example 8-1](#) are a good starting point, when autotuning is disabled.

Example 8-1 Sample SGA Initialization with Autotuning Disabled

```
db_cache_size = 820M
java_pool_size = 16M
large_pool_size = 16M
streams_pool_size = 16M
shared_pool_size = 144M
log_buffer = 10485760
```

The database cache or buffer pool holds recently accessed data from the database, to avoid repeated reading of the same data pages. In the case of Oracle Business Activity Monitoring database, database insert performance depends on fast primary key constraint checking, which is accomplished by virtue of the indexes created for each primary key being held in the database cache. The buffer pool hit ratio is a good indicator of a database cache that is too small. In general, the hit ratio should be above 95%. For an Oracle Business Activity Monitoring database, you should try to maintain a hit ratio at 99% or higher; this is generally is feasible with workloads that primarily insert incoming data to the database. With workloads having substantial report querying, 95% is a good target for the buffer pool hit ratio. You can obtain the buffer pool hit ratio using the following query:

```
SELECT NAME, PHYSICAL_READS, DB_BLOCK_GETS, CONSISTENT_GETS,
       1 - (PHYSICAL_READS / (DB_BLOCK_GETS + CONSISTENT_GETS))
       "Hit Ratio" FROM V$BUFFER_POOL_STATISTICS;
```

Oracle Business Activity Monitoring is also sensitive to the log buffer size. When the database has a log file sync wait, the log buffer holds any pending requests for commit processing. When there are a large number of log file sync waits, the log buffer may become full. When a request for the log buffer is unable to be fulfilled, the requester continues to make the request after short delay, called a log buffer retry.

Log buffer retries have a significant negative impact on performance. For workloads with a data input batch size in the thousands or for any workload with a significant number of log buffer retries, the log buffer size might need to be increased.

Use the following query to monitor the buffer retries in the database:

```
SELECT NAME, VALUE FROM V$SYSSTAT
       WHERE NAME = 'redo buffer allocation retries';
```

This metric is normally measured by collecting the value at the start of a workload, and again at the end of a workload, and then the difference between the two values is the number of log buffer retries required for running the workload. If the database is restarted at the beginning of a test workload, then the beginning value is 0, and only the end value is needed. Over about a 10 minute interval, the number of log buffer retries should be under 100 (in the ideal case, the value would be 0). Values in the 100-300 range only moderately impact performance, but larger values can have a sizable impact.

It is important to note that whenever you increase the size of an SGA component, such as the log buffer or the database cache, it is also important to increase the total SGA size by the same amount (the same amount increase as the component is increased).

Otherwise, the increased component takes the memory away from other objects in the SGA.

Note: The SGA size is rounded up to the next higher multiple of 4MB at database startup.

8.4 Database Re-Organization in the Presence of Deletion Activity

When Oracle Business Activity Monitoring is installed, the Oracle Business Activity Monitoring tablespace, `ORACLEBAM`, is created to hold datasets and metadata. When a Oracle Business Activity Monitoring user creates a dataset, an Oracle Database table is created to store objects in the dataset. This table has a system-generated primary key, named `SystemID`. When a Oracle Business Activity Monitoring user creates Oracle Business Activity Monitoring indexes on columns in a Oracle Business Activity Monitoring dataset, corresponding indexes are created on the corresponding table in the database.

The Oracle Business Activity Monitoring tablespace is created with automatic segment space management enabled. This helps to manage free space in segments mapped to Oracle Business Activity Monitoring tables, and enables database re-organization to take place without taking down the system. However, it does not completely eliminate the need for manual re-organization of the database, particularly after a significant accumulation of deletion activity

After a significant amount of deletion activity, the segments mapped to a Oracle Business Activity Monitoring dataset table may become sparsely filled with data, this condition is usually called internal fragmentation and leads to a very significant increase in the amount of I/O required to perform operations on the table. When there is internal fragmentation, any indexes created on the table, including the index for the primary key, will not be structured in an optimal, balanced configuration after a significant number of deletion operations. Thus, after a significant number of deletion operations, you should perform a database re-organize to improve the structure of the indexes and better pack the rows of the table in a smaller number of data pages. You could also write a script to delete the obsolete Oracle Business Activity Monitoring data and perform the reorganization as part of the deletion job.

The procedure for reorganizing segments for a Oracle Business Activity Monitoring dataset table is to drop the primary key for the table, drop indexes on the table, enable row movement for the table, shrink space on the table, deallocate unused space in the table, re-create the primary key for the table, re-create indexes for the table, and disable row movement for the table.

[Example 8-2](#) shows a sample script that reorganizes segments for two Oracle Business Activity Monitoring datasets, `Call_Detail`, and `Call_Agg` that are stored in database tables `_Call_Detail` and `_Call_Agg`.

Note: The reorganization procedure for an Oracle Business Activity Monitoring dataset table should be scheduled at regular intervals, such as once a week or once a month, or after significant deletion activity, or with deletion included in the script as noted.

Before you perform the reorganization procedure, you should be sure to perform a backup.

Example 8–2 Sample Oracle Business Activity Monitoring Table Reorganization

```
alter table "ORABAM"."_Call_Detail" drop primary key;
alter table "ORABAM"."_Call_Agg" drop primary key;
<drop indexes on "ORABAM"."_Call_Detail">
<drop indexes on "ORABAM"."_Call_Agg">
alter table "ORABAM"."_Call_Detail" enable row movement;
alter table "ORABAM"."_Call_Agg" enable row movement;

< Note: include deletion activity here, if it is added as part of the script>

alter table "ORABAM"."_Call_Detail" shrink space;
alter table "ORABAM"."_Call_Agg" shrink space;
alter table "ORABAM"."_Call_Detail" deallocate unused;
alter table "ORABAM"."_Call_Agg" deallocate unused;
alter table "ORABAM"."_Call_Detail" disable row movement;
alter table "ORABAM"."_Call_Agg" disable row movement;
alter table "ORABAM"."_Call_Detail" add primary key ("SysIterID");
alter table "ORABAM"."_Call_Agg" add primary key ("SysIterID");

<rebuild indexes on "ORABAM"."_Call_Detail">
<rebuild indexes on "ORABAM"."_Call_Agg">
```

Note: If Oracle Business Activity Monitoring tables are configured to run with row movement enabled, then you need to change the script shown in [Example 8–2](#) to remove the disable row movement commands.

8.5 Configuring Multiple Plan Monitor Services and Enterprise Links

Oracle Business Activity Monitoring users can configure multiple plan monitor services and multiple Enterprise Links to achieve higher performance.

See Also: *Oracle Business Activity Monitoring Installation Guide*

Monitoring Using Built-in Performance Tools

This appendix includes the following sections:

- [Summary of Oracle Application Server Built-in Performance Metrics](#)
- [Viewing Performance Metrics Using AggreSpy with Basic Installation](#)
- [Viewing Performance Metrics Using AggreSpy with Web Server](#)
- [Viewing Performance Metrics Using dmstool](#)
- [Viewing Performance Metrics Using AggreSpy \(for Standalone OC4J\)](#)
- [Using Built-in Performance Metrics on Windows Systems](#)

A.1 Summary of Oracle Application Server Built-in Performance Metrics

You can monitor performance using the Application Server Control Console **Performance** secondary tab, using the System MBean Browser from the JMX area of the **Administration** secondary tab, or by viewing the Oracle Application Server built-in performance metrics.

This appendix describes how to view the built-in performance metrics using the Oracle Application Server `AggreSpy` servlet or using the `dmstool` command.

[Table A-1](#) summarizes the built-in tools that allow you to view performance metrics.

Table A-1 Oracle Application Server Built-in Monitoring Commands

Command	Description
<code>AggreSpy</code>	<p><code>AggreSpy</code> is a pre-packaged servlet that reports performance metrics for an Oracle Application Server instance. You can only run <code>AggreSpy</code> when the home OC4J instance is running. By default the OC4J instance named <code>home</code> supports <code>AggreSpy</code>.</p> <p>Note: in some cases the home instance needs to be started to use <code>AggreSpy</code>.</p>
<code>dmstool</code>	<p>Allows you to monitor a specific performance metric, a set of performance metrics, or all performance metrics. Options allow you to specify a reporting interval to report the requested metrics. This command also allows you to show a text report listing all the built-in performance metrics available on the site. The <code>dmstool</code> command is located in the directory <code>\$ORACLE_HOME/bin</code> on UNIX systems and in <code>%ORACLE_HOME%\bin</code> on Windows systems.</p>

See Also: [Appendix C, "Performance Metrics"](#)

A.2 Viewing Performance Metrics Using AggreSpy with Basic Installation

The `AggreSpy` Servlet displays metrics for Oracle Application Server processes, including: OC4J, Oracle Process Manager and Notification Server, and other Oracle Application Server component processes.

A.2.1 Using the AggreSpy Display

`AggreSpy` organizes metrics into two areas: DMS Spies and Metric Tables.

- DMS Spies show the available metrics by parent process type and parent process number. By selecting individual DMS Spies, you can view, in text form, all metrics collected for the associated process.
- Metric Tables show the available metrics by metric table type and when multiple OC4Js are running include OC4J metrics from multiple OC4J instances. By selecting individual metric tables you can view, in table form, all metrics of a specified type. For example, metric tables allow you to view the metrics associated with OC4J Servlets and Oracle Process Manager and Notification Server processes.

DMS metric tables are identified by a name, such as `oc4j_servlet` and `opmn_process`. In `AggreSpy`, the term *metric table* refers to the built-in performance metric table names.

You can access performance metrics using `AggreSpy` from the following URL:

```
http://host:port/dmsoc4j/AggreSpy
```

where:

host is the host for the OC4J that provides the HTTP listener, for example, `tv.us.oracle.com`.

port is the OC4J provided HTTP listener port, for example 8888.

Note: You can only run *AggreSpy* when the home OC4J instance is running. By default, the OC4J instance named *home* supports *AggreSpy*.

[Figure A-1](#) shows a sample *AggreSpy* display. The display shows two frames, one containing a list of DMS Spies and DMS Metric Tables, and one showing selected values for the DMS Spies or the Metric Tables.

AggreSpy provides navigation and display options, including:

- Access DMS Spies and Metric Tables using the links in the left frame.
- Sort rows in metric tables by clicking on the column headings.
- Display a table containing the descriptions of a Metric Table metrics by clicking the Metric Definitions link shown on each metric table.

You need to refresh your browser to display built-in metric data after you start *AggreSpy*. When you first use *AggreSpy* many of the fields, and the complete list of DMS Spies may not contain all of the current Metric Tables. If you wait a short time, and then refresh the display, the data is available and *AggreSpy* shows the complete list of Metric Tables.

Note: The OC4J home instance must be running to use *AggreSpy*.

In the Basic install, the home instance starts up with the command, `opmnctl startall`, or by clicking **Start** using Application Server Control Console.

Figure A-1 AggreSpy Performance Metric Display

DMS Spies

[All DMS Spies](#)

home:OC4J:8888:6100 | [Text](#)
opmn:3500:6100 | [Text](#)

Metric Tables

[All Metric Tables](#)

[JDBC_ConnectionSource](#)
[JMSConnectionStats](#)
[JMSConsumerStats](#)
[JMSDestinationStats](#)
[JMSDurableSubscriberStats](#)
[JMSPersistenceStats](#)
[JMSProducerStats](#)
[JMSSessionStats](#)
[JMSStats](#)
[JMSStoreStats](#)
[JTAResource](#)
[JVM](#)
[actThreadReq](#)
[agg_oc4j_application](#)
[agg_oc4j_application_no_ra](#)
[agg_oc4j_application_per_iv](#)
[agg_oc4j_application_per_iv](#)
[agg_oc4j_ejb](#)
[agg_oc4j_ejb_module](#)
[agg_oc4j_ejb_module_no_r](#)

Table of Contents tvanraal-pc.us.oracle.com:8888/dmsoc4j/AggreSpy: OC4J:8888:6100

- [Spies](#)
- [Tables](#)

Spies

Process	Format	SpyType	Host	Port	Path	mid	instance
opmn:3500:6100	Text	opmnlocal	tvanraal-pc	6100	/connect		
home:OC4J:8888:6100	Text	oc4j	tvanraal-pc	8888	/dmsoc4j/Spy	1379667864	10g32.tvanraal-pc.us.oracle.com

[Top](#)

Tables

Name	numRows	numColumns
JDBC_ConnectionSource	1	23
JMSConnectionStats	104	7
JMSConsumerStats	100	18
JMSDestinationStats	18	53
JMSDurableSubscriberStats	8	13
JMSPersistenceStats	2	26
JMSProducerStats	2	11
JMSSessionStats	102	8
JMSStats	1	103
JMSStoreStats	18	55
JTAResource	1	80

A.3 Viewing Performance Metrics Using AggreSpy with Web Server

The AggreSpy Servlet displays metrics for Oracle Application Server processes, including: Oracle HTTP Server, OC4J, Oracle Process Manager and Notification Server, and other Oracle Application Server component processes.

Note: This section describes viewing performance metrics using AggreSpy using Oracle HTTP Server. Depending on the type of advanced installation that you choose, Oracle HTTP Server is installed on your system. If Oracle HTTP Server is not installed on your system, then the commands in this section will not work on your system.

This section covers the following topics:

- [Using the AggreSpy Display](#)
- [AggreSpy URL and Access Control with Web Server](#)
- [AggreSpy URL and Access Control with Web Server](#)
- [AggreSpy Limitation When Using Load Balancing With Multiple Instances](#)

A.3.1 Using the AggreSpy Display with Web Server

AggreSpy organizes metrics into two areas: DMS Spies and Metric Tables.

- DMS Spies show the available metrics by parent process type and parent process number. By selecting individual DMS Spies, you can view, in text form, all metrics collected for the associated process.
- Metric Tables show the available metrics by metric table type and when multiple OC4Js are running include OC4J metrics from multiple OC4J instances. By selecting individual metric tables you can view, in table form, all metrics of a specified type. For example, metric tables allow you to view the metrics associated with OC4J Servlets, Oracle HTTP Server Modules, and Oracle Process Manager and Notification Server processes.

Note: To view DMS metrics using AggreSpy, you may need to configure your browser to disable the use of a proxy for the localhost, if your system is configured to use proxies. By default Oracle Application Server only allows access for the localhost. See "[AggreSpy URL With a Proxy Server with Web Server](#)" on page A-6 for details.

DMS metric tables are identified by a name, such as `ohs_server` for the Oracle HTTP Server metrics. In AggreSpy, the term *metric table* refers to the built-in performance metric table names.

You can access performance metrics using AggreSpy from the following URL:

`http://host:port/dms0/AggreSpy`

where:

host is the Oracle HTTP Server host, for example, `tv.us.oracle.com`.

port is the Oracle HTTP Server listener port, for example `7777`.

Note: You can only run AggreSpy when the home OC4J instance is running. By default, the OC4J instance named `home` supports AggreSpy. Using an OracleAS Infrastructure, the home instance needs to be started to use AggreSpy, since by default the home instance is installed with OracleAS Infrastructure, but it is not started.

Figure A-1 shows a sample AggreSpy display. The display shows two frames, one containing a list of DMS Spies and DMS Metric Tables, and one showing selected values for the DMS Spies or the Metric Tables.

AggreSpy provides navigation and display options, including:

- Access DMS Spies and Metric Tables using the links in the left frame.
- Sort rows in metric tables by clicking on the column headings.
- Display a table containing the descriptions of a Metric Table's metrics by clicking the Metric Definitions link shown on each metric table.

You need to refresh your browser to display built-in metric data after you start AggreSpy. When you first use AggreSpy many of the fields, and the complete list of DMS Spies may not contain all of the current Metric Tables. If you wait a short time,

and then refresh the display, the data is available and AggreSpy shows the complete list of Metric Tables.

Note: The OC4J home instance must be running to use AggreSpy. When the home instance is down, requests to AggreSpy, `http://<host>:<port>/dms0/AggreSpy`, report an HTTP 500 Internal Server error.

In the J2EE install, the home instance starts up with the command, `opmnctl startall`, or by clicking **Start** using Application Server Control Console.

Figure A-2 AggreSpy Performance Metric Display

DMS Spies

[All DMS Spies](#)

HTTP_Server:OHS:2992:6100 | [Text](#)
 home:OC4J:12501:6100 | [Text](#)
 opmn:2152:6100 | [Text](#)

Metric Tables

[All Metric Tables](#)

[JDBC_ConnectionSource](#)
[JMSConnectionStats](#)
[JMSConsumerStats](#)
[JMSDestinationStats](#)
[JMSPersistenceStats](#)
[JMSSessionStats](#)
[JMSStats](#)
[JMSSStoreStats](#)
[JTAResource](#)
[JVM](#)
[agg_oc4j_application](#)
[agg_oc4j_application_no_rate](#)
[agg_oc4j_ejb](#)
[agg_oc4j_ejb_module](#)
[agg_oc4j_ejb_module_no_rate](#)
[agg_oc4j_ejb_no_rate](#)
[agg_oc4j_instance](#)
[agg_oc4j_instance_no_rate](#)
[agg_oc4j_jca_factory](#)
[agg_oc4j_jca_factory_no_rate](#)
[agg_oc4j_jca_pool](#)
[agg_oc4j_jca_pool_no_rate](#)
[agg_oc4j_jdbc_pool](#)
[agg_oc4j_jdbc_pool_no_rate](#)

Table of Contents 127.0.0.1:7201/dmsoc4j/AggreSpy: OC4J:12501

- [Spies](#)
- [Tables](#)

Spies

Process	Format	SpyType	Host	Port	Path	uid	instance
opmn:2152:6100	Text	opmlocal	tvanraal- pc	6100	/connect		
home:OC4J:12501:6100	Text	oc4j	127.0.0.1	7201	/dmsoc4j/Spy	814550108	10gR3.tvanraa pc.us.oracle.co
HTTP_Server:OHS:2992:6100	Text	ohs	127.0.0.1	7201	/dms0/Spy	814550107	10gR3.tvanraa pc.us.oracle.co

[Top](#)

Tables

Name	numRows	numColumns
JDBC_ConnectionSource	1	23
JMSConnectionStats	3	7
JMSConsumerStats	2	18
JMSDestinationStats	9	81
JMSPersistenceStats	5	56
JMSSessionStats	2	8
JMSStats	1	126
JMSSStoreStats	8	70
JTAResource	1	77
JVM	1	20
agg_oc4j_application	0	12
agg_oc4j_application_no_rate	0	8
agg_oc4j_ejb	0	10

A.3.2 AggreSpy URL With a Proxy Server with Web Server

If your browser is configured to use a proxy server, then to access AggreSpy on the localhost, you need to configure the browser to disable the use of proxies for the localhost. The exact steps required to disable the use of a proxy server for the localhost depends on the browser you use.

A.3.3 AggreSpy URL and Access Control with Web Server

By default, the `dms0/AggreSpy` URL is redirected and the redirect location is protected, allowing only the localhost (127.0.0.1) to access the AggreSpy Servlet.

To view metrics from a system other than the localhost you need to change the DMS configuration for the system that is running the Oracle Application Server that you want to monitor by modifying the file

`$ORACLE_HOME/Apache/Apache/conf/dms.conf` on UNIX, or
`%ORACLE_HOME%\Apache\Apache\conf\dms.conf` on Windows systems.

[Example A-1](#) shows a sample default configuration from `dms.conf`. This configuration limits AggreSpy to access metrics on the localhost (127.0.0.1). The port shown, 7200, may differ on your installation.

Example A-1 Sample dms.conf File for localhost Access for DMS Metrics

```
# proxy to DMS AggreSpy
Redirect /dms0/AggreSpy http://localhost:7200/dmsoc4j/AggreSpy
#DMS VirtualHost for access and logging control
Listen 127.0.0.1:7200
OpmnHostPort http://127.0.0.1:7200
<VirtualHost 127.0.0.1:7200>
    ServerName 127.0.0.1
```

By changing the `dms.conf` configuration to specify the host that provides, or serves DMS metrics, you can allow users on systems other than the localhost to access the DMS metrics from the location `http://host:port/dms0/AggreSpy`.

Caution: Modifying `dms.conf` has security implications. Only modify this file if you understand the security implications for your site. By exposing metrics to systems other than the localhost, you allow other sites to potentially view critical Oracle Application Server internal status and runtime information.

To view metrics from a system other than the localhost (127.0.0.1), do the following:

1. Modify `dms.conf` by changing the entries with the value for localhost "127.0.0.1" shown in [Example A-1](#) to the name of the server providing the metrics (obtain the server name from the `ServerName` directive in the `httpd.conf` file, for example `tv.us.oracle.com`).
2. [Example A-2](#) shows a sample updated `dms.conf` that allows access from a system other than the localhost (127.0.0.1).

Example A-2 Sample dms.conf File for Remote Host Access for DMS Metrics

```
# proxy to DMS AggreSpy
Redirect /dms0/AggreSpy http://tv.us.oracle.com:7200/dmsoc4j/AggreSpy
#DMS VirtualHost for access and logging control
Listen tv.us.oracle.com:7200
OpmnHostPort http://tv.us.oracle.com:7200
<VirtualHost tv.us.oracle.com:7200>
    ServerName tv.us.oracle.com
```

3. Restart, or stop and start the Oracle HTTP Server using Application Server Control Console or using the `opmnctl` command. For example,

```
%opmnctl restartproc process-type=HTTP_Server
```

or

```
%opmnctl stopproc process-type=HTTP_Server  
%opmnctl startproc process-type=HTTP_Server
```

See Also: *Oracle Application Server Security Guide* for information on Oracle HTTP Server access control

A.3.4 AggreSpy Limitation When Using Load Balancing With Multiple Instances

AggreSpy does not work as expected when using Oracle Application Server with multiple instances. When the Oracle HTTP Server `mod_oc4j` component load balances OC4J requests across multiple instances, AggreSpy may report results for systems that are not the localhost (127.0.0.1).

Note: In this case it is recommended that you use `dmstool` instead of AggreSpy.

A.4 Viewing Performance Metrics Using dmstool

The `dmstool` command allows you to view a specific performance metric, a set of performance metrics, or all performance metrics for an Oracle Application Server instance. The `dmstool` command also supports an option that allows you to set a reporting interval, specified in seconds, to report updated metrics every *t* seconds.

For example, you can monitor the performance of a specific servlet, JSP, EJB, EJB method, or database connection and you can request periodic snapshots of metrics specific to these components.

The format for using `dmstool` to display built-in performance metrics is:

```
% dmstool [options] metric metric ...
```

or

```
% dmstool [options] -list
```

or

```
% dmstool [options] -dump
```

[Table A-2](#) lists the `dmstool` command-line *options*. Following [Table A-2](#) this section presents examples that show sample usage with specific performance metrics. The `dmstool` command is located in the `$ORACLE_HOME/bin` directory on UNIX or in `%ORACLE_HOME%\bin` directory on Windows.

Note: You can use `dmstool` in scripts or in combination with other monitoring tools to gather performance data, to check application performance, or to build tools that modify your system based on the values of performance metrics.

See Also:

["Using dmstool to List the Names of All Metrics"](#) on page A-11

[Appendix C, "Performance Metrics"](#) for a list and description of the DMS metrics

A.4.1 Access Control for dmstool

By default, `dmstool` shows metrics only when it is run from the localhost (127.0.0.1). If you want to view metrics from an Oracle Application Server running on a remote host, then you need to use `dmstool` with the `-a` option, on the local host, and update the `dms.conf` file of the remote Oracle Application Server instance in the `$ORACLE_HOME/Apache/Apache/conf/` directory on UNIX or `%ORACLE_HOME%\Apache\Apache\conf\` directory on Windows.

The configuration changes required to control the access to metrics using `dmstool` are the same as those for accessing `dms0/AggreSpy`.

See Also: ["AggreSpy URL and Access Control with Web Server"](#) on page A-6

Table A-2 *dmstool* Command-line Options

Option	Description
<code>-a[address] opmn://host[:port]</code>	<p>By default, without the <code>-a</code> option, <code>dmstool</code> gets metrics from the Oracle Application Server instance with the same <code>\$ORACLE_HOME</code> as <code>dmstool</code>. When <code>dmstool</code> runs in the same <code>\$ORACLE_HOME</code> as the Oracle Process Manager and Notification Server, OPMN, the <code>-a</code> option is not required.</p> <p>You can specify <code>-a</code> with the <code>opmn://</code> prefix and the arguments shown to monitor the Oracle Application Server processes under OPMN control that produce DMS metrics (some OPMN controlled processes, for example Oracle Web Cache, do not expose DMS metrics).</p> <p>Where:</p> <p><code>host</code> is the domain name or IP address of the host on which the OPMN process is running.</p> <p><code>port</code> specifies the OPMN request port that supplies metrics. The request port is specified in <code>\$ORACLE_HOME/opmn/conf/opmn.xml</code>.</p> <p>For example, the following shows the specification in <code>opmn.xml</code> for a request port (request="6003"):</p> <pre><notification-server> <port local="6100" remote="6200" request="6003"/> . .</notification-server></pre> <p>Note, if you use <code>dmstool -a</code> to request metrics from a remote system, the system must be configured to provide metrics (by default you can access DMS metrics on the localhost).</p> <p>See Also: "AggreSpy URL and Access Control with Web Server" on page A-7</p>
<code>-c[ount] num</code>	<p>Specifies the number of times to retrieve values when monitoring metrics. If not specified, <code>dmstool</code> continues retrieving metric values until the process is stopped.</p> <p>The <code>-count</code> option is not used with the <code>-list</code> option.</p>
<code>-dump [format=xml]</code>	<p>Using <code>dmstool</code> with the <code>-dump</code> option reports all the available metrics on the standard output. Oracle recommends that you run with the <code>-dump</code> option periodically, such as every 15 to 20 minutes, to capture and save a record of performance data for your Oracle Application Server server.</p> <p>The <code>-dump</code> option also supports the <code>format=xml</code> query. Using this query at the end of the command line supplies the metric output in XML format.</p>
<code>-help</code>	List the <code>dmstool</code> command-line options.
<code>-i[nterval] secs</code>	<p>Specifies the number of seconds to wait between metric retrievals. The default is 5 seconds. The <code>interval</code> argument is not used with the <code>-list</code> option. The interval specified is approximate.</p> <p>Note: if the system load is high, the actual interval may vary from the interval specified using the <code>-interval</code> option.</p>

Table A-2 (Cont.) dmstool Command-line Options

Option	Description
<code>-l[ist] [-table]</code>	Generates a list of all metrics available. Use the <code>-list</code> option with the <code>-table</code> option to display a list of all the metric table names. Note, including metric names on the command-line is not valid when using the <code>-list</code> option with <code>dmstool</code> .
<code>-reset [-table <i>metric_table</i>]</code>	Resets the specified metrics or with the <code>-table</code> option, all of the metrics contained in the specified metric table. Event and phaseEvent metrics are reset to 0 (as if they were never updated). State metrics are reset to the current value (as if they started with the current value). Note: The <code>reset</code> option may reset information that Application Server Control Console uses to compute and report values.
<code>-table <i>metric_table</i></code>	Includes all the performance metrics for the specified metric table with the name, <i>metric_table</i> . See Appendix C, "Performance Metrics" or run <code>AggreSpy</code> for a list of metric table names.

A.4.2 Using dmstool to List the Names of All Metrics

Every Oracle Application Server performance metric has a unique name. Using `dmstool` with the `-list` option produces a list of all metric names. The `-list` output contains the metric names that you can use with `dmstool` to request monitoring information for a specific metric or set of metrics.

Using the following command, `dmstool` displays a list of all metrics available on the server:

```
% dmstool -list
```

This command displays a list of the available metrics.

See Also: [Appendix C, "Performance Metrics"](#)

A.4.3 Using dmstool to Report Values for Specific Performance Metrics

To monitor a specific metric or set of metrics, use `dmstool` and include the metric name on the command-line. For example, to monitor the time the JVM has been running, perform the following steps:

1. Use `dmstool` with the `-list` option to find the name of the metric that shows the JVM uptime:

```
% dmstool -list | grep JVM/upTime.value
/system1/OC4J:12502:6100/JVM/upTime.value
```

2. Use `dmstool` and supply the full metric name as an argument to show the metric value:

```
% dmstool /system1/OC4J:12502:6100/JVM/upTime.value
Wed Dec 21 15:37:08 PST 2005
/system1/OC4J:12502:6100/JVM/upTime.value 159312 msec
```

Using `dmstool`, the default repeat interval is 5 seconds, so this command shows the updated metric every 5 seconds. Use the `-count` option to limit the number of times `dmstool` reports values.

For example:

```
% dmstool /system1/OC4J:12502:6100/JVM/upTime.value -count 2
Wed Dec 21 15:39:38 PST 2005
/system1/OC4J:12502:6100/JVM/upTime.value 310031 msec
Wed Dec 21 15:39:43 PST 2005
/system1/OC4J:12502:6100/JVM/upTime.value 314516 msec
```

Note: In some cases, the full path of a metric name may contain a space. If the path contains a space, the space must be quoted on the `dmstool` command line, so that the shell sends the metric name to `dmstool` as a single argument.

A.4.4 Using dmstool With the Interval and Count Options

To monitor the requests completed for an application over an interval of one minute, use the following `dmstool` command, supplying metric names on the command-line:

```
% dmstool -i 60 -c 120 \
/system1/OC4J:3301:6003/oc4j/default/WEBs/processRequest.completed
```

This command reports 120 sets of output for the metric listed on the command line, while collecting data at intervals of 60 seconds:

```
Tue Oct 12 14:43:43 PDT 2004
/system1/OC4J:3301:6003/oc4j/default/WEBs/processRequest.completed 8576 ops

Tue Oct 12 14:44:43 PDT 2004
/system1/OC4J:3301:6003/oc4j/default/WEBs/processRequest.completed 8581 ops

Tue Oct 12 14:45:43 PDT 2004
/system1/OC4J:3301:6003/oc4j/default/WEBs/processRequest.completed 8588 ops
.
.
.
```

A.4.5 Using dmstool to Report All Metrics with Metric Values

Using `dmstool` with the `-dump` option displays all the metrics from an Oracle Application Server instance to the standard output.

The following command displays all available metrics:

```
% dmstool -dump
```

Oracle recommends that you run `dmstool` with the `-dump` option periodically, such as every 15 to 20 minutes, to capture and save a record of performance data. If you save performance data over time, this data can assist you if you need to analyze system behavior to improve performance or when problems occur.

A.4.6 Using dmstool to Report All Metrics with Metric Values in XML Format

When you need to process metric data, use the `format=xml` query on the `dmstool` command line to report all metric values in XML format.

The following command displays all available metrics using XML format:

```
% dmstool -dump format=xml
```

A.4.7 Using dmstool to Reset Metric Values

When you want to reset metric values, use the `reset` option on the `dmstool` command line to reset values for a set of metrics, or for all metrics in a specified metric table.

Using the `reset` option, Event and phaseEvent metrics are reset to 0, as if they were never updated, and State metrics are reset to the current value (as if they started with the current value).

The following command resets the specified metric:

```
% dmstool -reset /system1/OC4J:3000:6004/JVM/upTime.value
```

The following command resets the specified metric table:

```
% dmstool -reset /system1/OC4J:3000:6004/JVM/upTime.value
```

Note: The `reset` option may reset information that Application Server Control Console uses to compute and report values.

A.4.8 Using dmstool to View Metrics on a Remote Oracle Application Server System

Using `dmstool` with the `-a` option reports the metrics from a remote Oracle Application Server instance.

Note: By default the Oracle Application Server only allows `dmstool` to access metrics from the localhost. You need to modify `dms.conf` to support access from systems other than the localhost. See "[AggreSpy URL and Access Control with Web Server](#)" on page A-7 for information on DMS access control.

The following command displays all available metrics and metric values on the Oracle Application Server Instance, as specified with the `-a` option:

```
% dmstool -a opmn://system1:6003 -list
```

Using the `dmstool -a` option, supply an argument with the prefix `opmn://` and include the host name where you want to obtain metrics, and the OPMN request port number. The port specifies the OPMN request port that supplies metrics for Oracle Application Server which is specified the `request` attribute under the `<notification-server>` element in `$ORACLE_HOME/opmn/conf/opmn.xml` on UNIX and `%ORACLE_HOME%\opmn\conf\opmn.xml` on Windows.

See Also: "[AggreSpy URL and Access Control with Web Server](#)" on page A-7

A.5 Viewing Performance Metrics Using AggreSpy (for Standalone OC4J)

When you are using OC4J in standalone mode, without the Oracle Application Server, the AggreSpy Servlet allows you to access OC4J metrics.

When running OC4J standalone, access performance metrics using `AggreSpy` from the following URL:

```
http://myhost:myport/dms0/AggreSpy
```

Note: You can only run `AggreSpy` when OC4J is configured to support it, and OC4J is running. By default, OC4J supports `AggreSpy`.

[Table A-3](#) covers the `dmstool` option that only applies to OC4J standalone mode. In addition, the options shown in [Table A-2](#) also apply to `dmstool` (except the `-a` option with the `opmn://` prefix).

Table A-3 *dmstool Command-line Options (for Standalone OC4J only)*

Option	Description
<code>-a[address]</code> <code>host[:port][path],...</code>	For a standalone OC4J system, use the <code>-a</code> option. This specifies the <code>http://</code> protocol, where: <i>host</i> is the domain name or IP address of the host on which the Oracle HTTP Server is running and <i>port</i> specifies the associated port.

A.6 Using Built-in Performance Metrics on Windows Systems

Using Oracle Application Server on Windows systems, statistics collection needs to be enabled to view certain DMS metrics. If some DMS metrics report the value "0" for values that you expect to be other than 0, then statistics collection may be disabled on your system. To enable statistics collection on Windows systems where statistics collection is disabled, set the value of the following registry entry to 0.

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PerfProc\Performance\Disable  
Performance Counters
```

Note: Incorrectly editing the registry may severely damage your system. At the very least, you should back up any valued data on the computer before making changes to the registry.

Instrumenting Applications With DMS

The Oracle Dynamic Monitoring Service (DMS) enables application developers, support analysts, system administrators, and others to measure application specific performance information. This chapter describes DMS and shows a sample application that demonstrates how to instrument Oracle Application Server Java applications using DMS.

Note: Oracle Application Server provides a number of built-in metrics. Using DMS to instrument applications adds new metrics to the set of built-in metrics.

This chapter includes the following sections:

- [Introducing DMS Performance Metrics](#)
- [Adding DMS Instrumentation To Java Applications](#)
- [Validating and Testing Applications Using DMS Metrics](#)
- [Understanding DMS Security Considerations](#)
- [Conditional Instrumentation Using DMS Sensor Weight](#)
- [Dumping DMS Metrics To Files](#)
- [Resetting and Destroying Sensors](#)
- [DMS Coding Recommendations](#)
- [Using A High Resolution Clock To Increase DMS Precision](#)
- [Rolling Up DMS Data for Descendent Nouns](#)

See Also: [Appendix C, "Performance Metrics"](#)

B.1 Introducing DMS Performance Metrics

The Dynamic Monitoring Service (DMS) API allows you to add performance instrumentation to Oracle Application Server applications. At runtime OC4J collects performance information, called DMS metrics, that developers, system administrators, and support analysts use to help analyze system performance or monitor system status.

This section includes the following:

- [Instrumenting Applications With DMS Metrics](#)
- [Monitoring DMS Metrics](#)
- [Understanding DMS Terminology \(Nouns and Sensors\)](#)
- [DMS Naming Conventions](#)

Note: Oracle Application Server components, including OC4J, provide a number of predefined metrics. For a listing of the predefined metrics see [Appendix C, "Performance Metrics"](#).

B.1.1 Instrumenting Applications With DMS Metrics

DMS Instrumentation refers to the process you use when you insert DMS calls into application code. By using the DMS API you can enable your application to measure, collect, and save performance information.

To create DMS metrics you add DMS API calls that notify DMS when events occur, when important intervals begin and end, or when pre-computed values change their state. At runtime, DMS stores metrics in memory and allows you to save or view the metrics.

Oracle Application Server includes built-in DMS metrics. By adding DMS calls to your applications you can expand the set of built-in metrics. When you instrument your applications with DMS calls, you use the same API that the built-in metrics use. In addition, to save and display your metrics, you use the same monitoring tools that you use with built-in metrics.

Tip: ["Adding DMS Instrumentation To Java Applications"](#) on page B-8

B.1.2 Monitoring DMS Metrics

Monitoring DMS metrics refers to the process of retrieving performance metrics. When an application runs, DMS stores metrics in memory and allows you to show metrics on the console or to view metrics using a web browser.

Oracle Application Server provides several runtime tools for viewing and saving DMS metrics, including `dmstool` and the `AggreSpy` Servlet.

[Example B-1](#) shows a set of metrics output using `dmstool`.

Example B-1 Set of Sample `dmsDemo` Metrics Using `dmstool`

```
/dmsDemo [type=n/a]
/dmsDemo/BasicBinomial [type=MathSeries]
computeSeries.active:      0      threads
computeSeries.avg:    21.181818181818183      msecs
computeSeries.completed:  11      ops
computeSeries.maxActive:  1      threads
```

```

computeSeries.maxTime:    93      msec
computeSeries.minTime:   0       msec
computeSeries.time:     233     msec
lastComputed.value:     184756
loops.count:            604     ops
    
```

See Also: [Appendix A, "Monitoring Using Built-in Performance Tools"](#)


B.1.3 Understanding DMS Terminology (Nouns and Sensors)

This section introduces the terminology you need to understand to use DMS. [Figure B-1](#) illustrates the organization of a set of DMS metrics corresponding to the metrics in the demo application described in this chapter and the metrics shown in [Example B-1](#).

This section covers the following topics:

- [DMS Metrics](#)
- [DMS Sensors](#)
- [DMS Nouns](#)
- [DMS Rollup Nouns](#)
- [DMS Object Relationships](#)

Figure B-1 Organization of Sample Metrics From *dmsDemo* Application



B.1.3.1 DMS Metrics

DMS Metrics track performance information that developers, system administrators, and support analysts use to help analyze system performance or monitor system status.

B.1.3.2 DMS Sensors

DMS Sensors measure performance data and allow DMS to define and collect a set of metrics. Certain metrics are always included with a Sensor and other metrics are optionally included with a Sensor.

B.1.3.2.1 DMS PhaseEvent Sensors A **DMS PhaseEvent Sensor** measures the time spent in a specific section of code that has a beginning and an end. Use a PhaseEvent Sensor to track time in a method or in a block of code.

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

Table B–1 describes metrics available with a PhaseEvent Sensor.

Table B–1 DMS PhaseEvent Sensor Metrics

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

B.1.3.2.2 DMS Event Sensors A **DMS Event Sensor** is a Sensor that counts system events. Use a DMS Event Sensor to track system events that have a short duration, or where the duration of the event is not of interest but the occurrence of the event is of interest.

Table B–2 describes the metric that is associated with an Event Sensor.

Table B–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, where <code>sensor_name</code> is the name of the Event Sensor as specified in the DMS instrumentation API. Default: count is the default metric for an Event Sensor. No other metrics are available for an Event Sensor.

B.1.3.2.3 DMS State Sensors A **DMS State Sensor** is a Sensor to which you assign a precomputed value. State Sensors track the value of Java primitives or the content of a Java Object. The supported types include integer, double, long, and object. Use a State Sensor when you want to track system status information or when you need a performance metric that is not associated with an event. For example, use State Sensors to represent queue lengths, pool sizes, buffer sizes, or host names.

[Table B–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 B–3 DMS State Sensor Metrics

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

B.1.3.3 DMS Nouns

DMS Nouns (Nouns) organize performance data. Each Sensor, with its associated metrics is organized in a hierarchy according to Nouns. Nouns allow you to organize DMS metrics in a manner comparable to a directory structure in a file system. For example, Nouns can represent classes, methods, objects, queues, connections, applications, databases, or other objects that you want to measure.

A `Noun type` is a name that reflects the set of metrics being collected. For example, in the built-in metrics the `Noun type oc4j_servlet` represents the metrics collected for each servlet in each Web module within each J2EE application. And the `Noun type JVM` represents the set of metrics for each Java process (OC4J) currently running in the site.

Note: In [Appendix C, "Performance Metrics"](#), the `Noun type` is called the metric table name.

The `Noun` naming scheme uses a '/' as the root of the hierarchy, with each `Noun` acting as a container under the root, or under its parent `Noun`.

See Also: [Appendix C, "Performance Metrics"](#)

B.1.3.4 DMS Rollup Nouns

DMS Rollup Nouns are nouns that DMS generates when you include instrumentation to request a set of aggregate nouns. The rollup noun contains metrics from a set of Sensors in the descendent nouns of a specified noun type. A rollup noun also contains summary information.

See Also: ["Rolling Up DMS Data for Descendent Nouns"](#) on page B-21

B.1.3.5 DMS Object Relationships

This section describes the object relationships and attributes for DMS metrics, Sensors, and Nouns.

[Table B-4](#) describes the relationships between DMS objects. [Figure B-1](#) illustrates the relationships shown in [Table B-4](#) using a sample set of metrics.

Table B-4 DMS Object Relationships and Attributes

Object	Contains	Attributes
Noun	Sensors or other Nouns	Name, Noun Type, Parent
Sensor	Metrics	Name, Description, Sensor Type, Parent There are three Sensor Types: PhaseEvent, Event, and State.
Metric	Value	Name, Units designation

B.1.4 DMS Naming Conventions

Certain guidelines apply for defining DMS names. By following these guidelines, people viewing DMS metric reports can easily understand metrics across applications and across Oracle Application Server components.

Note: View the naming conventions as guidelines; for each convention there may be an exception. Try to be as clear as possible, if there is a conflict, you may need to make an exception.

This section covers the following topics:

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

B.1.4.1 General DMS Naming

DMS metric names consist of a Sensor name plus the "." character plus the metric. For example, the names: `computeSeries.time`, `loops.count`, and `lastComputed.value` are valid DMS metric names.

A Sensor name is a simple string, not including the "." or the derivation. For example `computeSeries`, `loops`, and `lastComputed` are Sensor names. A Sensor full name consists of the Sensor name, preceded by the name of its associated Noun, and a delimiter. For example, `/dmsDemo/BasicBinomial/computeSeries`,

`/dmsDemo/BasicBinomial/loops`, and
`/dmsDemo/BasicBinomial/lastComputed`.

A Noun name is a simple string, not including a delimiter. For example `BasicBinomial` is a Noun name. A Noun full name consists of the Noun name, preceded by the full name of its parent, and a delimiter. For example `/dmsDemo/BasicBinomial` is a full Noun name.

B.1.4.2 General DMS Naming Conventions and Character Sets

DMS names should be as compact as possible. Whenever possible, when you define Noun and Sensor names, avoid special characters such as white space, slashes, periods, parenthesis, commas, and control characters.

Table B-5 shows DMS replacement for special characters in names.

Table B-5 DMS Naming Special Character Replacement

Character	DMS Replacement Character
Space " " or Period "."	Underscore "_"
Control Character	Underscore "_"
"<"	"("
">"	")"
"&"	"^"
"" (double quote)	"" (backquote) That is, a backquote replaces a double quote.
" (single quote)	" (backquote). That is, a backquote replaces a single quote.

Note: Oracle Application Server includes a number of built-in metrics. The Oracle Application Server built-in metrics do not always follow the DMS naming conventions.

B.1.4.3 Noun and Noun Type Naming Conventions

A Noun name should be a name which identifies a specific entity of interest.

Noun types should have names which clearly reflect the set of metrics being collected. For example, `Servlet` is the type for a Noun under which the metrics that are specific to a given servlet fall.

Noun type names should start with a capital letter to distinguish them from other DMS names. All Nouns of a given type should contain the same set of sensors.

B.1.4.4 Sensor Naming Conventions

The following list outlines DMS Sensor naming conventions.

1. Sensor names should be descriptive, but not redundant. Sensor names should not contain any part of the Noun name hierarchy, or type, as this is redundant.
2. Sensor names should avoid containing the specification of the units for the individual metrics.

3. Where multiple words are required to describe a Sensor, the first word should start with a lowercase letter, and the following words should start with uppercase letters. For example `computeSeries`.
4. In general, using a "/" in a Sensor name should be avoided. 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, you need to use an alternative delimiter, such as "," or "_", which does not appear anywhere in the path; this allows the "/" to be properly understood as part of the Noun or Sensor name rather than as a delimiter.

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

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

and you can look this up using the string:

```
,oc4j,default,WEBs,defaultWebApp,JSPs,example/jsp/num/numguess.jsp,service
```

Where the delimiter is the "," character.

5. Event Sensor and PhaseEvent Sensor names should have the form *verbNoun* where *verb* and *Noun* are interpreted as defined by English grammar. For example, `activateInstance` and `runMethod`. When a PhaseEvent monitors a function, method, or code block, it should be named to reflect the task performed as clearly as possible.
6. The name of a State Sensor should be a Noun, possibly preceded by an adjective, which describes the semantics of the value which is tracked with this State. For example, `lastComputed`, `totalMemory`, `port`, `availableThreads`, `activeInstances`.
7. To avoid confusion, do not name Sensors with strings such as: ".time", ".value", or ".avg", that are the same as the default metrics or optional derivations for a Sensor, as shown in [Table B-1](#), [Table B-2](#), and [Table B-3](#).

B.2 Adding DMS Instrumentation To Java Applications

You can collect performance information in Java applications by adding DMS instrumentation to existing applications or by creating new applications that include DMS instrumentation.

The DMS samples shown in this chapter are supplied on the Oracle Technology Network Web site

<http://www.oracle.com/technology/tech/java/oc4j/demos/index.html>

The DMS `demo.zip` file includes a ready to deploy `.ear` file and source code with build instructions. The demo includes two servlets, `BasicBinomial.java` and `ImprovedBinomial.java`.

The `BasicBinomial` servlet shows how to use the DMS API to add DMS Sensors.

The `ImprovedBinomial` servlet expands on the `BasicBinomial` and illustrates measuring the improved code, as compared with the `BasicBinomial`.

`ImprovedBinomial` servlet also shows how to add more costly metrics that gather more detailed information.

Refer to the sample code for full details on the examples in this chapter.

To use DMS instrumentation, add DMS calls by performing the following steps:

- [Including DMS Imports](#)

- [Organizing Performance Data](#)
- [Defining and Using Metrics for Timing](#)
- [Defining and Using Metrics for Counting](#)
- [Defining and Using Metrics for Recording Status Information \(State Sensors\)](#)

B.2.1 Including DMS Imports

To use DMS you need to add DMS imports. The following example shows the imports that the sample application `BasicBinomial.java` requires.

```
import oracle.dms.instrument.DMSConsole;
import oracle.dms.instrument.Event;
import oracle.dms.instrument.Noun;
import oracle.dms.instrument.PhaseEvent;
import oracle.dms.instrument.State;
import oracle.dms.instrument.Sensor;
```

B.2.2 Organizing Performance Data

Define DMS Nouns to organize Sensors and their associated metrics. DMS Nouns organize Sensors in a tree hierarchy in a manner comparable to a directory structure in a file system, starting with a root at the top of the tree.

[Example B-2](#) shows a section of code using `Noun.create()` from the `BasicBinomial.java`.

In [Example B-2](#), `MathSeries` specifies the Noun type. The Noun type is a name that reflects the set of metrics being collected. For example, `MathSeries` represents the metrics collected for the sample application containing a Binomial series computation. `AggreSpy` displays Sensors using the same Noun type together.

It is good practice to only use Noun types for Nouns that directly contain Sensors. When a Noun contains only Nouns, as in the Noun `dmsDemo`, and does not directly contain Sensors, `AggreSpy` displays the Noun type as a metric table, with no metrics. [Example B-2](#) shows the `dmsDemo` Noun that includes a Noun, `BasicBinomial`, but no Sensors. When the Noun type is not included for such a Noun, `AggreSpy` does not display a metric table associated with the Noun.

Note: Start Noun type names with a capital letter to distinguish them from other DMS names.

Example B-2 Using `Noun.create` To Organize Sensors

```
private Noun binRoot;           // Container for Binomial series DMS metrics.
Noun base = Noun.create("/dmsDemo");
binRoot = Noun.create(base, "BasicBinomial", "MathSeries");
```

See Also: ["DMS Naming Conventions"](#) on page B-6

B.2.2.1 Choosing Noun Types

In general, nouns should not be of the same noun type as any of their ancestor or descendent nouns. Usually, this is easy to code, and provides a logical hierarchy for nouns of the same type at the same level. For example, in the `dmsDemo` application, there is a second servlet, `ImprovedBinomial`, and there is the `BasicBinomial` servlet. In this case, the instrumentation uses the noun of type `MathSeries` for both.

This noun is created under `/dmsDemo` in the same hierarchy level for both servlets. Adhering to this practice makes the generated metric tables easier to understand. It also prevents some minimal information loss in the reporting process.

B.2.3 Defining and Using Metrics for Timing

To create metrics that measure the duration of a segment of code, define and use a `PhaseEvent` Sensor using the following steps:

- [Defining `PhaseEvent` Sensors](#)
- [Using `PhaseEvent` Sensors](#)

B.2.3.1 Defining `PhaseEvent` Sensors

[Example B-3](#) shows the DMS calls that declare and create the `computeSeries` `PhaseEvent` Sensor. This code defines a DMS metric named `/dmsDemo/BasicBinomial/computeSeries.time`.

`PhaseEvent` Sensors support a set of optional metrics, along with the default metric `.time` (representing the time, as measured between the `PhaseEvent start()` and the `PhaseEvent stop()` calls). You can derive optional metrics with `PhaseEvent` Sensors individually or as a complete set. [Table B-1](#) shows the available metrics for a `PhaseEvent` Sensor. The `binComp.deriveMetric(Sensor.all)` call in [Example B-3](#) causes all the supported optional metrics to be computed and reported.

Note: Using the method `deriveMetric(Sensor.all)` is recommended for adding optional metrics. Using this method with `Sensor.all` adds all metrics; this is good practice since the list of optional metrics could change in a future Oracle Application Server release. In addition, the metrics are efficient to compute and are often useful in evaluating performance.

Example B-3 Defining `PhaseEvent` Sensors

```
private PhaseEvent binComp; // Time to compute Binomial series.
.
.
.
binComp = PhaseEvent.create(binRoot, "computeSeries",
    "Time to compute a Binomial series");
binComp.deriveMetric(Sensor.all);
```

B.2.3.2 Using `PhaseEvent` Sensors

To use a `PhaseEvent` Sensor, an application calls the `start()` method to indicate the beginning of a phase and subsequently calls the `stop()` method to indicate the completion of the phase.

[Example B-4](#) shows a code segment from `BasicBinomial.java` that uses the `start()` and `stop()` methods for the `/dmsDemo/BasicBinomial/computeSeries.time` metric. The long value named `token` that is returned from the `PhaseEvent start()` method must be passed to the corresponding `PhaseEvent stop()` method. This value is a timestamp representing the start time. Passing this value to the `stop()` method allows DMS to compute the `PhaseEvent` duration.

Note: To assure that `PhaseEvents` are stopped, each `PhaseEvent` `start()` method, together with the code to be measured should be in a `try` block with the `PhaseEvent` `stop()` method in a corresponding `finally` block, as shown in [Example B-4](#).

Example B-4 Using `start()` and `stop()` With `PhaseEvent` Sensors

```
long token = 0; // DMS
try {
    token = binComp.start(); // DMS
    BigInteger bins[] = bin(length);
    out.println("<H2>Binomial series for " + length + "</H2>");
    for (int i = 0; i < length; i++)
        out.println("<br>" + bins[i]);
}
finally {
    binComp.stop(token); // DMS
    out.close();
}
```

[Example B-4](#) shows code instrumented such that each time a phase starts, it is stopped (since the `stop` method is placed in the `finally` clause). This prevents runaway `Phase` Sensors; however, this can result in the time required to throw an exception possibly contributing to phase statistics. To prevent exception handling from impacting a `PhaseEvent`, use the `abort()` method, as shown in [Example B-5](#).

[Example B-5](#) shows a code sample where a `Phase` that is not successfully stopped will be aborted. The `abort` call removes the statistics corresponding to the corresponding start, and these statistics do not contribute to metric calculations.

Example B-5 Using `abort()` with `PhaseEvent` Sensors

```
PhaseEvent pe = heavyPhase(param);
long token1 = 0;
long token2 = 0;
boolean stopped = false;
try {
    token1 = binComp.start();
    if (pe != null) token2 = pe.start();
    BigInteger bins[] = bin(length);
    out.println("<H2>ImprovedBinomial series for " + length + "</H2>");
    for (int i = 0; i < length; i++)
        out.println("<br>" + bins[i]);
    if (pe != null) pe.stop(token2);
    binComp.stop(token1);
    stopped = true;
}
finally {
    if (!stopped) {
        if (pe != null) pe.abort(token2);
        binComp.abort(token1);
    }
}
```

B.2.4 Defining and Using Metrics for Counting

To create metrics that count the occurrences of an event, define and use an `Event` Sensor as follows:

- [Defining Event Sensors](#)

- [Using Event Sensors](#)

B.2.4.1 Defining Event Sensors

[Example B-6](#) shows the DMS calls that define an Event Sensor. This code allocates a counter and defines a DMS metric named `/dmsDemo/BasicBinomial/loops.count`.

Example B-6 Defining Event Sensors

```
private Event binLoop; // Loops needed for Binomial series.
.
.
.

binLoop = Event.create(binRoot, "loops", "Iterations to compute series");
```

B.2.4.2 Using Event Sensors

DMS increments a counter when an application calls the `occurred()` method for an Event Sensor. [Example B-7](#) shows the `occurred()` call for an Event Sensor that increments the `/dmsDemo/BasicBinomial/loops.count` metric.

Example B-7 Using occurred() With Event Sensors

```
binLoop.occurred();
```

B.2.5 Defining and Using Metrics for Recording Status Information (State Sensors)

DMS captures status information with State Sensors. State Sensors track the value of Java primitives or the content of a Java Object. The supported types include integer, double, long, and object, as specified in the third argument to the `create()` method. When a Java primitive State Sensor is updated with the wrong type, DMS attempts to convert the supplied value to the correct type. For Object type State Sensors, DMS stores a reference to the Object and by default and calls `toString()` on the object when the DMS value is sampled.

To create metrics that record status information, define and use a State Sensor as follows:

- [Defining State Sensors](#)
- [Using State Sensors](#)

B.2.5.1 Defining State Sensors

State Sensors support a default metric value, as well as optional metrics. You can define the `minValue` and `maxValue` optional metrics with State Sensors only if the State Sensor represents a numeric Java primitive (of type integer, double, or long). [Table B-3](#) shows the available metrics for a State Sensor. [Example B-3](#) shows how to enable optional metrics.

[Example B-8](#) shows the DMS calls that declare and create a State Sensor. This code defines a DMS metric named `/dmsDemo/BasicBinomial/lastComputed.value`.

Example B-8 Defining State Sensors

```
private State binLast; // Value of the last computed element in series.
.
```



```

.
.
binLast = State.create(binRoot, "lastComputed", State.OBJECT, "",
    "Value of last computed series element");

```

When you define a State Sensor, use an empty string in the fourth argument to the `create()` method if no units are associated with the State Sensor, otherwise use a string listing the appropriate units (see [Example B-8](#)). State Sensors are created without an initial value. If you need to check whether a State Sensor has been initialized, use the `isInitialized()` method.

If you want your State Sensor to store the string value of an object, and not store a reference to the object, use the `setCopy()` method with the value `TRUE`. This tells the State Sensor to store the result of calling `toString()` on an object rather than using a reference to the object for the metric value.

B.2.5.2 Using State Sensors

When an application calls a State Sensor's `update()` method, DMS updates the value of the State Sensor. [Example B-9](#) shows the `update()` call for a State Sensor that updates the `/dmsDemo/BasicBinomial/lastComputed.value` metric.

Example B-9 Using update() With State Sensors

```
binLast.update(bins[k-1].toString());
```

B.3 Validating and Testing Applications Using DMS Metrics

You should test and verify the accuracy of the metrics that you add to Java applications.

This section includes the following:

- [Validating DMS Metrics](#)
- [Testing DMS Metrics For Efficiency](#)

B.3.1 Validating DMS Metrics

Use the `dmstool` and the other available DMS monitoring tools to verify and test new metrics.

Try to validate the following for new metrics:

- Do expected metrics appear in the display? Test this by examining the code to make sure that all the metric names added using DMS instrumentation appear in your display or saved set of metrics.
- Do unexpected metrics appear in the display? Verify that you have only added the metrics that you planned to add.
- Are the metric values you see within reasonable ranges? Usually, upper and lower bounds for metrics can be established. You then test that the reported values for metrics do not exceed the expected bounds.

For example, a "size of pool" metric should never report a negative value.

- Make sure that new metrics are needed. For example, if you add a `PhaseEvent` that always measures an event of very short duration, consider changing the metric to an `Event` metric, or remove the metric.

- Make sure that new metrics are accurate. For most applications using DMS metrics, accuracy is more important than the performance cost of adding the DMS instrumentation. New DMS metrics should provide reliable and useful information.

Testing for accuracy can be difficult; however, if an alternate means of measuring a particular metric is available then use it to verify metric values. For example, if you submit a known number of requests to a server and measure total time for the experiment, then you predict correct values for the relevant metrics and compare them with the actual monitored values. As another example, you can verify an `Event Sensor count` metric by examining records that you write to a log file or to the console.

Check for timing inaccuracies that may apply for the metrics. Timing inaccuracies may be caused when low-resolution clocks time metrics for an interval of short duration. For example on Windows systems, the default Java clock advances only once every 15 milliseconds. DMS metrics reported for brief events on these systems must be analyzed with care. Consider using the high resolution clock to address this issue.

See Also: ["Using A High Resolution Clock To Increase DMS Precision"](#) on page B-17

B.3.2 Testing DMS Metrics For Efficiency

The use of DMS metrics has some influence on application performance. When adding metrics, note the following:

- The processing required for computing and storing metrics can slow down the execution of an application. DMS is fast, but it does have some required overhead cost. In addition, DMS cannot prevent developers from using the DMS API inefficiently. Therefore, before adding DMS instrumentation, establish reasonable expectations. After completing the implementation, measure the actual costs and compare them to your expectations. Be prepared to make changes to the instrumentation to reduce overhead costs until the measurements agree with expectations.
- DMS provides the `DMSConsole.getSensorWeight()` method to help you control the use of metrics. The central setting is an advisory measurement level that DMS does not enforce. To control which metrics to include, at runtime, the code must test the value for `SensorWeight` to determine whether to make DMS calls.
- When integrating DMS instrumentation with an existing package or when implementing a new feature, you should consider insulating a previously working system. For example, you could include an option to enable and disable new DMS metrics.
- Worrying about performance too soon often leads to costly design and implementation errors. According to Donald Knuth, "Premature optimization is the root of all evil".
- You should run your performance tests with and without DMS enabled. If your tests show unacceptable results with DMS enabled, then you may want to re-design or re-implement metrics.

B.4 Understanding DMS Security Considerations

DMS metrics do not support user based access to DMS reports. When you define and use a DMS metric, the metric is available to any administrator that has access to DMS metrics. This means when you add DMS metrics, it is good practice to avoid placing customer sensitive information in the metrics.

When you add DMS instrumentation, the following users have access to the DMS metrics that you create:

- Applications running in the same OC4J instance can access the DMS metrics.
- All users that have access to the `dmstool` command, or the `AggreSpy` Servlet have access to the metrics (by default this is limited to Administrators).

See Also:

- ["AggreSpy URL and Access Control with Web Server"](#) on page A-7
- ["Access Control for dmstool"](#) on page A-9

B.5 Conditional Instrumentation Using DMS Sensor Weight

Use the DMS Sensor weight feature to conditionally limit your instrumentation. With Sensor weight, you specify that applications execute expensive instrumentation only when the Sensor weight is set to a particular value. Using this feature enables you to include expensive metrics that you may only need for debugging.

[Example B-10](#) shows how to use `DMSConsole.getSensorWeight()` to test the value of the Sensor weight, and optionally define and use a metric.

The Sensor weight is set globally using the `oracle.dms.sensors` property on the command-line. Set this property using the OC4J startup options. Supported values for this property include: `none`, `normal`, `heavy`, and `all`.

Example B-10 Using SensorWeight for Conditional Instrumentation

```

/* DMS Method
 *
 * If the SensorWeight is high enough, return a phase with the
 * parameter in the name. Otherwise, return null.
 */
PhaseEvent heavyPhase(String param) {
    PhaseEvent pe = null;
    if (DMSConsole.getSensorWeight() > DMSConsole.NORMAL) {
        Noun base = Noun.create(binRoot, param, "MathSeries");
        pe = PhaseEvent.create(base, "computeSeries",
                               "Time to compute a Binomial series");
        pe.deriveMetric(Sensor.all);
    }
    return pe;
}

```

B.6 Dumping DMS Metrics To Files

In a Java application, use the following method to dump DMS metrics to a file.

The following code allows you to append or replace the contents of the specified file with the current metrics:

```
DMSConsole cons2 = new DMSConsole();
DMSConsole.dump("dmsmathseries.log", true, true);
```

The first argument specifies the file path name, the second argument specifies the output format, and the third argument specifies if the output is appended to the file or replaces the contents of the file.

B.7 Resetting and Destroying Sensors

The Sensor abstract class provides methods to control PhaseEvent, Event, and State Sensors. The `reset()` method resets a Sensor's metrics to initial values. The `getResetTime()` method determines if a Sensor has been reset. The `destroy()` method removes a Sensor from DMS and releases references to its underlying resources.

Note: Do not use these methods to reset or destroy built-in metrics. The `reset()` and `destroy()` methods are intended for use with metrics that you create. Application Server Control Console, and other Oracle Application Server administrative facilities could report unexpected values or have unexpected behavior if you use these methods on internal, built-in metrics.

B.8 DMS Coding Recommendations

The following list includes coding recommendations for working with DMS.

1. There is a global name space for DMS metrics. When you create a new Noun Sensor (PhaseEvent, Event, or State), its full name must not conflict with names in use by Oracle built-in metrics, or by other applications. It is therefore a good idea to have a root Noun for your application that contains the application's full name. This prevents name space collisions.

See Also: ["General DMS Naming"](#) on page B-6

2. Be sure all PhaseEvents are stopped. If the code block to be measured is not in a `try` block, then put it in a `try` block that includes PhaseEvent's `start()`. Put the PhaseEvent's `stop()` in a `finally` block. Alternatively, make use of the `abort()` method in the `finally` block, as shown in [Example B-5](#).

See Also: ["Using PhaseEvent Sensors"](#) on page B-10

3. Use the DMS naming conventions.

See Also: ["DMS Naming Conventions"](#) on page B-6

4. Avoid creating any DMS Sensor or Noun more than once. The DMS API allows this, and avoids creation of multiple objects, but DMS performs lookups for each subsequent creation attempt. Thus, whenever possible, you should define Sensors and Nouns during static initialization, or in the case of a Servlet, in the `init()` method.
5. Assign a type for each Noun that contains Sensors. If no type is assigned, the type is given the value "n/a" (not available). Nouns with the type specified as "n/a" are not shown in the AggreSpy display.

6. Only use PhaseEvents to measure a section of code that is expensive to execute, and takes a significant time to execute under some conditions. In the case where the code never takes significant time to execute, use an Event metric, or remove the PhaseEvent.
7. The DMS API calls are threadsafe; they provide sufficient synchronization to prevent races and access bugs.

B.8.1 Isolating Expensive Intervals Using PhaseEvent Metrics

Carefully consider the requirements for new metrics when you add DMS instrumentation. It is important to add a sufficient number of metrics to validate that your code is behaving as desired.

Try to observe the following guidelines when you add DMS metrics:

1. Add PhaseEvent Sensors only to provide an overview of the time the system spends in your block of code or module. You do not need to collect performance data for every method call, or for every distinct phase of your code or module.
2. When your code calls external code that you do not control, and that you expect could take a significant amount of time, add a PhaseEvent Sensor to track the start and the completion of the external code.

Following these guidelines for adding PhaseEvent metrics provides the following benefits:

- Helps to limit the amount of information that DMS collects.
- Allows those analyzing the system to prove that a module gives the expected runtime performance.
- Ensures that people viewing DMS metrics can validate runtime performance without seeing an overwhelming amount of data.
- Allows those analyzing system performance to separate and track your module from other system modules that are either expensive or failure prone.

B.9 Using 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 such as OC4J. Optionally, DMS supports a high resolution clock to increase the precision of performance measurements and lets you select the units for reporting time intervals. You can use a high resolution clock when you need to time phase events more accurately than is possible using the default clock or when the system's default clock does not provide the resolution needed for your requirements.

Note: The resolution of the default clock and of the high resolution clock is system dependent. On some systems the default clock may not provide sufficient resolution for timing requirements. In particular, on Windows platforms, many users request greater precision than the default clock provides, because it advances only once every 15 milliseconds. DMS metrics reported for brief events on these systems must be analyzed with care. Consider using the high resolution clock to address this issue.

This section covers the following topics:

- [Configuring DMS Clocks for Reporting Time for OC4J \(Java\)](#)
- [Configuring DMS Clocks for Reporting Time for Oracle HTTP Server](#)

B.9.1 Configuring DMS Clocks for Reporting Time for OC4J (Java)

For Java processes, the default clock uses `java.lang.System.currentTimeMillis()`. Selecting the high resolution clock changes this call for all applications running on the process where the clock is changed. You set the DMS clock and the reporting units globally using the `oracle.dms.clock` and `oracle.dms.clock.units` properties, which control process startup options.

For example, to use the high resolution clock with the default units, set the following property on the Java command line for OC4J.

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

Caution: Using the high resolution clock, the default units are different than the value that Application Server Control Console expects (msecs). If you need the Application Server Control Console displays to be correct when using the high resolution clock, then you need to set the units property as follows:

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

[Table B-6](#) shows supported values for the `oracle.dms.clock` property.

[Table B-7](#) shows supported values for the `oracle.dms.clock.units` property.

See Also:

Table B-6 *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	Specifies that DMS use the high resolution clock. DMS accesses the high resolution clock using JNI (the JNI calls depend on the clocks available on the underlying operating system). The default value for the units for the HIGHRES clock is NSECS.

Note: On Windows platforms with a Pentium processor, DMS uses the `QueryPerformanceCounter` function to provide timing for the high resolution clock (HIGHRES). If you are running on a system without a Pentium processor, DMS uses the DMS C clock to provide timing for the high resolution clock. The DMS C clock has microsecond precision which offers a significant improvement over the default clock available with `System.currentTimeMillis()`.

Table B-7 *oracle.dms.clock.units Property Values*

Value	Description
MSECS	Specifies that the time be converted to milliseconds and reported as "msecs". Note: This is the default value for the default clock.
NSECS	Specifies that the time be converted to nanoseconds and reported as "nsecs". Note: This is the default value for the high resolution clock.
USECS	Specifies that the time be converted to microseconds and reported as "usecs".

Note the following when using the high resolution DMS clock:

- When you set the `oracle.dms.clock` and the `oracle.dms.clock.units` properties, any combination of upper and lower case characters is valid for the value that you select (case is not significant). For example, any of the following values are valid to select the high resolution clock: `highres`, `HIGHRES`, `HighRes`.
- DMS checks the property values at startup. When you set the clock with a value that does not match those listed in [Table B-6](#), then DMS uses the default clock. If the `oracle.dms.clock` property is not set, DMS also uses the default clock.
- If the specified clock units property value does not match those listed in [Table B-7](#), then DMS uses the default units for the specified clock. If the `oracle.dms.clock.units` property is not set, DMS uses the default units for the specified the clock.

[Table B-8](#) lists the platform specific environment variables settings for supported platforms. To use the high resolution DMS clock, the environment variables need to be set appropriately. The high resolution clock uses the DMS C library. On UNIX systems, this requires `libdms2.so` to be in the specified environment variable path. On Windows systems this requires `yod.dll` to be in the `PATH` environment. If a nanosecond clock is not available, high resolution timings use a microsecond clock.

Table B-8 *Library Path Environment Variables for Supported Platforms*

Platform	Environment Variable
AIX	LIBPATH \$ORACLE_HOME/lib/libdms2.so is required in the path LD_LIBRARY_PATH \$ORACLE_HOME/lib/libdms2.so is required in the path

Table B–8 (Cont.) Library Path Environment Variables for Supported Platforms

Platform	Environment Variable
HP-UX	SHLIB_PATH \$ORACLE_HOME/lib/libdms2.so is required in the path LD_LIBRARY_PATH \$ORACLE_HOME/lib/libdms2.so is required in the path
Linux	LD_LIBRARY_PATH \$ORACLE_HOME/lib/libdms2.so is required in the path
Tru64 UNIX	LD_LIBRARY_PATH \$ORACLE_HOME/lib/libdms2.so is required in the path
Solaris	LD_LIBRARY_PATH \$ORACLE_HOME/lib/libdms2.so is required in the path
Windows 2000	%ORACLE_HOME%\Apache\Apache\yod.dll must be in the PATH
Windows 2003	%ORACLE_HOME%\Apache\Apache\yod.dll must be in the PATH
Windows XP	%ORACLE_HOME%\Apache\Apache\yod.dll must be in the PATH

B.9.2 Configuring DMS Clocks for Reporting Time for Oracle HTTP Server

The default clock for measuring Oracle HTTP Server performance has a resolution of microseconds (usecs). You can optionally select a higher resolution clock to monitor C processes running under Oracle HTTP Server. To use the High Resolution clock under Oracle HTTP Server, you need to set configuration options in `httpd.conf`, or specify environment variables on the command line.

[Table B–9](#) lists the environment variables that control the Oracle HTTP Server DMS clock. [Table B–10](#) describes the `httpd.conf` configuration options that control the Oracle HTTP Server DMS clock. If you set both the command line options and the `httpd.conf` configuration options, the configuration options override the values set on the command line.

Table B–9 OHS DMS Clock Environment Variables

Environment Variable	Description
DMS_CLOCK	Specifies the clock to use for DMS timing. The values are interpreted the same as with <code>oracle.dms.clock</code> . Valid Values: DEFAULT, HIGHRES
DMS_CLOCK_UNITS	Specifies the units for reporting DMS timing values. The values are interpreted the same as with <code>oracle.dms.clock.units</code> . Valid Values: MSECS, NSECS, USECS Default Value: USECS

Table B–10 OHS DMS Clock Configuration Parameters

Parameter	Description
DmsClock	Specifies the clock for HTTP listener processes started by OHS, as the <code>oracle.dms.clock</code> property does for Java processes. Valid Values: DEFAULT, HIGHRES

Table B-10 (Cont.) OHS DMS Clock Configuration Parameters

Parameter	Description
DmsClockUnits	Specifies the time units for HTTP listener processes started by OHS, exactly as the oracle.dms.clock.units property is for Java processes. Valid Values: MSECS, NSECS, USECS Default Value: USECS

Note: On Windows platforms with a Pentium processor, DMS uses the `QueryPerformanceCounter` function to provide timing for the high resolution clock (HIGHRES). If you are running on a system without a Pentium processor, DMS uses the DMS C clock to provide timing for the high resolution clock. The DMS C clock has microsecond precision which offers a significant improvement over the default clock available with `System.currentTimeMillis()`.

For example, if you want to use the high resolution clock and use the same units to show times for Java processes running under OC4J and for `mod_oc4j` running under Oracle HTTP Server, update the Oracle HTTP Server `httpd.conf` file to include the following parameters and values:

```
DmsClock=HIGHRES
DmsClockUnits=MSECS
```

Also, include the following values as startup options for the OC4J process:

```
-Doracle.dms.clock=HIGHRES
-Doracle.dms.clock.units=MSECS
```

Using these options DMS uses a high resolution clock for all the Oracle HTTP Server processes that it monitors, for the Java OC4J processes that it monitors, and DMS reports values using the milliseconds units (msecs).

Caution: Using the high resolution clock for the Oracle HTTP Server, the default units for the high resolution clock are NSECS on most platforms. If you need to use Application Server Control Console, it expects USECS for the units. If you need the Application Server Control Console displays to be correct when using the high resolution clock, then you need to set the units property as follows:

```
DmsClock=HIGHRES
DmsClockUnits=USECS
```

B.10 Rolling Up DMS Data for Descendent Nouns

Oracle Application Server 10g Release 3 (10.1.3.1.0) includes the DMS Rollup feature that lets you specify metric aggregation. You can use the Rollup feature to specify metric aggregation during DMS instrumentation; `rollup` is specified to apply to descendents of a specified noun type. You can specify whether the rollup should only apply to direct descendents or to all descendents. [Example B-11](#) shows code that generates a DMS tree, as represented in [Figure B-2](#). Each noun of type `myContainer` contains the `percentageFull`, `close`, and `open` Sensors (see [Figure B-2](#)).

Note: The code in [Example B-11](#) generates a noun tree hierarchy that violates the guidance described in, "[Choosing Noun Types](#)" on page B-9. In this example, it makes sense for some nouns to have descendents and ancestors of the same noun type. The rollup feature described in this section can collect data which might otherwise be lost.


Example B-11 DMS sample code creating noun hierarchy of metrics

```
// Create DMS Noun hierarchy for metrics.
Noun home = Noun.create(Noun.getRoot(), "Home", "myContainer");
Noun containers = Noun.create(home, "Containers", "myContainer");
Noun closets = Noun.create(containers, "Closets", "myContainer");
Noun bedrooms = Noun.create(closets, "Bedrooms", "myContainer");
Noun br1 = Noun.create(bedrooms, "BR1", "myContainer");

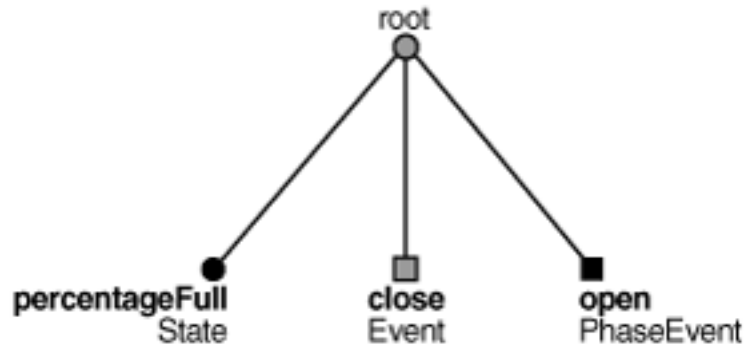
// Create a closet Noun and create Sensors for it.
Noun c1 = Noun.create(br1, "C1", "myContainer");
State percent = State.create(br1, "percentageFull", State.INTEGER, "percent",
"percentage full");
Event close = Event.create(br1, "close", "container closed");
PhaseEvent open = PhaseEvent.create(br1, "open", "open container");

// Derive metrics for State and PhaseEvent Sensors
percent.deriveMetric(Sensor.all);
open.deriveMetric(Sensor.all);
```

Figure B-2 Containers DMS Hierarchy Showing Tree Containing Metrics



[Figure B-3](#) shows a tree with a set of descendent containers. The nouns C1 and C2 under the bedrooms BR1 and BR2 are of type myContainer (see [Figure B-3](#) for a description of myContainer metrics).

Figure B–3 Noun *myContainer* showing Sample Sensors

Using the rollup feature, DMS lets you aggregate a summary for descendent Nouns. For example, you can add the rollup call to a *bedrooms* noun, as shown in [Example B–11](#). To aggregate *myContainer* type metrics under BR1, use the following call:

```
br1.rollup("myContainer", Noun.DIRECT);
```

This call creates a rollup noun named *myContainer_rollup* under */Home/Containers/Closets/Bedrooms/BR1*. The rollup noun contain the same sensors as the associated noun, including: *percentageFull*, *close*, and *open*.

DMS rollup metrics let you rollup the sensors in all descendent nouns of the given types or only those in the direct descendent nouns. Specifying `Noun.DIRECT` in the rollup call aggregates only direct descendent nouns of the specified type. To aggregate the metrics from all descendent nouns of type *myContainer* instead, use a call such as the following including `Noun.ALL`:

```
closets.rollup("myContainer", Noun.ALL);
```

Rollup metrics include aggregate summary information for their contents. [Table B–11](#) shows the available derived rollup metrics for each Sensor type.

Table B–11 Rollup Metrics Included Derived Metrics

Metric	Description
PhaseEvent	The derived metrics for a PhaseEvent rollup metric include the following: <ul style="list-style-type: none"> ■ <code>time</code>: the sum of <code>time</code> metrics. ■ <code>completed</code>: the sum of the <code>completed</code> metrics. ■ <code>maxTime</code>: the maximum of the <code>maxTime</code> metrics. ■ <code>minTime</code>: the minimum of the <code>minTime</code> metrics: ■ <code>avg</code>: the average time computed for all Sensors. ■ <code>active</code>: the sum of the <code>active</code> metrics.
Event	The derived metrics for a Event rollup metric include the following: <ul style="list-style-type: none"> ■ <code>sum</code>: the total of all <code>count</code> metrics. ■ <code>avg</code>: the average of all <code>count</code> metrics.

Table B-11 (Cont.) Rollup Metrics Included Derived Metrics

Metric	Description
State	The derived metrics for a State rollup metric include the following: <ul style="list-style-type: none"> ■ sum: the total of all value metrics. ■ avg: the average of all value metrics. ■ maxValue: the maximum of the maxValue metrics. ■ minValue: the minimum of the minValue metrics.
descendents	The rollup noun includes a <code>descendents</code> state sensor that reports whether the rollup covers only direct descendents or all descendents.
rolled	The rollup noun includes a <code>rolled</code> state sensor, which reports the number of nouns that are rolled up.
refresh	The rollup noun includes a <code>refresh</code> phase event, which reports the time spent aggregating the metrics for this rollup noun.

[Example B-12](#) shows sample metrics created for the `myContainer` rollup noun under `/Home/Containers/Closets`.

Example B-12 Test

```
myContainer_rollup
  descendent.value: all
  percentageFull.sum 40 percent
  percentageFull.avg 10.0 percent
  percentageFull.min 1 percent
  percentageFull.max 29 percent
  close.sum: 3
  close.avg: 0.75
  open.time: 871 msec
  open.completed: 4 ops
  open.maxTime: 722 msec
  open.minTime: 23 msec
  open.avg: 217.7 msec
  open.active: 0
  rolled.value: 4 nouns
  refresh.maxActive: 1 threads
  refresh.active: 0 threads
  refresh.avg: 0.2857142857142857 msec
  refresh.maxTime: 1 msec
  refresh.minTime: 0 msec
  refresh.completed: 7 ops
  refresh.time: 2 msec
```

Note that the metrics are similar to the `myContainer` metrics. The rollup metrics have several key differences, as follows:

1. The rollup noun contains the `descendent`, `rolled`, and `refresh` metrics (see [Table B-11](#) for details).
2. The `percentageFull` State contains `sum` and `avg` metrics rather than the `value` metric. The name of each metric reflects its content.
3. The `close` Event contains `sum` and `avg` metrics rather than the `count` metric. The name of each metric reflects its content.
4. The `open` PhaseEvent does not contain a `maxActive` metric as it would have no meaning in this context.

See Also: *Oracle Application Server DMS API Reference Javadoc*

Performance Metrics

This appendix lists built-in metrics that can help you analyze Oracle Application Server performance. The metrics fall into several distinct areas, such as Oracle HTTP Server, Oracle Containers for J2EE (OC4J). Each table in this chapter lists the metrics that are included in a corresponding Dynamic Monitoring Services metric table.

This appendix contains:

- [Oracle HTTP Server Metrics](#)
- [JVM Metrics](#)
- [JDBC Metrics](#)
- [mod_plsql Metrics](#)
- [Oracle Process Manager and Notification Server - OPMN Metrics](#)
- [DMS Internal Metrics](#)

C.1 Oracle HTTP Server Metrics

The tables, [Table C-1](#) through [Table C-5](#) describe the Oracle HTTP Server metrics.

[Table C-1](#) describes the HTTP server metrics. The metric table name is `ohs_server`.

Table C-1 HTTP Server Metrics (ohs_server)

Metric	Description	Unit
<code>busyChildren.value</code>	Number of busy Child processes	processes
<code>childFinish.count</code>	Number of child processes that finish	processes
<code>childStart.count</code>	Number of child processes that start	processes
<code>connection.active</code>	Number of connections currently open	threads
<code>connection.avg</code>	Average time spent servicing HTTP connections	usecs
<code>connection.completed</code>	Number of times an HTTP connection was established.	ops
<code>connection.maxTime</code>	Maximum time spent servicing any HTTP connection	usecs
<code>connection.minTime</code>	Minimum time spent servicing any HTTP connection	usecs
<code>connection.time</code>	Total time spent servicing HTTP connections	usecs
<code>error.count</code>	Number of HTTP errors	count
<code>get.count</code>	Number of GET requests	count
<code>handle.active</code>	Child servers currently in the handle processing phase	threads
<code>handle.avg</code>	Average time spent in module handler	usecs
<code>handle.completed</code>	Number of times the handle processing phase has completed	ops
<code>handle.maxTime</code>	Maximum time spent in module handler	usecs
<code>handle.minTime</code>	Minimum time spent in module handler	usecs
<code>handle.time</code>	Total time spent in module handler	usecs
<code>internalRedirect.count</code>	Number of times a module redirected a request to a new, internal URI	ops
<code>lastConfigChange.value</code>	Time and date configuration was last modified	time
<code>numChildren.value</code>	Total number of processes for request handling	processes
<code>numMods.value</code>	Number of loaded modules	ops
<code>post.count</code>	Number of POST requests	ops
<code>readyChildren.value</code>	Number of processes ready to handle a request	processes
<code>request.active</code>	Child servers currently in the request processing phase	threads
<code>request.avg</code>	Average time required to service an HTTP request	usecs
<code>request.completed</code>	Number of HTTP request completed	ops
<code>request.maxTime</code>	Maximum time required to service an HTTP request	usecs
<code>request.minTime</code>	Minimum time required to service an HTTP request	usecs
<code>request.time</code>	Total time required to service HTTP requests	usecs
<code>responseSize.value</code>	Size of response	bytes

C.1.1 Oracle HTTP Server Child Server Metrics

[Table C-2](#) describes the child server metrics.

The metric table name is `ohs_child`.

Table C–2 Oracle HTTP Server Child Server Metrics (ohs_child)

Metric	Description	Unit
pid.value	Child's process identifier	
slot.value	Child's slot identifier	
status.value	Current status of child	
time.value	Time spent by this child in processing most recent request	
url.value	URL of most recent request	

C.1.2 Oracle HTTP Server Responses Metrics

The Oracle HTTP Server responses metrics are included in the metric table type `ohs_responses`. This metric table includes one metric containing the count, number of times the response was generated, for each HTTP response type.

For example, `Success_OK_200.count: 28 ops`.

C.1.3 Oracle HTTP Server Virtual Host Metrics

[Table C–3](#) shows the Oracle HTTP Server virtual host metrics.

The metrics table type is `ohs_vhostSet`.

The `ohs_virtualHost` metric table type contains information on virtual host names and locations, and request and response metrics.

Table C–3 Oracle HTTP Server Virtual Host Metrics (ohs_virtualHost)

Metric	Description	Unit
request.active	Number of requests currently being processed by this host	threads
request.avg	Average time spent processing requests for this virtual host	usecs
request.completed	Number of requests processed by this virtual host	ops
request.maxTime	Maximum time spent processing any single request for this virtual host	usecs
request.minTime	Minimum time spent processing any single request for this virtual host	usecs
request.time	Total time spent processing requests for this virtual host	usecs
responseSize.value	Size of response	bytes
vhostType.value	Type of virtual host	

C.1.4 Aggregate Module Metrics

[Table C–4](#) shows the Oracle HTTP Server module metrics.

The metric table type is `ohs_module`.

Table C–4 Oracle HTTP Server Modules Metrics

Metric	Description	Unit
numMods.value	Number of loaded modules	

C.1.5 HTTP Server Module Metrics

There is one set of metrics for each module loaded into the server.

The metric table name is `ohs_module`.

Table C-5 Oracle HTTP Server Modules/mod_*.c Metrics (ohs_module)

Metric	Description	Unit
<code>decline.count</code>	Number of requests declined	ops
<code>handle.active</code>	Number of requests currently being handled by this module	requests
<code>handle.avg</code>	Average time required for this module	usecs
<code>handle.completed</code>	Number of requests handled by this module	ops
<code>handle.maxTime</code>	Maximum time required for this module	usecs
<code>handle.minTime</code>	Minimum time required for this module	usecs
<code>handle.time</code>	Total time required for this module	usecs

C.1.6 Oracle HTTP Server mod_oc4j Metrics

Table C-6 shows the `mod_oc4j` Failure Causes metrics. This table represents the categorization of errors that return an `INTERNAL_SERVER_ERROR` to the client.

The metric table name is `mod_oc4j_request_failure_causes`.

Table C-6 HTTP Server mod_oc4j Request Failure Causes Metrics

Metric	Description	Unit
<code>IncorrectReqInit.count</code>	The total number of times an internal error occurred. There could be a number of reasons, including: <code>mod_oc4j</code> not finding a connection endpoint, configuration errors, and others.	ops
<code>Oc4jUnavailable.count</code>	The total number of times that an <code>oc4j</code> JVM could not be found to service requests.	ops
<code>UnableToHandleReq.count</code>	The total number of times <code>mod_oc4j</code> declined to handle a request.	ops

Table C-7 shows the `mod_oc4j` Mount Point metrics. There is one mount point metric table for each mount point specified in `mod_oc4j.conf`. This table includes a set of metrics for each mount point specified, with each set grouped under the `mntPtId`. Where `id` is an integer that is automatically generated during module initialization.

The metric table name is `mod_oc4j_mount_pt_metrics`.

Table C-7 HTTP Server mod_oc4j Mount Point Metrics

Metric	Description	Unit
<code>Destination.value</code>	Specifies the destination name. For example, with: <code>Oc4jMount /j2ee/* home</code> The <code>Destination.value</code> would be <code>home</code>	String
<code>ErrReq.count</code>	Specifies the total number of requests, both session and non-session, that <code>mod_oc4j</code> failed to route to an <code>OC4J</code> .	ops
<code>ErrReqNonSess.count</code>	Specifies the total number of non session requests that <code>mod_oc4j</code> failed to route to an <code>oc4j</code> process.	ops
<code>ErrReqSess.count</code>	Specifies the total number of session requests that <code>mod_oc4j</code> failed to route to an <code>OC4J</code> process.	ops

Table C-7 (Cont.) HTTP Server mod_oc4j Mount Point Metrics

Metric	Description	Unit
Failover.count	Shows the total number of requests that have had a failover (which means that the had an error while talking to a JVM and switched over to another JVM).	ops
Name.value	Specifies the echo of the value specified as the path for Oc4jMount directive in mod_oc4j.conf. DMS changes certain characters, including: '/' and '*' to '_'. To preserve the actual path names specified, an internal table containing a mapping between mntPtId and the actual path name is created during mod_oc4j initialization. For example, with: Oc4jMount /j2ee/* home Name.value would be /j2ee/*	String
NonSessFailover.count	Specifies the total number of failovers for nonsession requests. Shows the number of requests that have had a failover (which means that the request had an error while talking to a JVM and switched over to another JVM).	ops
SessFailover.count	Specifies the total number of failovers for session requests. Shows the number of requests that have had a failover (which means that the request had an error while talking to a JVM and switched over to another JVM).	ops
SucReq.count	Specifies the total number of requests, both session and non-session, that mod_oc4j successfully routed to an OC4J instance.	ops
SucReqNonSess.count	Specifies the total number of non session requests that mod_oc4j successfully routed to an OC4J process.	ops
SucReqSess.count	Specifies the total number of session requests that mod_oc4j successfully routed to an OC4J process.	ops

Table C-8 shows the mod_oc4j Destination Metrics. This table includes a set of metrics for a specific destination. Each destination can have multiple mount points. There is one mntPts subtree for each mount point specified in mod_oc4j.conf.

The metric table name is mod_oc4j_destination_metrics.

Table C-8 HTTP Server mod_oc4j Destination Metrics

Metric	Description	Unit
ErrReq.count	Specifies the total number of requests, both session and non-session, that mod_oc4j failed to route to an OC4J.	ops
ErrReqNonSess.count	Specifies the total number of non session requests that mod_oc4j failed to route to an OC4J process.	ops
ErrReqSess.count	Specifies the total number of session requests that mod_oc4j failed to route to an OC4J process.	ops
Failover.count	Shows the total number of requests that have had a failover (which means that the had an error while talking to a JVM and switched over to another JVM..	ops
JVMCnt.value	Specifies the total number of routable OC4J JVMs that belong to this destination.	Number of JVMs
Name.value	Specifies the echo of the value specified as destination for Oc4jMount directive in mod_oc4j.conf, a single destination may appear several times in mod_oc4j.conf. Example: Oc4jMount /j2ee/* home,oc4jinstance2 Name.value would be home,oc4jinstance2	String
NonSessFailover.count	Specifies the total number of failovers for non session requests.	ops
SessFailover.count	Specifies the total number of failovers.	ops
SucReq.count	Specifies the total number of requests, both session and non-session, that mod_oc4j successfully routed to an OC4J.	ops
SucReqNonSess.count	Specifies the total number of non session requests that mod_oc4j successfully routed to an OC4J process.	ops
SucReqSess.count	Specifies the total number of session requests that mod_oc4j successfully routed to an OC4J process.	ops

C.1.7 Oracle HTTP Server SSL Metrics

Table C–9 describes the OSSL metrics.

The metric table type is `ohs_oss1`.

Table C–9 OHS_OSSL Metrics

Metric	Description	Unit
<code>checkcrl.time</code>	SSL checkcrl was invoked	time
<code>closessl.time</code>	SSL connection was closed	time
<code>connectssl.time</code>	SSL connection was established	time
<code>dataReceive.value</code>	OSSL Data received	kilobytes
<code>dataSent.value</code>	OSSL Data Sent	kilobytes
<code>entercache.time</code>	SSL entercache was invoked	time
<code>getcache.time</code>	SSL getcache was invoked	time
<code>handshake.time</code>	SSL handshake was invoked	time
<code>receive.time</code>	an encrypted message was received	time
<code>receiveErrors.count</code>	an error occurred in receive	ops
<code>send.time</code>	an encrypted message was sent	time
<code>sendErrors.count</code>	an error occurred in send	ops
<code>setfixup.time</code>	SSL setfixup was invoked	time

C.2 JVM Metrics

Table C–10 shows the JVM metrics. There is one set of metrics for each Java process (OC4J) running in the site.

The metric table type is `JVM`.

Table C–10 JVM Metrics (JVM)

Metric	Description	Unit
<code>activeThreadGroups.value</code>	The number of active thread groups in the JVM	integer
<code>activeThreadGroups.minValue</code>	The minimum number of active thread groups in the JVM	integer
<code>activeThreadGroups.maxValue</code>	The maximum number of active thread groups in the JVM	integer
<code>activeThreads.value</code>	The number of active threads in the JVM	threads
<code>activeThreads.minValue</code>	The minimum number of active threads in the JVM	threads
<code>activeThreads.maxValue</code>	The maximum number of active threads in the JVM	threads
<code>upTime.value</code>	Up time for the JVM	msecs
<code>freeMemory.value</code>	The amount of heap space free in the JVM	kilobytes
<code>freeMemory.minValue</code>	The minimum amount of heap space free in the JVM	kilobytes
<code>freeMemory.maxValue</code>	The maximum amount of heap space free in the JVM	kilobytes
<code>totalMemory.value</code>	The total amount of heap space in the JVM	kilobytes
<code>totalMemory.minValue</code>	The minimum amount of total heap space in the JVM	kilobytes
<code>totalMemory.maxValue</code>	The maximum amount of total heap space in the JVM	kilobytes

C.2.1 JVM Properties Metrics

Oracle Application Server creates a metric to track the value of each Java Property available through a call to `System.getProperties()` on any Java process. For each Java Property, a metric is created under the `/JVM/Properties` noun.

For example, each process should have a metric that contains the value of the `java.version` system property named, `/JVM/Properties/java_version.value`. The system converts property name components with a period, '.' to '_'.

If, during the life of a process, a property is deleted from the JVM system properties, the corresponding metric is deleted. If the value changes, this is reflected in the metric value the next time it is accessed. If a new property is added to the system properties, a new metric is created.

Note: The JVM Properties metrics are only available for viewing using the `Spies text` link in `AggreSpy`, or using the `dmstool` command to display metrics.

Table C–11 *JVM/Properties - JVM System Properties Metrics*

Metric	Description	Unit
A metric is created for each system property. Each property name has any of the "." characters in the name replaced with "_".	Contains the value of the Java system property.	String

C.3 JDBC Metrics

The following tables list the Oracle Application Server JDBC metrics.

C.3.1 JDBC Driver Metrics

[Table C–12](#) shows the JDBC driver metrics. There is one set of JDBC Driver metrics per JVM.

The metric table type is `JDBC_Driver`.

Table C–12 */JDBC/Driver - JDBC_Driver Metrics*

Metric	Description	Unit
<code>ConnectionCloseCount.count</code>	Total number of connections that have been closed.	ops
<code>ConnectionCreate.active</code>	Current number of threads creating connections.	ops
<code>ConnectionCreate.avg</code>	Average time spent creating connections.	msecs
<code>ConnectionCreate.completed</code>	Number of times this PhaseEvent has started and ended.	ops
<code>ConnectionCreate.maxTime</code>	Maximum time spent creating connections.	msecs
<code>ConnectionCreate.minTime</code>	Minimum time spent creating connections.	msecs
<code>ConnectionCreate.time</code>	Time spent creating connections.	msecs
<code>ConnectionOpenCount.count</code>	Total number of connections that have been opened.	ops

C.3.2 JDBC Data Source Metrics

[Table C–13](#) shows the JDBC datasource metrics. There is one set of data source metrics per data source.

The metric table type is `JDBC_DataSource`.

Table C–13 */JDBC/data-source-name - JDBC_Data Source Metrics*

Metric	Description	Unit
ConnectionCloseCount.count	Total number of connections that have been closed.	ops
ConnectionCreate.active	Current number of threads creating connections.	ops
ConnectionCreate.avg	Average time spent creating connections.	msecs
ConnectionCreate.completed	Number of times this PhaseEvent has started and ended.	ops
ConnectionCreate.maxTime	Maximum time spent creating connections.	msecs
ConnectionCreate.minTime	Minimum time spent creating connections.	msecs
ConnectionCreate.time	Time spent creating connections.	msecs
ConnectionOpenCount.count	Total number of connections that have been opened.	ops

C.3.3 JDBC Driver Specific Connection Metrics

[Table C–14](#) shows the JDBC driver connection metrics. There is one set of JDBC Connection metrics per connection.

The metric table type is `JDBC_Connection`.

Table C–14 */JDBC/Driver/CONNECTION - JDBC Driver Connection Metrics*

Metric	Description	Unit
CreateNewStatement.avg	Average time spent creating a new statement.	msecs
CreateNewStatement.completed	Number of times a request for a statement failed to be satisfied from the cache.	ops
CreateNewStatement.maxTime	Maximum time spent creating a new statement.	msecs
CreateNewStatement.minTime	Minimum time spent creating a new statement.	msecs
CreateNewStatement.time	Time spent creating a new statement (this does not include the time required to parse the statement. For information on the metric that includes the parse time see <code>Execute.Time</code> in Table C–17).	msecs
CreateStatement.avg	Average time spent getting a statement from the statement cache.	msecs
CreateStatement.completed	Number of times a request for a statement was satisfied from the cache.	ops
CreateStatement.maxTime	Maximum time spent getting a statement from the statement cache.	msecs
CreateStatement.minTime	Minimum time spent getting a statement from the statement cache.	msecs
CreateStatement.time	Time spent getting a statement from the statement cache.	msecs
JDBC_Connection_URL	Url specified for the connection	
JDBC_Connection_Username	User name used for the connection	
LogicalConnection.value	If this is a physical connection, then this refers to its logical connection, if any.	
StatementCacheHit.count	Statement found in cache	ops
StatementCacheMiss.count	Statement not found in cache	ops

C.3.4 JDBC Data Source Specific Connection Metrics

[Table C–15](#) shows the JDBC data source metrics. There is one set of JDBC data source specific connection metrics per data source per connection.

The metric table type is `JDBC_Connection`.

Table C–15 */JDBC/data-source-name/CONNECTION - JDBC Datasource Connection Metrics*

Metric	Description	Unit
CreateNewStatement.avg	Average time spent creating a new statement.	msecs
CreateNewStatement.completed	Number of times a request for a statement failed to be satisfied from the cache.	ops
CreateNewStatement.maxTime	Maximum time spent creating a new statement.	msecs
CreateNewStatement.minTime	Minimum time spent creating a new statement.	msecs
CreateNewStatement.time	Time spent creating a new statement (this time does not include the time required to parse the statement. For information on the metric that includes the parse time see <code>Execute.Time</code> in Table C–18).	msecs
CreateStatement.avg	Average time spent getting a statement from the statement cache.	msecs
CreateStatement.completed	Number of times a request for a statement was satisfied from the cache.	ops
CreateStatement.maxTime	Maximum time spent getting a statement from the statement cache.	msecs
CreateStatement.minTime	Minimum time spent getting a statement from the statement cache.	msecs
CreateStatement.time	Time spent getting a statement from the statement cache.	msecs
JDBC_Connection_Url	Url specified for the connection	
JDBC_Connection_Username	User name used for the connection	
LogicalConnection.value	If this is a physical connection, then this refers to its logical connection, if any.	
StatementCacheHit.count	Statement found in cache	
StatementCacheMiss.count	Statement not found in cache	

C.3.5 JDBC Connection Source Metrics

[Table C–16](#) shows the JDBC connection source metrics.

The metric table type is `JDBC_ConnectionSource`.

Table C–16 *JDBC Connection Source Metrics*

Metric	Description	Unit
CacheFreeSize.count	Number of free slots in the connection cache.	ops
CacheFreeSize.maxValue	Maximum number of free slots in the connection cache.	connections
CacheFreeSize.minValue	Minimum number of free slots in the connection cache.	connections
CacheFreeSize.value	Number of free slots in the connection cache.	connections
CacheGetConnection.active		threads
CacheGetConnection.avg	Average time spent getting a connection from the cache.	msecs
CacheGetConnection.completed	Number of times this PhaseEvent has started and ended.	ops
CacheGetConnection.maxTime	Maximum time spent getting a connection from the cache.	msecs
CacheGetConnection.minTime	Minimum time spent getting a connection from the cache.	msecs
CacheGetConnection.time	Time spent getting a connection from the cache or not.	msecs
CacheHit.count	Number of times a request for a connection has been satisfied from the cache.	ops
CacheMiss.count	Number of times a request for a connection failed to be satisfied from the cache.	ops
CacheSize.value	Number of physical connections in the cache	ops

C.3.6 JDBC Driver Statement Metrics

[Table C–17](#) shows the JDBC statement metrics. There is a set of JDBC statement metrics per connection per statement.

The metric table type is `JDBC_Statement`.

Note: The JDBC statement metrics are only available for JDBC connections that have enabled statement caching, and set the property `oracle.jdbc.DMSStatementCachingMetrics` to the value `true`. When JDBC statement caching is disabled, you can make the JDBC statement metrics available by setting the property `oracle.jdbc.DMSStatementMetrics` to `true`. To improve performance and to avoid collecting expensive metrics, by default these properties are both set to `false`.

Table C–17 */JDBC/Driver/CONNECTION/STATEMENT JDBC Statement Metrics*

Metric	Description	Unit
<code>Execute.time</code>	The time this statement has spent executing the SQL including the first fetch and the time required to parse the statement.	msecs
<code>Fetch.time</code>	The time this statement has spent in other fetches.	msecs
<code>SQLText.value</code>	The SQL being executed.	

C.3.7 JDBC Data Source Statement Metrics

[Table C–18](#) shows the JDBC statement metrics. There is a set of statement metrics per data source per connection per statement.

The metric table type is `JDBC_Statement`.

Note: The JDBC statement metrics are only available for JDBC connections that have enabled statement caching and set the property `oracle.jdbc.DMSStatementCachingMetrics` to the value `true`. When JDBC statement caching is disabled, you can make the JDBC statement metrics available by setting the property `oracle.jdbc.DMSStatementMetrics` to `true`. To improve performance and to avoid collecting expensive metrics, by default these properties are set to `false`.

Table C–18 */JDBC/data-source-name/CONNECTION/STATEMENT JDBC Statement Metrics*

Metric	Description	Unit
<code>Execute.time</code>	The time this statement has spent executing the SQL including the first fetch and the time required to parse the statement.	msecs
<code>Fetch.time</code>	The time this statement has spent in other fetches.	msecs
<code>SQLText.value</code>	The SQL being executed.	

C.3.8 JDBC Connection Pool Stats Metrics

Table C–19 shows the JDBC connection pool stats metrics.

The metric table type is `jdbc_connection_pool_stats`.

Table C–19 *JDBC Connection Pool Stats Metrics*

Metric	Description	Units
<code>CloseConnectionCount.value</code>	The number of connections closed.	connections
<code>CreateConnectionCount.value</code>	The number of connections created.	connections
<code>FreePoolSize.maxValue</code>	The number of available connections in the pool.	connections
<code>FreePoolSize.minValue</code>	The upper bound of number of available connections in the pool.	connections
<code>FreePoolSize.value</code>	The upper bound of number of available connections in the pool	connections
<code>FreePoolSizeUpperBound.value</code>	The upper bound of number of available connections in the pool	connections
<code>PoolSize.maxValue</code>	The total number of connections in the pool (used and available)	connections
<code>PoolSize.minValue</code>	The total number of connections in the pool (used and available)	connections
<code>PoolSize.value</code>	The total number of connections in the pool (used and available)	connections
<code>PoolSizeLowerBound.value</code>	The lower bound of total number of connections in the pool	connections
<code>PoolSizeUpperBound.value</code>	The upper bound of total number of connections in the pool	connections
<code>UseTime.time</code>	The time spent using a connection	time
<code>WaitTime.time</code>	The time spent waiting for a connection to be available	time
<code>WaitingThreadCount.maxValue</code>	The number of threads waiting for a connection	count
<code>WaitingThreadCount.minValue</code>	The number of threads waiting for a connection	count
<code>WaitingThreadCount.value</code>	The number of threads waiting for a connection	count

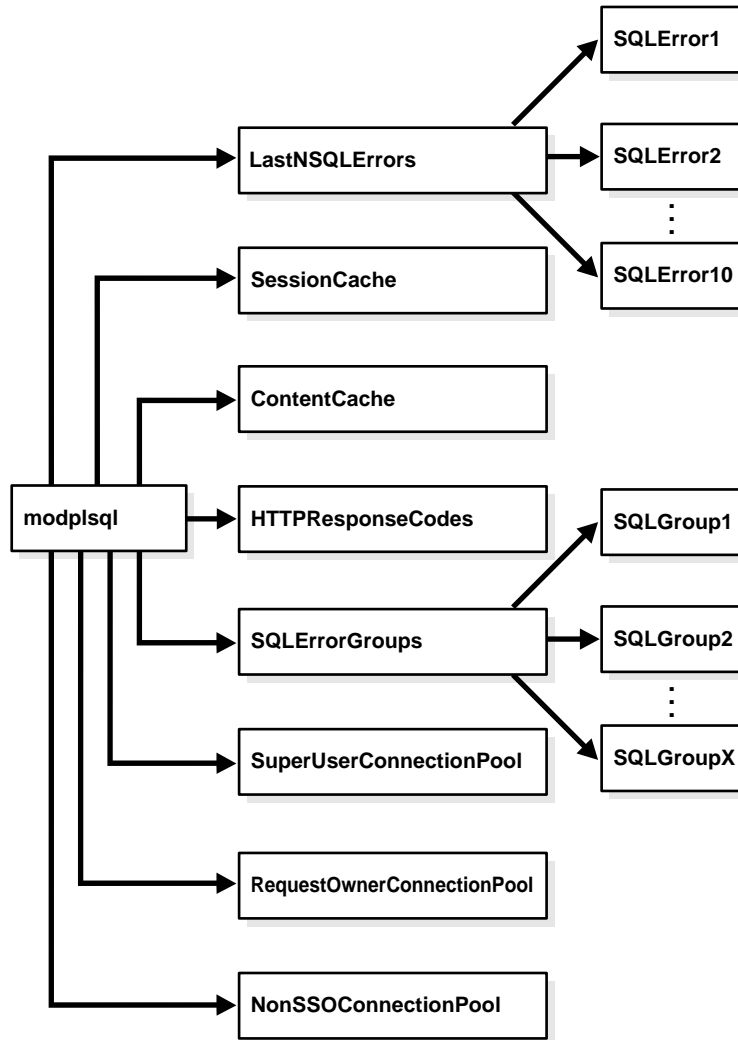
C.4 mod_plsql Metrics

This section describes the Oracle Application Server `mod_plsql` metrics.

Figure C–1, "mod_plsql Metric Tree" shows the structure of the `mod_plsql` metrics.

The tables in this section describe the relevant metrics.

Figure C-1 mod_plsql Metric Tree



The /modplsql/HTTPResponseCodes Metrics lists the response codes returned by mod_plsql.

The metric table name is modplsql_HTTPResponseCodes. This metric table includes one metric containing the count, number of times the response was generated, for each HTTP response type.

[type=modplsql_HTTPResponseCodes]

For example, the http404.count metric holds a count of the "HTTP 404: Not found" response codes.

Table C-20 lists the set of metrics for the mod_plsql session cache.

The metric table name is modplsql_Cache.

Table C-20 mod_plsql/SessionCache Metrics

Metric	Description	Unit
cacheStatus.value	Status of the cache. This can be either enabled or disabled.	status
newMisses.count	Number of session cache misses (new)	ops

Table C–20 (Cont.) mod_plsql/SessionCache Metrics

Metric	Description	Unit
staleMisses.count	Number of session cache misses (stale)	ops
hits.count	Number of session cache hits	ops
requests.count	Number of requests to the session cache	ops

Table C–21 lists the set of metrics for the mod_plsql content cache.

The metric table name is modplsql_ContentCache.

Table C–21 mod_plsql/ContentCache Metrics

Metric	Description	Unit
cacheStatus.value	Status of the cache, either enabled or disabled.	
newMisses.count	Number of content cache misses (new)	ops
staleMisses.count	Number of content cache misses (stale)	ops
hits.count	Number of content cache hits	ops
requests.count	Number of requests to the content cache	ops

The SQLErrorGroups metrics show the predefined groupings of SQL errors. For each group, the metrics in Table C–22 are recorded.

The metric table name is modplsql_SQLErrorGroup:

```
/modplsql/SQLErrorGroups/group [type=modplsql_SQLErrorGroup]
```

The *group* is based on the groupings in the Oracle Database Error Messages guide. For example, the metric name Ora24280Ora29249 represents the grouping Ora-24280 to Ora-29249. Each SQL error that occurs as a result of executing a request is put into the appropriate group based on its error code. If you are getting a high number of the same errors, you should investigate what is causing the problem, using the Oracle Database Error Messages guide for further details on the error message.

Table C–22 mod_plsql/SQLErrorGroups Metrics

Metric	Description	Unit
lastErrorDate.value	Date of the last request to cause the SQL error	date
lastErrorRequest.value	Last request to cause the SQL error	url
lastErrorText.value	SQL error text of the last error	error
error.count	Number of errors that have occurred within the group	ops

The LastNSQLErrors statistics show the last 10 SQL errors that have occurred while executing requests. These are updated in a round robin fashion. For each error, the metrics in Table C–23 are recorded.

The metric table name is modplsql_LastNSQLError:

```
/modplsql/LastNSQLErrors/<SQL Error Slot> [type=modplsql_LastNSQLError]
```

If you are getting a large number of the same errors, you should investigate what is causing the problem. Refer to the Oracle Database Error Messages guide for further details of the error represented by the errorText.value metric.

Table C-23 *mod_plsql/LastNSQLErrors Metrics*

Metric	Description	Unit
errorDate.value	Date the request caused the SQL error	date
errorRequest.value	Request causing the SQL error	url
errorText.value	SQL error text	error

[Table C-24](#) lists the set of metrics for the Non-SSO connection pool.

The metric table name is modplsql_DatabaseConnectionPool:

```
/modplsql/NonSSOConnectionPool [type=modplsql_DatabaseConnectionPool]
```

Table C-24 *mod_plsql/NonSSOConnectionPool Metrics*

Metric	Description	Unit
connFetch.maxTime	Maximum time to fetch a connection from the pool	usecs
connFetch.minTime	Minimum time to fetch a connection from the pool	usecs
connFetch.avg	Average time to fetch a connection from the pool	usecs
connFetch.active	Child servers currently in the pool fetch phase	threads
connFetch.time	Total time spent fetching connections from the pool	usecs
connFetch.completed	Number of times a connection has been requested from the pool	ops
newMisses.count	Number of connection pool misses (new)	ops
staleMisses.count	Number of connection pool misses (stale)	ops
hits.count	Number of connection pool hits	ops

[Table C-25](#) lists the set of metrics for the request owner connection pool.

The metric table name is modplsql_DatabaseConnectionPool:

```
/modplsql/RequestOwnerConnectionPool [type=modplsql_DatabaseConnectionPool]
```

Table C-25 *mod_plsql/RequestOwnerConnectionPool Metrics*

Metric	Description	Unit
connFetch.maxTime	Maximum time to fetch a connection from the pool	usecs
connFetch.minTime	Minimum time to fetch a connection from the pool	usecs
connFetch.avg	Average time to fetch a connection from the pool	usecs
connFetch.active	Child servers currently in the pool fetch phase	threads
connFetch.time	Total time spent fetching connections from the pool	usecs
connFetch.completed	Number of times a connection has been requested from the pool	ops
newMisses.count	Number of connection pool misses (new)	ops
staleMisses.count	Number of connection pool misses (stale)	ops
hits.count	Number of connection pool hits	ops

[Table C-26](#) lists the set of metrics for the super user connection pool.

The metric table name is modplsql_DatabaseConnectionPool:

```
/modplsql/SuperUserConnectionPool [type=modplsql_DatabaseConnectionPool]
```

Table C–26 *mod_plsql/SuperUserConnectionPool Metrics*

Metric	Description	Unit
<code>connFetch.maxTime</code>	Maximum time to fetch a connection from the pool	usecs
<code>connFetch.minTime</code>	Minimum time to fetch a connection from the pool	usecs
<code>connFetch.avg</code>	Average time to fetch a connection from the pool	usecs
<code>connFetch.active</code>	Threads currently in the pool fetch phase	threads
<code>connFetch.time</code>	Total time spent fetching connections from the pool	usecs
<code>connFetch.completed</code>	Number of times a connection has been requested from the pool	ops
<code>newMisses.count</code>	Number of connection pool misses (new)	ops
<code>staleMisses.count</code>	Number of connection pool misses (stale)	ops
<code>hits.count</code>	Number of connection pool hits	ops

C.5 Oracle Process Manager and Notification Server - OPMN Metrics

This sections lists the Oracle Process Manager and Notification Server (opmn) metrics.

This section includes the following:

- [OPMN_PM Metric Table](#)
- [OPMN_OC4J_PROC Table](#)
- [OPMN_HOST_STATISTICS Metric Table](#)
- [OPMN_IAS_INSTANCE Metric Table](#)
- [OPMN_IAS_COMPONENT Table](#)
- [OPMN ONS Metrics](#)
- [OPMN_APPCTX Table](#)

C.5.1 OPMN_PM Metric Table

The `opmn_pm` metric table is the root of the process manager subtree for the OPMN DMS metrics. The metrics in this metric table contain statistics about OPMN requests. An OPMN request is a command that has been issued to OPMN from a client, for example DCM, to perform an operation on one or more OPMN managed processes.

Requests can have one of three possible results:

- Success – success means OPMN handles the request successfully.
- Partial Success – partial Success means OPMN only handles part of the request successfully. For example, if a client wants OPMN to start three OC4J processes, and only two are successfully started, the request result is partial success.
- Failure – failure means the request fails.

[Table C–27](#) shows the metric table type `opmn_pm`.

Table C-27 OPMN_PM Metrics

Metric	Description	Unit
jobWorkerQueue.value	Specifies the number of jobs in the OPMN worker queue	ops
lReq.count	Specifies the number of local HTTP requests which OPMN handles	ops
procDeath.count	Specifies the number of processes which die after the process manager starts them	ops
procDeathReplace.count	Specifies the number of processes which are restarted after the process manager detects they are dead	ops
reqFail.count	Specifies the number of HTTP requests which fail	ops
reqPartialSucc.count	Specifies the number of HTTP requests which partially succeed	ops
reqSucc.count	Specifies the number of HTTP requests which succeed	ops
rReq.count	Specifies the number of remote HTTP requests which OPMN handles	ops
workerThread.value	Specifies the number of worker threads	threads

C.5.2 OPMN_OC4J_PROC Table

[Table C-28](#) shows the OPMN OC4J proc metrics that provides information on the OC4J process.

The metric table type is `opmn_oc4j_proc`.

Table C-28 OPMN_OC4J_proc Metrics

Metric	Description	Unit
oc4jinstance.value		
oc4jIsland.value	This is a backward compatibility metric.	

C.5.3 OPMN_HOST_STATISTICS Metric Table

The OPMN host statistics metric table provides information on the host running the OPMN process.

[Table C-29](#) shows the metric table type `opmn_host_statistics`.

Table C-29 OPMN_HOST_STATISTICS Metrics

Metric	Description	Unit
cpuIdle.value	Specifies the number of milliseconds the cpu(s) have been idle since an unspecified time.	milliseconds
freePhysicalMem.value	Specifies the amount of free physical memory on the host machine.	kilobytes
numProcessors.value	Specifies the number of processors available on the host machine.	integer
timestamp.value	Specifies the time that host statistics are taken. The timestamp is the number of milliseconds from an unspecified time.	milliseconds from an unspecified time
totalPhysicalMem.value	Specifies the total physical memory available on the host machine.	kilobytes

C.5.4 OPMN_IAS_INSTANCE Metric Table

The OPMN IAS instance subtree shows the Oracle Application Server instance node name.

[Table C-30](#) shows the metric table type `opmn_ias_instance`.

Table C–30 OPMN_IAS_INSTANCE Metrics

Metric	Description	Unit
iasCluster.value	Specifies the Oracle Application Server cluster name for the Oracle Application Server instance.	String

C.5.5 OPMN_IAS_COMPONENT Table

The OPMN IAS component subtree represents an Oracle Application Server component. The OPMN IAS component subtree includes several metric tables containing component information.

Table C–31 shows the metric table type `opmn_process_type`.

Table C–31 OPMN_PROCESS_TYPE Metrics

Metric	Description	Unit
moduleId.value	Specifies the values of attribute module-IDs, as specified in the process-type tag in the <code>opmn.xml</code> configuration file.	String

Table C–32 shows the metric table type `opmn_process_set`.

Table C–32 OPMN_PROCESS_SET Metrics

Metric	Description	Unit
numProcConf.value	Specifies the number, or maximum number, of processes configured for this process set.	String (integer)
numProcs.value	Number of process that exist for this process set	
IsService.value	Process set is configured as a service	String
reqFail.count	Specifies the number of HTTP requests which fail for this process set.	ops
reqPartialSucc.count	Specifies the number of HTTP requests which partially succeed for this process set.	ops
reqSucc.count	Specifies the number of HTTP requests which succeed for this process set	ops
restartOnDeath.value	Specifies whether, when a process dies, OPMN should restart the process.	String (boolean)

Table C–33 shows the metric table type `opmn_process`.

Table C–33 OPMN_PROCESS Metrics

Metric	Description	Unit
cpuTime.value	Shows the amount of CPU time used by the process.	CPU msecs
heapSize.value	Shows the heap size of the process.	Kilobytes
iasCluster.value	Shows the Oracle Application Server cluster name for the process	String
iasInstance.value	Shows the Oracle Application Server instance name for the process	String
indexInSet.value	Shows the process index in the process set. This value is only valid for OPMN managed processes, for OPMN unmanaged processes, this value has no meaning, and the value is always 0.	String (integer)

Table C–33 (Cont.) OPMN_PROCESS Metrics

Metric	Description	Unit
memoryUsed.value	The amount of memory used by the process. This metric is calculated in an operating system specific manner. On UNIX, this is the process image memory used value. This is all the memory in use by the process. On Windows, this is the working set memory used value. This is the same value that is reported by the Task Manager under the mem usage column. The working set is the set of memory pages touched recently by the threads in the process. If free memory in the system is over a certain threshold, pages are left in the working set of a process, even if they are not in use. When free memory falls below a certain threshold, pages are trimmed from the working sets. If needed, pages are soft-faulted back into the working set before they leave main memory.	
pid.value	The process ID for the process.	
privateMemory.value	The private memory of the process.	Kilobytes
sharedMemory.value	The shared memory for the process	Kilobytes
startTime.value	The start time of the process.	msecs
status.value	The status of the process. The status can have the following values: <ul style="list-style-type: none"> ▪ NONE – New process slot, no operations have been applied yet (no status). ▪ Init – process has been started, opmn is waiting for initialization to complete. ▪ Alive – process is fully started. ▪ Stop – process stop operation is in progress. ▪ Stopped – process has been fully stopped. ▪ Bounce – non-terminating process restart is in progress. ▪ Restart – process stop operation is in progress, prior to a new start being issued. ▪ InitFail – failure before init timeout reached, a stop and start will be attempted in the retry limit has not been reached. ▪ BounceFail – non-terminating process restart failed, as stop and start will be attempted if the retry limit has not been reached. 	String
type.value	The type of the process. See Table C–31 for information on process types.	
uid.value	The OPMN assigned ID for the process.	
upTime.value	The uptime for the process.	msecs

[Table C–34](#) shows the metric table type `opmn_connect`.

Table C–34 OPMN_CONNECT Metrics

Metric	Description	Unit
desc.value	Shows the port description, if available	String
host.value	Shows the host name	String (host name)
protocol.value		
port.value	Shows the port number	String (port number)

C.5.6 OPMN ONS Metrics

The Oracle Process Manager and Notification Server ONS subtree contains Oracle Notification System (ONS) information.

[Table C–35](#) shows the metric table type `opmn_ons`.

Table C–35 OPMN_ONS Metrics

Metric	Description	Unit
<code>notifProcessed.value</code>	The number of notifications processed by ONS.	ops
<code>notifProcessQueue.value</code>	The number of notifications in the process queue.	ops
<code>notifReceived.value</code>	The number of notifications received by ONS.	ops
<code>notifReceiveQueue.value</code>	The number of notifications in the receive queue.	ops
<code>workerThread.value</code>	The number of worker threads.	String (threads)

[Table C–36](#) shows the `local_port` metrics. The `../ons/local_port` subtree shows information about the ONS local port.

The metric table type is `opmn_connect`

Table C–36 OPMN ONS LOCAL_PORT Metrics

Metric	Description	Unit
<code>desc.value</code>	Port description	String
<code>host.value</code>	Host name	String
<code>port.value</code>	Port number	String

[Table C–37](#) shows the `remote_port` metrics. The `../ons/remote_port` subtree shows information about the ONS remote port.

The metric table type is `opmn_connect`

Table C–37 OPMN ONS REMOTE_PORT Metrics

Metric	Description	Unit
<code>desc.value</code>	Port description	String
<code>host.value</code>	Host name	String
<code>port.value</code>	Port number	String

[Table C–38](#) shows the `request_port` metrics. The `../ons/request_port` subtree shows information about the ONS request port.

The metric table type is `opmn_connect`

Table C–38 OPMN ONS REQUEST_PORT Metrics

Metric	Description	Unit
<code>desc.value</code>	Port description	String
<code>host.value</code>	Host name	String
<code>port.value</code>	Port number	String

[Table C–39](#) shows the `opmn_ons_topo_entry` metrics.

Table C–39 *OPMN ONS TOPO Entry Metrics*

Metric	Description	Unit
protocol.value	ons protocol version	
port.value	port value	
ip.value	ip address	

C.5.7 OPMN_APPCTX Table

[Table C–40](#) shows the opmn_appctx metrics.

Table C–40 *OPMN APPCTX Metrics*

Metric	Description	Unit
rtid.value		
routable.value		
state.value		

C.6 DMS Internal Metrics

[Table C–41](#) shows the DMS internal clock metrics.

Table C–41 *DMS-Internal Clock Metrics*

Metric	Description	Unit
logicalTime.value	The current time as measured with the DMS clock.	ticks
measuredFrequency.value	Number of clock ticks per second - measured.	ticks
measuredResolution.value	Time between ticks as measured with this clock.	
name.value		
overheadPerCall.value	The average duration of a call to get the time with this clock.	
reportedFrequency.value	The number of ticks per second the clock time is reported in.	ticks
requestedUnits.value	The string description of the units that times are reported in.	

[Table C–42](#) shows the DMS internal log metrics.

Table C–42 *DMS-Internal Log Metrics*

Metric	Description	Unit
initLogging.count		ops
messagesLogged.count		ops
status.value		

[Table C–43](#) shows the DMS internal measurement metrics.

Table C–43 *DMS-Internal Measurement Metrics*

Metric	Description	Unit
createNoun.count		ops
createSensor.count		ops
destroyNoun.count		ops

Table C-43 (Cont.) DMS-Internal Measurement Metrics

Metric	Description	Unit
destroySensor.count		ops
lastTreeNodeID.value		
sampleMetric.count		ops
sensorWeight.value		
treeNodes.maxValue		
treeNodes.value		

Table C-44 shows the DMS internal collector metrics.

Table C-44 DMS-Internal Collector Metrics

Metric	Description	Unit
logger.count		ops
logger.logged		ops
responseGenerateTime.active		threads
responseGenerateTime.avg		
responseGenerateTime.completed		
responseGenerateTime.maxActive		
responseGenerateTime.maxTime		
responseGenerateTime.minTime		
responseGenerateTime.time		

Table C-45 shows the DMS internal transtrace metrics.

Table C-45 DMS-Internal Transtrace Metrics

Metric	Description	Unit
expireMessages.avg		
expireMessages.completed		
expireMessages.maxActive		
expireMessages.maxTime		
expireMessages.minTime		
expireMessages.time		
messageCount.value		
pendingMessageCount.value		
s_debugEnabled.value		
s_dumpEnabled.value		
s_ecidEnabled.value		
s_transTraceEnabled.value		
storeSize.value		

OC4J Performance Metrics

This appendix covers the following metrics:

- [JTA Resource Metrics](#)
- [JCA Metrics](#)
- [OC4J J2EE Application Metrics](#)
- [OC4J JMS Metrics](#)
- [OC4J Task Manager Metrics](#)
- [Java Object Cache JOC Metrics](#)

D.1 JTA Resource Metrics

Table D-1 shows the JTA resource metrics.

The metric table type is `JTAResource`.

Table D-1 /oc4j/JTAResource Metrics

Metric with JSR-77 JTA Resource	Description	Unit
ActiveCount jtaresource_active	Total count of active transactions. A consistently high value can indicate a heavy load on a server.	ops
AverageCommitTime jtaresource_averageCommit	Average commit time of all transactions. This is the average of the <code>jtaresource_performTransaction</code> values, however, this is a mean average so there may be spikes in the system indicating other issues as well.	msecs
CommittedCount jtaresource_committed	Total count of transactions which have committed.	ops
HeuristicCommittedCount jtaresource_heuristicCommitted	Total count of heuristically committed transactions. A high value suggests the system or application may not be automated enough, for example, too much system administration in general or inadequate handling of transaction architecture, or a particular issue has occurred which required extensive administration. This is due to a subordinate <code>TransactionManager</code> and not a resource manager being rolledback while in the prepared state.	ops
HeuristicCount jtaresource_heuristic	Total count of all heuristically rolledback and committed transactions. See comments for <code>heuristicCommittedCount</code> and <code>heuristicRolledbackCount</code> .	ops
HeuristicMixedExceptionCount jtaresource_heuristicMixedException	Total count of <code>HeuristicMixedExceptions</code> encountered. A high value can indicate a high number of potentially non-ACID outcomes resulting from inconsistent or inappropriate administrative intervention.	ops
HeuristicRollbackExceptionCount jtaresource_heuristicRollbackException	Total count of <code>HeuristicRollbackExceptions</code> encountered. A high value suggests the system or application may not be automated enough, for example, too much system administration in general or inadequate handling of transaction architecture, or a particular issue has occurred which required extensive administration. Unlike the <code>rolledbackDueToAdminCount</code> metric which indicates administrative rollback at the root transaction manager level while a transaction is active, this is due to either a subordinate <code>TransactionManager</code> or resource manager being rolledback while in the prepared state.	ops
HeuristicRolledbackCount jtaresource_heuristicRolledback	Total count of heuristically rolledback transactions. A high value here suggests the system or application may not be automated enough for example, too much system administration in general or inadequate handling in transaction architecture, or a particular issue has occurred which required extensive administration. This is due to a subordinate <code>TransactionManager</code> and not a resource manager being committed while in the prepared state.	ops
IllegalStateExceptionCount jtaresource_illegalStateException	Total count of <code>IllegalStateExceptions</code> encountered. A high value should be rare and should only be possible as a result of prior administrative intervention.	ops
PerformTransaction jtaresource_performTransaction	Time from begin to end of the transaction. This is useful as a high-level indicator of performance issues, however, as it is a measure from the begin to the end of the transaction only, anything that occurs within this time could be the cause. For example, application logic, database activity, JMS activity, or transaction processing.	msecs
RollbackCompletion jtaresource_rollbackCompletion	Time required for a rollback completion. A high value indicates delays in the rollback calls on resource managers which may be a result of network latency or resource manager issues.	

Table D-1 (Cont.) /oc4j/JTA Resource Metrics

Metric with		
JSR-77 JTA Resource	Description	Unit
RollbackExceptionCount jtaresource_rollbackException	Total count of RollbackExceptions encountered. A high value can indicate an issue in the system (for example, a database is down) which results in performance degradation. This can result from both a direct internal system failure or from the application calling <code>setRollbackOnly</code> for some reason. Suggest examining the fine-grained rollback cause counts and logs for the cause of the rollbacks as well as looking into application code which calls <code>setRollbackOnly</code> .	ops
RolledbackCount jtaresource_rolledback	Total count of transactions which have rolledback. A high value can indicate an issue in the system (for example a database is down) which results in performance degradation.	ops
RolledbackDueToAdminCount jtaresource_rolledbackDueToAdmin	Total count of transactions that have rolledback due to administrative action. A high value here suggests the system or application may not be automated enough. For example, too much system administration in general or inadequate handling of transaction architecture, or a particular issue has occurred which required extensive administration.	ops
RolledbackDueToAppCount jtaresource_rolledbackDueToApp	Total count of transactions that have rolledback due to the application calling <code>setRollbackOnly</code> or <code>rollback</code> explicitly. A high value here can occur for any reason, but most often occurs due to some handled exception within an application, for example, <code>SQLException</code> during database update. Suggest looking into application code which calls <code>setRollbackOnly</code> or <code>rollback</code> to see why it is doing so.	ops
RolledbackDueToResourceCount jtaresource_rolledbackDueToResource	Total count of transactions that have rolledback due to and error in an enlisted resource. A high value here can indicate an issue with one or more resource managers, for example, database or the network connection between OC4J and these resources.	ops
RolledbackDueToTimedOutCount jtaresource_rolledbackDueToTimedOut	Total count of transactions that have rolledback due to timeout. A high number can indicate any number of issues are causing the transaction, or activity within transactional bounds, to take too long or the timeout value specified is not flexible enough. Suggest looking into what activities within the transactions involved, which may be of a certain type or application, are taking up time or adjust the <code>transaction-timeout</code> value in the <code>transaction-manager.xml</code> configuration file.	ops
SecurityExceptionCount jtaresource_securityException	Total count of SecurityExceptions encountered. A high value, or any value greater than 0, can indicate an issue with the identity on the thread executing this.	ops
SinglePhaseCommitCompletion jtaresource_singlePhaseCommitCompletion	Time required for a single-phase commit completion. A single phase commit involves committing of a single resource only and therefore no 2PC costs, for example logging, are incurred and a large value here generally indicates an issue with the resource which is being committed, for example network latency to the database, suggest looking closer at the metrics of the resource involved in the commit.	

Table D–1 (Cont.) /oc4j/JTAResource Metrics

Metric with	Description	Unit
JSR-77 JTA Resource		
SystemExceptionCount jtaresource_systemException	Total count of <code>SystemExceptions</code> encountered. A high value here should never occur and indicates a serious failure in the system. Suggest analyzing OC4J and resource manager logs.	
TransactionSuspended jtaresource_transactionSuspended	Time transaction has been suspended. A high value here indicates that the transactions are being held in a suspended state waiting for return a call in a different or no transactional context, generally from an EJB method call, or during propagation of a transaction context. Suggest analyzing the application to determine what activity is taking place during the suspend or whether there is a network latency in the case of propagation.	
TwoPhaseCommitCompletion jtaresource_ twoPhaseCommitCompletion	Time required for a two-phase commit completion. A high value here indicates delays in the prepare and commit calls on resource managers or the transaction record logging in OC4J. Suggest looking closer at the metrics of the resources involved in the prepare and commit as well as performance settings in for transaction record logging in OC4J found in the <code>transaction-manager.xml</code> configuration file.	

D.2 JCA Metrics

Table D–2 shows the JCA metrics.

The metric table type is `jca_connection_stats`.

Table D–2 oc4j/application/OracleASjms/JCAMetrics Metrics

Metric	Description	Unit
<code>closeCount.count</code>	Number of connection handles closed count.	ops
<code>createCount.count</code>	Number of connection handles created count.	ops
<code>poolName.value</code>	Name of connection pool value.	pool name
<code>useTime.avg</code>	Time spent using a connection.	time
<code>useTime.completed</code>	Time spent using a connection.	ops
<code>useTime.maxTime</code>	Time spent using a connection.	time
<code>useTime.minTime</code>	Time spent using a connection.	time
<code>useTime.time</code>	Time spent using a connection.	time
<code>waitTime.avg</code>	Time spent waiting for a connection to be available.	time
<code>waitTime.completed</code>	Time spent waiting for a connection to be available.	ops
<code>waitTime.maxTime</code>	Time spent waiting for a connection to be available.	time
<code>waitTime.minTime</code>	Time spent waiting for a connection to be available.	time
<code>waitTime.time</code>	Time spent waiting for a connection to be available.	time

Table D–3 shows the JCA connection pool stats metrics.

The metric table type is `jca_connection_pool_stats`.

Table D-3 `/oc4j/jca_connection_pool_stats` Metrics

Metric	Description	Unit
<code>closeCount.count</code>	Number of ManagedConnections closed.	ops
<code>createCount.count</code>	Number of ManagedConnections created.	ops
<code>errorCount.count</code>	Number of connection error events.	ops
<code>expiredCount.count</code>	Number of expired connections removed from pool.	ops
<code>freePoolSize.maxValue</code>	Number of free connections in the pool.	connections
<code>freePoolSize.minValue</code>	Number of free connections in the pool.	connections
<code>freePoolSize.value</code>	Number of free connections in the pool.	connections
<code>inactivityTimeout.value</code>	Configuration parameter: timeout for unused connections.	time
<code>inactivityTimeoutCheck.value</code>	Configuration parameter: when to check for unused connections.	
<code>initial-capacity.value</code>	Configuration parameter: number of connections to be pre-created by the pool.	ops
<code>invalidCount.count</code>	Number of invalid connections removed from pool.	ops
<code>maxPoolSize.value</code>	Configuration parameter: maximum number of connections.	connections
<code>minPoolSize.value</code>	Configuration parameter: minimum number of connections.	connections
<code>poolSize.maxValue</code>	Size of connection pool.	connections
<code>poolSize.minValue</code>	Size of connection pool.	connections
<code>poolSize.value</code>	Size of connection pool.	connections
<code>requestTimeoutCount.count</code>	Number of failed connection requests due to timeout.	ops
<code>scheme.value</code>	Scheme configuration parameter: connection pooling scheme.	
<code>useTime.avg</code>	Time spent using a connection	time
<code>useTime.completed</code>	Time spent using a connection	ops
<code>useTime.maxTime</code>	Time spent using a connection	time
<code>useTime.minTime</code>	Time spent using a connection	time
<code>useTime.time</code>	Time spent using a connection	time
<code>waitTime.avg</code>	Time spent waiting for a connection to be available	time
<code>waitTime.completed</code>	Time spent waiting for a connection to be available	ops
<code>waitTime.maxTime</code>	Time spent waiting for a connection to be available	time
<code>waitTime.minTime</code>	Time spent waiting for a connection to be available	time
<code>waitTime.time</code>	Time spent waiting for a connection to be available	time
<code>waitTimeout.value</code>	Configuration parameter: timeout waiting for a connection in <code>fixed_wait</code> scheme	
<code>waitingThreadCount.active</code>	Number of threads waiting for connection.	ops
<code>waitingThreadCount.avg</code>	Number of threads waiting for connection active.	
<code>waitingThreadCount.completed</code>	Number of threads waiting for connection active.	ops
<code>waitingThreadCount.maxActive</code>	Number of threads waiting for connection active.	threads
<code>waitingThreadCount.maxTime</code>	Number of threads waiting for connection active.	time
<code>waitingThreadCount.minTime</code>	Number of threads waiting for connection active.	time
<code>waitingThreadCount.time</code>	Number of threads waiting for connection active.	time

D.3 OC4J J2EE Application Metrics

This section lists the OC4J J2EE application related metrics.

This section covers the following metrics:

- [Web Module Metrics](#)
- [Web Context Metrics](#)
- [OC4J Servlet Metrics](#)
- [OC4J JSP Metrics](#)
- [OC4J EJB Metrics](#)
- [OC4J OPMN Info Metrics](#)
- [OC4J Work Management Pool Metrics](#)

D.3.1 Web Module Metrics

There is one set of metrics for each Web module within each J2EE application.

[Table D-4](#) shows the web module metrics.

The metric table type is `oc4j_web_module`.

Table D-4 *OC4J/application/WEBs Metrics*

Metric	Description	Units
<code>parseRequest.active</code>	Current number of threads trying to read/parse AJP or HTTP requests	
<code>parseRequest.avg</code>	Average time spent to read/parse requests	msecs
<code>parseRequest.completed</code>	Number of web requests that have been parsed	ops
<code>parseRequest.maxActive</code>	Maximum number of threads trying to read/parse AJP or HTTP requests	threads
<code>parseRequest.maxTime</code>	Maximum time spent to read/parse requests	msecs
<code>parseRequest.minTime</code>	Minimum time spent to read/parse requests	msecs
<code>parseRequest.time</code>	Total time spent to read/parse requests from the socket	msecs
<code>processRequest.active</code>	Current number of threads servicing web requests	
<code>processRequest.avg</code>	Average time spent servicing web requests	msecs
<code>processRequest.completed</code>	Number of web requests processed by this application	ops
<code>processRequest.maxActive</code>	Maximum number of threads servicing web requests	threads
<code>processRequest.maxTime</code>	Maximum time spent servicing a web request	msecs
<code>processRequest.minTime</code>	Minimum time spent servicing a web request	msecs
<code>processRequest.time</code>	Total time spent servicing this application's web requests	msecs
<code>resolveContext.active</code>	Current number of threads trying to create/find the servlet context	
<code>resolveContext.avg</code>	Average time spent to create/find the servlet context	msecs
<code>resolveContext.completed</code>	Count of completed context resolves	ops
<code>resolveContext.maxActive</code>	Maximum number of threads trying to create/find the servlet context	threads
<code>resolveContext.maxTime</code>	Maximum time spent to create/find the servlet context	msecs
<code>resolveContext.minTime</code>	Minimum time spent to create/find the servlet context	msecs
<code>resolveContext.time</code>	Total time spent to create/find the servlet context. Each web module (WAR) maps to a servlet context	msecs

D.3.2 Web Context Metrics

Table D-5 shows the web context metrics. There is one set of web context metrics for each Web context module within each J2EE application.

The metric table type is `oc4j_context`.

Table D-5 *OC4J/application/Webs/context Metrics*

Metric	Description	Unit
<code>resolveServlet.time</code>	Total time spent to create/locate servlet instances (within the servlet context). This includes the time for any required authentication.	msecs
<code>resolveServlet.completed</code>	Total Number of lookups for a servlet by OC4J	ops
<code>resolveServlet.minTime</code>	Minimum time spent to create/locate the servlet instance (within the servlet context)	msecs
<code>resolveServlet.maxTime</code>	Maximum time spent to create/locate the servlet instance (within the servlet context)	msecs
<code>resolveServlet.avg</code>	Average time spent to create/locate the servlet instance (within the servlet context)	msecs
<code>sessionActivation.active</code>	Number of active sessions	ops
<code>sessionActivation.time</code>	Total time in which sessions have been active	msecs
<code>sessionActivation.completed</code>	Number of session activations	ops
<code>sessionActivation.minTime</code>	Minimum time a session was active	msecs
<code>sessionActivation.maxTime</code>	Maximum time a session was active	msecs
<code>sessionActivation.avg</code>	Average session lifetime	msecs
<code>service.time</code>	Total time spent servicing requests. The service metrics for the servlet include any time spent in the calls to the database. If you need to determine just the <code>oc4j</code> service time, subtract the appropriate execution times.	msecs
<code>service.completed</code>	Total number of requests serviced	ops

D.3.3 OC4J Servlet Metrics

Table D-6 shows the servlet metrics. There is one set of servlet metrics for each servlet in each Web module within each J2EE application.

The metric table type is `oc4j_servlet`.

Table D-6 *OC4J/application/WEBS/context /SERVLETS/servlet Metrics*

Metric	Description	Units
<code>service.active</code>	Current number of threads servicing this servlet	threads
<code>service.avg</code>	Average time spent in servicing the servlet	msecs
<code>service.completed</code>	Total number of calls to service()	
<code>service.maxActive</code>	Maximum number of threads servicing this servlet	threads
<code>service.maxTime</code>	Maximum time spent on a servlet's service() call	ops
<code>service.minTime</code>	Minimum time spent on a servlet's service() call	msecs
<code>service.time</code>	Total time spent on the servlet's service() call	msecs

D.3.4 OC4J JSP Metrics

D.3.4.1 JSP Runtime Metrics

Table D-7 shows the JSP metrics. There is one set of JSP metrics for each Web context for each J2EE application.

The metric table type is `oc4j_jspExec`.

Table D-7 *OC4J/application/WEBs/context /JSP Metrics*

Metric	Description	Unit
<code>processRequest.time</code>	Time spent processing requests for JSPs Only used for Context/Application name	msecs
<code>processRequest.completed</code>	Number of requests for JSPs processed by this application	ops
<code>processRequest.minTime</code>	Minimum time spent processing requests for JSPs	msecs
<code>processRequest.maxTime</code>	Maximum time spent processing requests for JSPs	msecs
<code>processRequest.avg</code>	Average time spent processing requests for JSPs	msecs
<code>processRequest.active</code>	Current number of active requests for JSPs	ops

D.3.4.2 JSP Metrics

Table D-8 shows the JSP metrics. There is one set of metrics for each JSP in each Web module.

The metric table types are `oc4j_jsp(threadsafe=true)` and `oc4j_jsp(threadsafe=false)`.

To list these metrics using `dmstool`, enclose the metric table type in quotation marks.

For example:

```
dmstool -table "oc4j_jsp(threadsafe=true)"
```

Table D-8 *OC4J/application/WEBs/context /JSPjsp_name Metrics*

Metric	Description	Unit
<code>activeInstances.value</code>	Number of active instances. Only used when <code>threadsafe=false</code>	instances
<code>availableInstances.value</code>	Number of available (that is, created) instances. This value is only provided when <code>threadsafe=false</code> .	instances
<code>service.active</code>	Current number of active requests for the JSP	
<code>service.avg</code>	Average time spent servicing the JSP	msecs
<code>service.completed</code>	Number of requests for JSPs processed by this JSP	ops
<code>service.maxTime</code>	Maximum time spent servicing the JSP	msecs
<code>service.minTime</code>	Minimum time spent servicing the JSP	msecs
<code>service.time</code>	Time to serve a JSP (that is, actual execution time of the JSP)	msecs

D.3.5 OC4J EJB Metrics

D.3.5.1 OC4J EJB Session Bean Metrics

Table D-9 shows the EJB Session bean metrics showing information on each session bean.

The metric table type is `oc4j_ejb_session_bean`.

Table D–9 OC4J EJB Session Bean Metrics

Metric	Description	Unit
<code>session-type.value</code>	Provides information on the session type: <code>Stateless</code> or <code>Stateful</code>	String
<code>transaction-type.value</code>	Provides information on the transaction type: <code>Container</code> or <code>Bean</code>	String

D.3.5.2 EJB Entity Bean Metrics

Table D–10 shows the entity bean metrics. Oracle Application Server provides a set of these metrics for each type of bean in each EJB jar file in each J2EE application.

The metric table type is `oc4j_ejb_entity_bean`.

Table D–10 OC4J/application/EJBs/ejb-jar-module/ejb-name Metrics

Metric	Description	Unit
<code>transaction-type.value</code>	Possible values: <code>container</code> or <code>bean</code>	
<code>session-type.value</code>	Possible values: <code>stateful</code> or <code>stateless</code>	
<code>bean-type.value</code>	Possible values: <code>session</code> or <code>entity bean</code>	
<code>exclusive-write-access.value</code>	Possible values: <code>true</code> or <code>false</code>	
<code>isolation.value</code>	Possible values: <code>serializable</code> , <code>uncommitted</code> , <code>committed</code> , <code>repeatable_read</code> , <code>none</code> , <code>DB-determined</code> The value is <code>DB-determined</code> when the <code>isolation</code> attribute is omitted.	
<code>persistence-type.value</code>	Possible values: <code>container</code> or <code>bean</code> or	

D.3.5.3 EJB Method Metrics

Table D–11 shows the EJB method metrics. There is one set of EJB method metrics for each method within each type of EJB bean.

The metric table type is `oc4j_ejb_method`.

The `client.*` metrics show values for the actual implementation of the method. The `wrapper.*` metrics show values for the wrapper that was automatically generated for the method.

See Also: Chapter 6, "Advanced EJB Subjects" in *Oracle Containers for J2EE Enterprise JavaBeans Developer's Guide* for information on automatically generated wrappers.

Table D–11 OC4J/application/EJBs/ejb-jar-module/ejb-name/method-name Metrics

Metric	Description	Unit
<code>client.active</code>	Current number of threads accessing the actual implementation of this method	ops
<code>client.avg</code>	Average time spent inside the actual implementation of this method	msecs
<code>client.completed</code>	Number of requests for beans processed by this application	ops
<code>client.maxActive</code>	Maximum number of threads accessing the actual implementation of this method	ops
<code>client.maxTime</code>	Maximum time spent inside the actual implementation of this method	msecs
<code>client.minTime</code>	Minimum time spent inside the actual implementation of this method	msecs
<code>client.time</code>	Time spent inside the actual implementation of this method	msecs

Table D–11 (Cont.) OC4J/application/EJBs/ejb-jar-module/ejb-name/method-name Metrics

Metric	Description	Unit
ejbPostCreate.active	Current number of threads executing <code>ejbPostCreate</code>	ops
ejbPostCreate.avg	Average time spent in <code>ejbPostCreate</code>	msecs
ejbPostCreate.completed	Number of times this <code>ejbPostCreate</code> has been called	ops
ejbPostCreate.maxTime	Maximum time spent in <code>ejbPostCreate</code>	msecs
ejbPostCreate.minTime	Minimum time spent in <code>ejbPostCreate</code>	msecs
ejbPostCreate.time	Time spent in the <code>ejbPostCreate</code> method (entity beans)	msecs
trans-attribute.value	Transaction attribute. Possible values: <code>NotSupported</code> , <code>Supports</code> , <code>RequiresNew</code> , <code>Mandatory</code> , and <code>Never</code>	
wrapper.active	Current number of threads accessing the automatically generated wrapper method	
wrapper.avg	Average time spent inside the automatically generated wrapper method	msecs
wrapper.completed	Number of requests for beans processed by this application	ops
wrapper.maxActive	Maximum number of threads that access the wrapper	ops
wrapper.maxTime	Maximum time spent inside the automatically generated wrapper method	msecs
wrapper.minTime	Minimum time spent inside the automatically generated wrapper method	msecs
wrapper.time	Time spent inside the automatically generated wrapper method. Note: Not all the wrapper methods invoke the actual bean implementation at runtime (for example, <code>create</code> method in a stateless bean). This means that the time spent in the wrapper code could be less than the time spent in the bean implementation	msecs

D.3.5.4 EJB Stateless Bean Metrics

Table D–12 shows the EJB stateless bean metrics.

The metric table type is `oc4j_ejb_stateless_bean`.

Table D–12 OC4J EJB Stateless Bean Metrics

Metric	Description	Unit
pooled.count	Number of Pooled Instances count	count
pooled.maxValue	Number of Pooled Instances	ops
pooled.minValue	Number of Pooled Instances	ops
pooled.value	Number of Pooled Instances	ops
ready.count	Number of Ready Instances	count
ready.maxValue	Number of Ready Instances	ops
ready.minValue	Number of Ready Instances	ops
ready.value	Number of Ready Instances	
session-type.value	Session type	

D.3.5.5 EJB Stateful Bean Metrics

Table D–13 shows the EJB stateful bean metrics.

The metric table type is `oc4j_ejb_stateful_bean`.

Table D-13 OC4J EJB Stateful Bean Metrics

Metric	Description	
passive.count	Number of Passivated Instances	count
passive.maxValue	Number of Passivated Instances	ops
passive.minValue	Number of Passivated Instances	ops
passive.value	Number of Passivated Instances	ops
ready.count	Number of Ready Instances	count
ready.maxValue	Number of Ready Instances	ops
ready.minValue	Number of Ready Instances	ops
ready.value	Number of Ready Instances	ops
session-type.value	Session type	
transaction-type.value	Transaction	

D.3.5.6 EJB Message-Driven Bean Metrics

Table D-14 shows the message-driven bean metrics.

The metric table type is `oc4j_ejb_message-driven_bean`.

Table D-14 OC4J EJB Message-driven Bean Metrics

Metric	Description	Units
applicationExceptionCount.count	Number of application exceptions thrown	count
failedMessageDeliveryCount.count	Number of failed message deliveries	count
messageDelivery.avg	Message delivery attempts	time
messageDelivery.completed	Message delivery attempts	ops
messageDelivery.maxTime	Message delivery attempts	time
messageDelivery.minTime	Message delivery attempts	time
messageDelivery.time	Message delivery attempts	time
messageEndpointCount.value	Number of message endpoints	ops
messageEndpointType.value	Message endpoint type	Class name
pooled.count	Number of Pooled Instances	count
pooled.maxValue	Number of Pooled Instances	ops
pooled.minValue	Number of Pooled Instances	ops
pooled.value	Number of Pooled Instances	ops
ready.count	Number of Ready Instances	count
ready.maxValue	Number of Ready Instances	ops
ready.minValue	Number of Ready Instances	ops
ready.value	Number of Ready Instances	ops
startTime.value	The MDB available time for service	time
successfulMessageDeliveryCount.count	Number of successful message deliveries	count
systemExceptionCount.count	Number of SystemExceptions thrown	count
transaction-type.value	Transaction value	

D.3.6 OC4J OPMN Info Metrics

Table D–15 shows the OC4J OPMN information metrics.

The metric table type is `oc4j_opmn`.

Table D–15 OC4J OPMN Information Metrics

Metric	Description	Unit
<code>default_application_log.value</code>	Specifies the default application log file path.	
<code>ias_cluster.value</code>	Specifies the Oracle Application Server cluster name.	String
<code>ias_instance.value</code>	Specifies the Oracle Application Server instance name.	String
<code>jms_log.value</code>	Specifies the JMS log file path.	String
<code>oc4j_instance.value</code>	Specifies the OC4J instance ID.	String
<code>oc4j_island.value</code>	Specifies the OC4J island ID.	String
<code>opmn_group.value</code>	Specifies the OPMN group ID.	String
<code>opmn_sequence.value</code>	Specifies the OPMN sequence ID.	String
<code>rmi_log.value</code>	Specifies the RMI log file path name.	String
<code>server_log.value</code>	Specifies the application server log file path.	String

D.3.7 OC4J Work Management Pool Metrics

Table D–16 shows the OC4J Work management pool metrics.

The metric table type is `oc4j_workManagementPool`.

Table D–16 OC4J Work Management Pool Metrics

Metric	Description	Unit
<code>idleThreadCount</code>	Number of idle threads in the pool. This is a current thread pool state metric.	threads
<code>keepAlive</code>	Time before idle threads are removed from available pool. This is a configuration value metric	milliseconds
<code>maxPoolSize</code>	Maximum number of threads in the pool. This is a configuration value metric	threads
<code>maxQueueSize</code>	Maximum queue size. This is a configuration value metric.	work_requests
<code>minPoolSize</code>	Minimum number of threads in the pool. This is a configuration value metric.	threads
<code>queueFullEvent</code>	Number of work submission failures due to full queue. This is a current thread pool state metric.	ops
<code>queueSize</code>	Current queue size. This is a current thread pool state metric.	work_requests
<code>totalThreadCount</code>	Total number of threads in the pool. This is a current thread pool state metric.	threads
<code>workStartDuration</code>	Duration between work accepted and work started events. This is a current thread pool state metric. Waiting time is defined as the time period between the work submission is accepted and the execution of the work starts. This metric measures the duration between a work request submission is accepted by the pool and the time when a thread is allocated from the thread pool to run the work. If a thread is readily available, this would measure the processing overhead of the threadpool in finding an available thread and setting up the proper context for processing the work. If all available threads are busy handling other work requests, this time would also include the queuing time.	ops

D.4 OC4J JMS Metrics

OC4J JMS metrics are organized into metric tables and fall into two categories:

- **JMS API-level metrics:** collected on objects visible to the JMS API (for example, connections, sessions, producers, consumers, and browsers). JMS API-level metrics are collected and maintained only for Web and EJB clients (application clients also collect API-level metrics, but do so in their own JVM; these metrics are not available on the OC4J JMS server).
- **JMS Server-level metrics:** collected by the OC4J JMS server and maintained independent of client-state. JMS Server-level metrics are collected and maintained for all types of clients: Application, Web, and EJB.

Each OC4J JMS metric table (metric table type) contains metrics for instances of the same type; different instances have unique names. For each instance in a metric table, a set of metrics is collected. The names for metrics in each instance are unique IDs that OC4J JMS generates.

Instances may have one or more metrics whose value is the name of another metric instance. For example, the JMS session instances contain metrics that point to the parent containing JMS connection instance. You can use the pointers to navigate through the metrics.

A parent metric instance usually includes a counter metric indicating the number of child metrics of a certain type that have been created. Child metric instances may appear and disappear as the underlying objects are created and destroyed; the counter keeps track of the total number of such instances that were created during the lifetime of the parent.

Note: Oracle Application Server JMS metrics are available only for OC4J JMS (thus, metrics are not available for OJMS).

See Also: *Oracle Containers for J2EE Services Guide* for more information on OC4J JMS

D.4.1 JMS Metric Tables

OC4J JMS metrics are divided into three types, based on how they are updated:

1. **CTOR Metrics:** Metrics that are set in the constructor or initialization routine of the associated JMS object, and are never changed during the lifetime of the object.
2. **Normal Metrics:** Object level state metrics that are updated as soon as the associated state of the JMS object changes.
3. **Lazy Metrics:** these state metrics are updated lazily, that is, not as soon as the underlying metric value changes, but only periodically (these are typically server store metrics and are updated each time the store is cleaned up of expired messages).

[Table D-17](#) shows a summary of the organization of the OC4J JMS metric tables.

Table D–17 OC4J JMS Metric Tables

JMS Metric Table Type	Parent Table Type	Number of Instances	Description
JMSConnectionStats	JMSStats	1 per JMS connection	Statistics for the JMS connections active in this server
JMSDestinationStats	JMSStats	1 per permanent JMS destination	Statistics for each permanent JMS destination known to the OC4J JMS server
JMSDurableSubscriberStats	JMSStats	1 per JMS durable subscriber	Statistics for each JMS durable subscription known to this server
JMSMessageBrowserStats	JMSSessionStats	1 per JMS queue browser	Statistics for the JMS queue browsers in this server
JMSMessageConsumerStats	JMSSessionStats	1 per JMS message consumer	Statistics for the JMS consumers active in this server
JMSMessageProducerStats	JMSSessionStats	1 per JMS message producer	Statistics for the JMS producers active in this server
JMSPersistenceStats	JMSDestinationStats	1 per server-side persistent destination	Statistics for operations on the persistence file for each persistent destination
JMSRequestHandlerStats	JMSStats	1 per remote JMS connection	Statistics for the request handler thread servicing a remote JMS connection.
JMSSessionStats	JMSConnectionStats	1 per JMS session	Statistics for the JMS sessions active in this server
JMSStats	none	1	Statistics for the OC4J JMS Server
JMSStoreStats	JMSDestinationStats JMSTemporaryDestinationStats	1 per server-side message store	Statistics for each message store (one per queue, one per subscription per topic) on the OC4J JMS server
JMSTemporaryDestinationStats	JMSStats	1 per temporary JMS destination	Statistics for each temporary JMS destination known to the OC4J JMS server

D.4.2 JMS Stats Metric Table

[Table D–18](#) shows the JMS Stats metrics.

The metric table type is `JMSStats`.

Table D–18 JMSStats Metric Table

Metric	Description	Update	Unit
<code>activeConnections</code>			
<code>activeHandlers</code>			
<code>address</code>	The hostname(s) from which the JMS server accepts remote connections	ctor	string
<code>closeConnection</code>			
<code>closeConsumer</code>			
<code>commit</code>			
<code>connections</code>			
<code>createConsumer</code>			
<code>deqMessage</code>			
<code>enqMessage</code>			
<code>host.value</code>	The explicit hostname on which the OC4J JMS server is running.	ctor	string

Table D-18 (Cont.) JMSStats Metric Table

Metric	Description	Update	Unit
listMessages			
messageCommitted			
messageCount			
messageDequeued			
messageDiscarded			
messageEnqueued			
messageExpired			
messagePagedIn			
messagePagedOut			
messageRecovered			
messageRolledback			
oc4j.jms.checkPermissions			
oc4j.jms.debug	Value of the oc4j.jms.debug OC4J JMS control knob	ctor	bool
oc4j.jms.forceRecovery	Value of the oc4j.jms.forceRecovery OC4J JMS control knob	ctor	bool
oc4j.jms.j2ee14			
oc4j.jms.listenerAttempts	Value of the oc4j.jms.listenerAttempts OC4J JMS control knob	ctor	int
oc4j.jms.maxOpenFiles	Value of the oc4j.jms.maxOpenFiles OC4J JMS control knob	ctor	int
oc4j.jms.messagePoll	Value of the oc4j.jms.messagePoll OC4J JMS control knob	ctor	msecs
oc4j.jms.noDms	Value of the oc4j.jms.noDms OC4J JMS control knob	ctor	bool
oc4j.jms.noJmx			
oc4j.jms.pagingThreshold			
oc4j.jms.printStackTrace			
oc4j.jms.reconnectAttempts			
oc4j.jms.reconnectWait			
oc4j.jms.rememberALLXids			
oc4j.jms.saveAllExpired	Value of the oc4j.jms.saveAllExpired OC4J JMS control knob	ctor	bool
oc4j.jms.serverPoll	Value of the oc4j.jms.serverPoll OC4J JMS control knob	ctor	msecs
oc4j.jms.socketBufsize	Value of the oc4j.jms.socketBufsize OC4J JMS control knob	ctor	int
peekMessage			
pendingMessageCount			
port.value	The TCP/IP port on which the JMS server listens for incoming connections	ctor	int
registerConnection			
rollback			
startTime	System.currentTimeMillis() when the OC4J JMS server was started	ctor	msecs

Table D–18 (Cont.) JMSStats Metric Table

Metric	Description	Update	Unit
stats			
storeSize			
taskManagerInterval	The scheduling interval of the OC4J task manager (and the scheduling interval for the OC4J JMS expiration task)	ctor	msecs

D.4.3 JMS Request Handler Stats

Table D–19 shows the JMS Request Handler Stats.

The metric table type is `JMSRequestHandlerStats`.

Table D–19 JMSRequestHandlerStats Metrics

Metric	Description	Update	Unit
address.value	The hostname from which the remote connection originates (may be an implicit, special address)	ctor	string
connectionID.value	The ID of the <code>JMSConnectionStats</code> instance	ctor	string
host.value	The explicit hostname from which the remote connection originates	ctor	string
port.value	The TCP/IP port from which the remote connection originates	ctor	int
startTime.value	<code>System.currentTimeMillis()</code> when the request handler was started	ctor	string

D.4.4 JMS Connection Stats

Table D–20 shows the JMS Connection Stats.

The metric table type is `JMSConnectionStats`.

Table D–20 JMSConnectionStats Metrics

Metric	Description	Update	Unit
address.value	The implicit hostname of the remote JMS server host for this connection as specified in the connection factory used to create this connection; set only for non-local connections.	ctor	string
clientID.value	The administratively configured (for ctor) or programmatically set (for normal) clientID for this connection	ctor/normal	string
domain.value	The JMS domain ("queue", "topic", or "unified") of this connection	ctor	string
exceptionListener.value	The stringified name of the current exception listener for this connection	normal	string
host.value	The explicit hostname of the remote JMS server host for this connection; set only for non-local connections	ctor	string
isLocal.value	"true" if and only if the JMS connection is local to the OC4J JMS server in the same JVM	ctor	boolean
isXA.value	"true" if and only if the connection is in XA mode	ctor	boolean
port.value	The remote JMS server port for this connection; set only for non-local connections	ctor	int
startTime.value	<code>System.currentTimeMillis()</code> when this connection was created	ctor	msecs
user.value	The user identity for this connection	ctor	string
method-name	An interval timer metric (PhaseEvent Sensor) for every major method call in this connection object.	normal	

D.4.5 JMS Session Stats

Table D–21 shows the JMS Session Stats.

The metric table type is `JMSSessionStats`.

Table D–21 *JMSSessionStats Metrics*

Metric	Description	Update	Unit
<code>acknowledgeMode.value</code>	The acknowledge mode of this session. The valid modes are: <code>AUTO_ACKNOWLEDGE</code> , <code>CLIENT_ACKNOWLEDGE</code> , <code>DUPS_OK_ACKNOWLEDGE</code> , and <code>SESSION_TRANSACTED</code> .	ctor	string
<code>domain.value</code>	The JMS domain ("queue", "topic", or "unified") of this session	ctor	string
<code>isXA.value</code>	"true" if and only if the session is in XA mode	ctor	boolean
<code>sessionListener.value</code>	The stringified name of the current distinguished listener for this session	normal	string
<code>startTime.value</code>	<code>System.currentTimeMillis()</code> when this session was created	ctor	msecs
<code>transacted.value</code>	"true" if and only if the session is transacted	ctor	boolean
<code>txid.value</code>	The integer count of the current local transaction associated with this session; the counter is increment each time a local transaction is committed/rolledback; not set for non-transacted session	normal	int
<code>xid.value</code>	The Xid of the current distributed transaction associated with this session; set to a null/empty string when in a local transaction mode; not set if the session never participates in a global transaction	normal	string
<i>method-name</i>	An interval timer metric (PhaseEvent Sensor) for every major method call in this session object	normal	

D.4.6 JMS Message Producer Stats

[Table D–22](#) shows the JMS Producer Stats.

The metric table type is `JMSProducerStats`.

Table D–22 *JMSProducerStats Metrics*

Metric	Description	Update	Unit
<code>deliveryMode.value</code>	The current delivery mode of this producer. The valid delivery mode values are: <code>PERSISTENT</code> and <code>NON_PERSISTENT</code> .	normal	string
<code>destination.value</code>	The name of the identified destination for this producer; null/empty for an unidentified producer	ctor	string
<code>disableMessageID.value</code>	The value is <code>true</code> when message IDs are disabled for the producer	normal	boolean
<code>disableMessageTimestamp.value</code>	The value is <code>true</code> when message timestamps are disabled for the producer	normal	boolean
<code>domain.value</code>	The JMS domain ("queue", "topic", or "unified") of this producer	ctor	string
<code>priority.value</code>	The current priority of this producer	normal	int
<code>startTime.value</code>	<code>System.currentTimeMillis()</code> when this producer was created	ctor	msecs
<code>timeToLive.value</code>	The current <code>timeToLive</code> of this producer	normal	msecs
<i>method-name</i>	A phase timer (PhaseEvent Sensor) metric for every major method call in this producer object	normal	

D.4.7 JMS Message Browser Stats

[Table D–23](#) shows the JMS Browser Stats.

The metric table type is `JMSBrowserStats`.

Table D–23 JMSBrowserStats Metrics

Metric	Description	Update	Unit
<code>destination.value</code>	The name of the destination for this browser	ctor	string
<code>selector.value</code>	The message selector for this browser; null/empty string if unspecified	ctor	string
<code>startTime.value</code>	<code>System.currentTimeMillis()</code> when this browser was created	ctor	msecs
<i>method-name</i>	An interval timer metric (PhaseEvent Sensor) for every major method call in this browser object; calls to "hasMoreElements" and "nextElement" are made on individual enumeration objects, but counted as PhaseEvents in the browser object to simplify data collection, multiple enumerations can be active on the same browser	normal	

D.4.8 JMS Message Consumer Stats

[Table D–24](#) shows the JMS Message Consumer Stats.

The metric table type is `JMSMessageConsumerStats`.

Table D–24 JMSMessageConsumerStats

Metric	Description	Update	Unit
<code>destination.value</code>	The name of the destination for this consumer	ctor	string
<code>domain.value</code>	The JMS domain ("queue", "topic", or "unified") of this consumer	ctor	string
<code>messageListener.value</code>	The stringified name of the current message listener for this consumer	normal	string
<code>name.value</code>	The name of the durable subscriber for this consumer; set only for durable topic subscriptions	ctor	string
<code>noLocal.value</code>	The <code>noLocal</code> setting of a subscription; set only for topic consumers	ctor	boolean
<code>selector.value</code>	The message selector for this consumer; null/empty string if unspecified	ctor	string
<code>startTime.value</code>	<code>System.currentTimeMillis()</code> when this consumer was created	ctor	msecs
<i>method-name</i>	An interval timer metric (PhaseEvent Sensor) for every major method call in this consumer object	normal	

D.4.9 JMS Durable Subscription Stats

[Table D–25](#) shows the JMS Durable Subscription Stats.

The metric table type is `JMSDurableSubscriptionStats`.

Table D–25 JMSDurableSubscriptionStats Metrics

Metric	Description	Update	Unit
<code>clientID.value</code>	The <code>clientID</code> associated with this durable subscriptions	ctor	string
<code>destination.value</code>	The name of the topic for this durable subscription	ctor	string
<code>isActive.value</code>	"true" if and only if the durable subscription is currently active (being used by a consumer)	normal	boolean
<code>name.value</code>	The user-provided name of the durable subscription	ctor	string
<code>noLocal.value</code>	The <code>noLocal</code> flag for this durable subscription	ctor	boolean
<code>selector.value</code>	The JMS message selector for this durable subscription	ctor	string

D.4.10 JMS Destination Stats

[Table D–26](#) shows the JMS Destination Stats metrics

The metric table type is `JMSDestinationStats`.

Table D–26 JMSDestinationStats Metrics

Metric	Description	Update	Unit
domain.value	JMS domain, "queue" or "topic", of the destination	ctor	string
name.value	The configured name of the destination. As defined in <code>jms.xml</code>	ctor	string
locations.value	A comma-delimited list of JNDI names bound to the destination. As defined in <code>jms.xml</code>	ctor	string
<i>method-name</i>	An interval timer metric (PhaseEvent Sensor) for every major method call in the destination object	normal	

D.4.11 JMS Temporary Destination Stats

Table D–27 shows the JMS Temporary Destination Stats.

The metric table type is `JMSTemporaryDestinationStats`.

Table D–27 JMSTemporaryDestinationStats Metrics

Metric	Description	Update	Unit
connectionID.value	The ID of the <code>JMSConnectionStats</code> instance from which this temporary destination was created	ctor	string
domain.value	JMS domain, for example "queue" or "topic", of the destination	ctor	string
<i>method-name</i>	An interval timer metric (PhaseEvent Sensor) for every major method call in the destination object	normal	

D.4.12 JMS Store Stats

Table D–28 shows the JMS StoreStats metric table.

The metric table type is `JMSStoreStats`.

Table D–28 JMSStoreStats Metric

Metric	Description	Update	Unit
destination.value	A pretty-printed name of the JMS destination associated with this message store	ctor	string
messageCount.value	Total number of messages contained in this store	lazy	int
messageDequeued.count	Total number of message dequeues (transacted or otherwise)	normal	ops
messageDiscarded.count	Total number of message discarded after the rollback of an enqueue	normal	ops
messageEnqueued.count	Total number of message enqueues (transacted or otherwise)	normal	ops
messageExpired.count	Total number of message expirations	normal	ops
messagePagedIn.count	Total number of message bodies paged in	normal	ops
messagePagedOut.count	Total number of message bodies paged out	normal	ops
messageRecovered.count	Total number of messages recovered (either from a persistence file, or after the rollback of a dequeue)	normal	ops
pendingMessageCount.value	Total number of messages part of an enqueue/dequeue of an active transaction	lazy	int
storeSize.value	Total size, in bytes, of the message store.	lazy	bytes
<i>method-name</i>	An interval timer metric (PhaseEvent Sensor) for every major method call in the message store object	normal	

The following identity holds:

$\text{messageCount} = \text{messageRecovered} + \text{messageEnqueued} - \text{messageDequeued} - \text{messageDiscarded} - \text{messageExpired}$

If a message is both enqueued and dequeued in the same transaction, the `messageEnqueued` and `messageDequeued` events occur, but the `messageRecovered` and `messageDiscarded` events do not.

D.4.13 JMS Persistence Stats

Table D–29 shows the JMS Persistence Stats.

The metric table type is `JMSPersistenceStats`.

Table D–29 *JMSPersistenceStats Metrics*

Metric	Description	Update	Unit
<code>destination.value</code>	A pretty-printed name for the JMS destination associated with this persistence file	ctor	string
<code>holePageCount.value</code>	The number of 512b pages currently free in this file	normal	int
<code>isOpen.value</code>	"true" iff the persistence file descriptor is currently open (for LRU caching)	normal	boolean
<code>lastUsed.value</code>	<code>System.currentTimeMillis()</code> when this persistence file was last used (for LRU caching)	normal	msecs
<code>persistenceFile.value</code>	The absolute path name of the persistence file used for this persistent destination. This value differs depending on the operating system where OC4J is running.	ctor	string
<code>usedPageCount.value</code>	The number of 512b pages currently in use in this file	normal	int
<i>method-name</i>	An interval timer metric (PhaseEvent Sensor) for every major method call in the persistence file object	normal	

D.5 OC4J Task Manager Metrics

Table D–30 shows the OC4J Task Manager metrics.

The metric table type is `oc4j_task`.

Table D–30 *OC4J_taskManager Metrics*

Metric	Description	Unit
<code>interval.value</code>	Shows how often the task should run. The task manager executes all the tasks in a round-robin fashion. If the interval is zero, then the task manager executes the task when it is selected in the round robin.	msecs (Milliseconds)
<code>run().active</code>	Number of active threads.	threads
<code>run().avg</code>	Average time for the taskmanager to run the task	msecs
<code>run().completed</code>	Number of times the taskmanager has run the task.	ops
<code>run().maxActive</code>	Maximum number of active tasks.	threads
<code>run().maxTime</code>	Maximum time for the task to run.	msecs
<code>run().minTime</code>	Minimum time for the task to run.	msecs
<code>run().time</code>	Total time spent running the task manager	msecs

D.6 Java Object Cache JOC Metrics

Table D–31 shows the top level Java Object Cache metrics.

The metric table type is `joc`.

Note: The JOC metrics are only visible when the DMS element of the `javacache.xml` configuration file is set to the value `true`.

Table D–31 JOC Java Object Cache Metrics

Metric	Description	Unit
<code>disk_size.value</code>	Total number of bytes of disk consumed by objects in the cache	bytes
<code>memory_object_count.value</code>	Total number of objects in the cache	bytes
<code>memory_size.value</code>	Total number of bytes of memory consumed by objects in the cache	bytes
<code>response_q_size.value</code>	Response Queue size	ops
<code>task_count.value</code>	Total number of async tasks	ops
<code>time_q_size.value</code>	Time Queue size	ops
<code>worker_thread_count.value</code>	Total number of worker threads	threads

Table D–32 shows the java object cache region metrics.

The metric table type is `java_cache_region`.

Table D–32 Java Cache Region Metrics

Metric	Description	Unit
<code>disk_count.value</code>	Total number of objects in the region on disk	
<code>disk_size.value</code>	Total number of bytes of disk consumed by objects in the region	
<code>disk_average_load_time.value</code>	The average load time for objects in the region	
<code>memory_average_load_time.value</code>	The average load time for objects in the region	
<code>memory_object_access_count.value</code>	Total number of access of objects in the region	
<code>memory_object_count.value</code>	Total number of objects in the region	
<code>memory_size.value</code>	Total number of bytes of memory consumed by objects in the region	

Index

A

access logging, 6-3

active threads

- setting the percentage of threads to be tasked to process incoming threads, 7-14

activities

- idempotent, 7-3
- setting the maximum amount in memory within the same request, 7-14
- setting the maximum completion time for an idempotent service, 7-16
- setting the maximum time for a service to successfully complete an activity, 7-21

AggreSpy

- access control, A-7
- performance monitoring, A-2, A-4
- URL, A-7
- using, A-2, A-4
- using with standalone OC4J, A-13

Application Server Control

- monitoring Oracle Application Server with, 2-2

archive

- setting a time for checking the BPEL archive, 7-20

audit events

- controlling the amount, 7-12

audit trail

- setting, 7-12
- setting details, 7-11
- setting the logging level, 7-12

audit_details table

- definition, 7-10
- for storing strings larger than the threshold setting, 7-11

audit_trail table

- controlling the amount of audit events logged by a process, 7-12
- definition, 7-10
- for storing strings within the threshold setting, 7-11

auditDetailThreshold property

- definition, 7-11
- values, 7-12

auditLevel property

- controlling the amount of audit events logged by a process, 7-12

definition, 7-12

values, 7-12

B

BPEL

- activities

 - impacted by durable processes, 7-2
 - impacted by in-flight database storage, 7-3

- archive

 - setting a time for checking, 7-20

- dehydration store database

 - database parameters tuning, 7-26
 - redo logs performance tuning, 7-26
 - tablespace tuning, 7-26

- delivery cache

 - monitoring the size of, 7-13

- domain configuration property settings

 - definition, 7-2

- domain level

 - performance properties, 7-11
 - performance properties that cannot be edited, 7-11

- dspMaxThreads property

 - allocates WorkerBean threads, 7-5

- idempotent activity

 - definition, 7-3

- in-flight database storage

 - definition, 7-3

- invocations

 - descriptions of, 7-3

- JTA transactions

 - timeout value, 7-23

- OC4J performance tuning

 - JTA transaction timeout, 7-23

- one-way invocations

 - definition, 7-3
 - stored in the delivery cache, 7-13

- performance tuning

 - See tuning Oracle BPEL, 7-1

- pick activity

 - impacting durable processes, 7-2
 - in-flight database storage, 7-3

- process configuration property settings

 - definition, 7-2

- process level

- performance properties, 7-5
- receive activity
 - impacting durable processes, 7-2
 - in-flight database storage, 7-3
- tables
 - impacted by instance data growth, 7-10
- threading model
 - one-way invocation, 7-4
 - overview, 7-4
 - relationship with connecting pooling, 7-5
 - request-response invocation, 7-4
- transient processes
 - definition, 7-2
- tuning
 - dehydration store database performance tuning, 7-26
- wait activity
 - impacting durable processes, 7-2
 - in-flight database storage, 7-3
- worker queue connection pool
 - setting the minimum size, 7-22
- WorkerBean threads, 7-5
- bpelClasspath property
 - definition, 7-12
 - values, 7-12
- bpel.xml file
 - for setting process level performance properties, 7-5
- built-in performance metrics, 2-2

C

- cache
 - high water mark, 7-16
 - low water mark, 7-17
 - monitoring cache statistics in Oracle BPEL Control, 7-17
 - removing in-flight instance from, 7-18
 - warning about impacting performance, 7-19
- cached messages
 - viewing the number of, 7-13
- capacity, 1-6
- ci_id_range table
 - saving the instance ID range, 7-16
- com.evermind.server.ejb.TimeoutExpiredException from EJB, 3-16
- completionPersistLevel property
 - definition, 7-5
 - example of use, 7-6
 - values, 7-6
- completionPersistPolicy property
 - definition, 7-6
 - example of use, 7-7
 - values, 7-7
- concurrency
 - defined, 1-2
 - limiting, 1-6
- connection pooling
 - relationship with threading, 7-5
 - setting the invoker queue minimum size, 7-18

- setting the worker queue minimum size, 7-22
- size must be greater than or equal to sum of dspMaxThreads property value, 7-25
- contention, 1-4
 - defined, 1-2
- CPUs
 - insufficient, 1-4
 - optimizing use of, 7-15
- cube_instance table
 - definition, 7-10
 - growth impacted by completionPersistLevel property, 7-6
 - growth impacted by completionPersistPolicy property, 7-6
- cube_scope table
 - definition, 7-10
 - growth impacted by completionPersistLevel property, 7-6
 - growth impacted by completionPersistPolicy property, 7-6

D

- data source configuration, 7-24
- database
 - log file sync, 8-2
 - redo log files, 8-2
 - SGA tuning, 8-3
- database parameters
 - DB_CACHE_SIZE, 7-26
 - DB_FILE_MULTIBLOCK_READ_COUNT, 7-26
 - JOB_QUEUE_PROCESSES, 7-26
 - LOG_BUFFER, 7-26
 - PROCESSES, 7-26
 - SESSION_CACHED_CURSORS, 7-26
 - SHARED_POOL_SIZE, 7-26
 - tuning, 7-26
 - UNDO_RETENTION, 7-26
- database tables
 - audit_details, 7-11
 - audit_trail, 7-12
 - ci_id_range, 7-16
 - cube_instance, 7-6
 - cube_scope, 7-6
 - dlv_message, 7-13
 - document, 7-18, 7-19
 - growth impacted by completionPersistLevel property, 7-5
 - growth impacted by completionPersistPolicy property, 7-6
 - impacted by instance data growth, 7-10
 - work_item, 7-6
- datasourceJndi property
 - definition, 7-13
 - values, 7-13
- datasources
 - inactivity-timeout option, 3-6
 - initial-limit option, 3-6
 - max-connections attribute, 3-15
 - min-connections option, 3-6, 3-16

- num-cached-statements attribute, 3-8
- stmt-cache-size attribute, 3-8
- data-sources.xml file
 - configuring the data source entry, 7-24
 - location of, 7-24
- DB_CACHE_SIZE parameter
 - tuning, 7-26
- DB_FILE_MULTIBLOCK_READ_COUNT parameter
 - tuning, 7-26
- dehydration store database
 - in-flight database storage, 7-3
 - performance tuning, 7-26
- delivery cache
 - for storing one-way invocations, 7-13
- delivery service database tables
 - dlv_message, 7-13
 - invoke_message, 7-13
- deliveryPersistPolicy property
 - definition, 7-13
 - values, 7-13
 - warning about changing this property, 7-13
- directives
 - See also* httpd.conf directives
- dispatcher agent
 - setting the number of second between triggers of, 7-14
- dispatcher threads
 - setting the maximum number of, 7-14
 - setting the minimum number of, 7-15
- dlv_message table
 - definition, 7-10
 - for saving incoming requests, 7-13
- dlv_subscription table
 - definition, 7-10
- DMS
 - coding tips, B-16
 - conditional instrumentation, B-15
 - Event sensors, B-4
 - using, B-10
 - getSensorWeight, B-15
 - instrumentation
 - definition of, B-2
 - using, B-8
 - metrics
 - definition of, B-4
 - dumping to files, B-15
 - monitoring metrics, B-2
 - naming conventions, B-6
 - nouns, B-3, B-5
 - naming conventions, B-7
 - using, B-9
 - PhaseEvent sensors, B-4
 - using, B-10
 - rollup
 - descendents, B-24
 - refresh, B-24
 - rolled, B-24
 - rollup nouns, B-6
 - security, B-15
 - sensors, B-3
 - definition of, B-4
 - destroying, B-16
 - resetting, B-16
 - State sensors, B-4
 - using, B-11
 - terminology, B-3
 - testing metrics, B-14
 - validating metrics, B-13
- DMS metric tables, A-2, A-5
- dmstool
 - access control, A-8
 - address option, A-10, A-13
 - count option, A-10
 - dump option, A-10, A-12
 - format=xml option, A-10
 - interval option, A-10
 - list option, A-11
 - options, A-8
 - reset option, A-11
 - table option, A-11
 - using, A-8
- DNS
 - domain name server, 6-3
- document persistence threshold
 - setting, 7-18
- document table
 - storing large XML documents, 7-18
- document_ci_ref table
 - definition, 7-10
- document_dlv_msg_ref table
 - definition, 7-10
- documents
 - validating, 7-22
- domain data source JNDI name
 - setting, 7-13
- domain properties
 - auditDetailThreshold, 7-11
 - auditLevel, 7-12
 - bpelcClasspath, 7-12
 - datasourceJndi, 7-13
 - deliveryPersistPolicy, 7-13
 - dspAgentDelay, 7-14
 - dspInvokeAllocFactor, 7-14
 - dspMaxRequestDepth, 7-14
 - dspMaxThreads, 7-14
 - dspMinThreads, 7-15
 - expirationMaxRetry, 7-15
 - idempotentThreshold, 7-16
 - instanceKeyBlockSize, 7-16
 - instCacheHighWatermark, 7-16
 - instCacheLowWatermark, 7-17
 - instCachePolicy, 7-18
 - invokerQueueConnectionPoolMinSize, 7-18
 - largeDocumentThreshold, 7-18
 - minBPELWait, 7-19
 - optCacheOn, 7-19
 - optIdempotentRouting, 7-20
 - optSoapShortcut, 7-20
 - processCheckSecs, 7-20
 - relaxBpelAssignRules, 7-21

- slowPerfThreshold, 7-21
- statsLastN, 7-21
- syncMaxWaitTime, 7-22
- txDataSourceJndi, 7-22
- uddiLocation, 7-22
- validateXML, 7-22
- workerQueueConnectionPoolMinSize, 7-22
- domain transaction data source JNDI name setting, 7-22
- domains
 - improving performance and scalability, 7-14
- domain.xml file
 - for setting domain level performance properties, 7-11
- dspAgentDelay property
 - definition, 7-14
 - values, 7-14
- dspInvokeAllocFactor property
 - definition, 7-14
 - values, 7-14
- dspMaxRequestDepth property
 - definition, 7-14
 - values, 7-14
- dspMaxThreads property
 - connection pooling size must be greater than or equal to the sum of, 7-25
 - definition, 7-14
 - sum of InvokerBean and WorkerBean threads related to dspMaxThreads property value, 7-24
 - values, 7-15
- dspMinThreads property
 - definition, 7-15
 - values, 7-15
- durable processes
 - activities that impact, 7-2
 - definition, 7-2

E

- EJB configuration
 - InvokerBean threads, 7-24
 - Oracle BPEL Server, 7-23
 - WorkerBean threads, 7-24
- EJBs
 - metrics, D-8
- ErrorLog
 - directive, 6-4
- Event sensors, B-4
- eviction policy
 - setting, 7-18
- expiration calls
 - setting the maximum number of, 7-15
- expirationMaxRetry property
 - definition, 7-15
 - values, 7-15

F

- FAST_START_MTTT_TARGET parameter, 8-3

- functional demand, 1-6

G

- garbage collection
 - monitoring with the visual garbage collection (VC) tool, 7-17, 7-19
- garbage collection options, 3-4

H

- hash
 - defined, 1-2
- heap size
 - setting, 3-3
- HostNameLookups
 - directive, 6-3
- httpd.conf
 - directives
 - ErrorLog, 6-4
 - HostNameLookups, 6-3
 - KeepAlive, 6-3
 - KeepAliveTimeout, 6-3
 - ListenBacklog, 6-2
 - LogLevel, 6-4
 - MaxClients, 3-7, 6-2
 - MaxKeepAliveRequests, 6-3
 - MaxRequestsPerChild, 6-2
 - MaxSpareServers, 6-2
 - MinSpareServers, 6-2
 - StartServers, 6-2
 - ThreadsPerChild, 3-8
 - Timeout, 6-2

I

- idempotent property
 - definition, 7-7
 - example of use, 7-8
 - values, 7-7
- idempotent services
 - setting a routing shortcut for, 7-20
 - setting the maximum completion time, 7-16
- idempotentThreshold property
 - definition, 7-16
 - values, 7-16
- inactivity-timeout attribute, 3-6
- incoming messages
 - saving, 7-13
 - tuning the WorkerBean threads, 7-13
- incoming requests
 - saving to database tables, 7-13
- in-flight database storage
 - activities that impact, 7-3
 - pick activity, 7-3
 - receive activity, 7-3
 - wait activity, 7-3
- in-flight instances
 - removing from the cache, 7-18
 - setting the in-memory cache, 7-19
 - setting the maximum number to place in the cache

- before pruning, 7-16
 - setting the number to which the cache is pruned, 7-17
- INFO logging level, 3-10
- initial-limit attribute, 3-6
- in-memory activities
 - setting the maximum amount within the same request, 7-14
- in-memory cache
 - setting for in-flight instances, 7-19
- inMemoryOptimization property
 - definition, 7-8
 - example of use, 7-9
 - improving throughput, 7-8
 - values, 7-8
- instance cache
 - monitoring cache statistics in Oracle BPEL Control, 7-17
- instance data
 - impacting database table growth, 7-10
- instance IDs
 - controlling the instance ID range size, 7-16
- instanceKeyBlockSize property
 - definition, 7-16
 - values, 7-16
- instCacheHighWatermark property
 - definition, 7-16
 - values, 7-17
 - warning about changing, 7-16
- instCacheLowWatermark property
 - definition, 7-17
 - values, 7-17
 - warning about changing, 7-17
- instCachePolicy property
 - definition, 7-18
 - values, 7-18
- invocations
 - one-way, 7-3, 7-4
 - request-response, 7-4
 - two-way, 7-3
- invoke activities
 - using the nonBlockingInvoke property, 7-9
- invoke_message table
 - definition, 7-10
 - for saving incoming requests, 7-13
- invoker queue connection pooling
 - setting the minimum size of, 7-18
- InvokerBean threads
 - configuring, 7-24
- invokerQueueConnectionPoolMinSize property
 - definition, 7-18
 - values, 7-18

J

- J2EE
 - metrics, D-6
- Java options
 - XX+AggressiveHeap, 3-5
 - XX+DisableExplicitGC, 3-5

- Java Virtual Machine (JVM)
 - performance tuning, 7-25
- jca thread-pool, 3-15
- JDBC
 - metrics, C-7
 - stmt-cache-size, 3-8
- JOB_QUEUE_PROCESSES parameter
 - tuning, 7-26
- JSPs
 - metrics, D-7
- JTA transactions
 - reasons for transactions timing out, 7-23
 - two-way invocations, 7-4
- JVM
 - metrics, C-6
 - setting heap size, 3-3
- JVM garbage collection, 3-4
- JVM metrics
 - Properties metrics, C-6
- JVM system properties metrics, C-6

K

- KeepAlive directive, 3-7, 6-3
- KeepAlive httpd.conf directive, 6-3
- KeepAliveTimeout httpd.conf directive, 6-3

L

- large XML documents
 - impact on performance, 7-19
- largeDocumentThreshold property
 - definition, 7-18
 - values, 7-19
- latency
 - defined, 1-2
- ListenBacklog httpd.conf directive, 6-2
- listener threads
 - configuring, 7-14, 7-15
- lists
 - setting the size of request lists, 7-21
- load variances, 1-7
- local SOAP requests
 - setting a shortcut for, 7-20
- LOG_BUFFER parameter
 - tuning, 7-26
- LOG_CHECKPOINTS_TO_ALERT parameter, 8-3
- logging
 - access, 6-3
 - error, 6-4
 - performance and, 6-3
 - performance implications of, 6-3
- logging level
 - setting, 7-12
- logging levels
 - setting, 3-10
- LogLevel directive, 6-4
- logresolve utility, 6-3

M

MaxClients directive, 3-7, 6-2
max-connections attribute, 3-15
MaxKeepAliveRequests directive, 3-7
MaxKeepAliveRequests httpd.conf directive, 6-3
MaxRequestsPerChild httpd.conf directive, 6-2
MaxSpareServers httpd.conf directive, 6-2
memory
 JVM heap size, 3-3
metric table names, A-2, A-5
metric table types
 java_cache_region, D-21
 jca_connection_pool_stats, D-5
 jca_connection_stats, D-4
 JDBC_Connection, C-8
 jdbc_connection_pool_stats, C-11
 JDBC_ConnectionSource, C-9
 JDBC_DataSource, C-7
 JDBC_Driver, C-7
 JDBC_Statement, C-10
 JMSBrowserStats, D-17
 JMSConnectionStats, D-16
 JMSDestinationStats, D-18
 JMSDurableSubscriptionStats, D-18
 JMSMessageConsumerStats, D-18
 JMSPersistenceStats, D-20
 JMSProducerStats, D-17
 JMSRequestHandlerStats, D-16
 JMSSessionStats, D-17
 JMSStats, D-14
 JMSStoreStats, D-19
 JMSTemporaryDestinationStats, D-19
joc, D-20
JTAResource, D-2
JVM, C-6
 mod_oc4j_destination_metrics, C-5
 mod_oc4j_mount_pt_metrics, C-4
 mod_oc4j_request_failure_causes, C-4
 modplsql_Cache, C-12, C-13
 modplsql_DatabaseConnectionPool, C-14
 modplsql_HTTPResponseCodes, C-3, C-12
 modplsql_LastNSQLError, C-13
 modplsql_SQLErrorGroup, C-13
 oc4j_context, D-7
 oc4j_ejb_entity_bean, D-9
 oc4j_ejb_message-driven_bean, D-11
 oc4j_ejb_method, D-9
 oc4j_ejb_session_bean, D-9
 oc4j_ejb_stateful_bean, D-10
 oc4j_ejb_stateless_bean, D-10
 oc4j_jsp(threadsafe=false), D-8
 oc4j_jsp(threadsafe=true), D-8
 oc4j_jspExec, D-8
 oc4j_opmn, D-12
 oc4j_servlet, D-7
 oc4j_task, D-20
 oc4j_web_module, D-6
 oc4j_workManagementPool, D-12
 ohs_child, C-2
 ohs_module, C-3, C-4

 ohs_ossll, C-6
 ohs_responses, C-3
 ohs_server, C-2
 ohs_vhostSet, C-3
 ohs_virtualHost, C-3
 opmn_appctx, C-20
 opmn_connect, C-18, C-19
 opmn_host_statistics, C-16
 opmn_ias_instance, C-16
 opmn_oc4j_proc, C-16
 opmn_ons, C-18
 opmn_ons_topo_entry, C-19
 opmn_pm, C-15
 opmn_process, C-17
 opmn_process_set, C-17
 opmn_process_type, C-17
metrics
 acknowledgeMode.value, D-17
 activeConnections.time, D-14
 ActiveCount, D-2
 activeHandlers.time, D-14
 activeInstances.value, D-8
 activeThreadGroups.maxValue, C-6
 activeThreadGroups.minValue, C-6
 activeThreadGroups.value, C-6
 activeThreads.maxValue, C-6
 activeThreads.minValue, C-6
 activeThreads.value, C-6
 address.value, D-14, D-16
 applicationExceptionCount.count, D-11
 availableInstances.value, D-8
 AverageCommitTime, D-2
 bean-type.value, D-9
 busyChildren.value, C-2
 CacheFreeSize.value, C-9
 CacheGetConnection.avg, C-9
 CacheHit.count, C-9
 CacheMiss.count, C-9
 CacheSize.value, C-9
 cacheStatus.value, C-12, C-13
 checkcrl.time, C-6
 childFinish.count, C-2
 childStart.count, C-2
 client.active, D-9
 client.avg, D-9
 client.completed, D-9
 clientID.value, D-16, D-18
 client.maxActive, D-9
 client.maxTime, D-9
 client.minTime, D-9
 client.time, D-9
 CloseConnectionCount.value, C-11
 closeConnection.time, D-14
 closeConsumer.time, D-14
 closeCount.count, D-4, D-5
 closessl.time, C-6
 CommittedCount, D-2
 commit.time, D-14
 connection.active, C-2
 connection.avg, C-2

ConnectionCloseCount.count, C-7, C-8
 connection.completed, C-2
 ConnectionCreate.active, C-7, C-8
 ConnectionCreate.avg, C-7, C-8
 ConnectionCreate.completed, C-7, C-8
 ConnectionCreate.maxTime, C-7, C-8
 ConnectionCreate.minTime, C-7
 ConnectionCreate.time, C-7, C-8
 connectionID.value, D-16, D-19
 connection.maxTime, C-2
 connection.minTime, C-2
 ConnectionOpenCount.count, C-7, C-8
 connections, D-14
 connection.time, C-2
 connectssl.time, C-6
 connFetch.active, C-14, C-15
 connFetch.avg, C-14, C-15
 connFetch.completed, C-14, C-15
 connFetch.maxTime, C-14, C-15
 connFetch.minTime, C-14, C-15
 connFetch.time, C-14, C-15
 cpuIdle.value, C-16
 cpuTime.value, C-17
 CreateConnectionCount.value, C-11
 createConsumer.time, D-14
 createCount.count, D-4, D-5
 CreateNewStatement.avg, C-8, C-9
 CreateStatement.avg, C-8, C-9
 dataReceive.value, C-6
 dataSent.value, C-6
 decline.count, C-4
 default_application_log.value, D-12
 deliveryMode.value, D-17
 deqMessage.time, D-14
 desc.value, C-19
 Destination.value, C-4
 destination.value, D-17, D-18, D-19, D-20
 disableMessageID.value, D-17
 disableMessageTimestamp.value, D-17
 disk_average_load_time.value, D-21
 disk_Count.value, D-21
 disk_Size.value, D-21
 domain.value, D-16, D-17, D-18, D-19
 EJB, D-8
 ejbPostCreate.active, D-10
 ejbPostCreate.avg, D-10
 ejbPostCreate.completed, D-10
 ejbPostCreate.maxTime, D-10
 ejbPostCreate.minTime, D-10
 ejbPostCreate.time, D-10
 enqMessage.time, D-14
 entercache.time, C-6
 error.count, C-2, C-13
 errorCount.count, D-5
 errorDate.value, C-14
 errorRequest.value, C-14
 errorText.value, C-14
 ErrReq.count, C-4, C-5
 ErrReqNonSess.count, C-4, C-5
 ErrReqSess.count, C-4, C-5
 exceptionListener.value, D-16
 exclusive-write-access.value, D-9
 Execute.time, C-10
 expiredCount.count, D-5
 failedMessageDeliveryCount.count, D-11
 Failover.count, C-5
 Fetch.time, C-10
 freeMemory.maxValue, C-6
 freeMemory.minValue, C-6
 freeMemory.value, C-6
 freePhysicalMem.value, C-16
 freePoolSize.maxValue, D-5
 freePoolSize.minValue, D-5
 FreePoolSizeUpperBound.value, C-11
 freePoolSize.value, D-5
 getcache.time, C-6
 get.count, C-2
 handle.active, C-2, C-4
 handle.avg, C-2, C-4
 handle.completed, C-2, C-4
 handle.maxTime, C-2, C-4
 handle.minTime, C-2, C-4
 handle.time, C-2, C-4
 handshake.time, C-6
 heapSize.value, C-17
 HeuristicCommittedCount, D-2
 HeuristicCount, D-2
 HeuristicMixedExceptionCount, D-2
 HeuristicRollbackExceptionCount, D-2
 HeuristicRolledbackCount, D-2
 hits.count, C-13, C-14, C-15
 holePageCount.value, D-20
 host.value, C-19, D-14, D-16
 ias_cluster.value, D-12
 ias_instance.value, D-12
 iasCluster.value, C-17
 iasInstance.value, C-17
 idleThreadCount, D-12
 IllegalStateExceptionCount, D-2
 inactivityTimeoutCheck.value, D-5
 inactivityTimeout.value, D-5
 IncorrectReqInIt.count, C-4
 indexInSet.value, C-17
 initial-capacity.value, D-5
 internalRedirect.count, C-2
 interval.value, D-20
 invalidCount.count, D-5
 isActive.value, D-18
 isLocal.value, D-16
 isolation.value, D-9
 isOpen.value, D-20
 isXA.value, D-16, D-17
 J2EE, D-6
 JDBC_Connection_URL, C-8
 JDBC_Connection_Url, C-9
 JDBC_Connection_Username, C-8, C-9
 jms_log.value, D-12
 jobWorkerQueue.value, C-16
 JSP, D-7
 JVM, C-6

JVMCnt.value, C-5
 keepAlive, D-12
 lastConfigChange.value, C-2
 lastErrorDate.value, C-13
 lastErrorRequest.value, C-13
 lastErrorText.value, C-13
 lastUsed.value, D-20
 listMessages, D-15
 locations.value, D-19
 LogicalConnection.value, C-8, C-9
 lReq.count, C-16
 maxPoolSize, D-12
 maxPoolSize.value, D-5
 maxQueueSize, D-12
 memory_average_load_time.value, D-21
 memory_object_access_count.value, D-21
 memory_object_count.value, D-21
 memory_size.value, D-21
 memoryUsed.value, C-18
 messageCommitted, D-15
 messageCount, D-15
 messageCount.value, D-19
 messageDelivery.avg, D-11
 messageDelivery.completed, D-11
 messageDelivery.maxTime, D-11
 messageDelivery.minTime, D-11
 messageDelivery.time, D-11
 messageDequeued, D-15
 messageDequeued.count, D-19
 messageDiscarded, D-15
 messageDiscarded.count, D-19
 messageEndpointCount.value, D-11
 messageEndpointType.value, D-11
 messageEnqueued, D-15
 messageEnqueued.count, D-19
 messageExpired, D-15
 messageExpired.count, D-19
 messageListener.value, D-18
 messagePagedIn, D-15
 messagePagedIn.count, D-19
 messagePagedOut, D-15
 messagePagedOut.count, D-19
 messageRecovered, D-15
 messageRecovered.count, D-19
 messageRolledback, D-15
 minPoolSize, D-12
 minPoolSize.value, D-5
 mod_plsql, C-11
 moduleId.value, C-17
 Name.value, C-5
 name.value, D-18, D-19
 newMisses.count, C-12, C-13, C-14, C-15
 noLocal.value, D-18
 NonSessFailover.count, C-5
 notifProcessed.value, C-19
 notifProcessQueue.value, C-19
 notifReceived.value, C-19
 numChildren.value, C-2
 numMods.value, C-2, C-3
 numProcConf.value, C-17
 numProcessors.value, C-16
 oc4j_instance.value, D-12
 oc4j_island.value, D-12
 oc4j.jms.checkPermissions, D-15
 oc4j.jms.debug.value, D-15
 oc4j.jms.forceRecovery.value, D-15
 oc4j.jms.j2ee14, D-15
 oc4j.jms.listenerAttempts.value, D-15
 oc4j.jms.maxOpenFiles.value, D-15
 oc4j.jms.messagePoll.value, D-15
 oc4j.jms.noDms.value, D-15
 oc4j.jms.noJmx, D-15
 oc4j.jms.pagingThreshold, D-15
 oc4j.jms.printStackTrace, D-15
 oc4j.jms.reconnectAttempts, D-15
 oc4j.jms.reconnectWait, D-15
 oc4j.jms.rememberALLXids, D-15
 oc4j.jms.saveAllExpired.value, D-15
 oc4j.jms.serverPoll.value, D-15
 oc4j.jms.socketBufsize.value, D-15
 Oc4jUnavailable.count, C-4
 opmn_group.value, D-12
 opmn_sequence.value, D-12
 Oracle Application Server performance, C-1
 parseRequest.active, D-6
 parseRequest.avg, D-6
 parseRequest.completed, D-6
 parseRequest.maxActive, D-6
 parseRequest.maxTime, D-6
 parseRequest.minTime, D-6
 parseRequest.time, D-6
 passive.count, D-11
 passive.maxValue, D-11
 passive.minValue, D-11
 passive.value, D-11
 peekMessage, D-15
 pendingMessageCount, D-15
 pendingMessageCount.value, D-19
 PerformTransaction, D-2
 persistenceFile.value, D-20
 persistence-type.value, D-9
 pid.value, C-18
 pooled.count, D-10, D-11
 pooled.maxValue, D-10, D-11
 pooled.minValue, D-10, D-11
 pooled.value, D-10, D-11
 poolName.value, D-4
 PoolSizeLowerBound.value, C-11
 PoolSize.maxValue, C-11
 poolSize.maxValue, D-5
 poolSize.minValue, D-5
 poolSize.value, D-5
 port.value, C-19, D-15, D-16
 post.count, C-2
 priority.value, D-17
 privateMemory.value, C-18
 procDeath.count, C-16
 procDeathReplace.count, C-16
 processRequest.active, D-6, D-8
 processRequest.avg, D-6, D-8

processRequest.completed, D-6, D-8
 processRequest.maxActive, D-6
 processRequest.maxTime, D-6, D-8
 processRequest.minTime, D-6, D-8
 processRequest.time, D-6, D-8
 queueFullEvent, D-12
 queueSize, D-12
 readyChildren.value, C-2
 ready.count, D-10, D-11
 ready.maxValue, D-10, D-11
 ready.minValue, D-10, D-11
 ready.value, D-10, D-11
 receiveErrors.count, C-6
 receive.time, C-6
 registerConnection, D-15
 reqFail.count, C-16, C-17
 reqPartialSucc.count, C-16, C-17
 reqSucc.count, C-16, C-17
 request.active, C-2, C-3
 request.avg, C-2, C-3
 request.completed, C-2, C-3
 requestHandlers.count, D-15
 request.maxTime, C-2, C-3
 request.minTime, C-2, C-3
 requests.count, C-13
 request.time, C-2, C-3
 requestTimeoutCount.count, D-5
 resolveContext.active, D-6
 resolveContext.avg, D-6
 resolveContext.completed, D-6
 resolveContext.maxActive, D-6
 resolveContext.maxTime, D-6
 resolveContext.minTime, D-6
 resolveContext.time, D-6
 resolveServlet.avg, D-7
 resolveServlet.completed, D-7
 resolveServlet.maxTime, D-7
 resolveServlet.minTime, D-7
 resolveServlet.time, D-7
 response_q_size.value, D-21
 responseSize.value, C-2, C-3
 restartOnDeath.value, C-17
 rmi_log.value, D-12
 RollbackCompletion, D-2
 RollbackExceptionCount, D-3
 RolledbackCount, D-3
 RolledbackDueToAdminCount, D-3
 RolledbackDueToAppCount, D-3
 RolledbackDueToResourceCount, D-3
 RolledbackDueToTimedOutCount, D-3
 rReq.count, C-16
 run().active, D-20
 run().avg, D-20
 run().completed, D-20
 run().maxActive, D-20
 run().maxTime, D-20
 run().minTime, D-20
 run().time, D-20
 scheme.value, D-5
 SecurityExceptionCount, D-3
 selector.value, D-18
 sendErrors.count, C-6
 send.time, C-6
 server_log.value, D-12
 service.active, D-7, D-8
 service.avg, D-7, D-8
 service.completed, D-7, D-8
 service.maxActive, D-7
 service.maxTime, D-7, D-8
 service.minTime, D-7, D-8
 service.time, D-7, D-8
 SessFailover.count, C-5
 sessionActivation.avg, D-7
 sessionActivation.completed, D-7
 sessionActivation.maxTime, D-7
 sessionActivation.minTime, D-7
 sessionActivation.time, D-7
 sessionListener.value, D-17
 session-type.value, D-9, D-10, D-11
 setfixup.time, C-6
 sharedMemory.value, C-18
 SinglePhaseCommitCompletion, D-3
 SQLText.value, C-10
 staleMisses.count, C-13, C-14, C-15
 startTime.value, D-11, D-15, D-16, D-17, D-18
 opmn_process, C-18
 StatementCacheHit.count, C-8
 StatementCacheMiss.count, C-8, C-9
 stats, D-16
 status.value, C-18
 storeSize, D-16
 storeSize.value, D-19
 successfulMessageDeliveryCount.count, D-11
 SucReq.count, C-5
 SucReqNonSess.count, C-5
 SucReqSess.count, C-5
 SystemExceptionCount, D-4
 systemExceptionCount.count, D-11
 task_count.value, D-21
 taskManagerInterval.value, D-16
 time_q_size.value, D-21
 timestamp.value, C-16
 timeToLive.value, D-17
 totalMemory.maxValue, C-6
 totalMemory.minValue, C-6
 totalMemory.value, C-6
 totalPhysicalMem.value, C-16
 totalThreadCount, D-12
 transacted.value, D-17
 TransactionSuspended, D-4
 transaction-type.value, D-9, D-11
 trans-attribute.value, D-10
 TwoPhaseCommitCompletion, D-4
 txid.value, D-17
 type.value, C-18
 uid.value, C-18
 UnableToHandleReq.count, C-4
 upTime.value, C-6, C-18
 usedPageCount.value, D-20
 user.value, D-16

- useTime.avg, D-4, D-5
- useTime.completed, D-4, D-5
- useTime.maxTime, D-4, D-5
- useTime.minTime, D-4, D-5
- UseTime.time, C-11
- useTime.time, D-4, D-5
- vhostType.value, C-3
- waitingThreadCount.active, D-5
- waitingThreadCount.avg, D-5
- waitingThreadCount.completed, D-5
- waitingThreadCount.maxActive, D-5
- waitingThreadCount.maxTime, D-5
- WaitingThreadCount.maxValue, C-11
- waitingThreadCount.minTime, D-5
- WaitingThreadCount.minValue, C-11
- waitingThreadCount.time, D-5
- waitTime.avg, D-4, D-5
- waitTime.completed, D-4, D-5
- waitTime.maxTime, D-4, D-5
- waitTime.minTime, D-4, D-5
- waitTimeout.value, D-5
- WaitTime.time, C-11
- waitTime.time, D-4, D-5
- worker_thread_count.value, D-21
- workerThread.value, C-16, C-19
- workStartDuration, D-12
- wrapper.active, D-10
- wrapper.avg, D-10
- wrapper.completed, D-10
- wrapper.maxTime, D-10
- wrapper.minTime, D-10
- wrapper.time, D-10
- xid.value, D-17
- minBPELWait property
 - definition, 7-19
 - values, 7-19
- min-connections attribute, 3-6
- min-connections datasources option, 3-16
- min-instances attribute, 3-11
- MinSpareServers httpd.conf directive, 6-2
- mod_plsql metrics, C-11
- modplsql_Cache
 - metric table type, C-12, C-13
- modplsql_DatabaseConnectionPool
 - metric table type, C-14
- modplsql_HTTPResponseCodes
 - metric table type, C-12
- modplsql_LastNSQL_Error
 - metric table type, C-13
- modplsql_SQL_ErrorGroup
 - metric table type, C-13
- monitoring
 - performance statistics, 2-2

N

- naming conventions
 - DMS, B-6
- nonBlockingInvoke property
 - definition, 7-9

- example of use, 7-9
- setting if you have invoke activities in multiple
 - flow or flowN branches, 7-9
- values, 7-9
- nouns
 - creating, B-9
 - DMS, B-5
 - naming conventions, B-7
 - rollup, B-6
 - type, B-5
- num-cached-statements attribute, 3-8

O

- OC4J
 - monitoring performance statistics, 2-2
- OC4J performance tuning
 - data source configuration, 7-24
 - InvokerBean threads configuration, 7-24
 - Oracle BPEL Server EJB configuration, 7-23
 - WorkerBean thread configuration, 7-24
- OC4J thread pool
 - thread-pool element, 3-13
- oc4j_context
 - metric table type, D-7
- oc4j_ejb_entity_bean
 - metric table type, D-9
- oc4j_ejb_method
 - metric table type, D-9
- oc4j_jspExec
 - metric table type, D-8
- oc4j_servlet
 - metric table type, D-7
- oc4j_web_module
 - metric table type, D-6
- Oc4jCacheSize, 3-8
- onAlarm branch
 - setting the minimum wait time, 7-19
- opmn.xml file
 - configuring JVM parameters, 7-25
- optCacheOn property
 - definition, 7-19
 - values, 7-19
- optIdempotentRouting property
 - definition, 7-20
 - values, 7-20
- optSoapShortcut property
 - definition, 7-20
 - do not change, 7-20
 - values, 7-20
- Oracle BPEL Control
 - domain level performance properties that cannot
 - be edited, 7-11
 - for setting domain level performance
 - properties, 7-11
 - monitoring cache statistics, 7-17
 - monitoring instance cache statistics, 7-17
 - viewing statistics, 7-21
- Oracle BPEL Server
 - EJB configuration, 7-23

- InvokerBean threads configuration, 7-24
- WorkerBean threads configuration, 7-24
- Oracle Database
 - recommended for stress testing and production environments, 7-24
- Oracle HTTP Server
 - configuring with directives, 6-2
- oracle.j2ee.rmi.loadBalance property, 3-19
- orion-ejb-jar.xml file
 - configuring MDB J2EE listener threads, 7-14

P

- parameters
 - KeepAlive, 6-3
 - KeepAliveTimeout, 6-3
 - ListenBacklog, 6-2
 - MaxClients, 3-7, 6-2
 - MaxKeepAliveRequests, 6-3
 - MaxRequestsPerChild, 6-2
 - MaxSpareServers, 6-2
 - MinSpareServers, 6-2
 - Oc4jCacheSize, 3-8
 - StartServers, 6-2
 - ThreadsPerChild, 3-8
 - Timeout, 6-2
- performance
 - goals, 1-6
 - monitoring
 - native operating system, 2-3
 - network monitoring tools, 2-3
- performance properties
 - completionPersistLevel, 7-5
 - completionPersistPolicy, 7-6
 - idempotent, 7-7
 - inMemoryOptimization, 7-8
 - nonBlockingInvoke, 7-9
- persistent connections
 - KeepAlive directive, 6-3
- PhaseEvent sensors, B-4
- pick activity
 - setting the minimum wait time with the onAlarm branch, 7-19
- pool-cache-timeout attribute, 3-11
- processCheckSecs property
 - definition, 7-20
 - values, 7-21
- processes
 - durable, 7-2
 - transient, 7-2
- PROCESSES parameter
 - tuning, 7-26
- pruned instances, 7-17
 - retrieving, 7-16

R

- redo log file, 8-3
- redo log files, 8-2
- redo logs

- tuning, 7-26
- relaxBpelAssignRules property
 - definition, 7-21
 - do not use this property, 7-21
 - values, 7-21
- request lists
 - setting the size of, 7-21
- response time, 1-4
 - defined, 1-2
 - goal, 1-6
 - improving, 1-3
 - peak load, 1-7
- rollup
 - DMS, B-21

S

- scalability
 - defined, 1-2
- schema_md table
 - definition, 7-10
- server.xml file
 - setting the transaction timeout value, 7-23
- service time, 1-3, 1-4
 - defined, 1-2
- SESSION_CACHED_CURSORS parameter
 - tuning, 7-26
- SGA_MAX_SIZE parameter, 8-3
- SGA_TARGET parameter, 8-3
- SHARED_POOL_SIZE parameter
 - tuning, 7-26
- slow services
 - collecting statistics on, 7-21
- slowPerfThreshold property
 - definition, 7-21
 - values, 7-21
- SOAP requests
 - setting a shortcut for local requests, 7-20
- StartServers httpd.conf directive, 6-2
- State sensors, B-4
- statistics
 - viewing from Oracle BPEL Control, 7-21
- statsLastN property
 - definition, 7-21
 - values, 7-21
- syncMaxWaitTime property
 - definition, 7-22
 - values, 7-22
- system properties
 - DMS metrics, C-6

T

- tables
 - audit_details, 7-11
 - audit_trail, 7-12
 - ci_id_range, 7-16
 - cube_instance, 7-6
 - cube_scope, 7-6
 - dlv_message, 7-13

- document, 7-18, 7-19
- growth impacted by completionPersistLevel property, 7-5
- growth impacted by completionPersistPolicy property, 7-6
- work_item, 7-6
- tablespaces
 - tuning, 7-26
- task table
 - definition, 7-10
- think time
 - defined, 1-2
- thread pool options, 3-12
- thread-pool element, 3-13
- threads
 - InvokerBean threads configuration, 7-24
 - setting the maximum number of dispatcher threads, 7-14
 - setting the minimum number of dispatcher threads, 7-15
 - setting the percentage of active threads to be tasked to process incoming threads, 7-14
 - sum of InvokerBean and WorkerBean threads related to dspMaxThreads property value, 7-24
 - WorkerBean configuration, 7-24
- ThreadsPerChild directive, 3-8
- throughput
 - defined, 1-2
 - demand limiter and, 1-5
 - improving with the inMemoryOptimization property, 7-8
 - increasing, 1-4
- Timeout httpd.conf directive, 6-2
- TimeoutExpiredException
 - from EJB, 3-16
- transactions
 - reasons for transactions timing out, 7-23
 - setting the timeout value, 7-23
- tuning
 - Java Virtual Machine (JVM) performance tuning, 7-25
 - OC4J performance tuning, 7-23
- two-way invocations
 - definition, 7-3
 - JTA transactions, 7-4
- txDataSourceJndi property
 - definition, 7-22
 - values, 7-22

U

- uddiLocation property
 - definition, 7-22
- UNDO_RETENTION parameter
 - tuning, 7-26
- unit consumption, 1-6

V

- validateXML property
 - definition, 7-22
 - values, 7-22
- validating
 - incoming and outgoing documents, 7-22
- validation
 - setting, 7-22
- visual garbage collection (GC) tool
 - monitoring garbage collection, 7-17, 7-19

W

- wait activity
 - setting the minimum wait time, 7-19
- wait time
 - contention and, 1-4
 - defined, 1-2
 - parallel processing and, 1-3
 - setting the maximum time to wait for a result before returning, 7-22
- Windows
 - performance counters, A-14
- Windows metrics
 - enabling Performance Counters, A-14
- work_item table
 - definition, 7-10
 - growth impacted by completionPersistLevel property, 7-6
 - growth impacted by completionPersistPolicy property, 7-6
- WorkerBean threads
 - configuring, 7-24
 - tuning for incoming messages, 7-13
- workerQueueConnectionPoolMinSize property
 - definition, 7-22
 - values, 7-22

X

- XML document persistence threshold
 - setting, 7-18
- XML documents
 - impact of large documents on performance, 7-19
 - validating, 7-22
- xml format output for dmstool, A-10
- xml_document table
 - definition, 7-10
- Xms option, 3-3
- Xmx option, 3-3
- XX+AggressiveHeap option, 3-5
- XX+DisableExplicitGC option, 3-5
- XXAppendRatio option, 3-19