

Oracle® Enterprise Manager

Cloud Control Getting Started with Oracle Fusion Middleware
Management

12c Release 3 (12.1.0.3)

E24215-05

June 2013

Oracle Enterprise Manager Cloud Control Getting Started with Oracle Fusion Middleware Management 12c Release 3 (12.1.0.3)

E24215-05

Copyright © 2011, 2013, Oracle and/or its affiliates. All rights reserved.

Contributing Author: Namrata Bhakthavatsalam, Leo Cloutier, Genevieve D'Souza, Pradeep Gopal, Jacqueline Gosselin, Akshaya Govind, Alan Gray, Dan Hynes, Aravind Jayaraman, Dennis Lee, Pushpa Raghavachar

Contributor: Enterprise Manager Cloud Control Fusion Middleware Management Development Teams, Quality Assurance Teams, Customer Support Teams, and Product Management Teams

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

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

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

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

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

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

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

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

Contents

Preface	xxi
Audience	xxi
Documentation Accessibility	xxi
Related Documents	xxi
Conventions	xxi
 What's New in Getting Started with Oracle Fusion Middleware Management Manual?	xxiii
 Part I Managing Oracle Fusion Middleware	
 1 Introduction to Middleware Management	
1.1 Middleware Management with Enterprise Manager Cloud Control	1-1
1.2 Key Oracle Fusion Middleware Management Features	1-2
1.3 Managing Fusion Middleware with Fusion Middleware Control	1-3
 2 Discovering Middleware Targets	
2.1 Enabling Automatic Discovery of Fusion Middleware Targets	2-1
2.2 Discovering Targets Manually	2-8
2.2.1 Discovering a WebLogic 9.x or 10.x Domain Using Cloud Control	2-9
2.2.2 Discovering Multiple WebLogic Domains Using EM CLI	2-12
2.2.3 Discovering a Standalone Oracle HTTP Server	2-12
2.3 Discovering New or Modified Domain Members	2-13
2.3.1 Enabling Automatic Discovery of New Domain Members	2-13
2.3.2 Manually Checking for New or Modified Domain Members	2-13
2.3.3 Removing Targets	2-18
 3 Managing Middleware Targets	
3.1 Middleware Targets in Enterprise Manager	3-1
3.1.1 Oracle Fusion Middleware Components	3-2
3.1.2 Oracle Application Server Components	3-4
3.1.3 Non-Oracle Middleware Components	3-5
3.2 Monitoring Middleware Targets	3-5
3.2.1 Middleware Target Home Page	3-5

3.2.2	Target Home Page	3-7
3.2.3	Out-of-box Metrics.....	3-9
3.2.4	Analyzing Historical Performance.....	3-10
3.2.5	Setting Metric Thresholds for Alert Notifications.....	3-11
3.2.6	Monitoring Templates.....	3-11
3.2.7	Managing and Creating Blackouts	3-11
3.2.8	Extend Monitoring for Applications Deployed to WebLogic Server.....	3-12
3.2.9	Request Monitoring	3-12
3.2.9.1	Defining and Managing Transactions Being Monitored	3-13
3.3	Diagnosing Performance Problems.....	3-15
3.3.1	Using Home Pages to Diagnose Performance Issues	3-15
3.3.2	Diagnostics Snapshots.....	3-16
3.4	Administering Middleware Targets.....	3-17
3.5	Lifecycle Management	3-18
3.5.1	Managing Configurations	3-18
3.5.2	Compliance Management.....	3-19
3.5.3	Patch Management	3-20
3.5.4	Provisioning.....	3-20
3.5.4.1	Cloning from Test to Production Environments	3-20
3.5.4.2	Scaling Out Domains	3-21
3.5.4.3	Deploying / Undeploying Java EE Applications	3-21
3.6	Managing Service Levels	3-21
3.6.1	Service Dashboard	3-22
3.7	Job System	3-22
3.7.1	Log File Rotation.....	3-22
3.7.2	Log File Viewer	3-23
3.8	Topology Viewer.....	3-23
3.9	Support Workbench.....	3-23

4 Testing Application Load and Performance

4.1	Introduction to Application Replay	4-1
4.2	Testing Against Real-World Application Workloads.....	4-2
4.3	Capturing Application Workload Using RUEI	4-2
4.4	Synchronized vs. Non-Synchronized Replay	4-3
4.5	Prerequisites and Considerations.....	4-4
4.5.1	Using RUEI to Capture Application Workloads.....	4-4
4.5.2	Configuring Required User Privileges in Enterprise Manager.....	4-5
4.5.3	Using Synchronization.....	4-5
4.5.4	Setting up the Test System Database	4-6
4.5.5	Restarting the Database and Application Stack	4-6
4.5.6	Setting up the Capture Directory	4-7
4.6	Understanding the Capture and Replay Process	4-7
4.7	Creating Application Workload Captures	4-8
4.8	Monitoring the Capture Process.....	4-12
4.9	Replaying Application Workload Captures	4-13
4.9.1	Preparing to Replay Workload Captures.....	4-14
4.9.2	Understanding Replays and Replay Tasks	4-14

4.9.3	Resolving References to External Systems.....	4-14
4.9.4	Remapping URLs.....	4-15
4.9.5	Substituting Sensitive Data	4-15
4.9.6	Replaying Workload Captures	4-15
4.9.7	Analyzing Replay Results	4-20
4.10	Importing Replay Session Divergences into OpenScript.....	4-22
4.11	Troubleshooting	4-23

Part II Monitoring Exalogic Elastic Cloud, Exalytics Target, and Traffic Director

5 Monitoring Oracle Exalogic Elastic Cloud

5.1	Prerequisites to Discovering Oracle Exalogic Elastic Cloud	5-2
5.1.1	Importing Ops Center Certificate to the Oracle Management Agent Keystore	5-2
5.1.2	Critical Prerequisites For Oracle VM Manager Discovery	5-4
5.1.3	ZFS Plug-in and Oracle VM Manager Registration Preconfiguration Steps.....	5-5
5.2	Using the Exalogic Elastic Cloud Discovery Wizard.....	5-6
5.3	Upgrading Exalogic System Targets to the Version 12.1.0.3 Fusion Middleware Model	5-7
5.4	Displaying and Using the Exalogic Elastic Cloud Dashboard	5-7
5.5	Refreshing the Exalogic Elastic Cloud	5-8
5.6	Monitoring the Hardware Components of Exalogic Elastic Cloud.....	5-9
5.7	Viewing Application Deployments in Exalogic Elastic Cloud Targets	5-9
5.8	Viewing WebLogic Domains in Exalogic Elastic Cloud Targets	5-10
5.9	Viewing Coherence Clusters in Exalogic Elastic Cloud Targets.....	5-10
5.10	Viewing Hosts in Exalogic Elastic Cloud Targets.....	5-11
5.11	Exalogic Control Stack Monitoring	5-11
5.12	Visualizing Relationships Between Exalogic Software and Hardware Components....	5-12
5.13	Analyzing the Impact of Component Failures	5-12
5.14	Viewing Incidents and Status Changes Created for an Exalogic System in Ops Center as Incidents in Cloud Control	5-12
5.15	Creating Exalogic Reports	5-13
5.16	Viewing and Managing Exalogic Consumption Tracking	5-14
5.17	Viewing the vCPU Consumption Report.....	5-14
5.18	Viewing the Resource Consumption Trend.....	5-15
5.19	Exporting the vCPU Consumption Report	5-16
5.20	Configuring Exalogic Oracle Engineered System Healthchecks	5-16

6 Discovering and Monitoring an Exalytics Target

6.1	Discovering an Exalytics Target	6-1
6.1.1	Prerequisites for Discovering an Exalytics Machine Target.....	6-1
6.1.2	Discovering an Exalytics System Target	6-2
6.2	Monitoring an Exalytics System	6-2
6.2.1	Displaying and Using the Exalytics System Dashboard.....	6-3
6.2.2	Viewing the Generic Performance of an Exalytics System	6-3
6.2.2.1	Viewing the Performance Summary of an Exalytics System	6-3
6.2.2.2	Viewing the Alert History of an Exalytics System	6-4

6.2.2.3	Viewing Target Incidents of an Exalytics System.....	6-4
6.2.3	Viewing the Members Topology of an Exalytics System.....	6-5
6.2.3.1	Viewing a Configuration Topic	6-5
6.2.3.2	Creating a Custom Topology View	6-5
6.2.4	Viewing the Routing Topology of an Exalytics System	6-6
6.2.5	Refreshing an Exalytics System	6-6
6.2.6	Viewing the Job Activity of an Exalytics System	6-7
6.2.7	Administering Targets of an Exalytics System.....	6-7
6.2.7.1	Administering Target Access Privileges for an Exalytics System	6-7
6.2.7.2	Administering Target Blackouts for an Exalytics System	6-7
6.2.7.3	Removing an Exalytics System Target	6-8
6.3	Monitoring an Exalytics Machine.....	6-8
6.3.1	Displaying and Using the Exalytics Machine Homepage	6-8
6.3.2	Viewing the Generic Performance of an Exalytics Machine	6-9
6.3.2.1	Viewing the Target Alert History of an Exalytics Machine	6-9
6.3.2.2	Viewing the Target Incidents of an Exalytics Machine.....	6-9
6.3.3	Viewing the Members Topology of an Exalytics Machine	6-10
6.3.3.1	Viewing a Configuration Topic	6-10
6.3.3.2	Creating a Custom Topology View	6-11
6.3.4	Viewing the vCPU Consumption Report for an Exalytics Machine	6-11
6.3.4.1	Viewing the vCPU Consumption Report	6-12
6.3.4.2	Viewing the Resource Consumption Trend	6-12
6.3.4.3	Exporting the vCPU Consumption Report.....	6-13
6.3.5	Managing License Editions for an Exalytics Machine.....	6-13
6.3.5.1	Adding a License Edition to an Exalytics Machine Target.....	6-14
6.3.5.2	Creating a New License Edition for an Exalytics Machine	6-14
6.3.6	Viewing the Job Activity of an Exalytics Machine.....	6-14
6.3.7	Viewing the Collected Configuration of an Exalytics Machine	6-15
6.3.8	Administering the Exalytics Machine Target	6-15
6.3.8.1	Viewing the Target Monitoring Configuration of an Exalytics Machine.....	6-15
6.3.8.2	Administering Target Access Privileges for an Exalytics Machine.....	6-16
6.3.8.3	Administering Target Blackouts for an Exalytics Machine	6-16
6.3.8.4	Removing an Exalytics Machine Target.....	6-16

7 Oracle Traffic Director

7.1	Adding a Traffic Director to an Exalogic Target	7-2
7.2	About Traffic Director Configuration.....	7-2
7.2.1	Using the Traffic Director Configuration Page	7-2
7.2.2	Adding Traffic Director Target Configuration.....	7-3
7.2.2.1	Finding Configurations and Instances	7-4
7.2.2.2	Discovered Targets.....	7-5
7.2.2.3	Viewing Results	7-5
7.3	About Traffic Director Instance	7-6
7.4	About Traffic Director Refresh Flow.....	7-6
7.4.1	Adding New Targets to Newly Added Configurations	7-7
7.4.2	Adding New Targets for Newly Added Instances of Configurations.....	7-7
7.4.3	Deleting Targets of Configurations That Have Been Removed.....	7-7

7.4.4	Deleting Targets of Instances That Have Been Removed	7-7
-------	--	-----

Part III Monitoring Oracle WebLogic Domain and Oracle GlassFish Server

8 Monitoring WebLogic Domains

8.1	Updating the Agent Truststore	8-1
8.1.1	Importing a Demo WebLogic Server Root CA Certificate.....	8-2
8.1.2	Importing a Custom Root CA Certificate.....	8-2
8.2	Changing the Default AgentTrust.jks Password Using Keytool	8-2
8.3	Collecting JVM Performance Metrics for WebLogic Servers.....	8-2
8.3.1	Setting the PlatformMBeanServerUsed Attribute.....	8-3
8.3.2	Activating Platform MBeans on WebLogicServer 9.x to 10.3.2 versions.....	8-3

9 Overview of Oracle GlassFish Server Management

9.1	Before Getting Started	9-1
9.1.1	GlassFish Roles and Privileges	9-1
9.1.2	Adding Domain Certificate to Activate Start and Stop Operations.....	9-2
9.2	Understanding the Oracle GlassFish Domain	9-2
9.2.1	How to Access an Oracle GlassFish Domain.....	9-4
9.2.2	How to Add an Oracle GlassFish Domain To Be Monitored.....	9-5
9.2.3	Adding an Oracle GlassFish Domain: Finding and Assigning Targets.....	9-5
9.2.4	Adding an Oracle GlassFish Domain: Displaying Results	9-8
9.2.5	Refreshing an Oracle GlassFish Domain	9-9
9.3	Understanding the Oracle GlassFish Server Home Page.....	9-10
9.3.1	How to Access an Oracle GlassFish Server.....	9-11
9.4	Understanding the Oracle GlassFish Cluster Home Page	9-12
9.4.1	How to Access an Oracle GlassFish Cluster	9-13
9.5	Viewing Collected Configuration Data for Oracle GlassFish Members	9-13
9.6	Creating an Oracle GlassFish Server Configuration Comparison Template	9-14

Part IV Managing Oracle SOA

10 Overview of Oracle SOA Management

11 Discovering and Monitoring Oracle BPEL Process Manager

11.1	Supported Versions	11-1
11.2	Understanding the Discovery Mechanism.....	11-2
11.3	Understanding the Discovery Process	11-3
11.4	Setting Up Oracle Software Library	11-4
11.5	Discovering BPEL Process Manager	11-4
11.5.1	Deployed to Oracle Application Server	11-5
11.5.2	Deployed to Oracle WebLogic Managed Server	11-5
11.5.2.1	Discovering Oracle WebLogic Managed Server	11-6
11.5.2.2	Deployed to Oracle WebLogic Managed Server.....	11-6
11.5.3	Deployed to IBM WebSphere Application Server	11-8

11.5.3.1	Discovering IBM WebSphere Application Server	11-8
11.5.3.2	Deployed to IBM WebSphere Application Server	11-8
11.6	Configuring BPEL Process Manager	11-10
11.6.1	Specifying Details for Monitoring BPEL Process Manager	11-10
11.6.2	Registering BPEL Process Manager Credentials and Host Credentials	11-11
11.7	Troubleshooting BPEL Process Managers.....	11-11
11.7.1	Discovery Errors on Target Details Page	11-11
11.7.2	Discovery Errors on Review Page	11-12
11.7.3	Discovery Errors on Review Page	11-12
11.7.4	Display Errors on Processes Page.....	11-13
11.7.4.1	No Credentials Specified for Monitoring BPEL Process Manager.....	11-13
11.7.5	Retrieving the OPMN Port	11-14
11.7.6	javax.naming.NameNotFoundException Error	11-14
11.7.7	javax.naming.NamingException Error	11-14
11.7.8	javax.naming.NoInitialContextException Error.....	11-15
11.7.9	Error While Creating BPEL Infrastructure Services	11-16
11.7.10	Metric Collection Errors for BPEL Process Manager Partner Link Metrics	11-16
11.7.11	Agent Monitoring Metric Errors	11-17

12 Discovering and Monitoring Oracle Service Bus

12.1	Supported Versions	12-1
12.2	Understanding the Discovery Mechanism.....	12-2
12.3	Understanding the Discovery Process	12-2
12.4	Downloading One-Off Patches	12-3
12.5	Discovering Oracle Service Bus	12-4
12.5.1	Discovering OSB Deployed to WLS Not Monitored by Enterprise Manager.....	12-4
12.5.2	Discovering OSB Deployed to WLS Monitored by Enterprise Manager	12-5
12.6	Enabling Management Packs	12-6
12.7	Monitoring Oracle Service Bus in Cloud Control	12-6
12.7.1	Enabling Monitoring for OSB Services	12-6
12.8	Generating Oracle Service Bus Reports Using BI Publisher	12-7
12.9	Troubleshooting Oracle Service Bus	12-8
12.9.1	Required Patches Missing	12-8
12.9.2	System and Service	12-8
12.9.3	SOAP Test	12-9

13 Discovering and Monitoring the SOA Suite

13.1	New Features in This Release	13-1
13.2	List of Supported Versions	13-2
13.3	Monitoring Templates.....	13-2
13.4	Overview of the Discovery Process.....	13-3
13.5	Discovering the SOA Suite	13-3
13.5.1	Discovering the SOA Suite	13-4
13.5.2	Configuring the SOA Suite	13-4
13.6	Metric and Collection Settings	13-5
13.6.1	Viewing Application Dependency and Performance (ADP) Metrics	13-7
13.7	Configuring Instance Tracing	13-7

13.8	Setting Up and Using SOA Instance Tracing.....	13-7
13.8.1	SOA Instance Tracing at SOA Infrastructure Level.....	13-8
13.8.2	SOA Instance Tracing at SOA Composite Level	13-9
13.9	Monitoring Dehydration Store.....	13-10
13.9.1	Enabling Monitoring of the SOA Dehydration Store	13-10
13.9.2	Viewing the SOA Dehydration Store Data	13-11
13.10	Viewing the Service Topology	13-11
13.11	Publishing a Server to UDDI.....	13-12
13.12	Generating SOA Reports Using BI Publisher	13-12
13.13	Generating SOA Reports Using Information Publisher.....	13-14
13.14	Provisioning SOA Artifacts and Composites	13-14
13.15	Diagnosing Issues and Incidents	13-14
13.16	Verifying Target Monitoring Setup.....	13-15
13.16.1	Running Functionality-Level Diagnostic Checks	13-15
13.16.2	Running System-Level Diagnostic Checks	13-16
13.16.3	Repairing Target Monitoring Setup Issues	13-16
13.17	Searching Faults in the SOA Infrastructure	13-17
13.17.1	Overview of Faults and Fault Types in SOA Infrastructure	13-17
13.17.2	Overview of the Recovery Actions for Resolving Faults.....	13-18
13.17.3	Prerequisites for Searching, Viewing, and Recovering Faults	13-19
13.17.4	Searching and Viewing Faults	13-19
13.17.4.1	Setting Search Criteria	13-20
13.17.4.2	Finding Total Faults in the SOA Infrastructure	13-21
13.17.4.3	Limiting Faults Searched and Retrieved from the SOA Infrastructure.....	13-21
13.17.4.4	Searching Only Recoverable Faults	13-22
13.17.4.5	Searching Faults in a Particular Service Engine.....	13-22
13.17.4.6	Searching Faults by Error Message.....	13-22
13.17.4.7	Filtering Displayed Search Results	13-23
13.17.5	Recovering a Few Faults Quickly (Simple Recovery)	13-23
13.18	Creating and Tracking Bulk Recovery Jobs	13-23
13.18.1	Recovering Faults in Bulk (Bulk Recovery)	13-24
13.18.1.1	Navigating to the Bulk Recovery Jobs Page Directly	13-24
13.18.1.2	Navigating to the Create Bulk Recovery Page From SOA Infrastructure.....	13-25
13.18.1.3	Setting Fault Details for Recovering Faults in Bulk.....	13-25
13.18.1.4	Setting Recovery and Batch Details for Recovering Faults in Bulk.....	13-25
13.18.1.5	Scheduling Bulk Recovery Jobs to Run Once or Repeatedly	13-26
13.18.2	Tracking Bulk Recovery Jobs, and Viewing Their Results and Errors	13-27
13.18.3	Creating Bulk Recovery Jobs Using EMCLI and Web Services	13-28
13.18.3.1	Creating Bulk Recovery Jobs Using EMCLI	13-28
13.18.3.2	Viewing the Submitted Jobs and Outputs Using EMCLI.....	13-30
13.18.3.3	Creating Bulk Recovery Jobs through Web-Service	13-30
13.19	Recovering BPEL/BPMN Messages	13-30
13.20	Troubleshooting	13-31
13.20.1	Discovery	13-31
13.20.2	Monitoring.....	13-31
13.20.3	Instance Tracing.....	13-32
13.20.4	Recent Faults.....	13-32

13.20.5	Fault Management.....	13-32
13.20.5.1	Bulk Recovery	13-32
13.20.5.2	Fault Search and Recovery	13-34
13.20.5.3	Fault Management and Instance Tracing Errors.....	13-34
13.20.6	Application Dependency and Performance Integration.....	13-34
13.20.7	Information Publisher Reports	13-35
13.20.8	BI Publisher Reports.....	13-35
13.20.9	Systems and Services.....	13-36
13.20.10	BPEL Recovery	13-36
13.20.11	SOA License Issue.....	13-37
13.20.12	Dehydration Store Issue.....	13-37

Part V Managing Oracle Business Intelligence

14 Discovering and Monitoring Oracle Business Intelligence Instance and Oracle Essbase

14.1	Overview of Oracle Business Intelligence Targets You Can Monitor.....	14-1
14.1.1	Oracle Business Intelligence Instance	14-1
14.1.2	Oracle Essbase	14-2
14.2	Understanding the Monitoring Process.....	14-2
14.3	Discovering Oracle Business Intelligence Instance and Oracle Essbase Targets.....	14-4
14.3.1	Discovering Targets of an Undiscovered WebLogic Domain.....	14-4
14.3.2	Discovering New or Modified Targets of a Discovered WebLogic Domain	14-4
14.4	Monitoring Oracle Business Intelligence Instance and Essbase Targets	14-5
14.4.1	Performing General Monitoring Tasks.....	14-5
14.4.1.1	Viewing Target General and Availability Summary.....	14-6
14.4.1.2	Viewing Target Status and Availability History.....	14-7
14.4.1.3	Viewing Target Performance or Resource Usage	14-8
14.4.1.4	Viewing Target Metrics	14-9
14.4.1.5	Viewing or Editing Target Metric and Collection Settings	14-9
14.4.1.6	Viewing Target Metric Collection Errors.....	14-10
14.4.1.7	Viewing Target Health.....	14-10
14.4.1.8	Viewing Target Alert History	14-11
14.4.1.9	Viewing Target Incidents	14-11
14.4.1.10	Viewing Target Logs.....	14-12
14.4.1.11	Viewing Target Configuration and Configuration File	14-13
14.4.1.12	Viewing Target Job Activity.....	14-14
14.4.1.13	Viewing Target Compliance	14-14
14.4.2	Performing Target-Specific Monitoring Tasks	14-14
14.4.2.1	Viewing Oracle Business Intelligence Dashboard Reports	14-15
14.4.2.2	Viewing Oracle Business Intelligence Scheduler Reports	14-16
14.4.2.3	Viewing Oracle Business Intelligence Instance Key Metrics.....	14-17
14.4.2.4	Viewing Oracle Essbase Applications Summary.....	14-18
14.4.2.5	Viewing Oracle Essbase Application Data Storage Details.....	14-19
14.5	Administering Oracle Business Intelligence Instance and Essbase Targets.....	14-19
14.5.1	Performing General Administration Tasks.....	14-19
14.5.1.1	Starting, Stopping, or Restarting the Target.....	14-20

14.5.1.2	Administering Target Access Privileges	14-20
14.5.1.3	Administering Target Blackouts	14-20
14.5.1.4	Viewing Target Monitoring Configuration	14-21
14.5.2	Performing Target-Specific Administration Tasks	14-21
14.5.2.1	Viewing Oracle Business Intelligence Component Failovers	14-22
14.5.2.2	Editing Oracle Business Intelligence Monitoring Credentials	14-22

Part VI Using JVM Diagnostics and MDA Advisor

15 Introduction to JVM Diagnostics

15.1	Overview	15-1
15.1.1	Java Activity Monitoring and Diagnostics with Low Overhead	15-2
15.1.2	In-depth Visibility of JVM Activity	15-2
15.1.3	Real Time Transaction Tracing	15-2
15.1.4	Cross-Tier Correlation with Oracle Databases	15-2
15.1.5	Memory Leak Detection and Analysis	15-2
15.1.6	JVM Pooling	15-3
15.1.7	Real-time and Historical Diagnostics	15-3
15.2	New Features in this Release	15-3
15.3	Supported Platforms and JVMs	15-4
15.4	User Roles	15-4

16 Using JVM Diagnostics

16.1	Installing JVM Diagnostics	16-1
16.1.1	Monitoring a Standalone JVM	16-2
16.2	Setting Up JVM Diagnostics	16-2
16.2.1	Configuring the JVM Diagnostics Engine	16-3
16.2.2	Configuring JVMs and Pools	16-5
16.2.3	Register Databases	16-6
16.2.4	Configuring the Heap Analysis Hosts	16-8
16.2.5	Viewing Registered JVMs and Managers	16-9
16.3	Accessing the JVM Diagnostics Pages	16-9
16.4	Managing JVM Pools	16-10
16.4.1	Viewing the JVM Pool Home Page	16-10
16.4.2	Viewing the JVM Pool Performance Diagnostics Page	16-11
16.4.3	Viewing the JVM Pool Live Thread Analysis Page	16-12
16.4.4	Configuring a JVM Pool	16-14
16.4.4.1	Updating Pool Thresholds	16-14
16.4.5	Removing a JVM Pool	16-15
16.4.6	Add to Group	16-15
16.5	Managing JVMs	16-16
16.5.1	Viewing the JVM Home Page	16-16
16.5.2	Viewing the JVM Performance Diagnostics Page	16-18
16.5.3	Viewing the JVM Diagnostics Performance Summary	16-20
16.5.4	Viewing the JVM Live Thread Analysis Page	16-21
16.5.4.1	Cross Tier Analysis	16-23

16.5.4.2	JVM Diagnostics - Oracle Real Application Cluster Drill-Down	16-24
16.5.5	Viewing the JVM Live Heap Analysis Page	16-25
16.5.6	Working with Class Histograms	16-26
16.5.6.1	Saving a Class Histogram.....	16-27
16.5.6.2	Viewing Saved Histograms.....	16-27
16.5.6.3	Scheduling a Histogram Job	16-27
16.5.6.4	Comparing Class Histograms.....	16-27
16.5.6.5	Deleting Class Histograms.....	16-28
16.5.7	Taking a Heap Snapshot.....	16-28
16.5.8	Analyzing Heap Snapshots	16-30
16.5.8.1	Viewing Heap Usage by Roots.....	16-32
16.5.8.1.1	Top 40 Objects	16-33
16.5.8.1.2	Heap Object Information.....	16-33
16.5.8.1.3	Comparing Heap Snapshots	16-35
16.5.8.2	Viewing Heap Usage by Objects	16-35
16.5.8.3	Memory Leak Report	16-36
16.5.8.4	Anti-Pattern Report	16-36
16.5.9	Configuring a JVM	16-36
16.5.10	Removing a JVM.....	16-36
16.5.11	Add JVM to Group	16-37
16.6	Managing Thread Snapshots.....	16-37
16.6.1	Tracing Active Threads.....	16-37
16.7	Analyzing Trace Diagnostic Images.....	16-38
16.8	Uploading Trace Diagnostics Images	16-39
16.9	JVM Offline Diagnostics	16-40
16.9.1	Creating a Diagnostic Snapshot.....	16-41
16.9.2	Using the Diagnostic Snapshots Page.....	16-41
16.9.3	Analyzing a Diagnostic Snapshot	16-42
16.9.4	Viewing a Diagnostic Snapshot.....	16-42
16.10	Viewing JVM Diagnostics Threshold Violations.....	16-42
16.11	Using emctl to Manage the JVM Diagnostics Engine	16-43

17 Troubleshooting JVM Diagnostics

17.1	Cross Tier Functionality Errors.....	17-1
17.2	Trace Errors.....	17-5
17.3	Deployment Script Execution Errors	17-6
17.4	LoadHeap Errors.....	17-9
17.5	Heap Dump Errors on AIX 64 and AIX 32 bit for IBM JDK 1.6	17-9
17.6	Errors on JVM Diagnostics UI Pages.....	17-10
17.7	Frequently Asked Questions	17-10
17.7.1	Location of the JVM Diagnostics Logs.....	17-11
17.7.2	JVM Diagnostics Engine Status	17-11
17.7.3	JVM Diagnostics Agent Status	17-11
17.7.4	Monitoring Status	17-12
17.7.5	Running the create_jvm_diagnostic_db_user.sh Script	17-12
17.7.6	Usage of the Try Changing Threads Parameter.....	17-12
17.7.7	Significance of Optimization Levels	17-12

17.7.8	Custom Provisioning Agent Deployment.....	17-12
17.7.9	Log Manager Level.....	17-13
17.7.10	Repository Space Requirements.....	17-13

18 Using Middleware Diagnostics Advisor

18.1	Diagnosing Performance Issues with Oracle WebLogic Server.....	18-1
18.2	Diagnosing Performance Issues Using Middleware Diagnostics Advisor	18-2
18.3	Functioning of Middleware Diagnostics Advisor.....	18-2
18.4	Limiting the Scope of Middleware Diagnostics Advisor.....	18-3
18.5	Prerequisites	18-3
18.6	Enabling Middleware Diagnostics Advisor.....	18-4
18.7	Setting the Preferred Frequency for Purging Data	18-4
18.8	Enabling JMS Destination Metrics.....	18-5
18.9	Using Middleware Diagnostics Advisor to View and Diagnose Performance Issues...	18-5
18.10	Troubleshooting Issues Related to Middleware Diagnostics Advisor.....	18-6

Part VII Managing Oracle Coherence

19 Getting Started with Management Pack for Oracle Coherence

19.1	About Coherence Management	19-1
19.2	New Features.....	19-2
19.3	Understanding the Discovery Mechanism.....	19-3
19.3.1	Restarting a Node	19-4
19.3.2	Starting a JMX Management Node	19-6
19.3.2.1	Starting a JMX Management Node Using Shell.....	19-7
19.3.2.2	Starting a JMX Management Node Using WebLogic Console	19-7
19.3.3	Starting Other Nodes	19-7
19.3.3.1	Verifying the JMX Management Node Configuration.....	19-8
19.3.4	Using JVM Diagnostics with Coherence	19-8
19.3.5	Discovering Coherence Targets	19-8
19.3.6	Refreshing a Cluster	19-10
19.4	Enabling the Management Pack	19-10

20 Monitoring a Coherence Cluster

20.1	Cluster Level Pages.....	20-1
20.1.1	Cluster Level Home Page	20-1
20.1.1.1	General	20-2
20.1.1.2	Graphs.....	20-3
20.1.1.3	Cluster Management.....	20-3
20.1.1.4	Services.....	20-4
20.1.1.5	Applications	20-4
20.1.1.6	Cluster Level Operations.....	20-4
20.1.2	Cluster Level Node Performance Page.....	20-5
20.1.3	Cluster Level Cache Performance Page.....	20-6
20.1.4	Cluster Level Connection Performance Page	20-6
20.1.5	Cluster Level Administration Page.....	20-7

20.2	Detailed Pages	20-8
20.2.1	Node Home Page	20-8
20.2.2	Cache Home Page	20-10
20.2.3	Connection Manager Home Page.....	20-11
20.3	Performance Pages.....	20-12
20.3.1	Cache Performance Details Page	20-12
20.3.2	Connection Manager Performance Page.....	20-14
20.3.3	Connection Performance Page.....	20-15
20.3.4	Service Performance Page	20-15
20.3.5	Administration Pages.....	20-15
20.4	Log File Monitoring	20-15
20.5	Cache Data Management.....	20-17
20.5.1	View Explain Plan	20-18
20.5.2	View Trace	20-18
20.6	Reap Session Support	20-18
20.7	Push Replication Pattern.....	20-19
20.8	Transactional Cache Support	20-19
20.9	Integration with JVM Diagnostics	20-20
20.10	Viewing Performance Summary.....	20-21
20.11	Viewing Configuration Topology	20-21
20.12	Viewing Incidents	20-21
20.13	Troubleshooting Coherence	20-21
20.14	Best Practices	20-22
20.14.1	Mandatory Patches	20-22
20.14.2	Critical Metrics	20-22
20.14.3	Refresh Cluster	20-23
20.14.4	Custom Metrics	20-23
20.14.5	Monitoring Templates.....	20-27
20.14.6	Sample Code Snippets	20-27

Part VIII Using Identity Management

21 Getting Started with Identity Management

21.1	Benefits of the Using Identity Management Pack	21-1
21.2	Features of the Identity Management Pack.....	21-1
21.2.1	New Features for this Release	21-2
21.3	Monitoring Oracle Identity Management Components in Enterprise Manager	21-3

22 Prerequisites for Discovering Identity Management Targets

22.1	System Requirements	22-1
22.2	Installing Oracle Enterprise Manager Cloud Control 12c Release 2	22-2
22.3	Prerequisites for Discovering Identity Management Targets in Enterprise Manager...	22-2

23 Discovering and Configuring Identity Management Targets

23.1	Discovering Identity Management Targets.....	23-1
23.1.1	Discovering Identity Management 11g	23-1

23.1.2	Discovering Oracle Directory Server Enterprise Edition 6.x, 7.x, 11g	23-2
23.1.3	Discovering Oracle Access Manager Access Server 10.1.4.2 and 10.1.4.3.0	23-3
23.1.4	Discovering Oracle Access Manager Identity Server 10.1.4.2 and 10.1.4.3.0	23-4
23.1.5	Discovering Oracle Identity Federation Server 10.1.4.2 and 10.1.4.3.0	23-5
23.1.6	Discovering Oracle Identity Management Suite 10.1.4.2 and 10.1.4.3.0	23-6
23.1.7	Discovering Oracle Identity Manager Server 9.1.0.1	23-6
23.2	Collecting User Statistics for Oracle Internet Directory	23-7
23.3	Creating Identity Management Elements	23-8
23.3.1	Creating Identity and Access System	23-8
23.3.2	Creating Generic Service or Web Application Targets for Identity Management..	23-9
23.3.3	Creating a Service Dashboard Report.....	23-10

Part IX Discovering and Monitoring Non-Oracle Middleware

24 Discovering and Monitoring IBM WebSphere MQ

24.1	Introduction	24-1
24.2	Requirements.....	24-3
24.2.1	About Jar File Requirements for Local Monitoring	24-3
24.2.2	About Jar File Requirements for Remote Monitoring	24-3
24.3	Understanding Discovery	24-4
24.3.1	Discovery Parameters	24-4
24.3.2	Discovery Prerequisites for Local Agent	24-4
24.3.3	Discovery Prerequisites for Remote Agent	24-5
24.3.4	Queue Manager Cluster Discovery	24-5
24.3.5	Standalone Queue Manager Discovery	24-9
24.4	Monitoring	24-9
24.5	IBM WebSphere MQ Metrics	24-13

25 Discovering and Monitoring JBoss Application Server

25.1	JBoss Setup	25-1
25.2	JBoss Startup	25-1
25.3	Discovery.....	25-2
25.3.1	Understanding JBoss Local Discovery.....	25-2
25.3.2	Understanding JBoss Remote Discovery.....	25-2
25.3.3	Adding the JBoss Application Server	25-3
25.4	Monitoring JBoss JVM Metrics.....	25-6
25.5	About JVM Diagnostics.....	25-7
25.5.1	Deploying JVMD on JBoss.....	25-7
25.6	Troubleshooting	25-9

Part X Managing Oracle Data Integrator

26 Configuring and Monitoring Oracle Data Integrator

26.1	Setting Up ODI.....	26-1
26.2	Accessing Oracle Data Integrator	26-2

26.3	Managing Oracle Data Integrator.....	26-2
26.3.1	Starting, Stopping, and Restarting Oracle Data Integrator Agents.....	26-2
26.3.2	Managing Agent Status and Activities.....	26-3
26.3.3	Viewing Log Messages	26-3
26.4	Monitoring ODI Components.....	26-3
26.4.1	Viewing General Health of ODI Components	26-3
26.4.2	Viewing ODI Agents.....	26-5
26.4.3	Viewing Repository Details.....	26-6
26.4.4	Viewing Load Plan Executions/Sessions Statistics	26-8
26.5	Monitoring Run-Time Agents.....	26-9
26.6	Configuring Oracle Data Integrator Console.....	26-12
26.6.1	ODI Console Application Configuration Page.....	26-12
26.6.2	Configuring an Oracle Data Integrator Domain	26-13

Part XI Using Application Dependency and Performance

27 Introduction to Application Dependency and Performance

27.1	Overview	27-1
27.1.1	Managing Complex Java EE, SOA, OSB, and Portal Applications	27-2
27.1.2	Delivering a Service-Oriented View Across Environments	27-2
27.1.3	Avoiding Involvement of Java EE, SOA, OSB, Portal, and Application Experts	27-3
27.1.4	Eliminating Repetitive Do-It-Yourself (DIY) Manual Processes.....	27-3
27.1.5	ADP Solution.....	27-4
27.2	Architecture	27-4
27.2.1	ADP Java Agents	27-5
27.2.2	ADP Manager.....	27-6
27.2.2.1	ADP Manager and High Availability	27-6
27.2.3	ADP User Interface.....	27-6

28 Exploring Application Dependency and Performance

28.1	Exploring the User Interface.....	28-1
28.1.1	Accessing ADP	28-1
28.1.2	General ADP UI Elements.....	28-2
28.1.3	Drill Down in Operational Dashboard	28-2
28.1.4	Time Frame	28-3
28.1.5	Display Interval.....	28-3
28.1.5.1	Time Frame	28-3
28.1.5.2	Interval Context	28-3
28.1.5.3	Turning Off Time Frame Limitation.....	28-4
28.1.6	Graphs and Data Items	28-4
28.1.7	Custom Metrics	28-5
28.1.8	Functional View	28-5
28.1.9	Topology View	28-6
28.1.10	Architecture View.....	28-6
28.1.10.1	Accessing the Architecture View	28-7
28.1.10.1.1	Arrows in Architecture Views.....	28-7

28.1.10.1.2	Architecture View Summary	28-8
28.1.11	Metric Types	28-8
28.2	Exploring the Monitoring Tab	28-9
28.2.1	Monitoring ADP Entities	28-10
28.2.1.1	Monitoring SOA Suite 11g Performance	28-10
28.2.1.2	Monitoring OSB Performance	28-10
28.2.1.3	Monitoring Java EE Application Performance	28-11
28.2.1.4	Monitoring ADF Application Performance	28-11
28.2.2	WebLogic	28-12
28.2.2.1	Monitoring WebLogic Portal Performance	28-12
28.2.2.2	Adding WebLogic Domain To Be Monitored By Existing ADP Manager	28-12
28.2.2.3	Removing a WebLogic Domain From Monitoring	28-12
28.2.3	Oracle WebLogic Portals	28-13
28.2.3.1	Desktops	28-14
28.2.3.1.1	Display Portal Desktop - Desktop Structure Viewer	28-15
28.2.3.2	Portlet Drill Down	28-15
28.2.3.3	Pageflow Viewer	28-16
28.2.3.4	Books	28-16
28.2.3.5	Pages	28-16
28.2.3.6	Portlets	28-17
28.2.4	Oracle BPEL Processes	28-17
28.2.4.1	Delay Analysis View	28-18
28.2.4.2	Metadata View	28-19
28.2.4.3	Partner Links View	28-19
28.2.4.4	Partner Link Type Role View	28-19
28.2.4.5	Partner Link Bindings View	28-20
28.2.4.6	Modeled Entities View	28-20
28.2.4.7	Topology View	28-20
28.2.4.8	Node Hierarchy	28-20
28.2.5	Oracle ESB	28-21
28.2.5.1	Service Details View	28-22
28.2.5.2	Service Parent Details View	28-22
28.2.5.3	Service Definition View	28-22
28.2.5.4	Service Operations View	28-22
28.2.5.5	Operation Routing Rules View	28-23
28.2.6	Oracle WebCenter	28-23
28.2.6.1	ADF Task Flows	28-23
28.2.6.1.1	User-Defined Taskflows	28-24
28.2.6.1.2	Web 2.0 Service	28-24
28.2.6.2	JSF Pages	28-24
28.2.6.3	Portlets	28-25
28.2.6.4	Monitoring WebCenter Performance	28-25
28.2.7	Processes	28-25
28.2.7.1	Node Hierarchy	28-26
28.2.7.1.1	Delay Analysis View	28-27
28.2.7.1.2	Events View	28-27
28.2.7.2	Persistent Containers	28-27

28.2.7.2.1	Entity EJB Activity Table.....	28-28
28.2.7.2.2	Entity EJB Cache Table	28-28
28.2.7.2.3	Entity EJB Transactions Table.....	28-29
28.2.7.2.4	Entity EJB Locking Table.....	28-29
28.2.7.3	Instrumentation	28-29
28.2.8	Web Services.....	28-30
28.2.9	Pageflows	28-31
28.2.10	Services	28-31
28.2.10.1	HTTP	28-31
28.2.10.2	EJBs	28-32
28.2.10.3	JDBC.....	28-32
28.2.11	WSRP Producers	28-32
28.2.11.1	WSRP Summary	28-33
28.2.11.2	WSRP Topology	28-34
28.2.11.3	Display Portal Desktop.....	28-34
28.2.12	Integration.....	28-34
28.2.12.1	Health	28-35
28.2.12.1.1	Execute Queues.....	28-35
28.2.12.1.2	Async Dispatchers.....	28-36
28.2.12.1.3	Sync Dispatchers	28-37
28.2.12.1.4	JMS Destinations	28-37
28.2.12.1.5	Stateless Containers	28-38
28.2.12.1.6	Persistent Containers	28-39
28.2.12.2	Performance.....	28-41
28.2.12.3	Channels	28-42
28.2.12.4	Subscribers.....	28-42
28.2.13	Applications.....	28-42
28.2.13.1	Services.....	28-44
28.2.13.2	Dependencies	28-44
28.2.13.3	Deployments	28-45
28.2.13.4	Workshop Projects.....	28-46
28.2.13.5	Web Applications	28-46
28.2.13.6	Stateless Beans.....	28-47
28.2.13.7	Stateful Beans	28-47
28.2.13.7.1	Stateful EJB Cache	28-47
28.2.13.7.2	Stateful EJB Transactions.....	28-48
28.2.13.7.3	Stateful EJB Locking.....	28-48
28.2.13.8	Entity Beans.....	28-49
28.2.13.8.1	Entity EJB Activity.....	28-49
28.2.13.8.2	Entity EJB Cache	28-50
28.2.13.8.3	Entity EJB Transactions	28-50
28.2.13.8.4	Entity EJB Locking.....	28-50
28.2.13.9	Message Driven Beans	28-51
28.2.13.9.1	Message Driven EJB Activity	28-51
28.2.13.9.2	Message Driven EJB Transactions.....	28-52
28.2.14	Oracle WebLogic Resources.....	28-53
28.2.15	Oracle Resources.....	28-53

28.2.16	Custom Metrics	28-54
28.2.17	Status	28-54
28.2.18	Service Component Architecture (SCA).....	28-55
28.2.18.1	Components	28-55
28.3	Exploring the Configuration Tab.....	28-56
28.3.1	Database Configuration	28-56
28.3.2	Resource Configuration	28-56
28.3.3	Service Level Objective Configuration	28-56
28.3.3.1	Creating a New SLO	28-57
28.3.3.2	Defining SLO Parameters	28-57
28.3.3.3	SLO Blackout Configuration.....	28-58
28.3.3.4	Creating and Maintaining SLO Blackouts	28-58
28.3.3.5	Propagating Threshold Violation Events.....	28-59
28.3.4	Event Integration	28-60
28.3.5	Custom Metric Configuration.....	28-60
28.4	Exploring the Registration Tab	28-60
28.4.1	Using RMI Configuration for Managers	28-61
28.4.2	Adding a New Manager (RMI Configuration).....	28-61
28.4.3	Editing a Previously Configured Manager (RMI Configuration)	28-61
28.4.4	Removing or Disabling a Previously Configured Manager	28-62
28.5	Using emctl to Manage the ADP Diagnostics Engine	28-62

29 ADP Methodology

29.1	ADP Methodology Activities	29-2
29.1.1	Mapping Business SLAs to Performance SLOs.....	29-2
29.1.2	Specifying Target Performance Characteristics	29-3
29.1.3	Improving Performance.....	29-3
29.1.3.1	Characterizing Baseline Performance.....	29-3
29.1.3.2	Identifying Performance Bottlenecks	29-4
29.1.3.3	Removing Performance Bottlenecks	29-4
29.1.3.4	Setting SLOs on Key Metrics.....	29-4
29.2	Mapping Business SLAs to Performance SLOs	29-5
29.3	Characterizing Baseline Performance	29-6
29.4	Identifying Performance Bottlenecks	29-7
29.4.1	Determining System Level Performance.....	29-7
29.5	Setting SLOs on Key Metrics	29-8
29.6	Conclusion	29-10

30 Frequently Asked Questions About Application Dependency and Performance

30.1	Can I Erase the darchive Directory?	30-1
30.2	How Do I Undeploy the Agent?	30-1

A ADP Configuration Directories and Files

A.1	Configuration Directories	A-1
A.1.1	Directory Structure	A-1

A.1.2	Config Directory	A-2
A.1.3	Deploy Directory	A-2
A.2	Acsera.properties File.....	A-2
A.2.1	Log Files Management.....	A-3
A.2.2	Multi-Domain Monitoring Configuration.....	A-3
A.2.3	ADP RMI Port Assignment.....	A-3
A.2.4	ADP Aggregation and Data Life Time Configuration	A-3
A.2.5	Aggregating Incoming Metrics On the Fly	A-4
A.2.6	Listing Applications to Be Monitored or Excluded From Monitoring	A-4
A.2.7	Firewall Mitigation (for Internal RMI Ports)	A-4
A.2.8	SLO Dampening.....	A-5
A.3	UrlMap.properties	A-5

B Support Matrix for Application Dependency and Performance

Index

Preface

This document describes how to use Oracle Enterprise Manager Cloud Control to monitor and manage middleware software, including Oracle Fusion Middleware and Oracle WebLogic Server.

Audience

This document is intended for those who monitor and manage both Oracle applications and custom Java EE applications that run on a combination of Oracle Fusion Middleware, as well as non-Oracle middleware software.

Documentation Accessibility

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

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documents

For more information, see the following documents in the Oracle Enterprise Manager Release 12c documentation set:

- *Oracle Enterprise Manager Lifecycle Management Guide*
- *Oracle Enterprise Manager Cloud Control Administrator's Guide*

For the latest releases of these and other Oracle documentation, check the Oracle Technology Network at:

<http://www.oracle.com/technetwork/documentation/index.html#em>

Oracle Enterprise Manager also provides extensive Online Help. Click **Help** at the top of any Enterprise Manager page to display the online help window.

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.

What's New in Getting Started with Oracle Fusion Middleware Management Manual?

This section lists the changes to this manual since Enterprise Manager 12c release 12.1.0.2.

- Added the following chapters:
 - [Chapter 6, "Discovering and Monitoring an Exalytics Target"](#)
 - [Chapter 7, "Oracle Traffic Director"](#)
 - [Chapter 9, "Overview of Oracle GlassFish Server Management"](#)
 - [Chapter 18, "Using Middleware Diagnostics Advisor"](#)
 - [Chapter 24, "Discovering and Monitoring IBM WebSphere MQ"](#)
 - [Chapter 25, "Discovering and Monitoring JBoss Application Server"](#)
 - [Chapter 26, "Configuring and Monitoring Oracle Data Integrator"](#)
- Deleted the following chapters:
 - Monitoring Business Applications
 - Exporting Data (Application Dependency and Performance)
- Renamed the following chapter:
 - 'Troubleshooting Application Dependency and Performance' to [Chapter 30, "Frequently Asked Questions About Application Dependency and Performance"](#)
- Added the following section:
 - [Section 2.2.3, "Discovering a Standalone Oracle HTTP Server"](#)
- Rewrote the following section:
 - [Section 3.4, "Administering Middleware Targets"](#)

Part I

Managing Oracle Fusion Middleware

The chapters in this part describe how you can discover and monitor Oracle Fusion Middleware targets, including Oracle WebLogic Server and deployed Java EE applications.

The chapters are:

- [Chapter 1, "Introduction to Middleware Management"](#)
- [Chapter 2, "Discovering Middleware Targets"](#)
- [Chapter 3, "Managing Middleware Targets"](#)
- [Chapter 4, "Testing Application Load and Performance"](#)

Introduction to Middleware Management

This chapter provides an introduction to how you can use Oracle Enterprise Manager Cloud Control to monitor and manage middleware software, including Oracle Fusion Middleware and Oracle WebLogic Server.

Note: The features documented in this manual pertain to Enterprise Manager for Oracle Fusion Middleware Plug-in version 12.1.0.3.

This chapter covers the following:

- [Middleware Management with Enterprise Manager Cloud Control](#)
- [Key Oracle Fusion Middleware Management Features](#)
- [Managing Fusion Middleware with Fusion Middleware Control](#)

1.1 Middleware Management with Enterprise Manager Cloud Control

Middleware is the software that enables your enterprise applications to run. Managing the underlying middleware technology can be difficult, and IT organizations often have to rely on a variety of specialized tools. This can lead to inefficiency and may introduce complexities and risks.

Enterprise Manager Cloud Control is the definitive tool for middleware management and allows you to manage both Oracle applications and custom Java EE applications that run on a combination of Oracle Fusion Middleware as well as non-Oracle middleware software.

Oracle Enterprise Manager Cloud Control is a Web browser-based, graphical user interface that you can use to monitor multiple Oracle Fusion Middleware Domains and Oracle WebLogic Domains. In fact, Cloud Control provides deep management solutions for Oracle technologies including Oracle packaged applications, Oracle Database and Oracle VM.

Enterprise Manager Cloud Control supports the discovery, monitoring and central management of the entire family of Oracle Fusion Middleware components, including:

- Oracle WebLogic Server domains, clusters and single server instances
- Clustered and standalone Java EE applications
- Web tier components, including Oracle HTTP Server and Oracle Web Cache
- Service-Oriented Architecture (SOA) components
- Oracle Identity Management

- Metadata Services repositories
- Oracle WebCenter
- Oracle Portal
- Oracle Business Intelligence Discoverer
- Oracle Forms Services
- Oracle Reports
- Directory Server Enterprise Edition
- Oracle Coherence
- Oracle Exalogic Elastic Cloud
- Oracle Application Server
- Java EE

Cloud Control also offers extensive support for non-Oracle technologies through more than two dozen heterogeneous management plug-ins and connectors including Microsoft MOM, IBM WebSphere Application Server, JBoss Application Server, EMC storage, F5 BIG IP, Check Point Firewall, and BMC Remedy.

A key benefit of Enterprise Manager Cloud Control is that unlike other Fusion Middleware management utilities - such as Fusion Middleware Control and the WebLogic Server Administration Console - you can monitor and manage multiple middleware targets, such as all of your WebLogic Server domains, from a single console.

You can also view real time as well as historic performance metrics collected from middleware targets. This enables you to monitor the availability and performance of Oracle Fusion Middleware software both in real time and from a historical perspective for trend analysis and diagnosing availability and performance problems.

Enterprise Manager Cloud Control also enables you to manage the infrastructure upon which the middle tier depends. You can manage underlying operating systems and hosts on which the middleware software is installed. You can also monitor the databases used by deployed applications, enabling you to diagnose application performance problems and identify the true root cause of the problem and the tier (middleware, database) on which it occurs.

The built-in topology viewer allows you to visualize and monitor your entire Oracle Fusion Middleware environment in a graphical display. Topologies can be viewed for a single SOA composite, an Oracle WebLogic Domain, or across multiple Oracle WebLogic Domains.

Management of Service-Oriented Architecture (SOA) components such as BPEL processes and infrastructure components such as Oracle Service Bus, is also supported. The infrastructure provides monitoring, fault management, configuration management, deployment and dependency views of wiring between components.

1.2 Key Oracle Fusion Middleware Management Features

Cloud Control provides full historical monitoring across the middleware tier, from WebLogic Server instance and the Java virtual machine (JVM) it runs within to the Oracle Fusion Middleware components running on the application server. It also provides full configuration and lifecycle management of middleware components, while the product's extensive performance monitoring and diagnostics capabilities enable troubleshooting issues anywhere within the middleware tier.

With Oracle Enterprise Manager Cloud Control, you can:

- Centrally manage multiple Oracle Fusion Middleware Domains and WebLogic Domains.
- Manage third party products such as IBM WebSphere Application Server, JBoss Application Server, Apache Tomcat and the Microsoft .NET Framework.
- Manage non-middleware software such as underlying operating systems and hardware on which the middleware software is installed. This allows administrators to correlate middleware performance with its underlying host performance.
- Manage database software and diagnose application performance problems and identify the true root cause of the problem and the tier (middleware, database) on which it occurs.
- Monitor the availability and performance of Oracle Fusion Middleware software in real time and from a historical perspective for trend analysis.
- Diagnose availability and performance problems.
- Monitor and trace important end-user requests from the client to the service endpoint across all the servers and applications associated with each transaction.
- Use Application Dependency and Performance (ADP) to analyze Java EE and SOA applications.
- Monitor Java applications and diagnose performance problems in production using JVM Diagnostics.
- Define Service Level Objectives (SLOs) in terms of out-of-box system-level metrics as well as end user experience metrics to accurately monitor and report on Service Level Agreement (SLA) compliance.
- Perform several critical tasks like:
 - Setting thresholds on performance metrics. When these thresholds are violated, e-mail and page notifications are sent.
 - Tracking configuration changes and comparing configurations between example test environment and production environment.

1.3 Managing Fusion Middleware with Fusion Middleware Control

Fusion Middleware Control organizes a wide variety of performance data and administrative functions into distinct, Web-based home pages for the cluster, domain, servers, components, and applications. The Fusion Middleware Control home pages make it easy to locate the most important monitoring data and the most commonly used administrative functions all from your Web browser.

Fusion Middleware Control is a part of the Oracle Fusion Middleware installation. With Fusion Middleware Control, you can:

- Manage a single Oracle Fusion Middleware Domain and a single WebLogic Domain.
- Monitor the availability and performance of Fusion Middleware software in real time mode.
- Perform routine administration tasks such as deploying applications, configuring parameters etc.

For more details, see the *Oracle Fusion Middleware 11g Administrator's Guide*.

Discovering Middleware Targets

In order to manage and monitor Oracle Fusion Middleware components such as WebLogic Server instances and clusters, as well as the Java EE applications that depend on them, Enterprise Manager Cloud Control must first "discover" the Fusion Middleware Domain containing these components.

Once discovered, the domain and the components within it can be promoted to "managed target" status. In this process, Management Agents are assigned to each target, enabling Enterprise Manager Cloud Control to collect the data needed to monitor the target.

Note: The Oracle Fusion Middleware versions for which discovery is supported in Oracle Enterprise Manager Cloud Control release 12.1.0.3 is available from the certification matrix located on My Oracle Support (<https://support.oracle.com>).

This chapter contains the following sections:

- [Enabling Automatic Discovery of Fusion Middleware Targets](#)
- [Discovering Targets Manually](#)
- [Discovering New or Modified Domain Members](#)

2.1 Enabling Automatic Discovery of Fusion Middleware Targets

Cloud Control provides the ability to automatically search for potential Oracle Fusion Middleware targets on host machines that are already being managed by Cloud Control. You can configure automatic discovery to run on a regular schedule, such as every 2 days. Discovery is run every 24 hours by default.

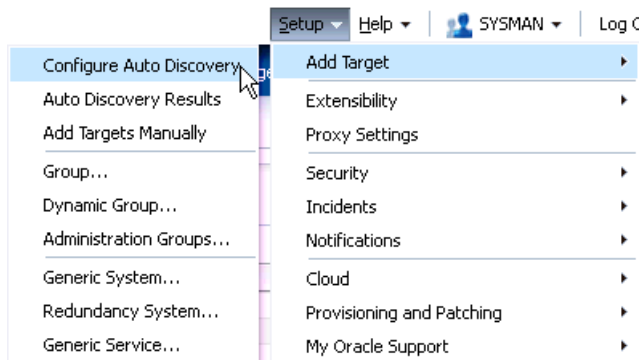
Once automatic discovery has been configured, you can check the Auto Discovery Results page to see what new Oracle Fusion Middleware targets have been discovered. You can then promote targets to managed target status by assigning a Management Agent to monitor and manage the target. To make things a bit easier, the Management Agent should already be pushed on to the target host. This enables you to configure auto-discovery on that host.

The following describes the process for configuring automatic discovery of Oracle Fusion Middleware targets only. However, Cloud Control actually enables you to configure discovery of a variety of Oracle target types.

Note: The automatic discovery feature is not supported for Oracle WebLogic Server version 8.x.

1. From the **Setup** menu, select **Add Target**, then select **Configure Auto Discovery**.

Figure 2–1 Configure Auto Discovery Option



2. On the Configure Auto Discovery page, select the Oracle Fusion Middleware link in the table to configure auto discovery for Oracle Fusion Middleware or click the icon in the **Configure Host Discovery** column to configure that Oracle Fusion Middleware row.

Figure 2–2 Configure Auto Discovery Page

Configure Auto Discovery

Instruction
Automatic discovery is a process that detects hosts, virtual servers and other targets so that they can be monitored and managed.

Configure Auto Discovery

- Setup discovery using IP Scan
- Setup discovery on Single Host
- Setup discovery on Multiple Hosts

→

Add Targets from Auto Discovery Results

- Add Non-Host Targets
- Add Discovered Hosts
- Ignore Discovered Targets

→

Configure Auto Discovery

	Type	Configure Host Discovery
Host Discovery		
Host and Oracle VM Manager Discovery using IP Scan	Agent-less	

	Type	Configure Auto Discovery
Auto Discovery		
All Discovery Modules	Agent-based	
Oracle Home Discovery	Agent-based	
Oracle Fusion Middleware	Agent-based	
Oracle Database, Listener and Automatic Storage Management	Agent-based	
Oracle Cluster and High Availability Service	Agent-based	
Discover Oracle Secure Backup Domain	Agent-based	

3. Set the schedule at which the discovery job will be run, in days. This schedule will be applied to all selected hosts. By default the job will run every 24 hours.

Figure 2–3 Schedule for Configuring Target Discovery**Configure Target Discovery for Target Types**

Use this page to configure target discovery and parameters for this discovery module.

Discovery Module: Oracle Fusion Middleware
 Target Types: Oracle WebLogic Domain

Schedule: ☒ Keep current discovery schedule on all hosts
☐ Use following discovery schedule for all target types on all hosts

Frequency Every Day(s)

View

Agent Host Name	Schedule	Discover	Configure
No data found.			

- Click **Add Host**. Select the host machines you want to include in the discovery.

Figure 2–4 Host Machines Available for Discovery

Search and Select: Targets - Oracle Enterprise Manager

Search and Select: Targets

Target Type: Host

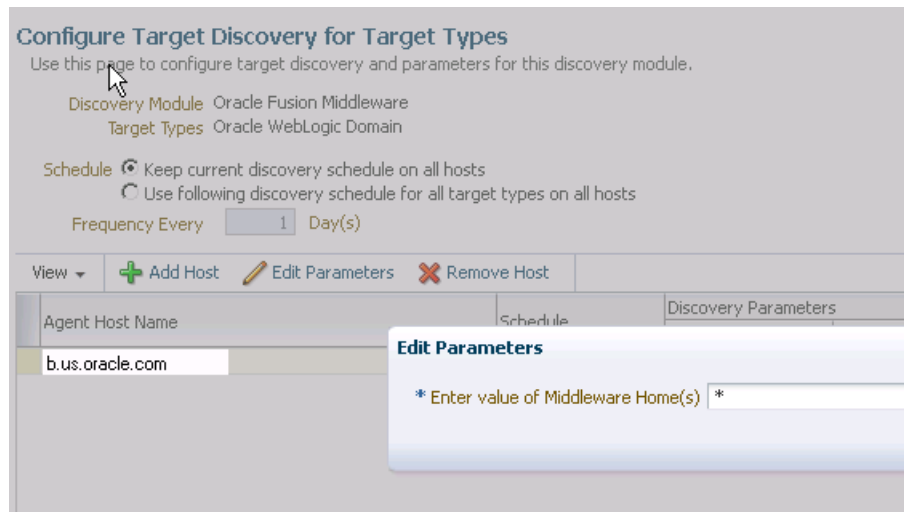
Target Name:

On Host:

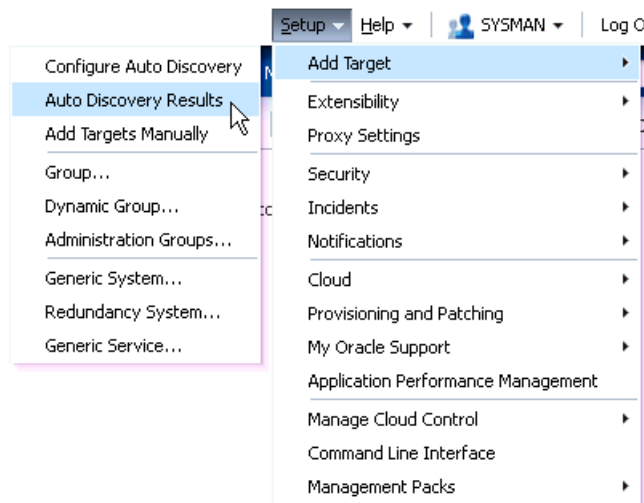
Target Name	Target Type	On Host	Status
a.us.oracle.com	Host	c.us.oracle.com	
b.us.oracle.com	Host	b.us.oracle.com	

- Select a host in the table, and then click **Edit Parameters** to specify the Middleware Homes to search for targets. The Middleware Home is the top-level directory for all Oracle Fusion Middleware products, created when Oracle WebLogic Server is installed.

Enter * to search all Middleware Homes, or specify the path to one or more Middleware Homes on the host, each separated by a comma.

Figure 2–5 Edit Parameters

6. Click the **OK** button located at the right of the screen. At this point, automatic discovery has been enabled, and the discovery job will run at the scheduled frequency.
7. Once automatic discovery has been enabled, you should check Cloud Control regularly to view the list of discovered targets. From the **Setup** menu, select **Add Target**, then select **Auto Discovery Results**.

Figure 2–6 Auto Discovery Results Option

8. Click the **Non-Host Targets** tab to view the discovered Oracle Fusion Middleware targets.

Figure 2–7 Auto Discovery Results Page

Auto Discovery Results

☒ **Instruction**
Review discovered unmanaged targets and promote targets to be managed by Enterprise Manager.

Configure Auto Discovery

- Setup discovery using IP Scan
- Setup discovery on Single Host
- Setup discovery on Multiple Hosts

→

Add Targets from Auto Discovery Results

- Add Non-Host Targets
- Add Discovered Hosts
- Ignore Discovered Targets

Host Targets (0) **Non-Host Targets (1)** Ignored Targets (0)

> Search

View ▾ Promote Rename Delete Ignore Refresh

Target Name	Target Type	Discovered On
WebTierabc	Oracle Home	Jul 16, 2012 5:03:06 f

9. Select a target, then click **Promote**.

If multiple targets of various types are listed, you can expand Search, then select the Target Type you are looking for (such as Oracle WebLogic Domain). Click **Search** to display the selected discovered target types.

Figure 2–8 Promote Tab

Auto Discovery Results

☒ **Instruction**
Review discovered unmanaged targets and promote targets to be managed by Enterprise Manager.

Configure Auto Discovery

- Setup discovery using IP Scan
- Setup discovery on Single Host
- Setup discovery on Multiple Hosts

→

Add Targets from Auto Discovery Results

- Add Non-Host Targets
- Add Discovered Hosts
- Ignore Discovered Targets

Host Targets (0) **Non-Host Targets (2)** Ignored Targets (0)

> Search

View ▾ **Promote** Rename Delete Ignore Refresh

Target Name	Target Type	Discovered On
/Farm_EMGC_DOMAIN/EMGC_DOMAIN b.us.oracle.co	Oracle WebLogic Domain	Jul 18, 2012 5:03:14
WebTierabc	Oracle Home	Jul 16, 2012 5:03:06

10. Supply or accept values for the following parameters:

Figure 2–9 shows the parameters that need to be provided.

Figure 2–9 Parameters**Add Fusion Middleware Farm: Find Targets**

Enterprise Manager can be configured to manage additional Oracle WebLogic Domains. For adding a secured Oracle then choose the proper JMX Protocol.

* Administration Server Host: b.us.oracle.com

* Port: 7022

* Username:

* Password:

* Unique Domain Identifier: Farm01

* Agent: b.us.oracle.com:1838

▼ **Advanced**

JMX Protocol: t3

Discover Down Servers: ☐

JMX Service URL:

External Parameters:

Discovery Debug File Name:

- **Administration Server Host**

Enter the host name on which the Administration Server is installed and running for the Oracle WebLogic Domain that you want to promote to a managed target, for example: myhost06.example.com

- **Port**

Enter the WebLogic Administration Server port. The default is 7001.

If the WebLogic Administration Server is secured using the Secure Sockets Layer (SSL) protocol, specify the location of the trusted keystore file. The keystore is a protected database that holds keys and certificates for an enterprise. See **Advanced Parameters**.

- **Enter the WebLogic Administration Server user name and password.**

If you want to discover the target only for monitoring purposes, then it is sufficient to provide a user name that has a monitoring role. If you want to monitor the target and also perform start/stop operations, then ensure that you provide a user name that has either an operator role or an administrator role.

Note: There is the potential of account locking issues if you enter the default WebLogic user name, and the account password is changed without updating the Enterprise Manager monitoring credentials for the Domain and Farm.

- **Unique Domain Identifier.**

Specify a Unique Domain Identifier. This value is used as a prefix to ensure farm names are unique in environments with the same domain name. By default, Enterprise Manager will pre-pend the name with "Farm", followed by a two-digit number, such as "Farm01".

- **Agent**

Host name for a Management Agent that will be used to discover the Fusion Middleware targets.

If a Management Agent exists on the WebLogic Administration Server host, the host name for this Management Agent will be supplied by default.

However, you can specify any Management Agent on any host that is managed by Cloud Control to perform the discovery.

Note: To access all supported features, Oracle recommends that you have a Management Agent local to each managed server in the domain and to use that Management Agent to monitor the WebLogic Servers on that local host machine. Though remote Management Agents can manage WebLogic Server targets, the local Management Agent is recommended.

Some features that are *not* supported when there is no local Management Agent:

- To patch a WebLogic Server, you need a local Management Agent on each WebLogic Server machine.
- If you want to use Oracle Support Workbench for a WebLogic Server target, then the target requires a local Management Agent.
- Cloning requires a local Management Agent at the source administration server machine and a local Management Agent at all destination machines.

Advanced Parameters

If the target domain is secured, expand the Advanced node to specify the protocol to use in the Protocol field. The default value is t3.

For additional details on discovering a domain secured using the Secure Sockets Layer (SSL) protocol, see section "C" in My Oracle Support Note 1093655.1. You can access My Oracle Support at the following URL:

<https://support.oracle.com/CSP/ui/flash.html>

■ JMX Protocol

Used to make a JMX connections to the Administration Server. For Secure domain JMX protocol - use t3s. If WebLogic domain is using a demo certificate, this certificate is automatically updated to monitoring and discovery agent. If a custom certificate is used, refer to [Chapter 8, "Monitoring WebLogic Domains"](#) for information on how to import a certificate.

■ Discover Down Servers

Use this check box to discover servers that are down. While adding Oracle Fusion Middleware WebLogic Domain targets to Enterprise Manager Cloud Control, you can now choose whether to add WebLogic Server targets that are discovered in a down state. This gives you more control in determining what to automatically add to Cloud Control for centralized management and monitoring.

To monitor down servers, their Listener Address must be set. Otherwise, these servers will have 'temp-host' as host value and the local agent cannot be determined for monitoring. Therefore, the servers will always be shown as down.

When servers come up, this WebLogic domain needs to be refreshed for populating the correct host name and monitoring.

■ JMX Service URL

Optionally supply the Java Management Extensions (JMX) Service URL that will be used to establish a JMX connection to the WebLogic Administration Server. For example:

```
service:jmx:t3://server.example.com:5555/jndi/weblogic.management.mbeanservers.domainruntime
```

If you do not specify this URL, Enterprise Manager will provide the Service URL based on the host port and protocol. If this is specified, the Administration server host and port information still must be provided in the input parameters.

- External Parameters

Optionally enter any system properties to be used by the Java process to connect to the WebLogic Administration Server in the Parameters field.

Supply space-separated name/value pairs. Preface each parameter with -D. For example:

```
-Dparam1=xxx -Dparam2=yyy -Dparam3=zzz
```

- Discovery Debug File Name

If problems occur while adding Middleware-related targets or refreshing domain membership, you can enable additional debugging information to quickly diagnose and resolve the issue. This file will be created in the discovery Agent's log directory.

11. Click **Continue**. Enterprise Manager will discover all Fusion Middleware targets within the domain.

12. Click **Close** in the Finding Targets dialog to automatically assign Management Agents to the discovered targets.

The Assign Agents page lists each Fusion Middleware target discovered and the Management Agent assigned to each. Agents are automatically assigned as follows:

- If a local Management Agent is installed on the discovered target host, that Agent will be assigned.
- If a local Management Agent cannot be found on the host, the Agent specified in the Targets page will be assigned.

Note that you can also manually assign Management Agents to specific targets, if desired.

13. Click **Add Targets** to assign Management Agents as listed in the Assign Agents page.

The Saving Target to Agent processing window appears, indicating how many total targets have been added and successfully saved. It will also indicate the number of targets that were unsuccessfully added.

14. Click **Close** in the processing window when finished. The Results page displays the targets and Agent assignments.

15. Click **OK** when finished. There may be a delay before these targets are visible and monitored. All the agents used for monitoring the targets must be up.

2.2 Discovering Targets Manually

This section covers the following:

- [Discovering a WebLogic 9.x or 10.x Domain Using Cloud Control](#)
- [Discovering Multiple WebLogic Domains Using EM CLI](#)

- Discovering a Standalone Oracle HTTP Server

2.2.1 Discovering a WebLogic 9.x or 10.x Domain Using Cloud Control

Oracle WebLogic Server release 9.x and 10.x domains and their respective components can be discovered using Enterprise Manager Cloud Control. A wizard guides you through the discovery process.

Note: Fusion Middleware 12c discovery also discovers dynamic servers, CAM (Common Administration Model) managed components, as well as a new target type called *Domain Application*.

Domain application represents the deployment of a target to a domain.

For Fusion Middleware 12c discovery, when Cloud Control discovers a Fusion Middleware WebLogic domain, it creates a Domain Application target for each Java EE application deployed in a Weblogic domain. For any Java EE application deployed on any number of servers, only one Domain Application target will be discovered.

Note: To discover a WebLogic Server domain, the Administration Server must be up because the Management Agent must make a JMX connection to it. If the Administration Server is down, discovery cannot occur.

Thereafter, to monitor the WebLogic Server domain, the Administration Server need not be up.

1. From the **Setup** menu, select **Add Target**, then select **Add Targets Manually**.
2. Select the **Add Non-Host Targets Using Guided Process (Also Adds Related Targets)** option.
3. Select **Oracle Fusion Middleware** as the target type.
4. Click **Add Using Guided Discovery**.
5. Supply or accept values for the following parameters:
 - **Administration Server Host**
Enter the host name on which the Administration Server is installed and running for the Oracle WebLogic Domain that you want to promote to a managed target, for example: myhost06.example.com
 - **Port**
Enter the WebLogic Administration Server port. The default value is 7001.
If the WebLogic Administration Server is secured using the Secure Sockets Layer (SSL) protocol, specify the location of the trusted keystore file. The keystore is a protected database that holds keys and certificates for an enterprise. See **Advanced Parameters**.
 - **Enter the WebLogic Administration Server user name and password.**
If you want to discover the target only for monitoring purposes, then it is sufficient to provide a user name that has a monitoring role. If you want to monitor the target and also perform start/stop operations, then ensure that

you provide a user name that has either an operator role or an administrator role.

- **Unique Domain Identifier.**

Specify a Unique Domain Identifier. This value is used as a prefix to ensure farm names are unique in environments with the same domain name. By default, Enterprise Manager will pre-pend the name with "Farm", followed by a two-digit number, such as, "Farm01".

- **Agent**

The host name for a Management Agent that will be used to discover the Fusion Middleware targets.

If a Management Agent exists on the WebLogic Administration Server host, the host name for this Management Agent will be supplied by default. However, you can specify any Management Agent on any host that is managed by Cloud Control to perform the discovery.

Note: To access all supported features, Oracle recommends that you have a Management Agent local to each managed server in the domain and to use that Management Agent to monitor the WebLogic Servers on that local host machine. Though remote Management Agents can manage WebLogic Server targets, the local Management Agent is recommended.

Some features that are *not* supported when there is no local Management Agent:

- To patch a WebLogic Server, you need a local Management Agent on each WebLogic Server machine.
- If you want to use Oracle Support Workbench for a WebLogic Server target, then the target requires a local Management Agent.
- Cloning requires a local Management Agent at the source administration server machine and a local Management Agent at all destination machines.

Advanced Parameters

If the target domain is secured, expand the Advanced node to specify the protocol to use in the Protocol field. The default value is t3.

For additional details on discovering a domain secured using the Secure Sockets Layer (SSL) protocol, see section "C" in My Oracle Support Note 1093655.1. You can access My Oracle Support at the following URL:

<https://support.oracle.com/CSP/ui/flash.html>

- **JMX Protocol**

Used to make a JMX connections to the Administration Server.

- **Discover Down Servers**

Use this check box to discover servers that are down. While adding Oracle Fusion Middleware WebLogic Domain targets to Enterprise Manager Cloud Control, you can now choose whether to add WebLogic Server targets that are discovered in a down state. This gives you more control in determining what to automatically add to Cloud Control for centralized management and monitoring.

To monitor down servers, their Listener Address must be set. Otherwise, these servers will have 'temp-host' as host value and the local agent cannot be

determined for monitoring. Therefore, the servers will always be shown as down.

When servers come up, this WebLogic domain needs to be refreshed for populating the correct host name and monitoring.

- Discover Application Versions

By default, each version of a deployed Java EE application will be discovered as a target. Therefore, with every new version, a new target will be discovered. Deselect this check box if you want only the active version of the application to be discovered.

- JMX Service URL

Optionally supply the Java Management Extensions (JMX) Service URL that will be used to establish a JMX connection to the WebLogic Administration Server. For example:

```
service:jmx:t3://server.example.com:5555/jndi/weblogic.management.mbeanservers.domainruntime
```

If you do not supply a value, Enterprise Manager will provide the Service URL based on the host port and protocol. If this is specified, the Administration server host and port information still must be provided in the input parameters.

- External Parameters

Optionally enter any system properties to be used by the Java process to connect to the WebLogic Administration Server in the Parameters field.

Supply space-separated name/value pairs. Preface each parameter with -D. For example:

```
-Dparam1=xxx -Dparam2=yyy -Dparam3=zzz
```

- Discovery Debug File Name

If problems occur while adding Middleware-related targets or refreshing domain membership, you can enable additional debugging information to quickly diagnose and resolve the issue.

6. Click **Continue**. Enterprise Manager will discover all Fusion Middleware targets within the domain.

7. Click **Close** in the Finding Targets dialog to automatically assign Management Agents to the discovered targets.

The Assign Agents page lists each Fusion Middleware target discovered and the Management Agent assigned to each. Agents are automatically assigned as follows:

- If a local Management Agent is installed on the discovered target host, that Agent will be assigned.
- If a local Management Agent cannot be found on the host, the Agent specified in the Targets page will be assigned.

Note that you can also manually assign Management Agents to specific targets, if desired.

8. Click **Add Targets** to assign Management Agents as listed in the Assign Agents page.

The Saving Target to Agent processing window appears, indicating how many total targets have been added and successfully saved. It will also indicate the number of targets that were unsuccessfully added.

9. Click **Close** in the processing window when finished. The Results page displays the targets and Agent assignments.
10. Click **OK** when finished. There may be a delay before these targets are visible and monitored. All the agents used for monitoring the targets must be up.

2.2.2 Discovering Multiple WebLogic Domains Using EM CLI

If you have multiple WebLogic domains that you want to manage through Enterprise Manager Cloud Control, you can use the Enterprise Manager Command Line Interface (EM CLI) `discover_wls` verb to discover them all at once, rather than discovering them one at a time using the discovery wizards.

The `discover_wls` verb can be used to discover WebLogic Server versions 7.x, 8.x, 9.x, and 10.x domains. The verb reads a file named `domain_discovery_file` that contains the information required to discover each domain.

See the *Enterprise Manager Command Line Interface* book for instructions on using the `discover_wls` verb.

2.2.3 Discovering a Standalone Oracle HTTP Server

Oracle Enterprise Manager can be configured to manage a standalone Oracle HTTP Server (OHS). A wizard guides you through the discovery process.

1. From the **Setup** menu, select **Add Target**, and then select **Add Targets Manually**.
2. On the Add Targets Manually page, select **Add Targets Using Guided Process (Also Adds Related Targets)**.
3. From the Target Types list, select **Oracle HTTP Server**, and then click **Add Using Guided Process...**
4. On the Add Standalone Oracle HTTP Server page, do the following:

- Specify the host name on which the OHS is installed. You can search for the host by clicking on the search icon.

When you click on the search icon, a Search and Select: Targets dialog box opens. Select a host name from the list, and then click **Select**.

- Specify the OHS port number.

The Agent URL gets generated automatically. You will not be able to edit this.

- Specify the target name.
- Select the Oracle Home location by clicking **Browse and Select**.

In the Remote File Browser dialog box, select the Oracle Home location, and click **OK**.

- Select the OHS version from the Version list.
- Click **Add Target**.

After the discovery succeeds, navigate to the **Target** menu, and select **All Targets**, or select **Middleware**. The discovered standalone Oracle HTTP Server target will be displayed.

2.3 Discovering New or Modified Domain Members

In the typical enterprise, Oracle WebLogic domains do not remain static. Instead, membership in the domain changes regularly: New Java EE applications are deployed, WebLogic Server instances are created or removed, clusters are added, and so on.

By default, Enterprise Manager Cloud Control is not automatically aware of changes made to Oracle WebLogic domains that have been configured as managed targets. However, the application does provide the ability to discover and uptake new or modified domain members.

This section covers the following:

- [Enabling Automatic Discovery of New Domain Members](#)
- [Manually Checking for New or Modified Domain Members](#)

2.3.1 Enabling Automatic Discovery of New Domain Members

You can enable a pre-defined Cloud Control job named "WebLogic Domain Refresh" to automatically discover new domain members and add them as managed targets.

Note: Whenever you perform the Refresh operation, the Administration Server must be up and the Discovery Agent must be able to connect to it using JMX.

1. From the **Targets** menu, select **Middleware**.
2. Click on the WebLogic Domain you want to enable the job for in the Middleware home page.
3. In the General region of the page, click the timestamp link next to the **WebLogic Domain Refreshed** property. The Refresh WebLogic Domain dialog opens.
4. Check the **Enable Automatic Refresh** box in the Refresh WebLogic Domain dialog, then click **OK**.

Once enabled, the job will check for new domain members once every 24 hours by default. To change the job settings, including the frequency at which it is run:

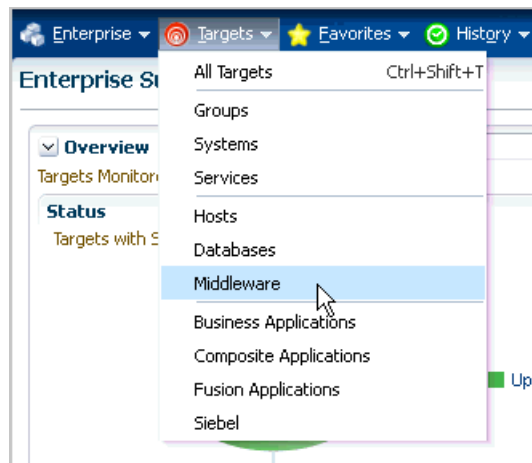
1. Click the **Jobs** tab.
2. Click the job title in the Job Activity page.
3. Click **Edit**.

2.3.2 Manually Checking for New or Modified Domain Members

You can use Cloud Control to check a domain for new or modified members on a periodic basis.


1. From the **Targets** menu, select **Middleware**.

Figure 2–10 Targets Menu

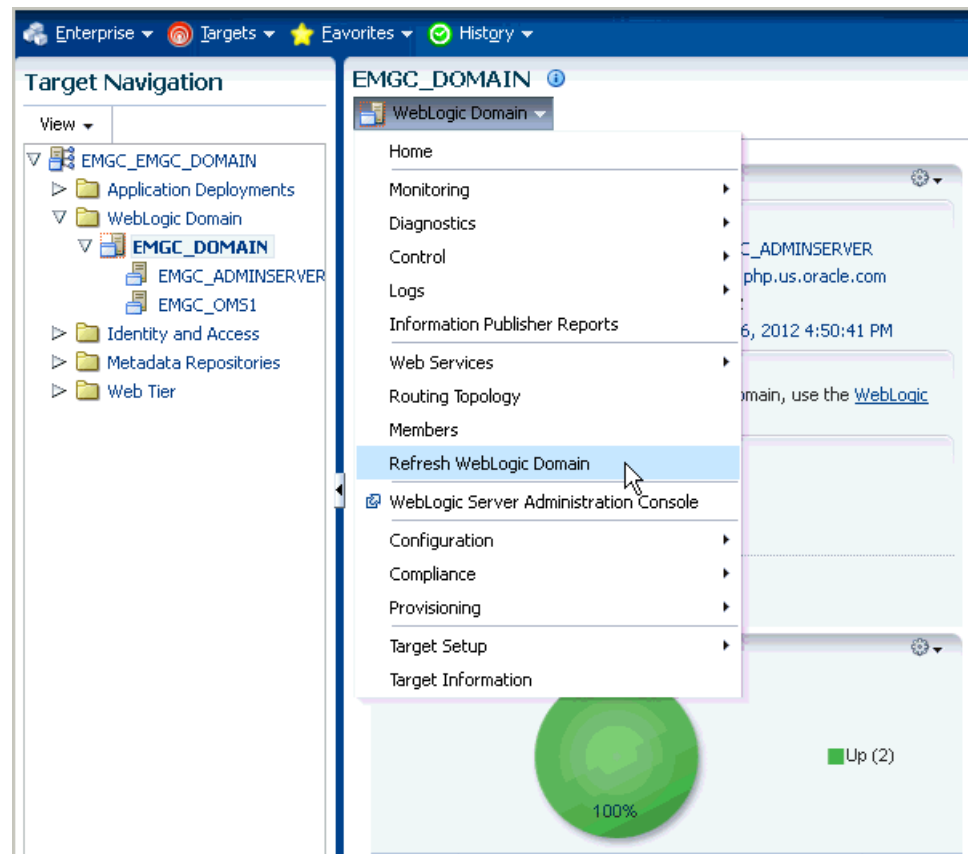


2. Click the WebLogic Domain you want to (enable the job for in the Middleware home page) refresh.

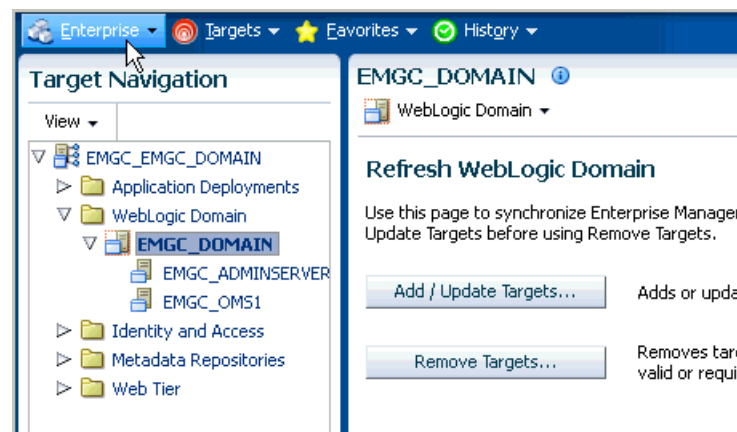
Figure 2–11 WebLogic Domain to Enable



3. From either the **Farm** or **WebLogic Domain** menu, select **Refresh WebLogic Domain**. The Refresh WebLogic Domain dialog opens.

Figure 2–12 Refresh WebLogic Domain Menu Option

4. Click **Add/Update Targets**. The Management Agent refreshes by connecting to the Administration Server. The Administration Server must be up for the refresh to occur. Click **Close** on the Confirmation page. Cloud Control will search the domain for new and modified targets.

Figure 2–13 Add/Update Option on Refresh WebLogic Domain Page

5. The Assign Agents page displays the Fusion Middleware targets discovered and the Management Agent assigned to each. Click **Add Targets** to assign Management Agents as listed in the Assign Agents page.

Agents are automatically assigned as follows:

- If a local Agent can be found on the target host, that Agent will be assigned.
- If a local Agent cannot be found on the host, the Agent specified in the Targets page will be assigned.

Note that you can also manually assign Agents to specific targets, if desired.

You can also modify the Domain Global Properties, for example, Contact, Cost Center, Lifecycle Status, and so on).

Figure 2–14 Assign Agents Page

Refresh WebLogic Domain EMGC_DOMAIN(Oracle WebLogic Domain): Assign Agents Cancel Add Targets

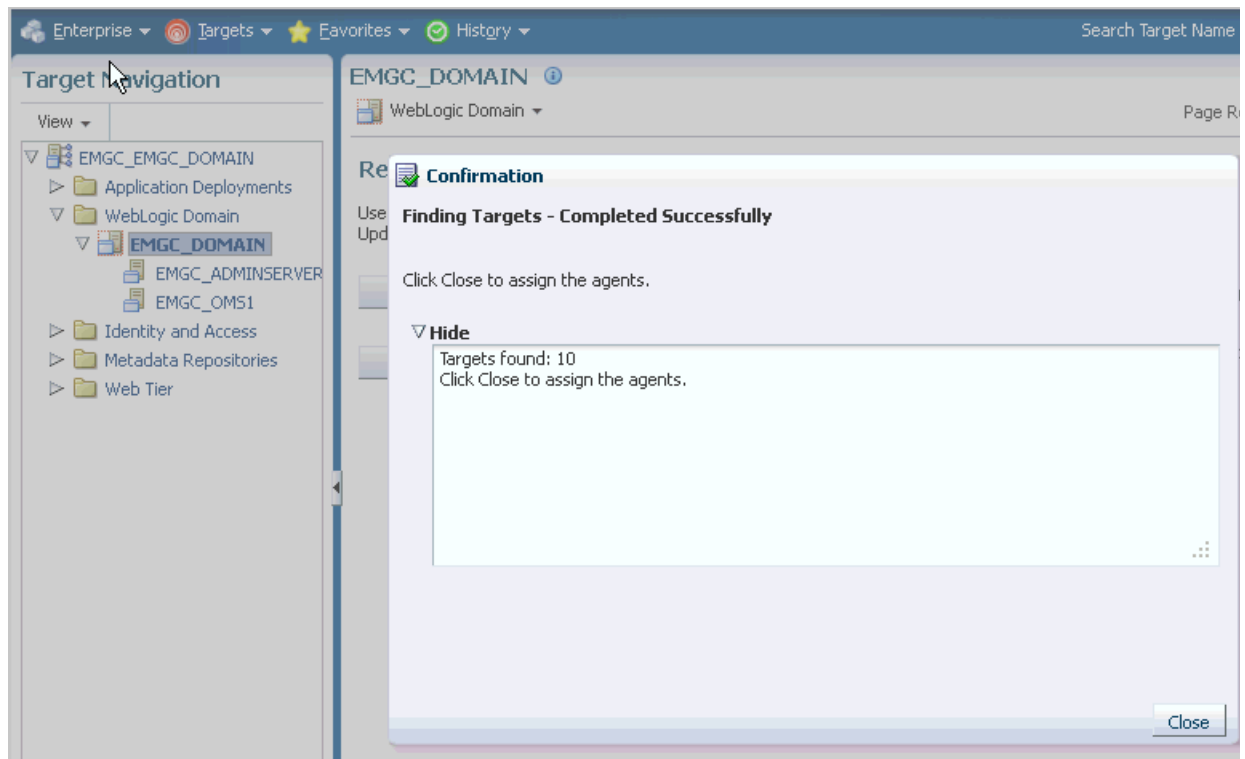
New Targets Found 0
Modified Targets Found 1
Targets Assigned To Local Agent 10

By default, each target is assigned a local agent that is on the same host as the target. If a local agent is not found, then the Agent monitoring the farm, <https://p.us.oracle.com:1838/emd/main/>, will be used.

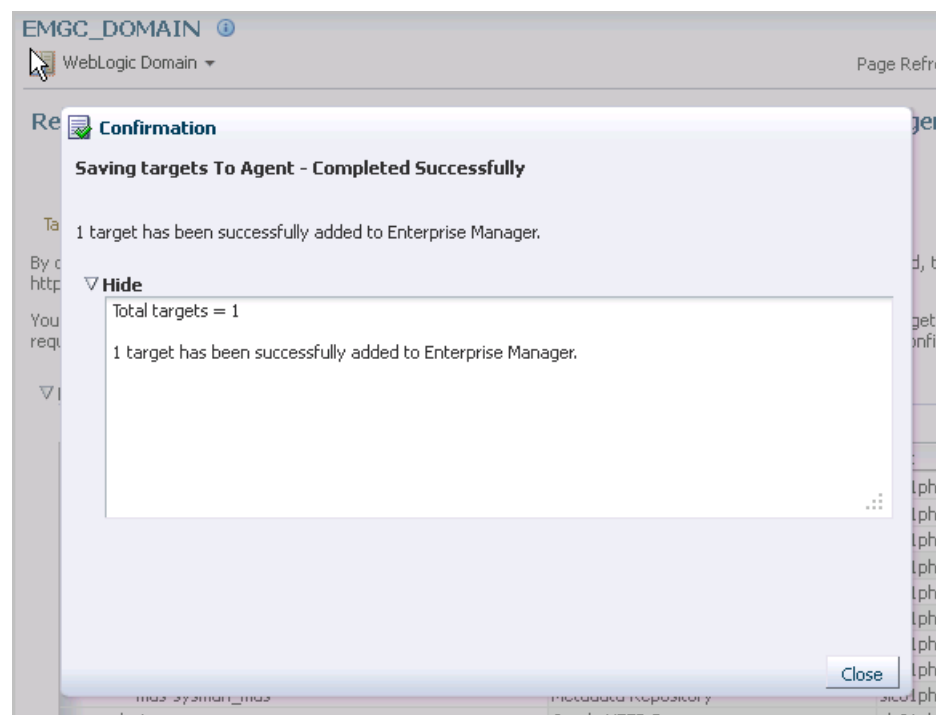
You can optionally override any agent assignment using the table below. All the agents used for monitoring the targets must be up. For the targets which require local agent, by default, only local agents are assigned. For such targets, if local agent is not found, the "Configured Agent" column is left blank.

▼ Hide Targets And Agent Assignments

Target Name	Target Type	Host	Configured Agent
▼ EMGC_EMGC_DOMAIN	Oracle Fusion Middleware Farm	b.us.oracle.com	b.us.oracle.c
▼ EMGC_DOMAIN	Oracle WebLogic Domain		
EMGC_ADMINSERVER	Oracle WebLogic Server		
▼ EMGC_OMS1	Oracle WebLogic Server		
emgc	Application Deployment		
empbs	Application Deployment		
oracle.security.apm(11.1.1.3.0)	Oracle Authorization Policy Manager		
mds-owsm	Metadata Repository		
mds-sysman_mds	Metadata Repository		
ohs1	Oracle HTTP Server		

Figure 2–15 Confirmation of Finding Targets

6. The Saving Targets to Agent processing window appears, indicating how many total targets have been added and successfully saved. It will also indicate the number of targets that were unsuccessfully added.

Figure 2–16 Confirmation Saving Targets to Agent

7. Click **Close** in the processing window when finished. The Results page displays the following options: Show Targets Details and Show Weblogic Domain Global Properties. The Show Targets Details page shows the targets and Agent assignments.

Note: If there were targets that were removed, you can go back to the Refresh WebLogic Domain page and click **Remove Targets** to remove the targets and any historical information in the Management Repository. See [Removing Targets](#).

Figure 2–17 Refresh WebLogic Domain Results Page



8. Click **OK** when finished. There may be a delay before these targets are visible and monitored. All the agents used for monitoring the targets must be up.

2.3.3 Removing Targets

Removing targets from the Management Repository:

- Identifies targets that are deleted from the Weblogic Domain, for example, WebLogic Servers, Clusters, Applications (both generic and custom), and any other System Components.
- Shows the list of targets which *might* have been deleted from the product, but Enterprise Manager cannot determine if they were deleted or not. For these targets, decide whether these targets should be deleted and mark them as such.
- Shows duplicate targets. For example, if for the same application deployment there is a custom and a generic target, the will shows the generic target which can be deleted.
- Shows the older versioned application deployments which can be deleted if a newer version of the same application is present.
- Lists all the down servers. You can decide to either blackout or delete these servers.

Managing Middleware Targets

This chapter describes how you can use Enterprise Manager to monitor Middleware software.

Note: Oracle provides a free self-paced course regarding the best practices on managing WebLogic and Service Oriented Architecture (SOA) applications and infrastructure. It consists of interactive lectures, videos, review sessions, and optional demonstrations, and lasts about two hours.

The *Oracle Enterprise Manager Cloud Control 12c: Best Practices for Middleware Management* self-study is available at <http://www.oracle.com/webfolder/technetwork/tutorials/tutorial/em/Oracle%20Enterprise%20Manager%20Cloud%20Control%2012c%20Middleware%20Management%20Best%20Practices%20Self-Study%20Course/player.html>

This chapter covers the following:

- [Middleware Targets in Enterprise Manager](#)
- [Monitoring Middleware Targets](#)
- [Diagnosing Performance Problems](#)
- [Administering Middleware Targets](#)
- [Lifecycle Management](#)
- [Managing Service Levels](#)
- [Job System](#)
- [Topology Viewer](#)
- [Support Workbench](#)

3.1 Middleware Targets in Enterprise Manager

After you have added a Middleware target (for example, Oracle Fusion Middleware, Oracle WebLogic Domain, Oracle Application Server, JBoss Application Server), you can view general information about the targets including their status and availability on the Middleware page. You can drill down into each target to get further details like how the target is performing, where it is deployed, the version, location of its home directory, and so on.

You can also view the number of critical, warning, and error alerts generated for the past 24 hours. These alerts indicate that a particular metric condition has been encountered. For example, an alert is triggered when a metric threshold is reached. Using these details, you can drill down to investigate the target and the problem that triggered the alert.

You can monitor the following components using Oracle Enterprise Manager Cloud Control:

- Oracle Fusion Middleware Components
- Oracle Application Server Components
- Non-Oracle Middleware Components

3.1.1 Oracle Fusion Middleware Components

A farm is a collection of components managed by Fusion Middleware Control. It can contain one Oracle WebLogic Domain, one Administration Server, one or more Managed Servers, and the Oracle Fusion Middleware components that are installed, configured, and running in the domain. You can monitor the following Oracle Fusion Middleware components using Enterprise Manager:

- **Oracle WebLogic Domains, Clusters, and Managed Servers:** A WebLogic Server domain is a logically related group of WebLogic Server resources that you manage as a unit. A domain includes one or more WebLogic Servers and may also include WebLogic Server clusters. Clusters are groups of WebLogic Servers instances that work together to provide scalability and high-availability for applications. With Oracle Enterprise Manager, you can monitor and manage the farm, domains, clusters, servers, and deployed applications.
- **Oracle SOA Suite:** The Oracle SOA Suite enables services to be created, managed, and orchestrated into SOA composite applications. Composite applications enable you to easily assemble multiple technology components into one SOA composite application. Oracle SOA Suite plugs into heterogeneous infrastructures and enables enterprises to incrementally adopt SOA. You can:
 - Automatically discover and model SOA components such as BPEL Process Manager, Oracle Service Bus, Service Engines, and so on.
 - Monitor the health and performance of the SOA components.
 - Trace the flow of an instance across all SOA Infrastructure applications.
 - Create systems, services, and aggregate services.
- **Oracle WebCenter:** The Oracle WebCenter is an integrated set of components with which you can create social applications, enterprise portals, collaborative communities, and composite applications, built on a standards-based, service-oriented architecture. It combines dynamic user interface technologies with which to develop rich internet applications, the flexibility and power of an integrated, multichannel portal framework, and a set of horizontal Enterprise 2.0 capabilities delivered as services that provide content, collaboration, presence, and social networking capabilities. Based on these components, Oracle WebCenter also provides an out-of-the-box, enterprise-ready customizable application, WebCenter Spaces, with a configurable work environment that enables individuals and groups to work and collaborate more effectively.
- **Oracle Web Tier:** This consists of:
 - **Oracle HTTP Server:** Oracle HTTP Server (OHS) is the underlying deployment platform for all programming languages and technologies that

Oracle Fusion Middleware supports. It provides a Web listener and the framework for hosting static and dynamic pages and applications over the Web. Based on the proven technology of the Apache 2.2.10 infrastructure, OHS includes significant enhancements that facilitate load balancing, administration, and configuration. It also includes a number of enhanced modules, or mods, which are extensions to the HTTP server that extend its functionality for other enterprise applications and services. You can:

- * Automatically discover and monitor Oracle HTTP Servers running within the application servers.
 - * View a list of metrics to gauge the server performance and virtual host performance.
 - * View the top URLs being accessed.
 - * Perform the enterprise configuration management tasks like viewing, comparing, and searching configuration information.
- **Oracle Web Cache:** Oracle Web Cache is a content-aware server accelerator, or reverse proxy, for the Web tier that improves the performance, scalability, and availability of Web sites that run on any Web server or application server, such as Oracle HTTP Server and Oracle WebLogic Server. Oracle Web Cache is the primary caching mechanism provided with Oracle Fusion Middleware. Caching improves the performance, scalability, and availability of Web sites that run on Oracle Fusion Middleware by storing frequently accessed URLs in memory. You can:
 - * Automatically discover and monitor OracleAS Web Cache instances running within application servers.
 - * View the metrics associated with this target to analyze their performance.
 - * Perform enterprise configuration tasks like viewing, comparing, and searching configuration information.
 - **Oracle Identity Management:** This is an enterprise identity management system that automatically manages users' access privileges within the resources of an enterprise. The architecture of Oracle Identity Management works with the most demanding business requirements without requiring changes to existing infrastructure, policies, or procedures. It provides a shared infrastructure for all Oracle applications. It also provides services and interfaces that facilitate third-party enterprise application development. These interfaces are useful for application developers who need to incorporate identity management into their applications.
 - **Oracle Portal:** This is a Web-based tool for building and deploying e-business portals. It provides a secure, manageable environment for accessing and interacting with enterprise software services and information resources. A portal page makes data from multiple sources accessible from a single location.
 - **Oracle Forms Services** is a middle-tier application framework for deploying complex, transactional forms applications to a network such as an Intranet or the Internet. With Oracle Forms Services, business application developers can quickly build comprehensive Java client applications that are optimized for the Internet without writing any Java code, and that meet (and exceed) the requirements of professional user communities. These Java client applications are Web-deployed applications available on demand for rapid processing of large amounts of data and rapid completion of complex calculations, analysis, and transactions.

- **Oracle Coherence** is a component of Oracle Fusion Middleware that enables organizations to predictably scale mission-critical applications by providing fast and reliable access to frequently used data. By automatically and dynamically partitioning data in memory across multiple servers, Oracle Coherence enables continuous data availability and transactional integrity, even in the event of a server failure. As a shared infrastructure, Oracle Coherence combines data locality with local processing power to perform real-time data analysis, in-memory grid computations, and parallel transaction and event processing. Oracle Coherence comes in three editions. You can:
 - Discover and manage a Coherence cluster and its various entities.
 - Monitor and configure various components such as nodes, caches, services, connections, and connection manager instances of a Coherence cluster.
 - Deploy and install a Coherence node based on the Provisioning Advisory framework.
- **Oracle Business Intelligence** is a complete, integrated solution that addresses business intelligence requirements. Oracle Business Intelligence includes Oracle Business Intelligence Reporting and Publishing, Oracle Business Intelligence Discoverer, and Oracle Business Intelligence Publisher. You can:
 - Manually discover Oracle BI Suite EE targets, and monitor their overall health.
 - Diagnose, notify, and correct performance and availability problems in Oracle BI Suite EE targets.
 - Access current and historical performance information using graphs and reports.
 - Perform enterprise configuration management tasks like viewing, comparing, and searching configuration information.
- **Oracle Universal Content Management System** provides a unified application for several different kinds of content management. It is an enterprise content management platform that enables you to leverage document management, Web content management, digital asset management, and records retention functionality to build and complement your business applications. Building a strategic enterprise content management infrastructure for content and applications helps you to reduce costs, easily share content across the enterprise, minimize risk, automate expensive, time-intensive and manual processes, and consolidate multiple Web sites onto a single platform for centralized management. Through user-friendly interfaces, roles-based authentication and security models, Oracle Universal Content Management empowers users throughout the enterprise to view, collaborate on or retire content, ensuring that all accessible distributed or published information is secure, accurate and up-to-date.

3.1.2 Oracle Application Server Components

You can monitor Oracle Application Server 10g components like Oracle Application Server Farms, Oracle Application Server Clusters, Oracle Application Servers, OC4J, Oracle HTTP Servers, Oracle Web Cache, Oracle Portal, Oracle Wireless, Oracle Forms Services, Oracle Reports Services, Oracle Business Intelligence, and Oracle Identity Management.

3.1.3 Non-Oracle Middleware Components

In addition to monitoring Oracle Middleware components, Enterprise Manager can also be used to monitor non-Oracle Middleware software. The third-party Middleware software that can be monitored includes the following:

- WebSphere Application Server
- WebSphere MQ
- JBoss Application Server
- Apache Tomcat
- Apache HTTP Server
- Microsoft Exchange Server
- Microsoft Internet Information Services
- Microsoft Active Directory
- Microsoft Commerce Server
- Microsoft BizTalk Server
- Microsoft Internet Security and Acceleration
- Microsoft .NET Framework

3.2 Monitoring Middleware Targets

Enterprise Manager organizes a wide variety of performance data and administrative functions into distinct, Web-based home pages for the domain, servers, components, and applications.

3.2.1 Middleware Target Home Page

Enterprise Manager provides centralized monitoring across domains, configuration management, provisioning, real time and historical performance analysis. You need to drill down to Oracle Fusion Middleware Control to perform administrative tasks and manage components.

The Middleware target home page provides an overview of middleware components configured as managed targets. You can use the Table tab to add or remove middleware targets, as well as set certain monitoring configuration properties for targets, from this page.

Note: Enterprise Manager no longer creates farm targets for Fusion Middleware release 12.1.0.4 domains.

The following details are displayed for middleware targets. Additional information can be displayed through the View Column menu:

- Type: The type of target being managed.
- Status: The availability of the target, if applicable. Note that some targets that represent a collection of components, such as a Fusion Middleware Farm, will not have a standalone status.
- Status Details: The availability of the middleware components associated with the target. The total number of components associated with the target, such as the

number of WebLogic Server instances associated with a WebLogic Domain, is shown outside of the parentheses.

- **Version:** The target version.
- **Compliance Score:** An overall evaluation of the target's compliance with policies defined in your enterprise, presented as a percentage of compliance. A compliance score of 100% indicates full compliance with a policy.

Heatmap

You can use the Graph tab to view the Middleware Targets Heatmap, a graphical representation of a set of targets depicted as boxes in the heatmap which are the root targets that are shown in the table tab. They can be grouped and optionally summarized by attributes like Version and Location. The size of the box represents the number of member targets. You can choose to color the boxes based on either the percent of members that are up or on average WebLogic CPU usage. Hover, click and double-click on graph elements to see more detail.

You can select **Area**, which limits which target types are displayed at a given time on the heatmap and table tabs. If you choose **All**, all target types are displayed. Other options include WebLogic Servers, WebLogic Full Hierarchy, GlassFish, JBoss, and so on. When an area is chosen that only contains one target type (such as the WebLogic Servers area), the heatmap displays that target type as the root node and does not display its parents.

You can use the **Show** drop-down menu to change to either of the two displays. When you choose CPU Usage for WebLogic Servers, the graph displays a heatmap with boxes that are root targets containing WebLogic servers. If a root target does not contain any WebLogic servers, it does not show up in the view. A box is displayed if the root target contains a WebLogic server that has a reported CPU value, not necessarily just a server. This means that servers that are down or blacked out are not counted in either the size of the box or the averaging. Also, if a server does not have a CPU value, which occurs if that metric has not yet been collected, or is blacked out, it will not be counted.

The box size is based on the number of targets and the box color is based on the metric value of all the targets it contains.

When you click a box, the selection details on the right display the number of WebLogic servers it contains, as well as the average CPU Usage value. Similarly, tooltips also display this information. The selection details also display a Properties section that shows the **Type** and **Target Version**, as well as a Monitoring and Diagnostics section that shows the **Descendant Target Incidents** and the **Configuration Changes**. The Properties section also displays any user-defined properties such as Contact, Location, Department, and so on, if they have been defined. You can click on the number for the incidents to display the Incidents Manager page where you can search, view, and manage exceptions and issues and you can track outstanding incidents and problems.

You can organize the display by using the **Organize First By** field and the **Then By** field, which allows you to choose a field on which to prioritize the display. Drilling allows you to focus on one section of the heatmap that was grouped using the Organize By menus. To do this, double-click on the header of the Organize First By box. This displays only the boxes that are in that box and hides all others. Click on the breadcrumbs to return to the original view. Summarizing turns the deepest Organize First By box into one box by summarizing all of the individual boxes it contains.

You can adjust the sliders to change color thresholds. Note that the color scale changes between the Percent of Targets Up and the CPU views. A lower value corresponds to

green, while larger numbers correspond to red. You can take advantage of the "Hide dark green boxes" option. The dark green corresponds to the darkest green on the color slider. For example, if your higher limit is 95 and you are showing Percent of Targets Up, all boxes with an average Percent of 95 or greater are hidden from view.

Searching Middleware Managed Targets

To minimize the number of targets displayed in the table and graph, and improve page performance and usability, use the Search function.

The Area drop-down list is used to specify a set of target types, for example Oracle WebLogic Servers. Only target types in that set will be returned when you click **Search**. Areas only appear in the drop-down list if you have access to at least one target in that area.

The search results are a hierarchy where the leaf nodes match all search fields. The leaf nodes are shown in context with parents who are only required to match the Area drop-down list's target types constraint.

Note: If you are searching for a single target and do not need hierarchy information, the **Search Target Name** option located in the upper right is available on most pages.

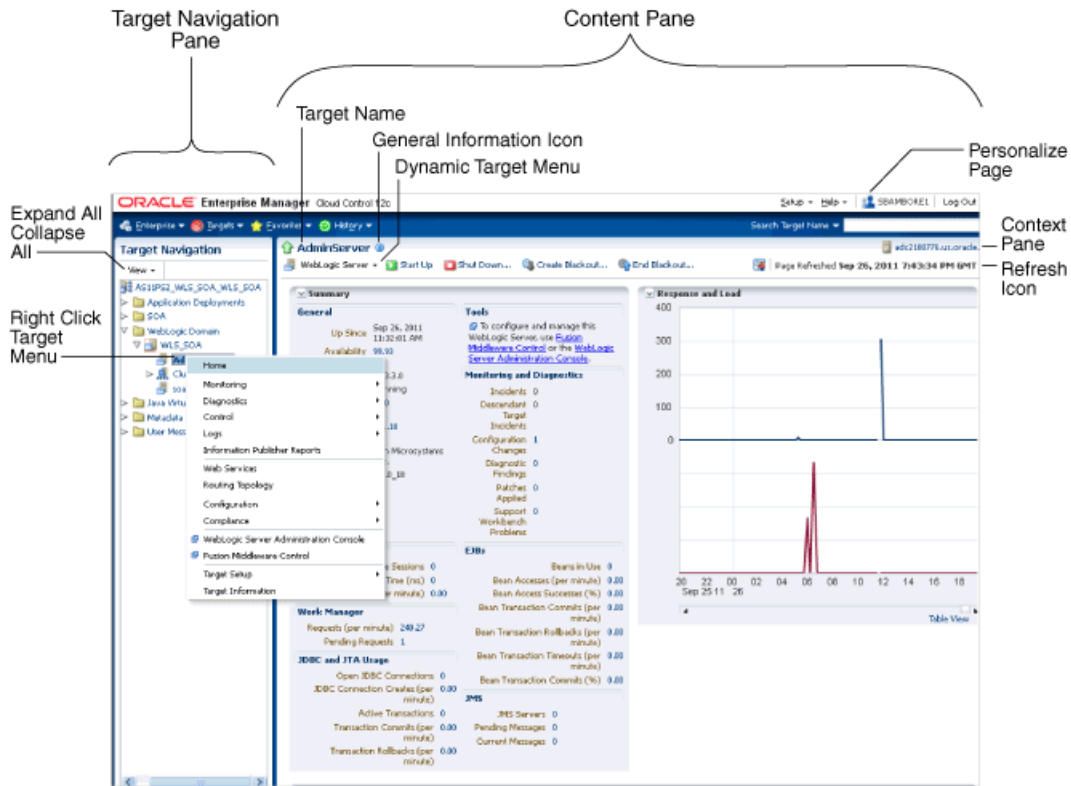
In addition:

- Though only the Area and Name appear when the Search region is collapsed, all the hidden search options still apply when you click **Search**.
- If you change the Area or some other search field, click **Search** to update the page contents. Clicking **Search** also saves your search criteria as the new default search the next time you visit the page.
- The Member Status Summary column in the table summarizes only the targets fetched by your search criteria. For example, if you decide to search the 'Oracle WebLogic Domains, Clusters, and Servers' area for targets with contact *Smith*, only domains, clusters, and servers matching *Smith* and their parents would be fetched and used to calculate the Member Status Summary column numbers. Targets which do not match Smith will not be shown or used in the summary column calculations.

3.2.2 Target Home Page

The Home pages make it easy to locate the most important monitoring data and the most commonly used administrative functions—all from your Web browser.

When you login into Enterprise Manager and select a Middleware target, the Home page for the target is displayed. For example, when you click on a WebLogic Server target in the Middleware page, the following screen is displayed.

Figure 3–1 WebLogic Server Home Page

This figure shows the target navigation pane on the left and the content page on the right. From the target navigation pane, you can expand the tree and select a component or an application. When you select a target, the target's home page is displayed in the content pane and that target's menu is displayed at the top of the page, in the context pane. You can also view the menu for a target by right-clicking the target in the navigation pane.

In the preceding figure, the following items are called out:

- **Target Navigation Pane** lists all of the targets in a navigation tree
- **Personalize Page Menu** provides a list of operations that you can perform on your environment such as, change your password and email address, choose your accessibility preferences, and so on. You can also review your Enterprise Manager privileges using the Entitlement Summary page.
- **Content Pane** shows the current page for the target. When you first select a target, that target's home page is displayed.
- **Dynamic Target Menu** provides a list of operations that you can perform on the currently selected target. The menu that is displayed depends on the target you select. The menu for a specific target contains the same operations as those in the Right-Click Target Menu.
- **Right-Click Target Menu** provides a list of operations that you can perform on the currently selected target. The menu is displayed when you right-click the target name in the target navigation pane. In the figure, even though the WebLogic Server is selected and its home page is displayed, the right-click target menu displays the operations for the selected target.
- **Target Name** is the name of the currently selected target.

- **Context Pane** provides the host name, the time of the last page refresh, the Refresh icon, and the Personalize Page icon.
- **View:** You can select options to Expand All / Collapse All, Scroll First, and Scroll Last in the navigation tree.
- **Refresh** icon indicates when the page is being refreshed. Click it to refresh a page with new data. (Refreshing the browser window refreshes the page but does not retrieve new data.)

From the Home page, you can also access the Fusion Middleware Control and WebLogic Server Administration Console by clicking on the appropriate link or selecting the appropriate menu item on the page.

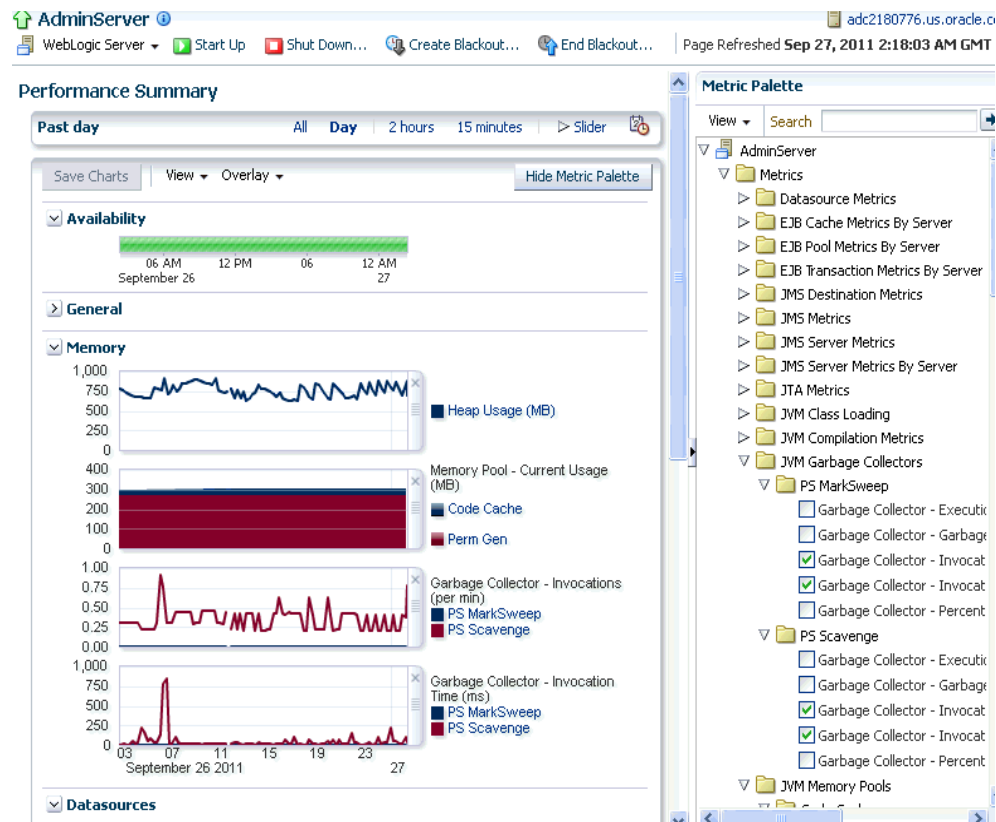
3.2.3 Out-of-box Metrics

Enterprise Manager provides a set of pre-defined performance metrics for each Middleware target. These metrics are available for domains, cluster, server, applications, web services, resources etc. The metric data is collected and stored in the Management Repository. For more details on the pre-defined metrics, see the *Oracle Fusion Middleware Metric Reference Guide*.

For example, Enterprise Manager can automatically monitor:

- The CPU or memory consumption of the application server, including detailed monitoring of individual Java Virtual Machines (JVMs) being run by Oracle WebLogic servers.
- Java EE application responsiveness from the application down through individual servlets and Enterprise JavaBeans (EJBs)
- Oracle HTTP Server session volumes, connection duration, and error rates
- Oracle Web Cache hit rates and volumes
- Top servlets based on number of requests, maximum processing time, and highest average processing time

The performance metrics provide details about the metric as a current real time value (30 seconds, 1 minute, or 5 minutes) or a previous value (past 24 hours, 7 days, or 31 days). The historical information is displayed as graphs and a table. By using graphs, you can easily watch for trends, and by using tables, you can examine details of past metric severity history. The out-of-box metrics can be viewed from the performance summary pages as shown below:

Figure 3–2 Performance Summary Page

You can change which charts are displayed on the performance page and then save the changes on a per-user, per-target-type basis. You can also save multiple customized versions of a performance page, giving each version a name. This will save time by allowing quick access to previously created version of the page. The Performance Summary feature allows you to create named chart views. The generic performance page is always shown in the context of one primary target. However, the performance of that target may be dependent on, or affect the performance of other targets. To explore these relationships you can chart metrics for multiple related targets on one performance page. The Performance Summary feature allows you to chart metrics for multiple related targets.

3.2.4 Analyzing Historical Performance

Enterprise Manager allows you to analyze historic metric data and perform trend analysis. In Fusion Middleware Control, you cannot analyze historical metric data and the real-time analysis is limited to a single domain. But in Enterprise Manager, the metrics are collected and stored in the Management Repository, so you can analyze the data well after the situation has changed. For example, you can use historical data and diagnostic reports to research an application performance problem that occurred days or even weeks ago.

You can even provide a customized time period for which the data should be retrieved from the Management Repository. You can customize the time period for:

- Pre-defined range of the last 24 hours, last 7 days, or last 31 days
- Customized range of any number of days, weeks, months, or years
- Any start date and end date (such that the duration is not greater than 99 years)

Note: You can analyze historical metric data with Enterprise Manager only. You cannot use Fusion Middleware Control to analyze historic performance.

3.2.5 Setting Metric Thresholds for Alert Notifications

Metric snapshots are named snapshots of a target's past performance. You can use metric snapshots to calculate thresholds based on deviations from this past performance. Thresholds are boundary values against which monitored metric values are compared. You can specify a threshold such that, when a monitored metric value crosses that threshold, an alert is generated. You can get critical alerts when a monitored metric has crossed its critical threshold or warning alerts when a monitored metric has crossed its warning threshold.

Enterprise Manager provides a comprehensive set of features that facilitates automated monitoring and generation of alerts. You can gather and evaluate diagnostic information for targets distributed across the enterprise, and an extensive array of Middleware performance metrics are automatically monitored against predefined thresholds. By selecting a metric, you can determine whether the thresholds have been defined for a particular metric. These thresholds are used as a mechanism to generate alerts. These alerts in turn are used to notify you whether a target is up or down, the response time is slow, and so on. Thus, you can monitor their overall performance.

You can set up corrective actions to automatically resolve an alert condition. These corrective actions ensure that routine responses to alerts are automatically executed, thereby saving you time and ensuring that problems are dealt with before they noticeably impact the users.

3.2.6 Monitoring Templates

You can also use monitoring templates to simplify the task of standardizing monitoring settings across your enterprise. You can specify the monitoring settings once and apply them to all Oracle Fusion Middleware targets. A Monitoring template defines all the parameters you would normally set to monitor Middleware target, such as:

- Target type to which the template applies
- Metrics (including user-defined metrics), thresholds, metric collection schedules, and corrective actions

When a change is made to a template, you can reapply the template across the affected targets in order to propagate the new changes. You can reapply monitoring templates as often as needed.

3.2.7 Managing and Creating Blackouts

Enterprise Manager comes with a bundle of performance and health metrics that enable automated monitoring of application targets in your environment. When a metric reaches the predefined warning or critical threshold, an alert is generated and the administrator is notified.

However, there are occasions when you want to perform maintenance work on your Middleware targets, and not want any alerts to be generated while you are bringing them down. In this case, you can schedule a blackout and suspend monitoring of the Middleware targets.

Blackouts allow you to suspend any data collection activity on one or more monitored targets, thus allowing you to perform scheduled maintenance on targets. If you continue monitoring during these periods, the collected data will show trends and other monitoring information that are not the result of normal day-to-day operations. To get a more accurate, long-term picture of a target's performance, you can use blackouts to exclude these special-case situations from data analysis. Enterprise Manager allows you to define new blackouts; view the status of existing blackouts; and edit, stop, and delete blackouts that are not required.

3.2.8 Extend Monitoring for Applications Deployed to WebLogic Server

Many administrators often require custom logic to be written to check for conditions specific to their application environments. Enterprise Manager allows integration of application instrumentation in the Enterprise Manager event monitoring infrastructure. If application developers expose application instrumentation using standards like JMX or Web Services operations, then you can build management plug-ins for the instrumentation using easy-to-use command line tools, and leverage the Enterprise Manager event monitoring system to monitor it. You do not have to edit any XML files or write any integration code to integrate such instrumentation. Follow these procedures to integrate application-defined instrumentation:

- Use Command Line Interfaces that analyze MBean interfaces for JMX and WSDL for Web Services and create management plug-ins
- Import Management Plug-in Archive in Enterprise Manager
- Deploy Management Plug-in to Management Agents
- Create Target-type instances for the target types defined in Management Plug-in Archive
- Leverage the Enterprise Manager event monitoring system including monitoring templates, corrective actions, historical and real time metric views, alerts, customization of notification rules, and methods on events generated from application instrumentation metrics.

Administrators are able to add performance metrics beyond those available out-of-box for JMX-instrumented applications deployed on Oracle WebLogic Server. Administrators can additionally monitor JMX-enabled applications by defining new target type that can be monitored using management plug-ins, and then use a command line tool `emjmxcli` to automate the generation of the target metadata and collection files. All JMX-enabled applications deployed to the WebLogic Server can be consolidated and monitored by a single management tool, Enterprise Manager.

3.2.9 Request Monitoring

Request Monitoring provides end-to-end visibility into requests and helps localize end-user performance problems based on the deployment model. You can monitor, understand, and visualize how servers interact with each other to deliver business end-user services requests. You can trace important end-user requests from the client to the service endpoint across all the servers and applications associated with each transaction. The trace data is used to model the call-path for these requests where call-paths represent the inter-server relationship and performance metrics for these requests. You can only monitor synchronous transactions running on WebLogic servers. With Request Monitoring, you can:

- Trace end-user transactions and capture the complete call-path of important requests.

- Reduce problem localization by diagnosing poorly performing requests and identifying the servers whose behavior (service time) deviates most from the norm.
- Reduce fault discovery time and set performance level objectives on request response times.
- Reduce fault-reason identification time by allowing the user to launch into Oracle Enterprise Manager JVM Diagnostics feature. This allows performance diagnostics to be done within the context of the problem experienced as well as historical statistics of the system when it was behaving normally.

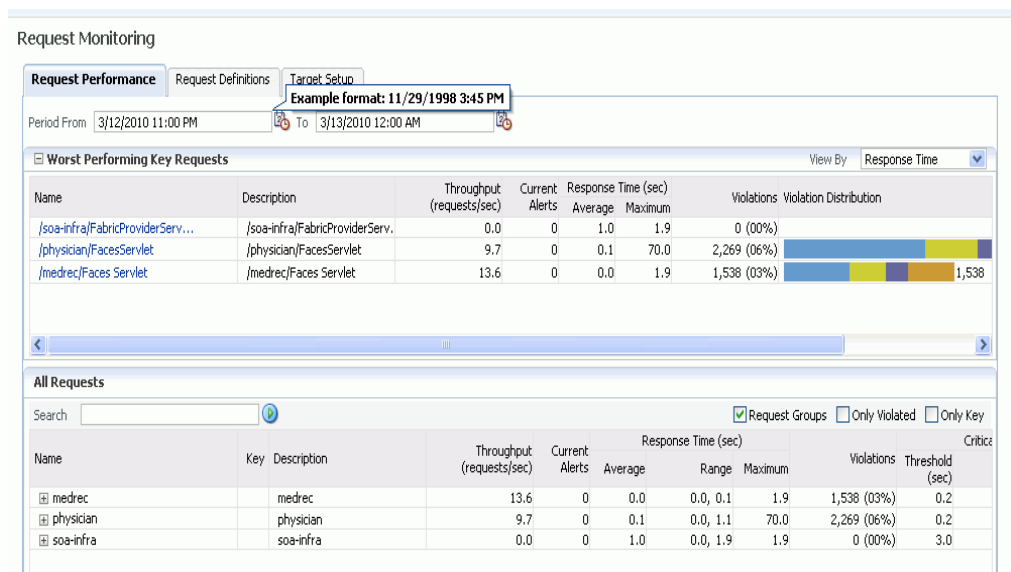
3.2.9.1 Defining and Managing Transactions Being Monitored

A request is a server entry-point that can be invoked by a Web Browser or an application. A group of request instances related to each other by some common attributes is known as a request. A collection of requests is known as a transaction group. A transaction can belong to one or more transaction groups. Violations occur when the critical and warning thresholds are exceeded. The administrator can mark one or more important requests as key requests. Alerts are generated for key requests only.

You can define requests that need to be monitored and tracked and group them for reporting purposes. A set of pre-defined requests and request groups are provided by default. You can create a new request, edit or delete an existing transaction. To define a request:

1. From the **Targets** menu, select **Middleware**.
2. From the **Middleware Features** menu, select **Request Monitoring**.
3. Click **Request Performance** to view the performance of the request during a specific time period.

Figure 3–3 Request Performance Page

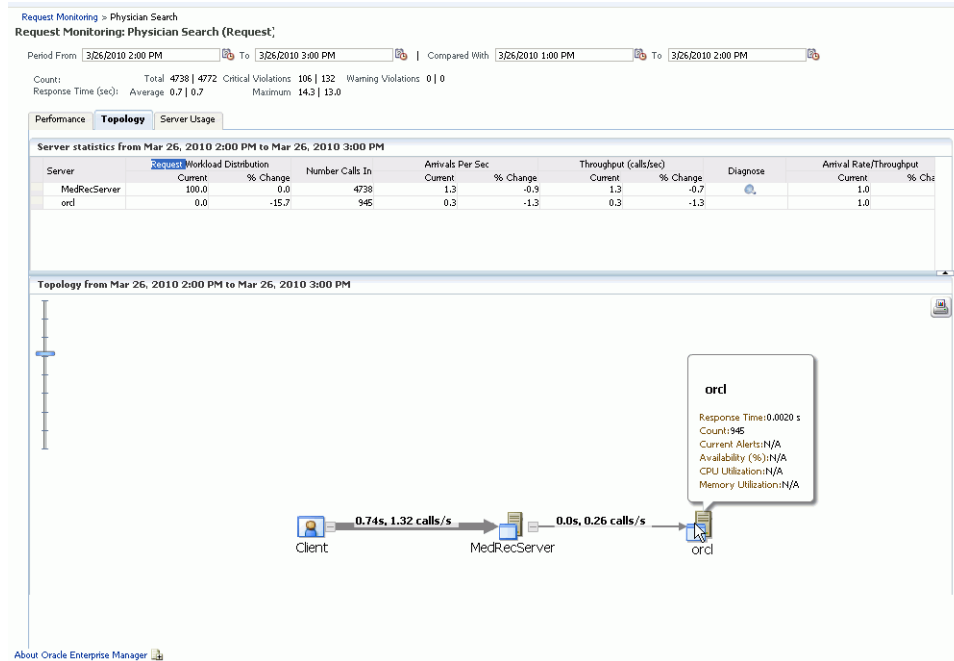


4. Click on a request to drill down to the Performance tab. The performance charts show the volume of requests over a time period and the average response time of all requests during this period. Two performance charts are displayed for the

current and the comparison period. The charts provide a historic response time analysis of the request over the selected period of time.

- Click on the **Topology** tab. The following screen is displayed.

Figure 3–4 Request Monitoring: Topology Page



The topology chart shows the aggregate server level call paths of the request. The chart shows the calls made between the clients and the clusters or servers. Click on the icon on a cluster to view all the servers that are part of the cluster. Place the cursor on a cluster or server to view the following details:

- **Response Time:** The average response time taken by the server to process all the calls for the selected period.
 - **Count:** The total number of calls that have been made to the server.
 - **Current Alerts:** These alerts show the status of the server or cluster and indicate whether it is up or down, blacked out, or unreachable.
- Click on the **Server Usage** tab. The following screen is displayed.

Figure 3–5 Server Usage Page

These charts show the workload distribution for the selected request in each server. The charts display the average of the percentage contribution of each server or cluster in the selected time period. You can see the workload distribution for good requests, requests that have violated the warning thresholds, and requests that have violated the critical thresholds.

3.3 Diagnosing Performance Problems

This section describes the methods and tools used to diagnose performance problems. You can:

- View the list of most active Servlets and JSPs and identify the ones that are causing the bottleneck.
- Analyze Java EE and SOA applications using Application Dependency and Performance.
- Use Java Diagnostics to diagnose performance problems in production.

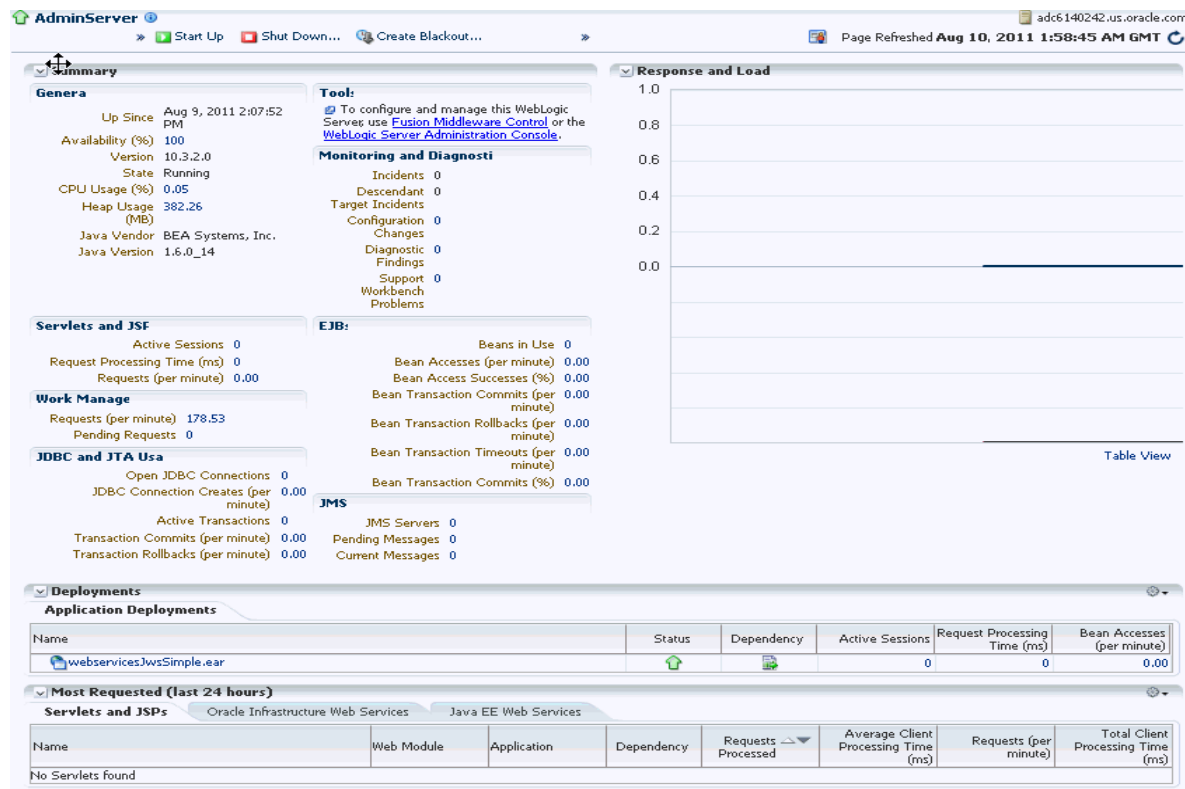
3.3.1 Using Home Pages to Diagnose Performance Issues

When you are troubleshooting performance problems, it can be helpful to know which servlets or JSPs are the most active. By viewing the Most Requested section on the WebLogic Server Home page, you can identify the most active Java servlets, JSPs, Web Services, or Java EE Services running on the WebLogic Server instance.

When you receive an alert notification, Enterprise Manager makes it easy for you to investigate the problem and take corrective actions wherever required. For example, notification of excessive CPU consumption by WebLogic Server may lead to investigation of the applications running on that instance. By using the Servlets and JSPs tab in the Most Requested section of the WebLogic Server Home page, you can quickly identify the highest volume or least responsive application. You can then drill

down and diagnose application's servlets, Java Server Pages (JSPs), or EJBs to identify the bottleneck.

Figure 3–6 WebLogic Server Home Page



3.3.2 Diagnostics Snapshots

A diagnostic snapshot consists of all the necessary data to diagnose an issue. The actual diagnostic snapshot data depends on what targets are included in generating the diagnostic snapshot. It also provides a collective snapshot of both JVM and WebLogic Server diagnostics and log data that can be exported or imported into other Cloud Control systems for analysis at a later date. This allows administrators to determine the root cause of problems and ensure that they do not occur again. These snapshots supplement the Fusion Middleware Support Workbench feature which now includes attaching diagnostic snapshots to Support Requests.

Diagnostic snapshots can be generated in the context of one or more Enterprise Manager targets like WebLogic Java EE Server, Java EE Application, Fusion Java EE Application, or Custom Java EE Application targets. These targets can be part of one single WebLogic Domain or multiple WebLogic Domains.

When generating the diagnostic snapshot, you can name the diagnostic snapshot, select the targets that should be used for generating the diagnostic snapshot, select the duration during which the data will be collected for the snapshot and also select an option to either import the generated diagnostic snapshot data into the same Enterprise Manager instance or export the generated diagnostic snapshot data into single or multiple files that can then be imported back into another Enterprise Manager instance (or the same Enterprise Manager instance) later.

3.4 Administering Middleware Targets

IT organizations typically have several WebLogic Domains - spanning test, stage, and production environments - to manage and administer on a regular basis.

Remembering details (such as URLs and credentials) for each of these domain's administration consoles can be difficult, and logging on to the appropriate console each time an administrative operation needs to be performed can be tedious.

Enterprise Manager Cloud Control addresses these challenges by exposing common WebLogic administration operations using its console directly; thereby, removing the need to drill down to the Oracle WebLogic Server Administration Console or to the Oracle Enterprise Manager Fusion Middleware Control console. Administration operations available directly from the Cloud Control console - beginning with Oracle Enterprise Manager Cloud Control 12c Release 3 and the Fusion Middleware Plug-in Release 12.1.0.4 - include the following:

- Locking a domain configuration using the Change Center prior to making configuration changes to prevent other administrators from making changes during their edit session. Administrators can continue to manage the changes using the Change Center by understanding which server instances need to be restarted for configuration changes to take effect, by releasing a lock, by activating changes, or by undoing changes.
- Viewing, configuring, and using MBeans for a specific Oracle WebLogic Server or Application Deployment target using the System MBean Browser.
- Creating, editing, deleting, controlling, or testing JDBC data sources.
- Recording configuration actions performed from within the Cloud Control console as a series of WebLogic Scripting Tool (WLST) commands, and then using WLST to replay the commands to help automate the task of configuring a domain.
- Configuring log file settings such as log file location, format of messages (for example, Oracle Diagnostic Logging - Text, Oracle Diagnostics Logging - XML), log level for both persistent loggers and active runtime loggers, and rotation policy (either size based or time based). Such settings are available for log files for the following Fusion Middleware target types: Oracle WebLogic Server, Application Deployment, SOA Infrastructure, Essbase Server, Directory Integration Platform Server, Oracle Virtual Directory, Oracle Reports Application, Oracle Reports Bridge, Oracle Reports Server, and Oracle Reports Tools.
- Performing selective tracing to gain more fine-grained logging data that is limited to a specific application name or other specific attributes of a request (for example, user name or client host).
- Starting, stopping, or restarting managed servers, clusters, domains or other Fusion Middleware components (for example, Oracle HTTP Server) immediately or scheduling the operation to occur at a future point in time. Enterprise Manager Cloud Control provides flexibility in terms of how to control such processes. Users can choose to control processes by using Node Manager, default scripts located in the domain home, or custom scripts. Relying on scripts to perform process control requires the boot.properties file to be properly configured. The process control feature is accessible from the target home page, target menu, or from the Job Library using the out-of-the-box Fusion Middleware Process Control job.

Middleware administration features (with the exception of the process control feature), require the Oracle Java Required Files (JRF) to be enabled for the WebLogic Domain to be administered. In addition, for the previously listed administration features (with the exception of the process control feature) the necessary domain credentials must be used. That is, once administrators access the administration functionality and

preferred credentials are not already specified for the domain, administrators must log in to the domain prior to performing such administrative operations.

3.5 Lifecycle Management

Enterprise Manager Cloud Control offers lifecycle management solutions that help you meet all lifecycle management challenges easily by automating time-consuming tasks related to discovery, initial provisioning and cloning, patching, configuration management, ongoing change management, and compliance management.

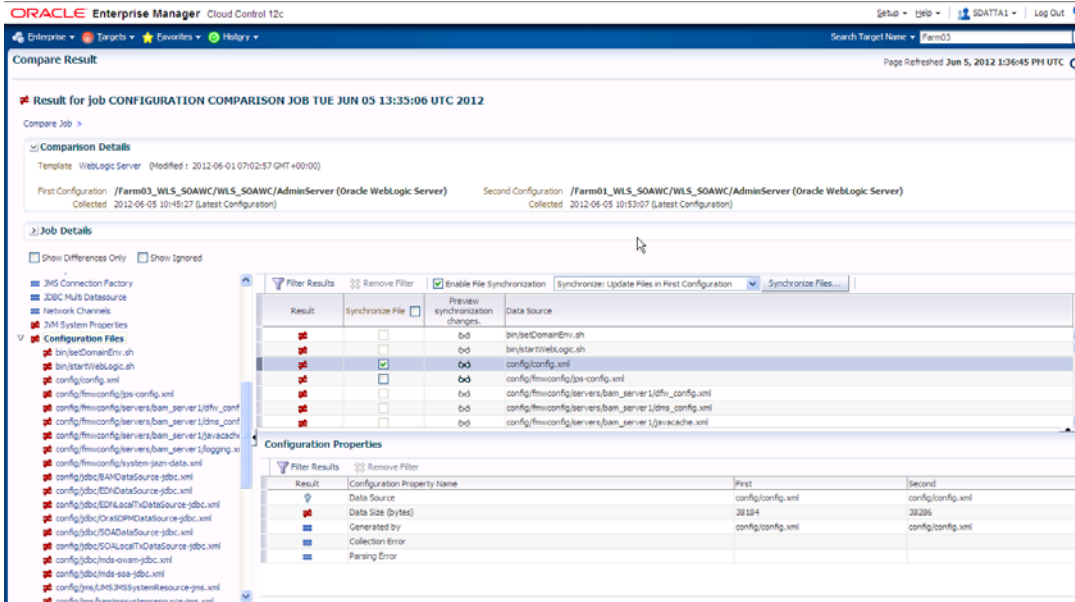
3.5.1 Managing Configurations

Enterprise Manager provides a suite of configuration management functions that can be performed on Middleware targets.

Oracle Management Agent collects configuration information about Oracle Fusion Middleware targets from their respective configuration files, and communicates this information over HTTP/HTTPS to Oracle Management Service, which stores it in the Management Repository. This information is periodically collected and updated while maintaining the audit of changes. Configuration files for Middleware targets are also collected. For example, for WebLogic Server, the `config.xml` configuration file is collected from the WebLogic Administration Server. The Enterprise Manager configuration management capabilities efficiently guide the users to desired configuration data in a particular component.

You can compare these configuration details and view the differences and similarities between the two instances of a Middleware target. You have the flexibility to compare two last collected configurations or two saved configuration files. You can also compare one configuration with multiple configurations or one configuration in the Management Repository with a saved configuration file. When a comparison operation results in differences that you do not require, you can synchronize the configurations so that one of configurations replaces the other one. This synchronization can be performed on demand based on the configuration files being compared.

Figure 3-7 Comparing Configuration Files



You can also compare configurations by using the default comparison templates. A comparison template is associated with a specific target type that determines the configuration item type and property that is to be compared. A template can specify rules or expressions that enable you parse comparison data and fine-tune comparisons. For example, you can specify rules that indicate which differences must trigger an alert and which differences must be ignored when the configuration is compared.

Using Enterprise Manager, you can search configurations across Middleware targets and find configuration anomalies - whether they are a mismatch of an install/patch version of Oracle Fusion Middleware software, or they are a mismatch of the software configuration data. You can perform more intelligent searches to identify all the components hosting a particular application or other resources. You can create and save more intelligent searches. For example, you can create a new search to retrieve all 10.3.5 WebLogic Server targets running on the Linux 64 bit platform that are using JDK 1.6.0_31.

In addition, for BPEL Process Manager targets, you can view the BPEL Processes, its different versions, and the suitcase files associated with each version. You can also compare the BPEL Process suitcase files of different versions and track the changes that were made to a version. shows you how the versions can be selected and compared. This allows you to identify the cause for improved or deteriorated performance due to a change in the BPEL Process suitcase file.

3.5.2 Compliance Management

Enterprise Manager Cloud Control offers the following compliance management solutions:

- The compliance results capability enables you to evaluate the compliance of Middleware targets and systems as they relate to your business best practices for configuration, security, and storage. In addition, compliance results provide advice of how to change configuration to bring your Middleware targets and systems into compliance.

- Using the compliance library, you can define, customize, and manage:
 - Compliance frameworks
 - Compliance standards
 - Compliance standard rules

By using these self-defined entities, you can test your environment against the criteria defined for your company or regulatory bodies.

3.5.3 Patch Management

Patching is one of the critical phases of the software lifecycle that helps you maintain the software over a period of time and keep it updated with bug fixes and latest features offered by the software vendor. However, in today's world, with numerous software deployments across your enterprise, patching becomes very complex and virtually impossible to manage. Considering these challenges, Enterprise Manager Cloud Control provides an integrated My Oracle Support Patches and Updates functionality that not only simplifies the patch operation but also offers timely and easy access to patch and support information during the patch planning phase. This helps you utilize the wealth of information from My Oracle Support to implement the best possible patch rollout for your organization. Enterprise Manager Cloud Control introduces a new concept called Patch Plans that enables you to create a collection of patches you want to apply to one or more targets. Each target can have a separate group of patches. In addition, you can save the deployment options of a patch plan as a patch template, and have new patch plans created out of the patch template.

3.5.4 Provisioning

Rather than spend resources on manually installing and configuring Oracle Fusion Middleware software, administrators would rather spend time and money on more strategic initiatives. To help achieve this, Enterprise Manager has automated common provisioning operations such as cloning Oracle SOA Suite 12c and scaling out an Oracle WebLogic Domain. Making such critical datacenter operations easy, efficient and scalable results in lower operational risk and lower cost of ownership. To access these provisioning operations, navigate to the Patching and Provisioning page in Enterprise Manager and click **Middleware Provisioning** and select either of the following:

- **Fusion Middleware Provisioning:** You can automate the cloning of WebLogic Domains and / or Middleware Homes either from a reference installation or from a profile present in the software library.
- **Fusion Middleware Domain Scale Up:** You can automate the scaling up or scaling out of a domain or cluster by adding a new managed server to an existing cluster or by cloning a managed server.

For more details on using these procedures, see the *Enterprise Manager Lifecycle Management Guide*.

3.5.4.1 Cloning from Test to Production Environments

Typically, creating a new environment to support SOA applications entails several manual, error prone installation and configuration steps. With Oracle Enterprise Manager this can be accomplished with very little effort and time using a predefined, customizable deployment procedure. This deployment procedure clones an existing SOA Suite environment to a new set of hardware per a hierarchical series of steps.

These predefined steps can be edited or disabled and new steps or custom scripts can be added to the deployment procedure to satisfy unique business needs.

The deployment procedure also supports secure host authentication using super user do (sudo) or pluggable authentication modules (PAM). While running the deployment procedure, administrators can specify configuration settings such as the domain name, credentials for the administration console, port values, and JDBC data resources. After the procedure completes, the newly created SOA environment is discovered and automatically added to the console for centralized management and monitoring.

3.5.4.2 Scaling Out Domains

To address growing business demands, modern data centers must augment and relocate resources quickly. Using Oracle Enterprise Manager, administrators can rapidly scale out a WebLogic Domain and Cluster with additional managed servers to accommodate an increase in application load.

3.5.4.3 Deploying / Undeploying Java EE Applications

You can deploy, undeploy, and redeploy Java EE applications (.war and .ear files) on a WebLogic Server. You can create a Java EE Application component in the Software Library and deploy multiple versions of an application, or roll-back to a previous version.

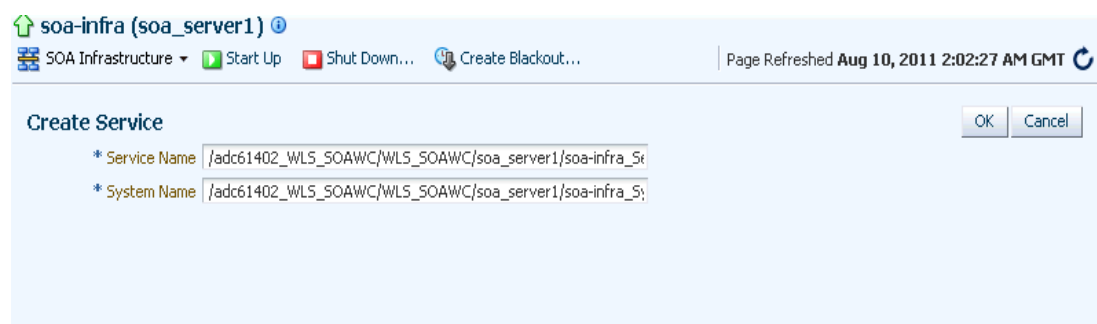
3.6 Managing Service Levels

Enterprise Manager allows you to create infrastructure services for Middleware targets such as Oracle BPEL Process Manager targets, Oracle Service Bus targets and Oracle SOA Composite and SOA Infrastructure instances.

An infrastructure service is a dependency service that is created to identify the infrastructure components on which the Middleware target depends. Here, the infrastructure components refer to hosts, databases, application servers, and so on that work together to host the Middleware target.

You can either create an infrastructure service with a new system or an existing system, or simply refresh an existing infrastructure service, if there is already one existing. By creating infrastructure services and systems, you can better manage your Middleware targets and also the components on which the Middleware targets depend.

Figure 3–8 Create Service for SOA Infrastructure



For example, once you create an infrastructure service for an Oracle SOA Infrastructure target, Enterprise Manager allows you to create an aggregate service for every process within that SOA Infrastructure target. An aggregate service is a logical

grouping of services, in this case, infrastructure services and availability services. Aggregate Services give you a bird's-eye view of the services that have been created for the SOA Infrastructure target and helps you monitor their availability, performance, and usage. Service availability can be composed of both metrics on the underlying target and service test results from period synthetic transaction execution.

You can define service level (measure of service quality) for a service. A service level is defined as the percentage of time during business hours a service meets specified availability, performance and business criteria.

A Service Level specifies the percentage of time a service meets the performance and availability criteria as defined in the Service Level Rule. By default, a service is expected to meet the specified criteria 85% of the time during defined business hours. You may raise or lower this percentage level according to service expectations. A service level measures service quality using two parameters: Expected and Actual Service Levels.

- **Expected Service Level:** A Service Level specifies the percentage of time a service meets the performance and availability criteria as defined in the Service Level Rule. By default, a service is expected to meet the specified criteria 85% of the time during defined business hours. You may raise or lower this percentage level according to service expectations.
- **Actual Service Level:** The Actual Service Level defines the baseline criteria used to define service quality.

3.6.1 Service Dashboard

The Service Dashboard provides a consolidated view of the critical aspects of the service including the status, availability, type of service, performance, and the SLAs that have been enabled for this service. It also shows the performance and usage metrics for the service, status of the key components, and any system incidents.

You can view all the information related to the service on a single page and assess the health of the service. You can customize the dashboard by adding or removing regions according to your requirements and make these changes available to all the users.

You can also personalize the dashboard and make changes that are visible only to you and not to the other users.

3.7 Job System

You can use Enterprise Manager job system to schedule tasks you want to automate. You can schedule a job for a target by selecting the **Control** menu option on the Home page. For example, for an Oracle WebLogic Server Domain, you can create a job to automatically start, stop, or restart the servers in the domain. You can view details about the jobs that are scheduled, running, suspended, or the ones that have a problem. You can also use jobs to automate the execution of the WLST (WebLogic Scripting Tool) scripts. See the *Enterprise Manager Cloud Control Administrator's Guide* for more details on the Job System and its functionality.

3.7.1 Log File Rotation

Oracle Application Server components generate log files containing messages that record all types of events, including startup and shutdown information, errors, warning messages, access information on HTTP requests, and additional information.

However, the information recorded in log files is voluminous, thus making it difficult to track what update was made at what time. Also because of the huge quantity of information updated periodically, the log files grow in size and occupy more space on the system over a period of time. The only way to manage these log files is to manually archive the contents to another file and store them in a different location.

Considering these impediments, Enterprise Manager has been enhanced with a log rotation feature that helps you manage the logs of Oracle Application Server components more effectively. In particular, you can use Enterprise Manager to:

- Schedule a job that automatically rotates a log at the scheduled date and time
- Manage space on your system by storing the rotated log files in a different directory

Enterprise Manager allows you to view the logs of a particular Oracle Application Server component type and select the ones that need to be rotated. Note that a log rotation job can also be part of a multi-task job.

When a log rotation job is executed, Enterprise Manager automatically stops the component whose logs have to be rotated. After it is stopped, the content from its existing log file is moved to another file that is distinguished with the timestamp when it was actually rotated. The original log file is kept empty for new log details to be populated. Once this is done, Enterprise Manager restarts the component.

Note: The log rotation feature cannot be used with the WebLogic Server.

3.7.2 Log File Viewer

You can centrally search logs generated by WebLogic and Oracle Fusion Middleware across all Oracle Fusion Middleware components. You can perform structured log searches based on log properties such as time, severity, or Execution Context ID (ECID). You can also download log files or export messages to a file. This feature provides ready access to log files no matter where they are stored on the file system.

3.8 Topology Viewer

Enterprise Manager provides a Topology Viewer which is a graphical representation of routing relationships across targets, components and elements. You can easily determine how requests are routed across components. For example, you can see how requests are routed from Oracle Web Cache, to Oracle HTTP Server, to a Managed Server, to a data source.

The Topology Viewer provides the basic navigation applications, such as zoom, pan, and fit-to-contents. You can change the source of data being viewed, the layout mode, and the flow direction between objects. Filters allow you to alter global properties of the topology diagram, such as the visibility of link labels or altering the link style. It enables you to easily monitor your environment. You can see which entities are up and which are down. You can also print the topology or save it to a .gif file. For more details, see the Enterprise Manager Online Help.

3.9 Support Workbench

Enterprise Manager Support Workbench enables you to investigate, report, and, in some cases, repair problems (critical errors). You can gather first-failure diagnostic data, obtain a support request number, and upload diagnostic data to Oracle Support.

The Support Workbench also recommends and provides easy access to Oracle advisors that help you repair SQL-related problems, data corruption problems, and more. You can use Support Workbench within Fusion Applications to:

- **View a Fusion Aggregated Diagnostic Summary:** Fusion applications are deployed across multiple systems, and incidents are therefore recorded in multiple Automatic Diagnostic Repository homes. You can get a quick summary of diagnostic data across all targets and Automatic Diagnostic Repository homes aggregated by the instance, product family, or cluster application.
- **Execute Tests to Diagnose a Problem:** You can link to Oracle Application Diagnostics Testing for a given Fusion Application, where you can execute additional tests to help diagnose the problem.
- **Create a Package and Upload it to Oracle Support:** You can create a package, add additional diagnostics and files, upload the package to Oracle Support, and then create a Service Request.

Testing Application Load and Performance

This chapter describes how you can perform load and performance testing of applications with real-world production workloads using the Application Replay feature of Enterprise Manager. With Application Replay you can capture application workloads on production systems, and then replay them against test systems while maintaining the precise timing, concurrency, and transaction order of the workload.

This chapter covers the following:

- [Introduction to Application Replay](#)
- [Testing Against Real-World Application Workloads](#)
- [Capturing Application Workload Using RUEI](#)
- [Synchronized vs. Non-Synchronized Replay](#)
- [Prerequisites and Considerations](#)
- [Understanding the Capture and Replay Process](#)
- [Creating Application Workload Captures](#)
- [Monitoring the Capture Process](#)
- [Replaying Application Workload Captures](#)
- [Importing Replay Session Divergences into OpenScript](#)
- [Troubleshooting](#)

4.1 Introduction to Application Replay

Application Replay enables realistic testing of planned changes to any part of the application stack from application server down to disk, by re-creating the production workload on a test system. Using Application Replay, you can capture a workload on the production system and replay it on a test system with the exact timing, concurrency, and transaction characteristics of the original workload. This enables you to fully assess the impact of the change, including undesired results, new contention points, or plan regressions. In addition, extensive analysis and reporting is provided to help identify any potential problems, such as new errors encountered and performance divergence. Types of changes that can be tested with Application Replay include application server upgrades, hardware updates, O/S changes, configuration changes, and so on. Capturing real-world production workload eliminates the need to develop simulation workloads or scripts, resulting in significant cost reduction and time savings. By using Application Replay, realistic testing of complex applications that previously took months using load simulation tools can now be completed in

days. As a result, you can rapidly test planned changes and adopt new technologies with a higher degree of confidence and at lower risk.

4.2 Testing Against Real-World Application Workloads

Today's enterprise application deployments are highly complex and, therefore, challenging to manage. They comprise multiple tiers, such as Web servers, application servers, and databases, running on multiple hosts. Their software architecture combines multiple independent components, such as client-side user interfaces, business logic and data access mechanisms, in addition to stateful client-server protocols typically built over HTTP.

Due to the complexity of these structures, predicting the behavior of the entire stack in a production environment is extremely difficult. Given the complexity of these deployments, and the absence of system-wide verification techniques, effective testing is critical to ensuring successful deployment after an infrastructure change.

The Application Replay feature provides a testing structure that works by first capturing the entire workload relevant to an application (as generated by the application's Web interface) at the production site.

The captured application workload is then moved to the test environment, where the replay driver infrastructure on one or more hosts, reproduce the captured workload, preserving its original properties, such as concurrency and request timings.

Finally, extensive performance and correctness data from all layers of the stack is collected and reported. This enables you to compare the replay with the original captured workload. In this way, any issues resulting from infrastructure changes that occurred during the replay can be identified, and appropriate troubleshooting action undertaken to prevent them from occurring in production. Moreover, it increases your confidence in a successful deployment.

Benefits

The use of *real* workloads offers a number of significant advantages over testing techniques based on synthetic workloads. In particular:

- It provides a system-wide perspective starting from the user's activity. This is in contrast to the traditional piecemeal testing of individual components that provides little information on their combined behavior and performance under a realistic workload.
- Rather than relying on pre-determined scenarios, the use of real workloads provides comprehensive testing, subjecting the system to real users operations. For Web applications, this not only means exploring all possible ways a user interacts with the system, but also all possible load conditions. This is necessary because systems behave quite differently under different workload characteristics (for example, the number of concurrent users).
- Far greater insight is obtained into possible errors. Test results include data for every layer of the stack, and these can be correlated across different layers. Besides performance, it also provides a means to verify correct execution, by checking for errors or unexpected server responses.

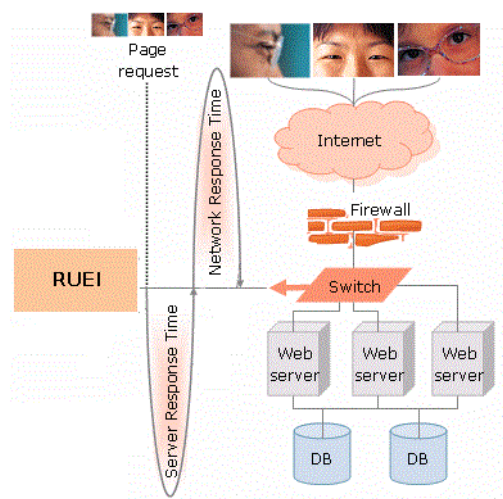
4.3 Capturing Application Workload Using RUEI

In order to capture Web application workloads, Application Replay uses Oracle Real User Experience Insight (RUEI). This is a Web-based utility to report on real-user traffic requested by, and generated from, your Web infrastructure. It measures the

response times of pages and user flows at the most critical points in your network infrastructure. It provides you with powerful analysis of your network and business infrastructure, while an insightful diagnostics facility allows application managers and IT technical staff to perform root-cause analysis.

Typically, RUEI is installed before the Web servers, behind a firewall in the DMZ. The data collection method is based on Network Protocol Analysis (NPA) technology. This data collection method is shown in [Figure 4-1](#).

Figure 4-1 How RUEI Collects Data



When an object is requested by a visitor, RUEI sees the request and starts measuring the time the Web server requires to present the visitor with the requested object. At this point, RUEI knows who requested the page (IP client), which object was requested, and from which server the object was requested (IP server).

When the Web server responds and sends the object to the visitor, RUEI sees that response, and stops timing the server response time. At this stage, RUEI can see whether there is a response from the server, whether this response is correct, how much time the Web server required to generate the requested object, and the size of the object.

RUEI is also able to see whether the object was completely received by the visitor, or if the visitor aborted the download (proof of delivery). Therefore, RUEI can determine the time it took for the object to traverse the Internet to the visitor, and can calculate the Internet throughput between the visitor and the server (connection speed of the visitor).

Further information about RUEI is available from the following location:

<http://www.oracle.com/us/products/enterprise-manager/index.html>

4.4 Synchronized vs. Non-Synchronized Replay

Application Replay supports the use of synchronization. This ensures that each replayed request sees exactly the same database state it saw during the capture process. In this way, the responses generated during capture and replay can be directly compared. In addition, data divergence between capture and replay is minimized. This is explained at greater length in ["Analyzing Replay Results"](#).

While the use of synchronized workload captures is clearly preferable, it is important to understand that:

- The use of the synchronization facility places additional requirements on your production and test environments. These are explained in ["Using Synchronization"](#).
- Non-synchronized workload captures require considerably less disk space. This is because information about the current database state does not need to be retained. Therefore, their use is recommended in test scenarios where differences between the responses generated in the capture and replay environments are not an issue. For example, in loading and stress testing.

4.5 Prerequisites and Considerations

This section describes the requirements that must be met, and the issues that should be considered, in order to use the Application Replay facility for workload capture and replay. It is *strongly* recommended that you carefully review this information before proceeding with a workload capture.

Important: It is *strongly* recommended that you review the Oracle Support Web site to obtain up-to-date information about supported RUEI, application server, and database versions, as well as patches, configurations, known issues, and workarounds.

This section covers the following:

- [Using RUEI to Capture Application Workloads](#)
- [Configuring Required User Privileges in Enterprise Manager](#)
- [Using Synchronization](#)
- [Setting up the Test System Database](#)
- [Restarting the Database and Application Stack](#)
- [Setting up the Capture Directory](#)

4.5.1 Using RUEI to Capture Application Workloads

In order to use RUEI to capture your application workloads, you must ensure that:

- RUEI version 12.1 (or higher) has been configured to monitor the required applications. See the Oracle Support Web site (<http://www.oracle.com/support/contact.html>) for information about required releases and hot fixes. Information about deployment options and requirements is available from the *Oracle Real User Experience Insight Installation Guide*.
- You have a valid user name and password combination. If necessary, contact your RUEI Administrator. Note that the user account must have Security Officer permission. For further information about roles and permissions, see the *Oracle Real User Experience Insight User's Guide*.
- You have the URL used to access the RUEI installation. If necessary, contact your RUEI Administrator.
- The configured RUEI logging and masking policies are consistent with the use of Application Replay. This is described in the following section.

RUEI Configuration for Application Replay

As mentioned above, you must ensure that the RUEI logging and masking policies are configured as follows:

1. Select **Configuration, Security, Masking, URL prefix masking**, and click the Default masking action setting. This must be set to "Logging".
2. Note that if you expect a high level of traffic during the workload capture, it is recommended that you select **Configuration, Security, Collector data retention policy**, and ensure that sufficient storage has been assigned for each application that is planned to be captured.

For further information on these configuration procedures, see Chapter 13 "Managing Security-Related Information" of the *Oracle Real User Experience Insight User's Guide*.

4.5.2 Configuring Required User Privileges in Enterprise Manager

The following Enterprise Manager privileges must be assigned to users of the Application Replay facility:

- ASREPLAY_VIEWER in order to view captures, replays, and replay tasks.
- ASREPLAY_OPERATOR in order to create, modify, or submit captures, replays, and replay tasks.

In addition to the above, users must also be assigned the `PERFORM_OPERATION_ANYWHERE` privilege.

In order for database users to run the Application Replay facility with database capture, the following privileges must be granted to the user:

```
GRANT ADMINISTER ANY SQL TUNING SET TO asreplay;
GRANT EXECUTE ON DBMS_LOCK TO asreplay;
GRANT EXECUTE ON DBMS_WORKLOAD_CAPTURE TO asreplay;
GRANT EXECUTE ON DBMS_WORKLOAD_REPLAY TO asreplay;
GRANT CREATE SESSION TO asreplay;
GRANT CREATE ANY DIRECTORY TO asreplay;
GRANT SELECT_CATALOG_ROLE TO asreplay;
GRANT BECOME USER TO asreplay;
GRANT DROP ANY DIRECTORY TO asreplay;
```

Note that in the above example, the database user is assumed to be called `asreplay`.

4.5.3 Using Synchronization

In order to make use of the synchronization functionality within Application Replay, you must ensure that:

- Oracle Fusion Middleware has been installed and configured. See the Oracle Support Web site (<http://www.oracle.com/support/contact.html>) for information about required releases and patch sets. For further information, see the *Oracle Fusion Middleware Installation Planning Guide* available at the following location:

http://download.oracle.com/docs/cd/E14571_01/install.1111/b32474/start.htm

- Oracle Database has been installed and configured. See the Oracle Support Web site (<http://www.oracle.com/support/contact.html>) for information about required releases and patch sets. For further information, see the relevant *Database Installation Guide* available at the following location:

http://www.oracle.com/pls/db112/portal.portal_db?selected=11&frame=

- The database user needs to be assigned the privileges described in the *Oracle Database Security Guide* available at the following location:
http://download.oracle.com/docs/cd/B19306_01/network.102/b14266/admusers.htm
- The user capturing and replaying application workloads has access to the host directory used to store capture files.
- The database user has EXECUTE privilege on DBMS_WORKLOAD_REPLAY.
- The files used during replay are on a shared file location that is accessible by all Replay Client systems using the same directory path. If synchronization has been enabled, this is also required for the appropriate database systems.
- It is recommended that the time zone and time settings for all host systems (such as RUEI, Oracle Fusion Middleware, and Oracle database) used during capture and replay should be the same to prevent clock skew.

Note that if you are using the database synchronization facility, be aware that only one database capture can be generated against a production environment at given time.

4.5.4 Setting up the Test System Database

Before a workload can be replayed, the logical state of the application data on the replay system should be similar to that of the capture system when replay begins. Therefore, you should have a strategy in place to restore the application server and database state on the test system. To restore the application server state, you should consult your application administrator. To restore the database state, consider using one of the following methods:

- Recovery Manager (RMAN) DUPLICATE command. For further information, see the *Oracle Database Backup and Recovery User's Guide*.
- Snapshot standby. For further information, see the *Oracle Data Guard Concepts and Administration*.
- Data Pump Import and Export. For further information, see the *Oracle Database Utilities*.

4.5.5 Restarting the Database and Application Stack

This step is not required, although restarting the database and application server components before starting capture ensures that ongoing and dependent transactions are allowed to be completed or rolled back before the capture begins. This minimizes divergence during replay, which is always desirable. On the other hand, if the database is not restarted before the capture begins, transactions that are in progress, or have yet to be committed, may not be fully captured in the workload. Therefore, ongoing transactions may not be replayed properly, because only the part of the transaction whose calls were captured will be replayed, resulting in data divergence.

Moreover, any subsequent transactions with dependencies on the incomplete transactions may also generate errors during replay. Before restarting the database and application server components, determine an appropriate time to shut down the production database before the workload capture when it is the least disruptive. For example, you may want to capture a workload that begins at 8:00 a.m. However, to avoid service interruption during normal business hours, you may not want to restart the database during this time. In this case, you should consider starting the workload

capture at an earlier time, so that the database can be restarted at a time that is less disruptive. Once the application is restarted, it is important to start the capture before any user sessions reconnect and start issuing any workload. Otherwise, transactions performed by these user sessions will not be replayed properly in subsequent replays, because only the part of the transaction whose calls were executed *after* the workload capture is started will be replayed.

4.5.6 Setting up the Capture Directory

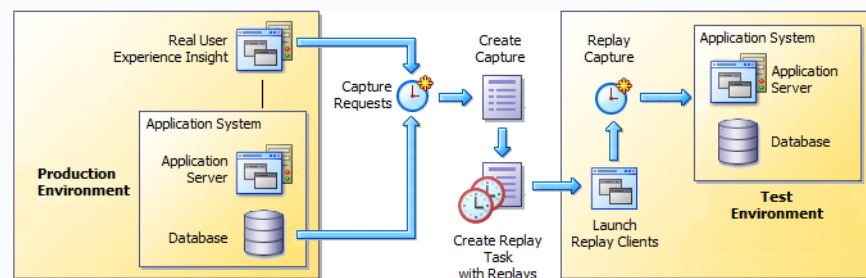
Determine and set up the directory where the captured workload will be stored. Before starting the capture, ensure that the directory is empty and has sufficient disk space to store the workload. If the directory runs out of disk space during a capture, the capture will be terminated.

To estimate the required disk space, it is recommended that you run a test capture on your workload for a short duration (typically, a few minutes), and then use this to extrapolate the space required for a full capture. To avoid potential performance issues, you should also ensure that the target replay directory is mounted on a separate file system.

4.6 Understanding the Capture and Replay Process

Figure 4–2 shows the architecture of the Application Replay facility.

Figure 4–2 Application Capture and Replay Architecture



The capture part of Application Replay operates within the context of a production environment. This deployment comprises Web and application servers, and a database. The Web-tier capture mechanism is provided by RUEI. It writes information about the monitored traffic to capture files. These contain HTTP requests, responses, and timings, along with all other data necessary to accurately reproduce the production workload against a test system. Once the capture is complete, the generated files constitute a complete representation of the entire production workload. In addition, the synchronization facility also captures the corresponding database workload for both capture analysis.

The replay part of Application Replay operates within the context of a test system. This comprises an application stack that runs the system configuration under test. One or more Replay Clients reproduce the captured workload, preserving its original properties, such as concurrency and request timings. Further, the Application Replay facility uses synchronization to ensure that each replayed request sees the exact application state it saw during capture so that the responses are directly comparable. Finally, it collects a wealth of performance and verification data from all layers of the stack, and allows you to compare the replay with the original capture upon which it is based.

Depending on the volume and concurrency of the workload capture, it may be necessary to deploy multiple Replay Clients, each assigned a portion of the workload. Recommendations about required Replay Clients based on the captured workload are available when scheduling a replay.

4.7 Creating Application Workload Captures

To create an application workload capture:

1. From the **Enterprise** menu, select **Quality Management**, then **Application Replay**.
2. Click **Captures**. The currently defined captures are listed. An example is shown in [Figure 4–3](#).

Figure 4–3 Application Replay Page

Application Replay							
> Overview							
Selected Item: Captures Item: Replay Tasks							
View Menu Create... Create Like... Edit... Delete Query By Example							
Previous 1-16 of 16 Next							
Select	Name	Status	System	Replay Tasks	Owner	Creation Date	Description
<input checked="" type="radio"/>	manula_ip2	Completed	EBS6170_system	1	JSMITH	02-Sep-2011 20:46:50	
<input type="radio"/>	ebs_atc_ip	Completed	EBS6170_system	1	PJONES	02-Sep-2011 16:40:55	
<input type="radio"/>	ruei12_form_https_cookei	Completed	EBS6170_system	1	PJONES	01-Sep-2011 20:52:26	
<input type="radio"/>	ebs_manual_ip	Completed	EBS6170_system	1	PJONES	01-Sep-2011 16:28:07	
<input type="radio"/>	manual_compare_http_form	Failed	EBS6170_system	0	JSMITH	01-Sep-2011 15:53:06	
<input type="radio"/>	EBS_https_Form_cookie	Completed	EBS6170_system	1	PJONES	01-Sep-2011 15:01:39	
<input type="radio"/>	EBS_Form_HTTPS_R12	Completed	EBS6170_system	1	PJONES	31-Aug-2011 23:34:21	
<input type="radio"/>	fod_synch2	Completed	EBS6170_system	1	JSMITH	31-Aug-2011 00:23:35	
<input type="radio"/>	fod_synch	Failed	EBS6170_system	0	JSMITH	30-Aug-2011 22:40:18	

3. Click **Create** or **Create Like**. The page shown in [Figure 4–4](#) appears.

Figure 4–4 Create Capture (Overview) Page

Application Replay

Overview System RUEI Application Database Storage Schedule Review

Create Capture : Overview Back Step 1 of 7 Next Cancel

* Name:

Description:

Capture Prerequisites

These prerequisites should be completed before performing a capture.

- ☒ Ensure there is sufficient free disk space on the selected host system to store the capture. You should consider performing a short duration capture, and using it as the basis for estimating the requirements for a full capture.
- ☒ If you intend to enable database capture for this capture (and enable database synchronization for its subsequent replay), ensure that you can restore the test system database state to match that of the database at the start of the capture. A successful replay requires application transactions accessing application data identical to that on the capture system. Common methods to restore application data state include point-in-time recovery, flashback, and import/export.

4. Specify a unique name for the new capture. Optionally, specify a brief description for the traffic to be captured. It is recommended that you include an indication of the purpose and scope of the capture. Carefully review the prerequisite

information, and click the acknowledgement check boxes to indicate that they have been met. When ready, click **Next**. The page shown in [Figure 4-5](#) appears.

Figure 4-5 Create Capture (System) Page

Application Replay

Overview **System** RUEI Application Database Storage Schedule Review

Create Capture : System Back Step 2 of 7 Next Cancel

Select the System target representing the application for which traffic is to be captured. The system members will be used to determine whether a database capture can be made and used later for synchronized replay. Note that if the selected system does not contain a supported version of Fusion Middleware and Oracle Database, the option to enable database capture will not be available in the subsequent Database step.

* System Name

Name	Type	Status
myshop_test_db	Database Instance	↑
pcc41689.us.myshop.com	Host	↑

- Click the **Select System** icon, and select the target that represents the applications for which traffic is to be captured. It is recommended that you review the status of the selected component, and ensure that it will be available throughout the planned capture. Note that if the selected target does not include supported versions of Oracle Fusion Middleware and Oracle Database components, the Create Capture (Database) Page (shown in [Figure 4-7](#)) is not available, and is skipped. When ready, click **Next**. The page shown in [Figure 4-6](#) appears.

Figure 4-6 Create Capture (RUEI Application) Page

Application Replay

Overview System **RUEI Application** Database Storage Schedule Review

Create Capture : RUEI Application Back Step 3 of 7 Next Cancel

Specify the URL and credential for the Real User Experience Insight (RUEI) instance to be used for the capture. After providing these details, click **Show Applications** to select the application(s) to be included in the capture.

* RUEI URL

* RUEI Username

* RUEI Password

Show Applications

Select	Name	Data Size (MB)	Sessions	Hits	Page Views
<input checked="" type="checkbox"/>	bookings	37	1212	4523678	1734
<input type="checkbox"/>	EBSR12	89	681	84564323	6534
<input checked="" type="checkbox"/>	Siebel	21	869	739135	9212
<input type="checkbox"/>	Buss-partner	91	5478	3568902	9371
<input checked="" type="checkbox"/>	CRM	45	732	5345657	78321
<input checked="" type="checkbox"/>	catalog	121	932	9249123	2952

- Specify the URL used to access the RUEI installation. This must be based on a secure (HTTPS) connection. Specify a valid user name and password combination. The specified user must have Security Officer permission. If necessary, contact your RUEI Administrator for this information.

Click **Show Applications** to view the applications currently being monitored by the specified RUEI deployment. Note that you can use the traffic information available for each application to determine its suitability for capture. In particular, when selecting the applications to be included in the capture, you should ensure that the applications are running, and traffic volumes and error levels are within acceptable bounds. When ready, click **Next**. The page shown in [Figure 4-7](#) appears.

Figure 4-7 Create Capture (Database) Page

Application Replay

Overview System RUEI Application **Database** Storage Schedule Review

Create Capture : Database Back Step 4 of 7 Next Cancel

Select whether database capture should be enabled during application capture and specify the necessary database credential. If database capture is enabled, select whether Automatic Workload Repository (AWR) data should be exported.

Database Capture Disabled Enabled **Information** Database capture is necessary to provide synchronized replay. Assuming replay starts from a database state as similar as possible to the state when capture started, replay synchronization attempts to order requests so that every replayed request operates on the same data as it did during capture. Disabling synchronization generates a replay in which the recorded timings are reproduced without taking into account the data dependencies.

Database Name: Oemrep_Database Database Host Name: pcc41689.us.myshop.com

Database Credential

Credential Preferred Credential Named Credential **New Credentials**

* Username: admin

* Password: [masked]

* Confirm Password: [masked]

Role: NORMAL

☒ Save As: NC_OEMREP_D_2011-04-27-045237

☐ Set As Preferred Credentials

Test

Database Host Credential

Credential Preferred Credential Named Credential **New Credentials**

* Username: admin

* Password: [masked]

* Confirm Password: [masked]

☒ Save As: NC_OEMREP_D_2011-04-27-045238

☐ Set As Preferred Credentials

Test

* Database Capture Intermediate Storage Location: /tmp/ **Information** If you have an Oracle cluster database, make sure this directory is on a shared file system and accessible by all database instances using the same directory path.

Automatic Workload Repository

AWR Data Export Disabled Enabled **Information** Exporting of AWR data during capture enables performance comparison of the database between capture and replay. Because AWR data exportation can have a significant impact on the production system, it should be scheduled at a time that will not negatively impact the production system.

Start: **Immediately** Later [calendar icon] (UTC-08:00) US Pacific Time

- Specify whether database capture should be enabled during application capture. This is required for synchronized replay. If enabled, specify the necessary database and host credentials, the file system location on the database host system used for intermediate capture storage, and whether Automatic Workload Repository (AWR) data should be exported. Note that if AWR export is enabled, you need to specify when exporting should begin. By default, it is performed immediately after capture is completed. When ready, click **Next**. The page shown in [Figure 4-8](#) appears.

Figure 4–8 Create Capture (Storage) Page

Application Replay

Overview System RUEI Application Database **Storage** Schedule Review

Create Capture : Storage Back Step 5 of 7 Next Cancel

Specify the file system location at which the capture data should be stored. Note that this includes not only the capture files themselves, but also the storage of the RUEI files from which they are derived, as well as any data synchronization information.

* Storage Host pcc41689.us.myshop.com

* Storage Location /tmp

Storage Host Credential

Credential ☐ Preferred Credential ☐ Named Credential ☒ New Credentials

* UserName admin

* Password

* Confirm Password

☒ Save As admin

☐ Set As Preferred Credentials

Test

8. Specify the host and file system location where the capture data should be stored. Note that this includes not only the capture files themselves, but also the storage of the RUEI files from which they are derived, as well as any data synchronization information.

Important: Capture files can require large amounts of disk space. Therefore, it is recommended that you perform a short capture, and then use that as the basis for calculating the required disk space for the planned capture. In addition, be aware that the generated capture files are in a proprietary format, and should *not* be modified.

Click the **Select Target** icon. A new window opens that allows you to view the available targets. Click a target to select it. Note that only one host target can be selected. You can use the **Target Type** menu to restrict the listing of targets to specific types. Note that it is also recommended that you review the status of the selected targets, and ensure that they will be available throughout the planned capture. Specify the credentials of the selected storage host. When ready, click **Next**. The page shown in [Figure 4–9](#) appears.

Figure 4–9 Create Capture (Schedule) Page

Application Replay

Overview System RUEI Application Database Storage **Schedule** Review

Create Capture : Schedule Back Step 6 of 7 Next Cancel

Schedule start time and duration for capture.

Start ☐ Immediately ☒ Later 27/04/2011 05:31:59 (UTC-08:00) US Pacific Time

Duration ☐ Indefinitely ☒ For 90 minutes ☐ Until

9. Specify the start and stop times for the planned capture. By default, capture starts immediately. Note that, by default, the capture will run for the next 15 minutes. It is *strongly* recommended that you carefully consider the capture's duration and, if scheduled to run indefinitely, regularly review the capture process to prevent the creation of excessively large capture files. When ready, click **Next**. The page shown in [Figure 4–10](#) appears.

Important: If the capture is configured to run indefinitely, it must be stopped manually from the Capture Page. It is *strongly* recommended you regularly check the size of the created capture to prevent running out of storage space.

Figure 4–10 Create Capture (Review) Page

Application Replay

Overview System RUEI Application Database Storage Schedule **Review**

Create Capture : Review Back Step 7 of 7 Next Submit Cancel

Review details before initiating capture.

Overview

Capture Name Siebel CRM
Capture Description North America and Europe Siebel CRM application.

System

System Name myshop_bookings_db

RUEI Application

RUEI URL https://myshop.com/ruei
RUEI Username admin
RUEI Applications bookings|SiebelCRM|catalog

Database

Enable Database Capture Yes
Database Name myshop_test_db
Database Username admin
Database Host Name pcc41689.us.myshop.com
Database Host Username admin
Database Capture Intermediate Storage Location /tmp
Enable AWR Data Export Yes
AWR Data Export Schedule Start Immediately

Storage

Storage Host pcc41689.us.myshop.com
Storage Host Username admin
Storage Location /tmp

Schedule

Start Later 27/04/2011 05:31:59 (UTC-08:00) US Pacific Time
Duration For 90 minutes

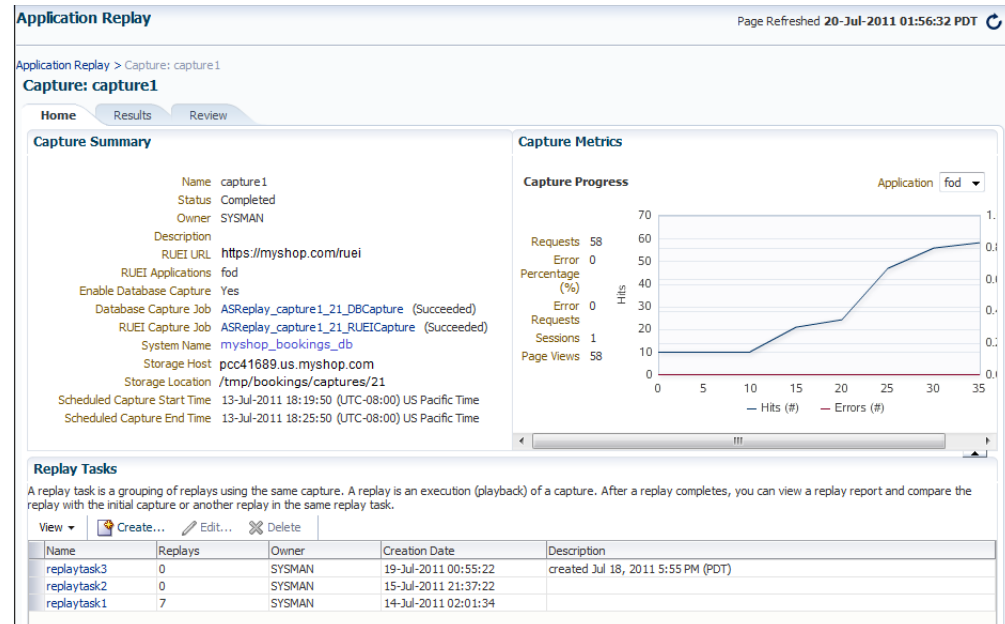
10. Review the planned capture's properties before launching it. If necessary, use the **Back** and **Next** buttons to amend the capture's properties. When ready to launch the new capture, click **Submit**.

4.8 Monitoring the Capture Process

Once a capture has been started, you can monitor the capture process to ensure that the intended traffic is being correctly captured, and that the application system is

working under normal conditions. An example of a capture progress report is shown in [Figure 4–11](#).

Figure 4–11 Capture Page



Be aware that there is a lead time of appropriately 10 minutes after the start of a capture before progress information about it becomes available. This is available from the Application Replay page ([Figure 4–3](#)). In either case, the characteristics of the capture are detailed in terms of its volume, performance, and errors. In this way, you can assess the quality of the capture, and its usefulness for testing purposes.

Note that the number of requests monitored during the capture process is particularly useful for assessing the capture's progress. In the event of an unusually high level of errors, you can use the Job page (available by clicking the RUEI Capture Job setting) to drill-down into specific errors.

4.9 Replaying Application Workload Captures

You can replay a workload capture against a test system. Besides issuing identical HTTP requests, the replay mechanism also mimics the characteristics of the capture in terms of concurrency and timing. This section provides information about the following parts of the replay process:

- [Preparing to Replay Workload Captures](#)
- [Understanding Replays and Replay Tasks](#)
- [Resolving References to External Systems](#)
- [Remapping URLs](#)
- [Substituting Sensitive Data](#)
- [Replaying Workload Captures](#)
- [Analyzing Replay Results](#)

4.9.1 Preparing to Replay Workload Captures

Proper planning of the workload replay ensures that the replay will be accurate. Replaying a workload capture requires the following steps:

- Ensure that the application data state on the test system is logically equivalent to that of the capture system at the start time of workload capture. This is described in ["Setting up the Test System Database"](#).
- All references to external systems have been resolved. This is explained in ["Resolving References to External Systems"](#).

4.9.2 Understanding Replays and Replay Tasks

It is important to understand that a replay is an execution (playback) of a workload capture. A replay task is a group of replays based on the same capture. After a replay is completed, you can view a replay report and compare the replay with the initial capture, or create another replay within the same replay task. Typically, the replays within a replay task perform the same purposes. For example, a database or host system configuration with multiple parameter changes.

It is recommended that replays are grouped into the same replay task in order to facilitate comparison. For example, those that relate to the testing of the same database upgrade patch.

4.9.3 Resolving References to External Systems

A captured workload may contain references to external systems, such as database links or external tables. It is critical that you reconfigure these external interactions to avoid impacting other production systems during replay. Typical external references that need to be resolved before replaying a workload are shown in [Table 4–1](#).

Table 4–1 *References to External Systems*

Type	Description
Database links	Typically, it is not desirable for the replay system to interact with other databases. Therefore, you should reconfigure all database links to point to an appropriate database that contains the data needed for replay.
External tables	All external files specified using directory objects referenced by external tables need to be available to the database and application server during replay. The content of these files should be the same as during capture, and the filenames and directory objects used to define external tables should also be valid.
Directory objects	You should reconfigure any references to directories on the production system by appropriately redefining the directory objects present in the replay system after restoring the database.
URLs	URLs/URIs that are stored in the database and application server need to be configured so that Web services accessed during the capture will point to the appropriate URLs during replay. If the workload refers to URLs that are stored in the production system, you should isolate the test system network during replay.
E-mails	To avoid resending E-mail notifications during replay, any E-mail server accessible to the replay system should be configured to ignore requests for outgoing E-mails.

Important: To avoid impacting other production systems during replay, it is *strongly* recommended that you run the replay within an isolated private network that does not have access to the production environment hosts.

4.9.4 Remapping URLs

URLs in the workload capture files need to be remapped to different values before replay within the test environment. For example, the Web application URL in every request needs to be remapped to that of the test system.

Note that wildcard characters are not supported within remapped URLs. All required domain and port numbers must be fully specified.

4.9.5 Substituting Sensitive Data

The RUEI installation monitoring your network traffic can be configured to omit the logging of sensitive information. This is called *masking*, and prevents passwords and other sensitive information from being recorded on disk. Further information on the use of this facility is available from the *Oracle Real User Experience Insight User's Guide*.

It is important to understand that Application Replay only supports the substitution of one value for each masked field. For example, if an application logon password field is masked, you will need to set up one common alternative logon password for all user accounts in the test system.

4.9.6 Replaying Workload Captures

To replay a workload capture using Enterprise Manager:

1. From the **Enterprise** menu, select **Quality Management**, then **Application Replay**. The page shown in [Figure 4–3](#) appears.
2. From the **Replay Tasks** section, click **Create** or **Create Like**. The page shown in [Figure 4–12](#) appears.

Figure 4–12 Create Replay Task Page

Application Replay

Create Replay Task OK Cancel

Create a new replay task. A replay task is a grouping of replays using the same capture. From the replay task home page, new replays can be created to replay the associated capture against a test system.

* Capture Name: myshop_bookings_db

* Name: CRM_Western_Sphere

Description: 90 mins Siebel CRM application (North American and west Europe).

3. Click the **Select Capture** icon. A new window opens that allows you to select the capture upon which the replay task should be based. Specify a unique name for the new replay task. Optionally, specify a brief description. It is recommended that you include an indication of the replay task's purpose and scope. When ready, click **OK**. You are returned to the page shown in [Figure 4–3](#).
4. Click the newly created replay task. The page shown in [Figure 4–13](#) appears.

Figure 4–13 Create Replay (Overview) Page

Application Replay

Create Replay Back Next Submit Cancel

Tasks

- Overview
- System
- Capture Storage Credential
- Replay Clients
- Synchronization
- URL Mappings
- Masked Data Substitution
- Additional Replay Parameters
- Schedule
- Review

Overview

* Name: Siebel CRM

Description: North America and Europe Siebel CRM application.

Replay Prerequisites

These prerequisites should be completed before replaying a capture.

- ☒ Restore database to match the state of the production database at the start of capture.
- ☒ Resolve any references to external systems.

External Production System Warnings

A replay should be performed in a completely isolated test environment. Make sure to resolve all references to external systems before starting a replay so that your production environment is not harmed. The following are examples of external references that an application might have. Your application may have additional or different external references.

- Web Services
- Authentication Servers (e.g. SSO, LDAP)
- Application URLs
- Database Links
- Directory Objects
- Streams

- Specify a name for the replay. It must be unique among the selected replay task. Optionally, specify a brief description for the traffic to be replayed. It is recommended that you include an indication of the purpose and scope of the replay. Carefully review the prerequisite information, and click the acknowledgement check boxes to indicate that they have been met. When ready, click **System**. The page shown in [Figure 4–14](#) appears.

Figure 4–14 Create Replay (System) Page

Application Replay

Create Replay Back Next Submit Cancel

Tasks

- ☒ Overview
- System
- Capture Storage Credential
- Replay Clients
- Synchronization
- URL Mappings
- Masked Data Substitution
- Additional Replay Parameters
- Schedule
- Review

System

Select the System target that represents the test application against which the associated capture is to be replayed. The selected system members determine whether database synchronization can be performed during replay. If the selected system does not contain supported versions of Fusion Middleware and Oracle Database targets, the subsequent Synchronization task will be disabled.

* Replay System: pcc41689.us.myshop.com

Name	Type	Status
myshop_test_db	Database Instance	↑
pcc41689.us.myshop.com	Host	↑

- Click the **Select System** icon, and select the targets that represent the test environment against which the capture should be replayed. It is recommended that you review the status of the selected components, and ensure that they will be available throughout the planned replay. When ready, click **Capture Storage Credential**. The page shown in [Figure 4–15](#) appears.

Figure 4–15 Create Replay (Capture Storage Credential) Page

Application Replay

Create Replay Back Next Submit Cancel

Tasks

- Overview
- System
- Capture Storage Credential**
- Replay Clients
- Synchronization
- URL Mappings
- Masked Data Substitution
- Additional Replay Parameters
- Schedule
- Review

Capture Storage Credential

Specify the capture storage host credential.

Storage Host pcc41689.us.myshop.com

Storage Location /tmp/bookings/captures/21

Storage Host Credential

Credential ☒ Preferred Credential ☐ Named Credential ☐ New Credentials

Preferred Credential Name Privileged Host Credentials

Credential Details

Attribute	Value
UserName	jsmith
Password	*****

[More Details](#)

7. Specify the credentials of the selected storage host. When ready, click **Replay Clients**. The page shown in Figure 4–16 appears.

Figure 4–16 Create Replay (Replay Clients) Page

Application Replay

Create Replay Back Next Submit Cancel

Tasks

- Overview
- System
- Capture Storage Credential
- Replay Clients**
- Synchronization
- URL Mappings
- Masked Data Substitution
- Additional Replay Parameters
- Schedule
- Review

Replay Clients

Specify the distribution of replay clients among hosts.

Estimated Replay Clients Needed 1

Replay Client Hosts

Using the provided estimate as a guide, specify the replay client hosts and the distribution of replay clients among them.

View + Add ✎ Edit... ✕ Delete

Host	Username	Replay Clients	Status
pcc41689.us.myshop.com	admin	1	

Replay File System

Specify a shared file directory location that is accessible from all replay client hosts. This location will be used for staging the capture to be replayed and storing the replay results.

* Replay File Location /tmp

Information

If you enable synchronized replay, make sure this directory is on a shared file system and accessible by all database instances using the same directory path.

8. Specify the Replay Clients that will be used to generate the workload on the test system. Note that the provided estimate should be used as a basis for scaling the planned replay. Specify the file system location to be used for storing the capture files and replay results. This location must be on a shared file system and accessible from the Replay Client hosts and database hosts (if synchronization is enabled) via exactly the same file directory path. When ready, click **Synchronization**. The page shown in Figure 4–17 appears.

Figure 4–17 Create Replay (Synchronization) Page

9. Specify whether database synchronization should be enabled during the replay. Note that, if enabled, you will need to provide relevant database and host credentials. When ready, click **URL Mappings**. The page shown in [Figure 4–18](#) appears.

Figure 4–18 Create Replay (URL Mappings) Page

10. Specify the URL mappings that should be used during the replay. That is, how the URLs encountered during capture should be substituted when replayed within the test environment. When ready, click **Masked Data Substitution**. The page shown in [Figure 4–19](#) appears.

Figure 4–19 Create Replay (Masked Data Substitution) Page

Application Replay

Create Replay Back Next Submit Cancel

Tasks

- Overview
- System
- Capture Storage Credential
- Replay Clients
- Synchronization
- URL Mappings
- Masked Data Substitution**
- Additional Replay Parameters
- Schedule
- Review

Masked Data Substitution

These sensitive data fields are masked by RUEI during capture. Specify the new replay values to substitute during replay.

Sensitive Data Field	Replay Value	URL	RUEI Application	
			Name	Type
passwordField	*****	http://myshop.testenv.com:8083/OA_HTML/OA.jsp	EBSR 12	EBS
l_password	*****	http://myshop.com/StoreFrontModule/faces/check	Catalog	ADF
SWEPassword	*****	http://sd4.corp/siebel.com/calcenter_enu/start.swe	Contact	Siebel
fidPassword	*****	http://myshop.com/basket/login	MyShop	MyShop

- RUEI can be configured to omit the logging of sensitive information (such as passwords, credit card details, and so on) from being recorded on disk. Because the values of these fields are not recorded, they need to be explicitly specified for replay. When ready, click **Additional Replay Parameters**. The page shown in Figure 4–20 appears.

Figure 4–20 Create Replay (Additional Replay Parameters) Page

Application Replay

Create Replay Back Next Submit Cancel

Tasks

- Overview
- System
- Capture Storage Credential
- Replay Clients
- Synchronization
- URL Mappings
- Masked Data Substitution
- Additional Replay Parameters**
- Schedule
- Review

Additional Replay Parameters

Replay Rate

Configure parameters controlling the rate at which replay progresses.

Run Replay ☐ Same Rate As Captured ☒ Custom Rate

Session Start Time Scale

Think Time Scale

Maintain Request Rate ☒

About Custom Replay Rate Parameters

Session Start Time Scale: Scales the elapsed time from when the capture started to when the session connects with the specified value. Controls the rate of logon activity during replay.

Think Time Scale: Scales the elapsed time between two successive user calls for the same session. Controls the replayed request rate.

Maintain Request Rate: If enabled, the system will correct the think time (based on the think_time_scale parameter) between calls when user calls take longer to complete during replay than during capture. Maintains capture request rate.

Advanced Replay Configuration

Add custom values for parameters used during replay.

Advanced Settings ☒ Standard ☐ Custom

- Specify whether the replay should progress at the same rate as in the original capture or at an alternative rate. In the case of the latter, you need to specify the appropriate session start time scale, think time scale, and whether the request rate should be maintained at the original rate. Expand the **Advanced Replay Parameters** section to specify whether the standard advanced settings should have

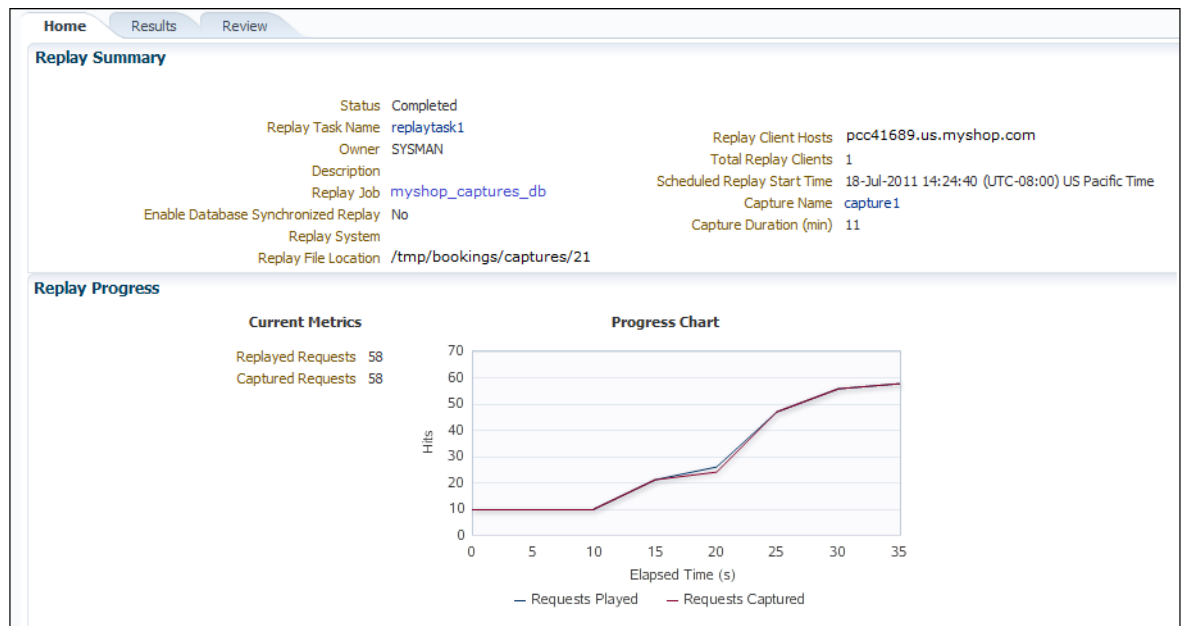
their standard values or custom values. When ready, click **Schedule**. The page shown in [Figure 4–21](#) appears.

Figure 4–21 Create Replay (Schedule) Page

13. Specify when the replay should start. When ready, click **Review**.
14. A summary of the planned replay is displayed. If necessary, use the **Back** and **Next** buttons to move between sections. When ready, click **Submit** to launch the replay.

4.9.7 Analyzing Replay Results

Detailed information about a selected replay is available by clicking it within the Application Replay Page. An example of a replay overview is shown in [Figure 4–22](#).

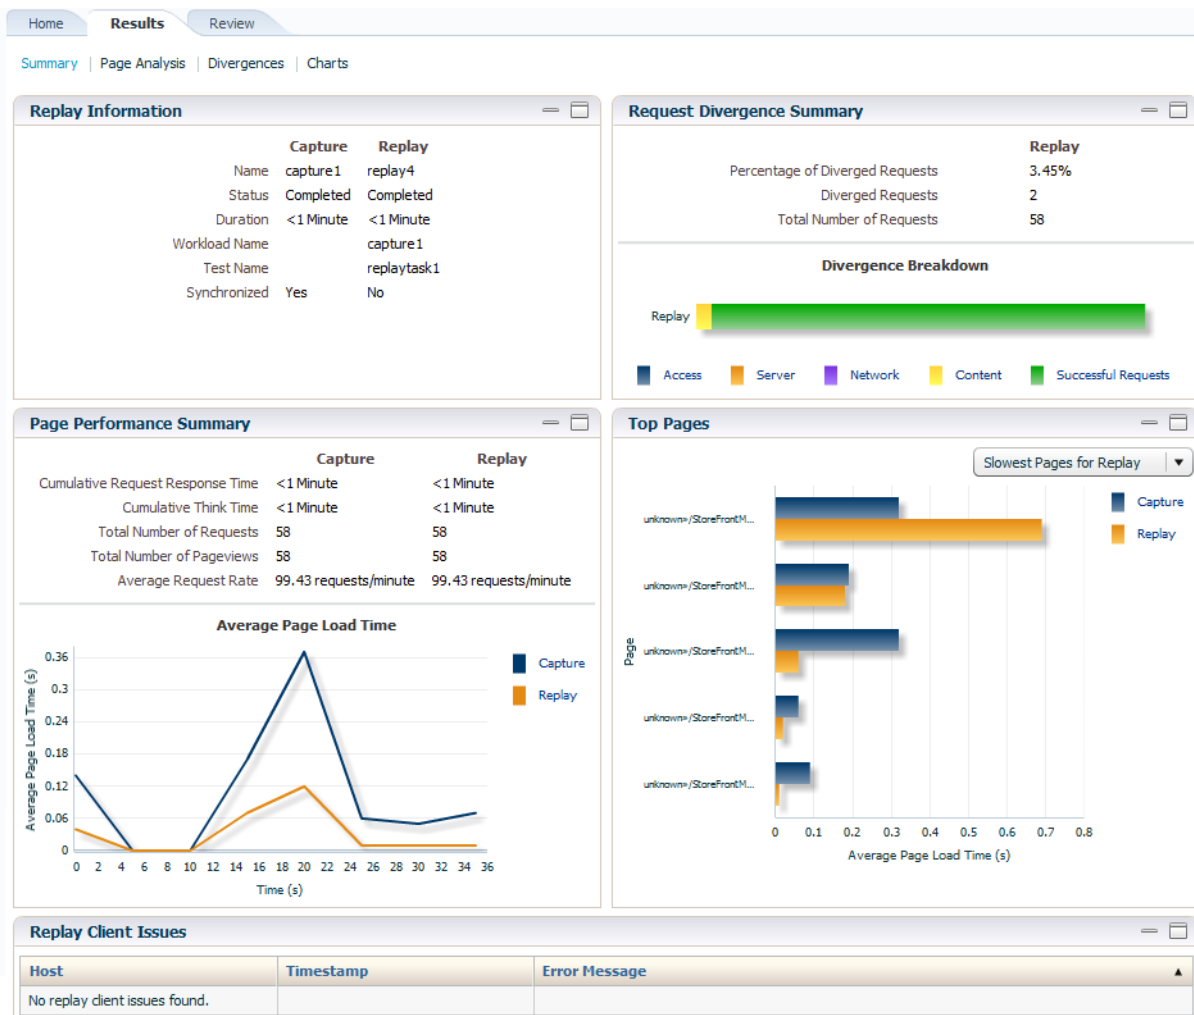
Figure 4–22 Example Replay Summary

It consists of three parts:

- **Home:** provides a overview of the replay, its associated replay task, and the capture upon which it based. The progress of the replay, and a comparison with the original capture, is also provided.
- **Results:** provides more detailed information about the request divergence. This includes a comparison of page performance during the original capture and the replay, and the application pages that experienced the highest level of divergence. An example is shown in [Figure 4–23](#).

Within this section, The **Page Analysis** section allows you to an analysis of each application page across selected metrics. The **Divergences** section allows you to view information restricted to specific divergence types (such as access, content, and so on). The **Charts** section allows you to view detailed replay information across specific metrics (such as average page load time).

- **Review:** provides information about the replay environment (such as the credentials, host, and replay clients), as well as the URL mappings and masked data substitutions used during the replay.

Figure 4–23 Example Replay Results Summary

4.10 Importing Replay Session Divergences into OpenScript

Using the Oracle Application Testing Suite's OpenScript module, you can further analyze divergence errors by viewing them by session. Currently Application Replay calculates divergence statistics and presents them on a per page basis. Enterprise Manager allows you to export this session data to a .zip file and import this data into OpenScript.

Creating a Session .zip File

The session .zip file generated from Enterprise Manager contains both capture and replay data.

To create a .zip file:

1. Navigate to the *Replay Divergence* Page.
2. Choose **View by Session** as the viewing **Mode**.
3. In the **Sessions** list, click on the Session ID that contains divergences (shown in the Divergences column). The *Session Divergence Detail* dialog appears.

4. In the *Pages* region, click **Fetch Page Details**.
5. In the *Export Session Data* region, click **Export**.

Once the export operation is complete, a link to the session data .zip file appears.

6. Click on the link to download the session data .zip file and save the file to your local system.

Importing a Session .zip File into OpenScript

Once you have created a .zip file containing the capture and replay data, you must import this data into OpenScript.

To import the contents of the .zip file:

1. From OpenScript, you next need to create a Load Testing script.
From the **File** menu, select **New**. The **New Project** dialog displays. From here, you can choose a wizard to create a new Load Testing script.
2. From the tree list, choose an application type and click **Next**. The New Script dialog appears.
3. Enter a script name and click **Finish**. OpenScript creates the new script.
4. From the OpenScript Tools menu, select Import Oracle Real User Experience Insight (RUEI) Session Log. The *Import RUEI Log* dialog appears.
5. Click **Browse...** The *Open RUEI Log* dialog appears.
6. Navigate to the session .zip file you created earlier, select the file and click **Open**.
7. Click **OK**.

4.11 Troubleshooting

This section provides guidance on dealing with the most common problems encountered when capturing and replaying workloads. In addition, it is recommended that you review the Oracle Support Web site for information about known issues and workarounds. It is available at the following location:

<https://support.oracle.com>

RUEI Installation

Ensure that the RUEI installation used to monitor the applications in the workload capture meets the following requirements:

- Check the Oracle Support Web site for information about supported versions and required hot fixes.
- RUEI has been configured to monitor the required applications. Information about deployment options and requirements is available from the *Oracle Real User Experience Insight Installation Guide*.
- You have a valid user name and password combination. If necessary, contact your RUEI Administrator. Note that the user account must have Security Officer permission. For further information about roles and permissions, see the *Oracle Real User Experience Insight User's Guide*.
- Full-Session Replay (FSR) has been enabled, and sufficient storage has been assigned, for each application that is planned to be captured. In addition, you should ensure that each application's data replay logging and masking settings are

compatible with the use of FSR. For information, see the *Oracle Real User Experience Insight User's Guide*.

- Ensure that your Web server has been configured to use static SSL certificates. This is necessary because RUEI does not support dynamic SSL certificates.
- If your application make use of jumbo frames, increase the RUEI capture length from its default 2kb to 64kb by issuing the following command on the RUEI Reporter system as the `root` user:

```
execsql config_set_profile_value wg System config CaptureLength replace 65536
```

Synchronization

If you are using synchronization replay, ensure that your capture and replay environments meet the requirements described in [Section 4.5.3, "Using Synchronization"](#).

Capture Checklist

In addition to the requirements indicated above, you should also ensure that:

- RUEI is correctly capturing all required traffic using the appropriate logging and masking policies. For information on verifying its configuration, see the *Oracle Real User Experience Insight User's Guide* available at the following location:

<http://www.oracle.com/technetwork/documentation/realuserei-091455.html>
- The database host user ID belongs to the same group as the Enterprise Manager Agent user account.

Replay Checklist

In addition to the requirements indicated above, you should also ensure that:

- All required URLs have been correctly remapped, as described in [Section 4.9.4, "Remapping URLs"](#). Check whether the test system has been configured as HTTP or HTTPS, the domain name of the Web server, and the relevant port number. In addition, verify that the full domain name is specified in the URL, and not just the host name.
- It is *strongly* recommended that you do not replay a captured workload in a production environment.
- Ensure that you have provided a substitute value for all sensitive data fields that were masked during capture. This is described in [Section 4.9.5, "Substituting Sensitive Data"](#).

Part II

Monitoring Exalogic Elastic Cloud, Exalytics Target, and Traffic Director

The chapters in this part describe how you can monitor Oracle Exalogic Elastic Cloud, Oracle Exalytics Target, and Oracle Traffic Director.

The chapters are:

- [Chapter 5, "Monitoring Oracle Exalogic Elastic Cloud"](#)
- [Chapter 6, "Discovering and Monitoring an Exalytics Target"](#)
- [Chapter 7, "Oracle Traffic Director"](#)

Monitoring Oracle Exalogic Elastic Cloud

Oracle Exalogic is an engineered hardware and software system designed to provide the ideal platform for the Oracle Fusion Middleware and business applications. By deploying Exalogic Elastic Cloud, you can increase the performance and efficiency of existing Linux, Solaris and Java applications and improve application and infrastructure reliability, scalability and availability.

The Oracle Exalogic Elastic Cloud is a system consisting of software, firmware and hardware, on which you can deploy Oracle business applications and Oracle Fusion Middleware or software. Exalogic dramatically improves performance of standard Linux, Solaris and Java applications and reduces costs across the application lifecycle.

Oracle Exalogic Elastic Cloud is modeled as a system target rather than a group target within Enterprise Manager Cloud Control. The following targets can comprise the Oracle Exalogic Elastic Cloud:

- Infiniband Network - An Infiniband Network has IB Switches as members. This is a one to many association.
- SOA Infrastructure
- WebLogic Domains
- Coherence Clusters
- OTD
- Hardware Targets (such as PDU, CISCO Switch, ZFS, and ILOMs)

For more information about adding assets to Exalogic which can then be used for monitoring on Enterprise Manager Cloud Control, refer to the documentation at the link below:

http://docs.oracle.com/cd/E27363_01/doc.121/e27511/asset_mgmt.htm#OPCFG4201

- Hosts

You can use the following sections in this chapter to learn more about how to monitor the Oracle Exalogic Elastic Cloud:

- [Prerequisites to Discovering Oracle Exalogic Elastic Cloud](#)
- [Using the Exalogic Elastic Cloud Discovery Wizard](#)
- [Upgrading Exalogic System Targets to the Version 12.1.0.3 Fusion Middleware Model](#)
- [Displaying and Using the Exalogic Elastic Cloud Dashboard](#)
- [Refreshing the Exalogic Elastic Cloud](#)

- [Monitoring the Hardware Components of Exalogic Elastic Cloud](#)
- [Viewing Application Deployments in Exalogic Elastic Cloud Targets](#)
- [Viewing WebLogic Domains in Exalogic Elastic Cloud Targets](#)
- [Viewing Coherence Clusters in Exalogic Elastic Cloud Targets](#)
- [Viewing Hosts in Exalogic Elastic Cloud Targets](#)
- [Visualizing Relationships Between Exalogic Software and Hardware Components](#)
- [Viewing Incidents and Status Changes Created for an Exalogic System in Ops Center as Incidents in Cloud Control](#)
- [Creating Exalogic Reports](#)
- [Viewing and Managing Exalogic Consumption Tracking](#)
- [Analyzing the Impact of Component Failures](#)

5.1 Prerequisites to Discovering Oracle Exalogic Elastic Cloud

There are several steps you must perform before you can discover Oracle Exalogic Elastic Cloud in Enterprise Manager Cloud Control. The following sections outline these steps in detail and provide you with the information you need to set up the Exalogic Elastic Cloud targets:

- [Importing Ops Center Certificate to the Oracle Management Agent Keystore](#)
- [Critical Prerequisites For Oracle VM Manager Discovery](#)
- [ZFS Plug-in and Oracle VM Manager Registration Preconfiguration Steps](#)

5.1.1 Importing Ops Center Certificate to the Oracle Management Agent Keystore

This section provides step-by-step instructions about how to export the certificate from the Enterprise Manager Ops Center keystore and how to import the certificate into the Enterprise Manager agent keystore. These steps are a prerequisite for Exalogic Elastic Cloud discovery in Enterprise Manager Cloud Control.

Use these steps to export the Ops Center Enterprise Controller trust certificate and import it to the Enterprise Manager Cloud Control Management Agent:

1. Change directory to `OPS_AGENT_HOME/oem-ec/security/jsse` where `OPS_AGENT_HOME` is the path to the Oracle Enterprise Manager Ops Center home. For example, `/etc/opt/sun/cacao2/instances/oem-ec/security/jsse`:

```
[root@localhost~]# cd OPS_AGENT_HOME/oem-ec/security/jsse
```

This directory change is on the VM where the Ops Center Enterprise Controller is running.

2. Export the certificate using `keytool` that corresponds to the JDK configured to run Ops Center

```
[root@localhost jsse]# $JAVA_HOME/jre/bin/keytool -export  
-alias cacao_agent -file oc.crt -keystore truststore  
-storepass trustpass
```

This command exports the certificate file (`oc.crt`) in the local directory, for example, `OPS_AGENT_HOME/oem-ec/security/jsse`. In the above command, *trustpass* is the default password for the Ops Center keystore unless changed. The Ops Center Certificate is stored in file `oc.crt`. You can import the Ops Center certificate

to the Oracle Management Agent TrustStore using the following steps. The Agent used for communication with Enterprise Manager Ops Center should be running on the same host instance. If it is not, transfer the certificate to the host instance where the agent that you want to configure is running.

3. Change directories to the agent instance home, where \$OMS_AGENT_HOME is the Oracle Management Agent home:

```
cd $OMS_AGENT_HOME/core/12.1.0.2.0/stage/sysman/config/montrust
```

4. Enter the following command:

```
$OMS_AGENT_HOME/agent_inst/bin/emctl secure add_trust_cert_to_jks -trust_certs_loc $CERT_LOC/oc.crt -password welcome -alias <alias_name>
```

If <alias_name> provided in the above command has already been used to import another certificate, you may see the following error:

```
keytool error: java.lang.Exception: Certificate not imported, alias <alias_name> already exists.
```

In this case, provide another <alias_name> in the command and try again.

Alternatively you can enter this command:

```
[root@localhost montrust]# :keytool -import -keystore $AGENT_HOME/agent_inst/sysman/config/montrust/AgentTrust.jks -alias <alias_name> -file <path to oc certificate file>
```

When prompted, enter the password *welcome* unless changed from the default, and click **Enter**.

If you follow this step for exporting the certificate, the path to the certificate file will be OPS_AGENT_HOME/oem-ec/security/jsse/oc.crt

5. To verify that the certificate has been imported correctly, run the following command:

```
[root@localhost montrust]# $JAVA_HOME/jre/bin/keytool -list -keystore $OMS_AGENT_HOME/core/12.1.0.2.0/stage/sysman/config/montrust/AgentTrust.jks
```

You should see the entry for *wlscertgencab* alias you created in the output.

6. After verifying that the certificate has been successfully imported, restart the agent by using the stop and start procedure in the following documentation:

http://docs.oracle.com/cd/E25054_01/doc.1111/e24473/emctl.htm#BABEFBDI

5.1.2 Critical Prerequisites For Oracle VM Manager Discovery

Warning: Exalogic configurations supporting virtualization must each have Oracle Virtual Machine Manager (OVMM) properly configured for read-only access by Enterprise Manager Cloud Control as a mandatory requirement as described in the following procedure. Failure to configure each Oracle Virtual Machine Manager for read-only access by Enterprise Manager Cloud Control will result in an unsupported configuration. Operations performed in this unsupported configuration may result in software outages and necessitate the re-installation of the Exalogic Control software.

If Oracle VM Manager was registered with Enterprise Manager Cloud Control prior to performing the Oracle VM Manager -- Enterprise Manager read-only configuration procedure above, then you must choose Synchronize Operation from the Oracle VM Manager Target menu in Enterprise Manager.

You must perform the following steps as a prerequisite in order to discover the Oracle VM Manager in read-only configuration.

Note: If Oracle VM Manager was registered with Enterprise Manager Cloud Control prior to performing the read-only configuration procedure below, then you must perform Synchronize Operation from the Oracle VM Manager Target menu in Enterprise Manager to re-register it in read-only mode. Please confirm that the Oracle VM Manager target is in read-only configuration as described at the end of section 6.1.2 of this document.

1. Login to Oracle VM Manager vServer as oracle user, and then perform the commands in the sequence below.
2. `cd /u01/app/oracle/ovm-manager-3/ovm_shell`
3. `sh ovm_shell.sh --url=tcp://localhost:54321 --username=admin --password=<ovmm admin user password>`
4. `ovm = OvmClient.getOvmManager ()`
5. `f = ovm.getFoundryContext ()`
6. `j = ovm.createJob ('Setting EXALOGIC_ID');`

The EXALOGIC_ID can be found in the em-context.info on dom0 located in the following file path location:

`/var/exalogic/info/em-context.info`

You must log in to dom0 as a root user to obtain this file. For example, if the em-context.info file content is *ExalogicID=Oracle Exalogic X2-2 AK00018758*, then the EXALOGIC_ID will be AK00018758.

7. `j.begin ();`
8. `f.setAsset ("EXALOGIC_ID", "<Exalogic ID for the Rack>");`
9. `j.commit ();`
10. `Ctrl/d`

You can validate the configuration following these steps by observing that several menus for the *manager/pool/server/guest vm* target are disabled. Specifically, you can validate that OVM Manager is in read-only mode by right-clicking the registered OVM manager under Infrastructure Cloud.

All the active operation menu items in the second and third tier of the list, such as Create Zone, Create Virtual Server Pool, and Manage Network, are disabled (grayed out) and are not available for selection. Conversely, if you choose an item from the pull-down menus that are displayed when you click either the VM Manager or VM Guest menu option, the same menu items may not be grayed out. However, if you select one of the active management options, Enterprise Manager displays an error message stating that the action cannot be performed, thus indicating that the OVM Manager is in read-only mode.

5.1.3 ZFS Plug-in and Oracle VM Manager Registration Preconfiguration Steps

For virtualized Exalogic configurations, there are several preconfiguration steps you must complete. The first configures the Enterprise Manager Cloud Control ZFS plug-in (optional), the second registers the Oracle VM Manager with Enterprise Manager Cloud Control (mandatory), and the third configures the Exalogic Guest Base Template and the Exalogic network.

For non-virtualized Exalogic configurations, the ZFS plug-in can be configured in step one outlined below as part of a recommended but optional configuration:

1. Configure the ZFS Plug-in

- It is recommended that the agent used for the ZFS configuration reside on the same host as the Exalogic Enterprise Controller. If an Enterprise Manager Cloud Control Agent is not present on the vServer, deploy one using the following documentation:

http://docs.oracle.com/cd/E24628_01/install.121/e22624/install_agent.htm#CACJEFJI

- Install and Configure the Enterprise Manager Cloud Control ZFS plug-in using the following documentation:

http://download.oracle.com/otn/java/oem/Oracle_GC_Plugin_Installguide.html

2. Register the Oracle VM Manager with Enterprise Manager Cloud Control

- Use the Enterprise Manager Cloud Control Agent deployed on the same vServer as the OVM Manager. If an Enterprise Manager Cloud Control Agent is not present on the vServer please deploy one using the following documentation:

http://docs.oracle.com/cd/E24628_01/install.121/e24089/part_installing_agent.htm

- In order to see Exalogic vServers in the Enterprise Manager 12c Exalogic navigation tree you must register the Oracle VM Manager with Enterprise Manager Cloud Control.

3. As a prerequisite for Enterprise Manager management of a virtual Exalogic configuration the Exalogic Guest Base Template and Exalogic network must be configured according to the following documentation:

http://docs.oracle.com/cd/E24628_01/doc.121/e35776/emex1.htm

5.2 Using the Exalogic Elastic Cloud Discovery Wizard

You can use the Exalogic Elastic Cloud Discovery wizard to discover and monitor an Exalogic target in Enterprise Manager.

The Exalogic Elastic Cloud Discovery process identifies the targets present in the Exalogic Elastic Cloud and maps them to Enterprise Manager targets, then adds Enterprise Manager targets as Exalogic Elastic Cloud system members.

To use the Exalogic Elastic Cloud Discover wizard, follow these steps:

1. In Enterprise Manager, navigate to the Systems page.
2. From the Add drop-down list, choose **Exalogic Elastic Cloud** and click **Add**.
Enterprise Manager displays the Discover Exalogic Elastic Cloud page which allows you to enter the parameters and values required to discover an Oracle Exalogic target.
3. Specify a unique **Name** for the Oracle Exalogic target you want to monitor in the Name field.
4. Select an **Agent** on one of the hosts in the Exalogic System to perform the discovery. For virtual Exalogic configurations, select the agent running on the same vServer as the Exalogic enterprise controller. If an Enterprise Manager Cloud Control Agent is not present on the vServer, deploy one using the following documentation:

http://docs.oracle.com/cd/E24628_01/install.121/e22624/install_agent.htm#CACJEFJI

For non-virtual Exalogic configurations, select an Agent on one of the hosts in the Exalogic System to perform the discovery.

If you choose an agent that is not part of the Exalogic System, an error message appears stating that no Exalogic Property File can be found and indicating that you must choose an Agent which is on an Exalogic System Host. For a virtual deployment the host may not have the ID.

5. Choose the **Deployment Type** from the drop-down. You can select either **Physical** or **Virtual**.
6. For virtualized Exalogic configurations, provide details about the Ops Center Enterprise Controller monitoring the Exalogic Elastic Cloud you want to discover. You must enter information in the **Ops Center Host Name** field, the **Ops Center Port** field, and the credentials in the **Username** and **Password** fields.

For deployment type physical (not-virtualized) Exalogic, you must provide the desired Enterprise Manager target name and a corresponding Enterprise Manager agent on an Exalogic rack host target.

Note: For Exalogic deployment types "physical" (not-virtualized) there is no hardware rack schematic in the hardware-tab page view as there is for a properly configured virtualized Exalogic target.

7. Click **Next**. Enterprise Manager displays the Discover Oracle Exalogic Targets: Discovered Targets page. The page displays all member targets (such as PDU, Switches, ILOM) that are discovered as part of the system and displays other discovered hardware targets in virtual discovery.

8. Click **Finish** to complete discovery. You can also choose **Back** to return to the Discover Exalogic Elastic Cloud page or **Cancel** to terminate the discovery process.

Enterprise Manager displays a confirmation that the Exalogic Elastic Cloud instance has been added and begins to monitor the Exalogic Elastic Cloud target. The new target is displayed on the Systems page.

5.3 Upgrading Exalogic System Targets to the Version 12.1.0.3 Fusion Middleware Model

You can upgrade the Exalogic System targets discovered in Enterprise Manager version 11 of Grid Control and Enterprise Manager version 12 of Cloud Control to the 12.1.0.3 Fusion Middleware model. Before you upgrade the Exalogic System targets, however, the Management Agent monitoring the Exalogic System target must also be upgraded and the Management Agent should have the version 12.1.0.3 Fusion Middleware Plug-in.

The upgrade job is available as a library job by choosing **Jobs** from the **Enterprise** menu and then selecting **Library**. Select the job name, **Upgrade Exalogic Systems To Fusion Middleware 12.1.0.3.0 Model**. The job can be run more than once. Run this job after upgrading the OMS and the Management Agent monitoring the Exalogic System target. If the job is successful, Enterprise Manager displays a confirmation message indicating that all Exalogic systems have been upgraded and the job is complete. If one or more Management Agents monitoring Exalogic System targets are not upgraded, the job will fail. Under such circumstances, first upgrade those Management Agents and then try submitting this job to upgrade the Exalogic System targets.

5.4 Displaying and Using the Exalogic Elastic Cloud Dashboard

Use the Software tab on the Exalogic Elastic Cloud Dashboard to display status information including alerts and key performance metrics of the following targets in the Exalogic Elastic Cloud:

- Application Deployments
- WebLogic Domains
- Coherence Clusters
- Hosts
- SOA
- OTD

The Software tab displays an Incident section in the table for each element, such as Application Deployments or WebLogic Domains, that shows alerts on which you can drill down for more information.

You can also use the Software tab to view the alarms that are generated in OpsCenter for an Exalogic Elastic Cloud system as incidents for an Exalogic system target. The Software tab displays alarms and incidents generated for OpsCenter as incidents in Cloud Control for an Exalogic System target so that you can view all the issues from both the systems at one place. This allows you to use the existing integration between Cloud Control and the ticketing system to manage the OpsCenter incidents.

The Software tab displays large numbers of software targets that are members of the Exalogic system. You can use the inter-region filtering option to conduct a page level search. For example, if the search is performed for an Oracle WebLogic Domain, the Applications region is also filtered to show all the applications deployed on the

WebLogic servers that are members of the searched WebLogic domain. Similarly the Host region is filtered to show hosts that are hosting the WebLogic servers that are part of the searched WebLogic domain.

You can also use the Exalogic Elastic Cloud Dashboard page to access the Hardware tab where you can view information about the hardware and infrastructure of the Exalogic Elastic Cloud. You can view information such as Temperature (by turning on the Temperature option at the top of the Exalogic Schematic) and Status, as well as a Hardware Schematic depicting the current state of the hardware.

The Hardware tab displays the number of hardware units in each category along with their status in the Overview section. It also shows the Incidents associated with the hardware. The Exalogic Hardware schematic shows the position of each hardware component on the rack. You can use the Incidents section to monitor hardware issues that arise and then drill down to the Incident Manager page by clicking the **Incident Summary** for more details about the incident.

To display and use the Exalogic Elastic Cloud Dashboard, follow these steps:

1. In Enterprise Manager, navigate to the Systems page. Filter the entries in the Search field by choosing **Exalogic Elastic Cloud** from the drop-down list and click **Go**. In the Search Results table, choose the Exalogic Elastic Cloud you want to view.

Enterprise Manager displays the Exalogic Elastic Cloud Home page where you can monitor the components on the Exalogic Dashboard.

2. You can view detailed information about each software component by choosing the component name from the Exalogic Elastic Cloud drop-down list. Only the software components are visible in the Exalogic drop-down list. Other hardware components can be seen in the Members tab.

Enterprise Manager displays the component page you selected. For example, select WebLogic Domains Summary page to see the charts showing the status of weblogic servers, Request Processing Time metric information, CPU Usage, Requests per minute, and Heap Usage data.

3. You can return to this page at any time by choosing Home from the Exalogic Elastic Cloud drop-down list.
4. You can display General Information about the Exalogic target by choosing Target Information from the Exalogic Elastic Cloud drop-down list.

5.5 Refreshing the Exalogic Elastic Cloud

You should consider refreshing the system whenever you upgrade the Exalogic Rack, for example, from a quarter to half. You should also consider refreshing when you discover Oracle VM Manager or ZFS targets after the Exalogic system is discovered. To refresh the Exalogic Elastic Cloud, follow these steps:

1. From the Targets menu, select **Exalogic**.
2. From the list of Exalogic targets, click the Exalogic Elastic Cloud target you want to view.
3. From the Exalogic Elastic Cloud menu, choose **Refresh Exalogic Elastic Cloud**.

Cloud Control displays a page where you must click **Refresh** to start the refresh action or **Cancel** to cancel. When the refresh action completes, Enterprise Manager displays a Confirmation page that shows the targets found and prompts you to press any button to continue.

4. Click **Close**.

When you close the Confirmation page, the Discover Exalogic Elastic Cloud: Discovered Targets page appears where you can view all targets associated with the system. You can then click **Add Targets** to save the targets.

5.6 Monitoring the Hardware Components of Exalogic Elastic Cloud

You can monitor the hardware components of an Exalogic Elastic Cloud by following these steps:

1. From the Enterprise Manager Systems page, click the Exalogic Elastic Cloud target from the table. Alternatively, you can choose **Exalogic** from the Targets menu.

Enterprise Manager displays the Exalogic Elastic Cloud Dashboard page.

2. Click the **Hardware** tab to view a schematic showing the hardware configuration of the Exalogic Elastic Cloud.

You can use the Hardware tab to view information such as Temperature (by turning on the **Temperature** option at the top of the Exalogic Schematic) and Status, as well as a Hardware Schematic depicting the current state of the hardware.

The Hardware tab displays the number of hardware units in each category along with their status in the Overview section. It also shows the Incidents associated with the hardware. The Exalogic Schematic depicts the hardware, such as Infiniband Switches, Storage Heads, Storage Disks, and Compute Nodes and employs a Legend that indicates the current state (Up, Down, Unallocated, and so on) of each.

You can use the Incidents section to monitor hardware issues that arise and then drill down to the Incident Manager page by clicking the **Incident Summary** for more details about the incident.

5.7 Viewing Application Deployments in Exalogic Elastic Cloud Targets

Use the Application Deployments page in the Exalogic Elastic Cloud target area to view details about the applications hosted on the hosts running on the Exalogic Elastic Cloud target.

To view application deployments in Exalogic Elastic Cloud targets, follow these steps:

1. In Enterprise Manager, navigate to the Systems page. Filter the entries in the Search field by choosing **Exalogic Elastic Cloud** from the drop-down list and clicking **Go**. In the Search Results table, choose the Exalogic Elastic Cloud you want to view.

Enterprise Manager displays the Exalogic Elastic Cloud Home page where you can monitor the status of the Exalogic target and its components.

2. Choose **Application Deployments** from the Exalogic Elastic Cloud menu.

Enterprise Manager displays the Application Deployments page.

3. You can choose to show All Domains or filter by specific domains by choosing the domain from the Show menu.
4. You can drill down to specific applications, targets, domains, or dependencies by clicking on its related value in each row.

5. You can filter the list of applications by choosing a value from the Status drop-down. You can select from Up, Down, Unknown, Blackout, and All.
6. You can change the column appearance of the table by clicking **View** and choosing which Columns to display, expanding or collapsing rows, or scrolling to the first or last row. You can also reorder columns.
7. You can use the Topology tab to display a pictorial view of the Application Deployments in various relational configurations.

5.8 Viewing WebLogic Domains in Exalogic Elastic Cloud Targets

You can use Enterprise Manager to view details about the domains hosted on the virtual machines running on Exalogic Elastic Cloud target. To view WebLogic domains in Exalogic Elastic Cloud targets, follow these steps:

1. In Enterprise Manager, navigate to the Systems page. Filter the entries in the Search field by choosing **Exalogic Elastic Cloud** from the drop-down list and click **Go**. In the Search Results table, choose the Exalogic Elastic Cloud you want to view.

Enterprise Manager displays the Exalogic Elastic Cloud Home page where you can monitor the status of the Exalogic target and its components.

2. Choose **WebLogics Domain** from the Exalogic Elastic Cloud menu. You can choose to view either a Summary of the WebLogic Domains or specific information about Members.

Enterprise Manager displays the related WebLogic Domain page.

3. On the Summary page you can view a chart that shows the status of the WebLogic Domains and displays the percentage of domains that are up and down. You can also view server information that shows the Server Status and alert and policy violation information for each. You can monitor charts that display metric information such as Request Processing Time and CPU Usage and you can drill down through these charts for more detailed information. Change the chart view to a table view by clicking **Table View** or **Chart View** beneath each table or chart. The Server table displays information about servers and domains showing host information and related metrics.
4. On the Members page, you can view the Status information along with alerts and policy violation and metric data for each WebLogic Server or Domain. Use the Performance Summary section to view metrics for each, such as Host and Cluster information and metrics such as Heap Usage and Request Processing Time.

5.9 Viewing Coherence Clusters in Exalogic Elastic Cloud Targets

You can use Enterprise Manager Cloud Control to view details about the Coherence targets that comprise the Exalogic Elastic Cloud target. To display Coherence Clusters in Cloud Control, follow these steps:

1. In Enterprise Manager, navigate to the Systems page. Filter the entries in the Search field by choosing **Exalogic Elastic Cloud** from the drop-down list and click **Go**. In the Search Results table, choose the Exalogic Elastic Cloud you want to view.

Enterprise Manager displays the Exalogic Elastic Cloud Home page where you can monitor the status of the Exalogic target and its components.

2. Choose **Coherence Clusters** from the Exalogic Elastic Cloud menu.

Enterprise Manager displays the Coherence Clusters page.

3. You can view a chart that shows the status of the Coherence Clusters and displays the percentage of clusters that are up and down.
4. You can drill down to specific values for each cluster such as Alerts and Policy Violations along with Node information.
5. You can filter the list of clusters by choosing a value from the Status drop-down. You can select from Up, Down, Unknown, Blackout, and All.
6. You can change the column appearance of the table by clicking **View** and choosing which Columns to display. You can also reorder columns.
7. The Coherence Clusters page displays two charts showing the Top Nodes With Lowest Available Memory and Caches With Lowest Hit To Get Ratio. You can drill down to specific node information by clicking on the Node name below the Top Nodes With Lowest Available Memory chart.
8. The Nodes table displays information about each Node, including Host and several metric values such as Memory Available, Gets, and Puts.
9. The Applications table displays information about applications such as Local Attribute Cache, Clustered Session Cache, and other metrics. You can drill down to specific information about each application by clicking the Application name.

5.10 Viewing Hosts in Exalogic Elastic Cloud Targets

You can view details about the host targets hosted on the physical and virtual machines running on the Exalogic Elastic Cloud target. To display the hosts information, follow these steps:

1. In Enterprise Manager, navigate to the All Targets page. Filter the entries in the Search field by choosing **Exalogic Elastic Cloud** from the drop-down list and click **Go**. In the Search Results table, choose the Exalogic Elastic Cloud you want to view.

Enterprise Manager displays the Exalogic Elastic Cloud Home page where you can monitor the status of the Exalogic target and its components.

2. Choose **Hosts** from the Exalogic Elastic Cloud menu.

Enterprise Manager displays the Hosts page.

3. You can view a chart that shows the status of the hosts and displays the percentage of hosts that are up and down.
4. You can view information about the Middleware Targets that lists the Type, Status, CPU Utilization percentage, Memory Utilization percentage, and Incident statistics along with Configuration Changes.
5. You can view charts showing the CPU Utilization percentage based on time and similarly, Memory Utilization based on time.

5.11 Exalogic Control Stack Monitoring

The Exalogic Elastic Cloud Software control software stack can be monitored by choosing Control Stack in the left hand navigation tree of the Exalogic target. Cloud Control displays a master-detail page of Exalogic Elastic Cloud Software virtual machines, corresponding hosts and control software element. Selecting each row master element results in a change in the detail of the page and corresponding key

performance indicator charts. This page provides a unified view of incidents related to the Exalogic Elastic Cloud Software control software stack.

5.12 Visualizing Relationships Between Exalogic Software and Hardware Components

Use Enterprise Manager Cloud Control to visualize relationships among Exalogic software and hardware components. To see the routing topology among these components, follow these steps:

1. From the Targets menu, select **Exalogic**.
2. From the list of Exalogic targets, click the Exalogic Elastic Cloud target you want to view.
3. From the Exalogic Elastic Cloud menu, select **Routing Topology**.
4. View the relationships between Exalogic software and hardware components.

5.13 Analyzing the Impact of Component Failures

You can analyze the impact of component failures in an Exalogic environment by following these steps:

1. From the Targets menu, select **Exalogic**.
2. From the list of Exalogic targets, click the Exalogic Elastic Cloud target you want to view.
3. From the navigation panel, click the Oracle Infiniband Network you want to view.
4. From the Infiniband Network menu, select **Topology**.
5. Identify any problems with a component by observing status and incident badges, or by showing metric values by way of the Annotations toolbar menu.
6. Identify components linked to the problem component in the diagram.

5.14 Viewing Incidents and Status Changes Created for an Exalogic System in Ops Center as Incidents in Cloud Control

This section describes how to configure Enterprise Manager Cloud Control in order to consolidate alarms/incidents from OpsCenter and consequently view all the issues from both systems in a single console. Furthermore, this configuration enables you to use the existing integration between Enterprise Manager Cloud Control and the ticketing system to manage the Ops Center incidents. Please note that once you start to manage OpsCenter incidents within Enterprise Manager Cloud Control, if for any reason you alter incidents in OpsCenter, those changes will not be reflected in Cloud Control.

The functionality of viewing Ops Center Alarms in Cloud Control is disabled by default. You can enable this functionality to view Ops Center Alarms in Cloud Control by performing the following steps:

1. From the **Exalogic System** target menu, select **Monitoring** and then choose **Metric and Collection Settings**.
2. Enable the Collection Schedule for the Ops Center Alarms metric by clicking **Collection Schedule**.

To view incidents and status changes created for an Exalogic system in Ops Center as incidents in Cloud Control, follow these steps:

1. From the **Targets** menu, choose **Exalogic**.
Enterprise Manager displays the Systems page.
2. From the Systems page, click the **Name** of the Exalogic Elastic Cloud in the Systems table.
Enterprise Manager displays the Exalogic Elastic Cloud Dashboard.
3. Click the **Hardware** tab to view alarms and incidents in Ops Center as an incident in Cloud Control for the Exalogic System target.

5.15 Creating Exalogic Reports

You can create the following Exalogic reports in Enterprise Manager:

- Oracle BI Instance Performance -- This report covers all BI instances running on the selected Exalogic rack. First a list of all WebLogic domains associated with the selected Exalogic rack is read, then all 'oracle_ias_farm' targets associated with those WebLogic domains are determined and then the list of all BI instances associated with those 'oracle_ias_farm' targets are identified.
- Oracle Internet Directory Performance -- This report covers all OID servers running on the selected Exalogic rack. First a list of all WebLogic domains associated with the selected Exalogic rack is read, then all 'oracle_ias_farm' targets associated with those WebLogic domains are determined and then the list of all OID servers associated with those 'oracle_ias_farm' targets are identified.
- Oracle Service Bus Performance -- This report covers all OSB servers running on the selected Exalogic rack. First a list of all WebLogic domains associated with the selected Exalogic rack is read and then all OSB server targets associated with those WebLogic domains through 'farm' association are identified. This report has two sections under each OSB server – Top Business Services and Top Proxy Services. These sections show the top five business services and the top five proxy services respectively. The top five services are determined based on the decreasing order of max value of throughput (for example, messages per minute) metric for each service during the selected period.
- SOA Infrastructure Performance -- This report covers all SOA Infra targets running on the selected Exalogic rack. First a list of all WebLogic domains associated with the selected Exalogic rack is read and then all SOA Infra targets associated with those WebLogic domains through 'farm' association are identified. The SOA Infra Engine related charts cover only the following four engine types - bpel, bpmn, mediator and workflow.
- Oracle WebLogic Server Performance -- This report covers all WebLogic servers running on the selected Exalogic rack. The list of all WebLogic domains associated with the selected Exalogic rack is read. The report is divided into sections based on different WebLogic domains. For each domain, the top five servers are shown in the charts for each metric. The top five servers for each metric are determined by decreasing order of max value for the metric for each server in the selected period.
- Compliance Summary -- This report give a brief summary about the compliance standards evaluated for the middleware targets like WebLogic server, domains, clusters, host, and so on, running on the Exalogic system. The report shows the list of top ten least compliant WebLogic domains and host and Compliance Standard Summary.

- **Compliance Details** -- This report give details on about the compliance standards evaluated for middleware targets like WebLogic domains, servers, clusters, hosts, and so on, running on the Exalogic system. The reports show the top ten least compliant WebLogic domains and host and a detailed list on compliance rules violated for the middleware targets.

The Reports menu item is visible in the Exalogic system menu if BIP is configured with Enterprise Manager. If BIP is not configured, the menu item does not appear.

To create Exalogic reports, follow these steps:

1. From the Targets menu, select **Exalogic**.
2. From the list of Exalogic targets, click the Exalogic Elastic Cloud target for which you want to create reports.
3. From the Exalogic Elastic Cloud menu, choose **Reports** and then select the report you want to create from the menu.

5.16 Viewing and Managing Exalogic Consumption Tracking

You can use resource utilization features in Enterprise Manager to view consumption tracking and set the license edition for each target of a Exalogic system. To view consumption tracking information, follow these steps:

Note: You can view the Consumption Report page only for virtual Exalogic machine targets.

1. From the Targets menu, select **Exalogic**.
2. From the list of Exalogic targets, click the Exalogic Elastic Cloud target for which you want to view consumption information.
3. From the Exalogic Elastic Cloud menu, choose **Resource Utilization**. and then select **Consumption Tracking**.

Cloud Control displays the vCPU Consumption page where you can perform the following tasks:

- View the resource consumption trend for the previous 24 hours by choosing **View Trends**.
- Export the data on the vCPU Consumption page for the selected Time Period, View, and Product by clicking **Export**.
- Click **Manage License Edition** to access a dialog box where you can set the license edition for each target of the current Exalogic system.
- View the details for the selected product in the Details table. The table displays the Collection Time/Host, vCPUs, Target Name, Target Type, Version, Cost Center, and Owner.

5.17 Viewing the vCPU Consumption Report

To view the vCPU Consumption Report for the software components running on an Exalogic Machine, do the following:

1. On the vCPU Consumption page, select one of the time periods in the **Time Period** drop-down list.

Note: When you select a time period, the Product table and Details table refreshes automatically for the chosen time period.

2. From the Product list, select the product or the products for which you want to view the consumption report. By default all the products are selected.
3. From the Product table, select a **Product** for which you want to view the consumption report.

Note: When you select a product, the Details table refreshes automatically to display the Consumption Report for the product.

4. The Details table displays the consumption details for the product or products that you have selected.

You can add more details as columns to the Details table by clicking **View**, selecting **Columns**, and then selecting the items you want to be displayed. You can select which columns you want to be visible and which columns to hide by clicking **Manage Columns**.

By default, the Details table displays the collection time and the number of total vCPUs consumed. Expand the collection time and host name to view all the following details:

- Collection time of the point when vCPU was collected
- Host name of the host on which software is running
- The Exalogic machine name
- The number of total vCPUs consumed by the software/target running at that time
- The target name
- The target type
- The owner of the target

5.18 Viewing the Resource Consumption Trend

To view the resource consumption trend for an Exalogic machine target, do the following:

1. On the vCPU Consumption page, from the Product drop-down list, select the product for which you want to view the resource consumption trend.

Select **All** if you want to view the resource consumption trend for all the targets.

Note: When you make a selection from the Product drop-down list, the Product table and the Details table automatically refreshes to display the details for the product or products that you have selected.

2. From the Time Period drop-down list, select the time period for which you want to view the resource consumption trend.
3. Click **View Trends**.

4. The Resource Consumption Trend window opens.
If you selected All products, then the graphical representation of the resource consumption for all the products will be displayed. You can choose to view trends for individual products as well, by clicking on the Product drop-down list, and selecting the product.
5. You can also choose to view the resource consumption trend in a table format by clicking **Table View**.
On the Table view page, you can customize the details you want to view in the table by clicking **View**.
6. Click **OK** to close the Resource Consumption trend window.

5.19 Exporting the vCPU Consumption Report

To export the vCPU consumption report for an Exalogic machine target, do the following:

1. On the vCPU Consumption page, select the following:
 - **Time Period**
Select a time period for which you want to view the vCPU consumption report for.
 - **View**
Select **Product**.
 - **Product**
Select the products for which you want to view the vCPU consumption report.
Select **All** if you want to view the vCPU consumption report for all the products.
2. Click **Export**.
3. On the Opening License dialog box, select the program with which you want to open the file. Microsoft Office Excel is the default program. Click **OK**.

5.20 Configuring Exalogic Oracle Engineered System Healthchecks

Health check targets are discovered independently of Exalogic. This target can be navigated by a menu item which is added as part of integration of Health Check with Exalogic. The HealthCheck menu item is hidden in Exalogic system and rack menu if the HealthCheck target is not discovered. Once the HealthCheck target is discovered, you must refresh the Exalogic Cloud, which will then enable the menu item. You can then navigate to the HealthCheck page from the Exalogic system and rack menu.

Each rack has its own HealthCheck target and you can navigate to its own HealthCheck target (in case of multi rack scenario) while the Exalogic system displays the HealthCheck target of the primary rack.

For more information about configuring Exalogic Oracle-engineered system healthchecks, see the *Oracle Enterprise Manager System Monitoring Plug-In Installation Guide for Oracle Engineered System Healthchecks*.

Discovering and Monitoring an Exalytics Target

This chapter contains the following sections:

- [Discovering an Exalytics Target](#)
- [Monitoring an Exalytics System](#)
- [Monitoring an Exalytics Machine](#)

6.1 Discovering an Exalytics Target

In order to manage and monitor Oracle Fusion Middleware components running on Exalytics, such as WebLogic domain, Oracle BI Foundation 11g, and Oracle TimesTen (in-memory database), Enterprise Manager Cloud Control must first discover the Exalytics machine containing these components.

Once discovered, the Exalytics machine and the components within it can be promoted to "managed target" status, enabling Enterprise Manager Cloud Control to collect the data needed to monitor the target.

This section contains the following topics:

- [Prerequisites for Discovering an Exalytics Machine Target](#)
- [Discovering an Exalytics System Target](#)

6.1.1 Prerequisites for Discovering an Exalytics Machine Target

Note: An Exalytics system consists of one or more Exalytics machines. The machine(s) can be physical, virtualized, or a mix of both. These prerequisites are for discovering an Exalytics machine target.

Before discovering an Exalytics machine target, meet the following prerequisites:

- A Management Agent for discovering targets must be deployed onto the physical Exalytics machine (host).

For discovering a virtual machine, make sure that at least one of the virtual Exalytics has a Management Agent running on it. However, all virtual machines require a Management Agent to be deployed on them for monitoring.

- To identify the Exalytics machine, ensure that you have the context info file which contains the Exalytics machine ID.

6.1.2 Discovering an Exalytics System Target

To discover an Exalytics system target in an Exalytics system, follow these steps:

1. From the **Setup** menu, select **Add Target**, and then select **Add Targets Manually**.
2. On the Add Targets Manually page, select **Add Targets Using Guided Process (Also Adds Related Targets)**.
3. From the Target Types drop-down list, select **Exalytics System**, and then click **Add Using Guided Process...**
4. On the Discover Exalytics System page, provide a name for the Exalytics System, and then click **Add Machine**.
5. Provide the following details:
 - **Name**
Provide a unique for the Oracle Exalytics machine target.
 - **Agent**
Enter or select the Management Agent to use for the Oracle Exalytics Machine discovery process.
 - **Deployment Type**
Select **Physical** or **Virtual** depending on if you want to discover a physical target or a virtual target.
 - **Host Name** (only for Virtual target)
Enter the host name or IP Address where Oracle Virtual Server is running
 - **User Name and Password**
For a physical Exalytics machine, provide credentials of the user which has privileges to run the imageinfo command.

For a virtual Exalytics machine, provide credentials to log in to Oracle Virtual Server (OVS). You should have privileges to run the imageinfo command.
6. A confirmation box appears. Click **OK**.

Note: Click **Add Targets** to save the targets as a Manageable Entity.

6.2 Monitoring an Exalytics System

You can use Enterprise Manager Cloud Control 12c to monitor and administer certain Exalytics systems. Proactively monitoring the status, performance, and health of Exalytics system targets enables you to use your Exalytics system more efficiently and react faster to potential issues.

By monitoring a system using Enterprise Manager, you obtain a complete and up to date overview of the status, availability, performance, and health of the system. Enterprise Manager displays complex target performance data in a simple form, using graphs and pie charts. It also keeps you informed about target metrics crossing their threshold levels, target alerts, and target incidents that require user action.

Using the Exalytics System console, you can do the following tasks:

- [Displaying and Using the Exalytics System Dashboard](#)

- [Viewing the Generic Performance of an Exalytics System](#)
- [Viewing the Members Topology of an Exalytics System](#)
- [Viewing the Routing Topology of an Exalytics System](#)
- [Refreshing an Exalytics System](#)
- [Viewing the Job Activity of an Exalytics System](#)
- [Administering Targets of an Exalytics System](#)

6.2.1 Displaying and Using the Exalytics System Dashboard

To access the Exalytics System dashboard, from the **Targets** menu, select **Exalytics**. On the Systems page, select the Exalytics System that you want to access.

The Exalytics System Dashboard displays status information including alerts and key performance metrics of the following Fusion Middleware targets that are managed and monitored by the Exalytics System:

- WebLogic domains
- Oracle BI Foundation 11g
- Oracle TimesTen (in-memory database)
- Hosts

By default, the table for each Fusion Middleware component displays the name of the component, the type, the status, incidents, the host the component is running on, and the port. You can add more columns by clicking **View**, selecting **Columns** and then choosing the options that you want displayed as columns in the table.

The dashboard also displays an Incidents and Problems section that shows alerts for targets that have problems. The table also displays the status of the target, the severity of the problem, the escalation level, and the time when the target was last updated.

You can view detailed information about the Exalytics System from the Exalytics System menu, such as performance summary, alert history, target incidents, members topology, routing topology, and job activity. You can also perform management and administrative tasks such as refreshing the system, removing the system, granting target access privileges, and administering target blackouts.

6.2.2 Viewing the Generic Performance of an Exalytics System

Using the Enterprise Manager Exalytics System console, you can perform the following generic performance tasks:

- [Viewing the Performance Summary of an Exalytics System](#)
- [Viewing the Alert History of an Exalytics System](#)
- [Viewing Target Incidents of an Exalytics System](#)

6.2.2.1 Viewing the Performance Summary of an Exalytics System

To view the performance summary of the Exalytics system, follow these steps:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, select the Exalytics System that you want to access.
2. On the Exalytics system home page, click on the **Exalytics System** menu, select **Monitoring**, and then select **Performance Summary**.

3. The following charts are displayed:

- Average query time (in seconds)
- Server queries (per second)
- Status of BI Instance
- Active Sessions
- Failed Requests
- Requests (per minute)

You can change the order of the charts by clicking on the **View** menu, and selecting **Reorder Charts**. This enables you to move a chart up, down, or to remove it.

4. On the Performance Summary page, you can do the following:

- You can view the performance of the Exalytics summary for the entire day, 2 hours, or for 15 minutes.
- You can set what you want to be displayed on the chart by clicking the **View** menu and selecting or deselecting the items.
- You can overlay the chart displayed with the chart from the previous chart, or you can overlay it with the performance summary chart of another Exalytics System.

To do this, click on **Overlay**, and then select an option.

- You can view the metric palette by clicking on **View Metric Palette**.

6.2.2.2 Viewing the Alert History of an Exalytics System

To view the alert history of an Exalytics system, from the Exalytics System menu displayed on the Exalytics system home page, select **Monitoring**, then select **Alert History**.

The Alert History page provides details about target metrics, such as the periods when a particular metric was beyond its critical threshold value, the periods when the metric could not be calculated, and so on. These details help you plan corrective measures for metric-related problems, before any severe damage or prolonged downtime can occur.

The Alert History page contains the following elements:

Table 6–1 Target Alert History

Element	Description
Metric	Parameter related to the performance of the target.
History	Condition of the metric at various times. The condition can have the values Critical, Warning, Clear, and No Data.

6.2.2.3 Viewing Target Incidents of an Exalytics System

To view the incidents related to the target, from the Exalytics System menu displayed on the target home page, select **Monitoring**, then select **Incident Manager**.

The Incident Manager page provides details about the various events, related to the target, that negatively impact any hardware or software component running on the Exalytics system. These events require user action. The details provided by this

section, such as the incident summary, severity, target, target type, and so on, are essential for troubleshooting.

6.2.3 Viewing the Members Topology of an Exalytics System

To view the members topology of an Exalytics system, follow these steps:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, select the Exalytics System that you want to access.
2. From the Exalytics System menu displayed on the Exalytics system home page, select **Members**, and then select **Topology**.
3. The Configuration Topology page for the System is displayed.

The topology viewer enables you to understand and visualize your IT environment in a graphical format. The inter-relationships between systems display results such as on what systems does this system rely, what systems will be affected if this system goes down, and how are the members of my system related?

On this page, you can perform the following tasks:

- [Viewing a Configuration Topic](#)
- [Creating a Custom Topology View](#)

6.2.3.1 Viewing a Configuration Topic

To view a configuration topic, from the **View** list, select any of the following:

- **Uses**
This view helps you determine the targets that the selected target depends on. If a target is having problems, this view can be useful in helping you determine whether its problems have been caused by another target it depends on.
- **Used By**
This view shows you the targets that depend on the selected target. This can be useful, for example, if you are planning on shutting down the target and need to know what other targets will be affected
- **System Members**
This view shows the members of the system.
- **Custom views** that have been defined by end users.

Note: The Uses, Used By, and System Members views are out-of-box topology views. They cannot be modified.

6.2.3.2 Creating a Custom Topology View

To create a custom topology view, do the following:

1. From the **Customize** menu, select **Create Custom View**.
2. Provide the name and description for the topology and select one of the Initial Contents:
 - **Copy Current View** to create a topology view similar to the one you are viewing.

- **Create Empty View** to create a topology view that starts with the root node.
3. Choose one of the following expose options:
 - Expose the custom view for all targets of the current target type. For example, if you are creating a topology view for a database target, the new view will be available for all database targets.
 - Expose the custom view for the current target only.
To share the view, click **Share this custom view with other users**.
 - Click **OK**.
 4. Reduce the unwanted information in the topology by highlighting the target and selecting **Hide Relationships...** in the **Customize** menu.
You can also display relationships that are not being displayed by selecting a target. From the **Customize** menu, select **Target**, then select **Show More Relationships to Target Type....**
 5. Click **OK**.

6.2.4 Viewing the Routing Topology of an Exalytics System

A topology diagram is an interactive diagram showing Enterprise Manager targets and their statuses, incidents, and other target information, as well as the relationships between targets.

To view the Exalytics Routing Topology page, do the following:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, select the Exalytics System that you want to access.
2. From the Exalytics System menu displayed on the target home page, select **Routing Topology**.

The Routing Topology page displays middleware targets in a single topology diagram, allowing you to see dependencies among targets and the impact of component failures.

You can use this page to visualize which software components are communicating with each other, and which hardware components they are using in order to understand the system. You can learn which components of a system are impacted by a problem with another component.

The topology graph expresses the relationships between the targets of the Exalytics system. The targets are annotated to indicate status and incidents. Tool tips show additional information about entities and relationships. You can pan and zoom the graph, as in popular map programs, and also expand and collapse desired areas of the graph. The graph can be filtered to show only desired nodes.

6.2.5 Refreshing an Exalytics System

To refresh an Exalytics System, do the following:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, select the Exalytics System that you want to access.
2. From the Exalytics System menu displayed on the target home page, select **Refresh Exalytics System**.
3. The Exalytics target discovery page opens.

You can rediscover Exalytics machine targets to add to the Exalytics system, or you can remove existing targets from the system.

To discover Exalytics machine targets, refer to [Discovering an Exalytics System Target](#).

6.2.6 Viewing the Job Activity of an Exalytics System

To view the past, currently running, and scheduled jobs related to a target, follow these steps:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, select the Exalytics System that you want to access.
2. From the Exalytics System menu displayed on the target home page, select **Job Activity**.
3. The Job Activity page displays target jobs related to target administrative tasks, such as starting the target, stopping the target, target blackouts, and so on.

Use the Job Activity page to search for a particular job and retrieve job details such as the owner, status, scheduled start time, and so on. You can also use the Job Activity page to perform target job administration tasks, such as creating, editing, suspending, and resuming a job.

6.2.7 Administering Targets of an Exalytics System

You can perform the following administration tasks on a target of an Exalytics system:

- [Administering Target Access Privileges for an Exalytics System](#)
- [Administering Target Blackouts for an Exalytics System](#)
- [Removing an Exalytics System Target](#)

6.2.7.1 Administering Target Access Privileges for an Exalytics System

To manage the access privileges for a target, follow these steps:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, select the Exalytics System that you want to access.
2. From the Exalytics System menu displayed on the Exalytics system home page, select **Target Setup**, then select **Administrator Access**.
3. Click **Add** to grant target access privileges to a role or an administrator.

Use the Access page to set target privileges for roles and administrators. The available privileges are View, Operator, and Full.

View only allows you to view the target in the console, whereas Operator allows you to view targets, and perform all administrative actions except deleting targets. Full allows you to view targets, and perform all administrative actions.

6.2.7.2 Administering Target Blackouts for an Exalytics System

To administer the blackouts for a target, follow these steps:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, select the Exalytics System that you want to access.
2. From the Exalytics System menu displayed on the Exalytics system home page, select **Monitoring**, then select **Blackouts**.

3. Blackouts suspend data collection on a monitored target. Blackouts are useful when you want to perform scheduled maintenance tasks on monitored targets.

Use the Blackouts page to search for existing target blackouts, edit existing blackouts, define new blackouts, and stop blackouts. You can also create and stop blackouts using the Exalytics System menu. To create or stop a blackout, from the Exalytics System menu, select **Control**, then select **Create Blackout** or **End Blackout**, respectively.

6.2.7.3 Removing an Exalytics System Target

To remove an Exalytics machine target from an Exalytics System, do the following:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, select the Exalytics System that you want to access.
2. On the Exalytics system home page, from the Exalytics System menu, select **Target Setup**, and then **Remove Target**.
3. A warning will be displayed, asking if you want to proceed. Click **Remove**.

6.3 Monitoring an Exalytics Machine

You can use Enterprise Manager Cloud Control 12c to monitor and administer certain Exalytics machine targets. Monitoring the status, performance, and health of Exalytics machine targets enables you to set up a more efficient Exalytics system.

By monitoring a target using Enterprise Manager, you obtain a complete and up to date overview of the status, availability, performance, and health of the target. Enterprise Manager displays complex target performance data in a simple form, using graphs and pie charts. It also keeps you informed about target metrics crossing their threshold levels, target alerts, and target incidents that require user action.

Using the Exalytics Machine console, you can do the following tasks:

- [Displaying and Using the Exalytics Machine Homepage](#)
- [Viewing the Generic Performance of an Exalytics Machine](#)
- [Viewing the Members Topology of an Exalytics Machine](#)
- [Viewing the vCPU Consumption Report for an Exalytics Machine](#)
- [Managing License Editions for an Exalytics Machine](#)
- [Viewing the Job Activity of an Exalytics Machine](#)
- [Viewing the Collected Configuration of an Exalytics Machine](#)
- [Administering the Exalytics Machine Target](#)

6.3.1 Displaying and Using the Exalytics Machine Homepage

To view the Exalytics Machine homepage, follow these steps:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, click on the expansion button of the Exalytics System that contains the Exalytics target that you want to access. A list of Exalytics targets is displayed.
2. Select the Exalytics target that you want. This takes you to the Exalytics Machine home page.

The Exalytics Machine home page displays the summary of the functioning of the Exalytics target. It shows the general information of the target, the overview of

incidents and problems, the status of the target, the compliance summary, jobs activity, target blackouts, and the number of configuration changes.

You can view more detailed information of the target by clicking on the Exalytics Machine menu and selecting the option that you want to view, such as target alert history, target incidents, members topology, vCPU consumption report, job activity and collected configuration. You can also manage and administer tasks such as administering target access privileges, administering target blackouts, and removing the target.

The Exalytics Machine homepage also contains a **Personalize Page** button on the top right corner. This enables you to customize the homepage by adding, removing, or configuring regions. You can also change the layout of the page.

6.3.2 Viewing the Generic Performance of an Exalytics Machine

Using the Enterprise Manager Exalytics System console, you can perform the following generic performance tasks:

- [Viewing the Target Alert History of an Exalytics Machine](#)
- [Viewing the Target Incidents of an Exalytics Machine](#)

6.3.2.1 Viewing the Target Alert History of an Exalytics Machine

To view the alert history of an Exalytics machine, do the following:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, click on the expansion button of the Exalytics System that contains the Exalytics target that you want to access. A list of Exalytics targets is displayed.
2. Select the Exalytics target that you want. This takes you to the Exalytics Machine home page.
3. From the Exalytics Machine menu displayed on the Exalytics machine home page, select **Monitoring**, then select **Alert History**.

from the Exalytics Machine menu displayed on the Exalytics machine home page, select **Monitoring**, then select **Alert History**.

The Alert History page provides details about target metrics, such as the periods when a particular metric was beyond its critical threshold value, the periods when the metric could not be calculated, and so on. These details help you plan corrective measures for metric-related problems, before any severe damage or prolonged downtime can occur.

The Alert History page contains the following elements:

Table 6–2 Target Alert History

Element	Description
Metric	Parameter related to the performance of the target.
History	Condition of the metric at various times. The condition can have the values Critical, Warning, Clear, and No Data.

6.3.2.2 Viewing the Target Incidents of an Exalytics Machine

To view the incidents related to the target, navigate to the Incidents section, by following these steps:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, click on the expansion button of the Exalytics System that contains the Exalytics target that you want to access. A list of Exalytics targets is displayed.
2. Select the Exalytics target that you want. This takes you to the Exalytics Machine home page.
3. From the Exalytics Machine menu displayed on the target home page, select **Monitoring**, then select **Incident Manager**.

The Incident Manager page provides details about the various events, related to the target, that negatively impact any hardware and software components running on the Exalytics machine. These events require user action. The details provided by this section, such as the incident summary, severity, target, target type, and so on, are essential for troubleshooting.

6.3.3 Viewing the Members Topology of an Exalytics Machine

To view the members topology of an Exalytics machine, follow these steps:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, click on the expansion button of the Exalytics System that contains the Exalytics target that you want to access. A list of Exalytics targets is displayed.
2. Select the Exalytics target that you want. This takes you to the Exalytics Machine home page.
3. From the Exalytics Machine menu displayed on the Exalytics machine target home page, select **Members**, and then select **Topology**.
4. The Configuration Topology page for the System is displayed.

The topology viewer enables you to understand and visualize your IT environment in a graphical format. The inter-relationships between systems are used to answer questions such as:

- On what targets does this target rely?
- What targets will be affected if this system goes down?
- How are the members of my target related?

On this page, you can perform the following tasks:

- [Viewing a Configuration Topic](#)
- [Creating a Custom Topology View](#)

6.3.3.1 Viewing a Configuration Topic

To view a configuration topic, from the **View** list, select any of the following:

- **Uses**
This view helps you determine the targets that the selected target depends on. If a target is having problems, this view can be useful in helping you determine whether its problems have been caused by another target it depends on.
- **Used By**
This view shows you the targets that depend on the selected target. This can be useful, for example, if you are planning on shutting down the target and need to know what other targets will be affected
- **System Members**

This view shows the members of the system.

- **Custom views** that have been defined by end users.

Note: The Uses, Used By, and System Members views are out-of-box topology views. They cannot be modified.

6.3.3.2 Creating a Custom Topology View

To create a custom topology view, do the following:

1. From the **Customize** menu, select **Create Custom View**.
 2. Provide the name and description for the topology and select one of the Initial Contents:
 - **Copy Current View** to create a topology view similar to the one you are viewing.
 - **Create Empty View** to create a topology view that starts with the root node.
 3. Choose one of the following expose options:
 - Expose the custom view for all targets of the current target type. For example, if you are creating a topology view for a database target, the new view will be available for all database targets.
 - Expose the custom view for the current target only.

To share the view, click Share this custom view with other users.

 - Click **OK**.
 4. Reduce the unwanted information in the topology by highlighting the target and selecting **Hide Relationships...** in the **Customize** menu.
- You can also display relationships that are not being displayed by selecting a target. From the **Customize** menu, select **Target**, then select **Show More Relationships to Target Type....**
5. Click **OK**.

6.3.4 Viewing the vCPU Consumption Report for an Exalytics Machine

To navigate to the Consumption Report page, follow these steps:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, click on the expansion button of the Exalytics System that contains the Exalytics target that you want to access. A list of Exalytics targets is displayed.
2. Select the Exalytics target that you want. This takes you to the Exalytics Machine home page.
3. On the Exalytics machine home page, click on the **Exalytics Machine** menu, select **Resource Utilization**, and then select **Consumption Report (All Exalytics)**.

Note: You can view the Consumption Report page only for virtual Exalytics machine targets.

You can perform the following tasks on the vCPU Consumption page:

- [Viewing the vCPU Consumption Report](#)

- [Viewing the Resource Consumption Trend](#)
- [Exporting the vCPU Consumption Report](#)

6.3.4.1 Viewing the vCPU Consumption Report

To view the vCPU Consumption Report for the software components running on an Exalytics Machine, do the following:

1. On the vCPU Consumption page, select one of the time periods in the Time Period drop-down list.

Note: When you select a time period, the Product table and Details table refreshes automatically for the chosen time period.

2. From the Product list, select the product or the products that you want to view the consumption report for. By default all the products are selected.
3. From the Product table, select a Product you want to view the consumption report for.

Note: When you select a product, the Details table refreshes automatically to display the Consumption Report for the product.

4. The Details table displays the consumption details for the product or products that you have selected.

You can add more details as columns to the Details table by clicking **View**, selecting **Columns**, and then selecting the items you want to be displayed. You can select which columns you want to visible and which columns to hide by clicking **Manage Columns**.

By default, the Details table displays the collection time and the number of total vCPUs consumed. Expand the collection time, and host name, to view all the following details:

- Collection time of the point when vCPU was collected
- Host name of the host on which software is running
- The Exalytics machine name
- The number of total vCPUs consumed by the software or the target running at that time
- The target name
- The target type
- The owner of the target

6.3.4.2 Viewing the Resource Consumption Trend

To view the resource consumption trend for an Exalytics machine target, do the following:

1. On the vCPU Consumption page, from the Product drop-down list, select the product you want to view the resource consumption trend for.

Select **All** if you want to view the resource consumption trend for all the targets.

Note: When you make a selection from the Product drop-down list, the Product table and the Details table automatically refreshes to display the details for the product or products that you have selected.

2. From the Time Period drop-down list, select the time period for which you want to view the resource consumption trend for.
3. Click on **View Trends**.
4. The Resource Consumption Trend window opens.
If you selected All products, then the graphical representation of the resource consumption for all the products will be displayed. You can choose to view trends for individual products as well, by clicking on the Product drop-down list, and selecting the product.
5. You can also choose to view the resource consumption trend in a table format by clicking **Table View**.
On the Table view page, you can customize the details you want to view in the table by clicking **View**.
6. Click **OK** to close the Resource Consumption trend window.

6.3.4.3 Exporting the vCPU Consumption Report

To export the vCPU consumption report for an Exalytics machine target, do the following:

1. On the vCPU Consumption page, select the following:
 - **Time Period**
Select a time period for which you want to view the vCPU consumption report for.
 - **View**
Select **Product**.
 - **Product**
Select the products you want to view the vCPU consumption report for.
Select **All** if you want to view the vCPU consumption report for all the products.
2. Click **Export**.
3. On the Opening License dialog box, select the program you want to open the file with. Microsoft Office Excel is the default program. Click **OK**.

6.3.5 Managing License Editions for an Exalytics Machine

To navigate to the Consumption Report page, follow these steps:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, click on the expansion button of the Exalytics System that contains the Exalytics target that you want to access. A list of Exalytics targets is displayed.
2. Select the Exalytics target that you want. This takes you to the Exalytics Machine home page.

3. On the Exalytics machine home page, click on the **Exalytics Machine** menu, select **Resource Utilization**, and then select **Consumption Report (All Exalytics)**.

Note: You can view the Consumption Report page only for virtual Exalytics machine targets.

You can perform the following tasks related to managing license editions on the vCPU Consumption page:

- [Adding a License Edition to an Exalytics Machine Target](#)
- [Creating a New License Edition for an Exalytics Machine](#)

6.3.5.1 Adding a License Edition to an Exalytics Machine Target

To add a license edition to a target, do the following:

1. On the vCPU Consumption page, click **Manage License Edition**.
2. In the Manage License Edition dialog box, expand the Search icon. Enter details of the target you want to add a license edition to. Click **Search**.
3. Select the target from the Target table, and then click **Set License Edition**.

Note: To view more details of the table, click **View**, and then select all the items you want to view as columns in the target table.

4. Click **Done**.

6.3.5.2 Creating a New License Edition for an Exalytics Machine

To create a new license edition for an Exalytics machine target, do the following:

1. On the vCPU Consumption page, from the View list, select **License Edition**.
2. Click **Manage License Edition**.
3. On the Manage License Edition page, select the target you want a create a new license for, and then click **Set License Edition**.
4. In the Set License Edition box, you will see a list of licenses that the target already has.

Click **Add**.

5. Enter the name of the new license you want to add to the target.

Click **OK**.

6.3.6 Viewing the Job Activity of an Exalytics Machine

To view the past, currently running, and scheduled jobs related to a target, follow these steps:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, click on the expansion button of the Exalytics System that contains the Exalytics target that you want to access. A list of Exalytics targets is displayed.
2. Select the Exalytics target that you want. This takes you to the Exalytics Machine home page.

3. From the Exalytics Machine menu displayed on the target home page, select **Job Activity**.
4. The Job Activity page displays target jobs related to target administrative tasks, such as starting the target, stopping the target, target blackouts, and so on.
5. The Job Activity page displays target jobs related to target administrative tasks, such as starting the target, stopping the target, target blackouts, and so on.

6.3.7 Viewing the Collected Configuration of an Exalytics Machine

To view the collected configuration data of an Exalytics machine target, do the following:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, click on the expansion button of the Exalytics System that contains the Exalytics target that you want to access. A list of Exalytics targets is displayed.
2. Select the Exalytics target that you want. This takes you to the Exalytics Machine home page.
3. From the Exalytics machine menu displayed on the Exalytics machine target home page, select **Configuration** and **Last Collected**.
4. The **Latest Configuration** page contains the following configuration data of the target:
 - Configuration items and properties
 - System configuration data as well as all system members and their configuration data
 - System and target relationships (immediate, member of, uses, used by, and so forth)
 - Custom configuration collection data

The browser window consists of left and right panes. The left pane is a tree hierarchy. The right pane consists of tabs that display information in tables. As you navigate in the tree, your selection dictates the contents in the right pane. Depending on the selection, tabs appear containing data such as properties and values, relationships, a hierarchical structure of a system and its members, and file contents in both a parsed and raw text format.

You can take any of several actions as you view a configuration in the browser. These actions are available from the **Actions** menu above the tabs. The tree hierarchy in the left pane also has context menus available.

6.3.8 Administering the Exalytics Machine Target

You can perform the following administration tasks on the Exalytics machine target:

- [Viewing the Target Monitoring Configuration of an Exalytics Machine](#)
- [Administering Target Access Privileges for an Exalytics Machine](#)
- [Administering Target Blackouts for an Exalytics Machine](#)
- [Removing an Exalytics Machine Target](#)

6.3.8.1 Viewing the Target Monitoring Configuration of an Exalytics Machine

To view the monitoring configuration details for a particular target, follow these steps:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, click on the expansion button of the Exalytics System that contains the Exalytics target that you want to access. A list of Exalytics targets is displayed.
2. Select the Exalytics target that you want. This takes you to the Exalytics Machine home page.
3. From the Exalytics Machine menu displayed on the target home page, select **Target Setup**, then select **Monitoring Configuration**.

The Monitoring Configuration page provides information about instance properties of the target, which provide internal details about target monitoring.

6.3.8.2 Administering Target Access Privileges for an Exalytics Machine

To manage the access privileges for a target, follow these steps:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, click on the expansion button of the Exalytics System that contains the Exalytics target that you want to access. A list of Exalytics targets is displayed.
2. Select the Exalytics target that you want. This takes you to the Exalytics Machine home page.
3. From the Exalytics Machine menu displayed on the target home page, select **Target Setup**, then select **Administrator Access**.
4. Click **Add** to grant target access privileges to a role or an administrator.

Use the Access page to set target privileges for roles and administrators. The available privileges are View, Operator, and Full.

View only allows you to view the target in the console, whereas Operator allows you to view targets, and perform all administrative actions except deleting targets. Full allows you to view targets, and perform all administrative actions.

6.3.8.3 Administering Target Blackouts for an Exalytics Machine

To administer the blackouts for a target, follow these steps:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, click on the expansion button of the Exalytics System that contains the Exalytics target that you want to access. A list of Exalytics targets is displayed.
2. Select the Exalytics target that you want. This takes you to the Exalytics Machine home page.
3. From the Exalytics Machine menu displayed on the target home page, select **Monitoring**, then select **Blackouts**.
4. Blackouts suspend data collection on a monitored target. Blackouts are useful when you want to perform scheduled maintenance tasks on monitored targets.

Use the Blackouts page to search for existing target blackouts, edit existing blackouts, define new blackouts, and stop blackouts. You can also create and stop blackouts using the Exalytics Machine menu. To create or stop a blackout, from the Exalytics System menu, select **Control**, then select **Create Blackout** or **End Blackout**, respectively.

6.3.8.4 Removing an Exalytics Machine Target

To remove an Exalytics machine target from an Exalytics System, do the following:

1. From the **Targets** menu, select **Exalytics**. On the Systems page, click on the expansion button of the Exalytics System that contains the Exalytics target that you want to access. A list of Exalytics targets is displayed.
2. Select the Exalytics target that you want. This takes you to the Exalytics Machine home page.
3. On the Exalytics system home page, select an Exalytics machine target.
4. On the Exalytics machine home page, from the Exalytics Machine menu, select **Target Setup**, and then **Remove. Target...**
5. A warning will be displayed, asking if you want to proceed. Click **Remove**.

Oracle Traffic Director

Oracle Traffic Director (Traffic Director) is a software-level load balancer, used to load-balance incoming HTTP connections among origin-servers (host:port pair) that host the actual content.

Traffic Director has the following functions:

- Reverse Proxy—Distributes incoming traffic among servers using load-balancing algorithms. The forwarding mechanism is based on URI and on handling sessions.
- Proxy Caching—Stores frequently accessed html pages.

You can use the Traffic Director to create a configuration that involves defining virtual-servers, listeners, origin-servers, and server-pools. This configuration is then deployed on a set of hosts. The instances of the same configuration form a configuration.

Note: For release 12.1.0.3, only two Traffic Director instances belonging to the same configuration are supported within Enterprise Manager.

The virtual-server is the main component that is configured with the load-balancing properties, for example, the servers to distribute traffic to (origin-servers and pools) to use, the IP-address and port on which to listen for requests (listener), and so on.

Hence, the typical hierarchy of a Traffic Director deployment is that of a Traffic Director configuration consisting of a set of instances deployed on hosts, and each instance has components like virtual-server, listener, proxy-cache along with origin-servers and origin-server-pools.

Following is the target list:

- Traffic Director Configuration
The Traffic Director Configuration target contains configuration and performance metrics pertaining to components like virtual-server, proxy-cache, tcp-proxy, origin-servers and server-pools, at the configuration-level.
- Traffic Director Instance
The Traffic Director Instance has performance metrics for the same components but at the instance level.

Traffic Director Instance target is used to monitor one instance of a Traffic Director Configuration running on a host. This target shows the performance information of the Instance running on that host.

Use the following sections in this chapter to learn more about Traffic Director.

- [Adding a Traffic Director to an Exalogic Target](#)
- [About Traffic Director Configuration](#)

- [About Traffic Director Instance](#)
- [About Traffic Director Refresh Flow](#)

For additional information refer to the *Oracle Traffic Director Administrator's Guide*.

7.1 Adding a Traffic Director to an Exalogic Target

Oracle Traffic Director:adding to Exalogic target to manually add the Traffic Director dashboard to an Exalogic target, perform the following steps:

1. From Enterprise Manager, select **Exalogic** from the **Targets** menu.
2. Click the link of the name of the Exalogic Elastic Cloud to which you want to add the Traffic Director dashboard.
3. On the Exalogic Elastic Cloud page, click the **Software** tab.
4. On the SYSMAN icon located at the top right of the page, select **Personalize Page**.
5. On the Editing Page, click **Add Content**. On the Add Content popup, move to the Traffic Director item and click the associated **Add** link. Click **Close**.
6. On the Editing page, click **Close**.

The Traffic Director region is now visible from the Exalogic Elastic Cloud dashboard (Software tab).

7.2 About Traffic Director Configuration

A configuration is used to create a configuration of Traffic Director instances all having similar functionality. The configuration mainly contains the description of the:

- Servers to direct the incoming traffic (Origin Servers and Server Pools)
- IP address and port to listen for incoming requests (Listener)

The following steps must be performed when creating a configuration:

- A configuration must be created, with one HTTP virtual-server. This virtual server is configured to accept requests for `www.oracle.com`.
- A HTTP Listener must be created whose port is set as 80.
- Two origin-servers are created with the aforementioned host:port pairs.
- Then a server pool is created with these two servers, and the virtual server is associated with this server pool.

Once the entities are created, this configuration is deployed on the hosts.

7.2.1 Using the Traffic Director Configuration Page

The performance metrics, performance information, or performance summary visible on the configuration home page, is at the configuration level that is, the performance is an aggregate of that entity's performance across the hosts on which the configuration instances are running. For example, the virtual server metrics visible on the configuration home page, are an aggregate of metrics at the instance level.

The Traffic Director Configuration page lists the following regions:

- **Summary**—Provides the general information regarding the configuration including how long the configuration has been up, its availability and the version

of configuration. In addition, you can go directly to configure and monitor the Traffic Director by using the link to the Traffic Director Admin Console.

Also, the Monitoring and Diagnostics section lists whether there are any incidents involving the Traffic Director. An incident is an event that represents an issue requiring resolution. Click the incident to determine what needs attention.

- Response and Load
- Performance
- Instances
- Virtual Servers and Origin Servers
- Failover Group—Shows all the failover groups in the configuration. However, failover groups in the Instance Target home page shows only the failover groups of which the instance is a part.
- Exalogic—Shows information about Traffic Director Instances associated to the Exalogic Elastic Cloud target on which the region has been added.

For seeing detailed information about the Traffic Director Instances/Configurations shown in this region, click the links to navigate to the respective target home pages.

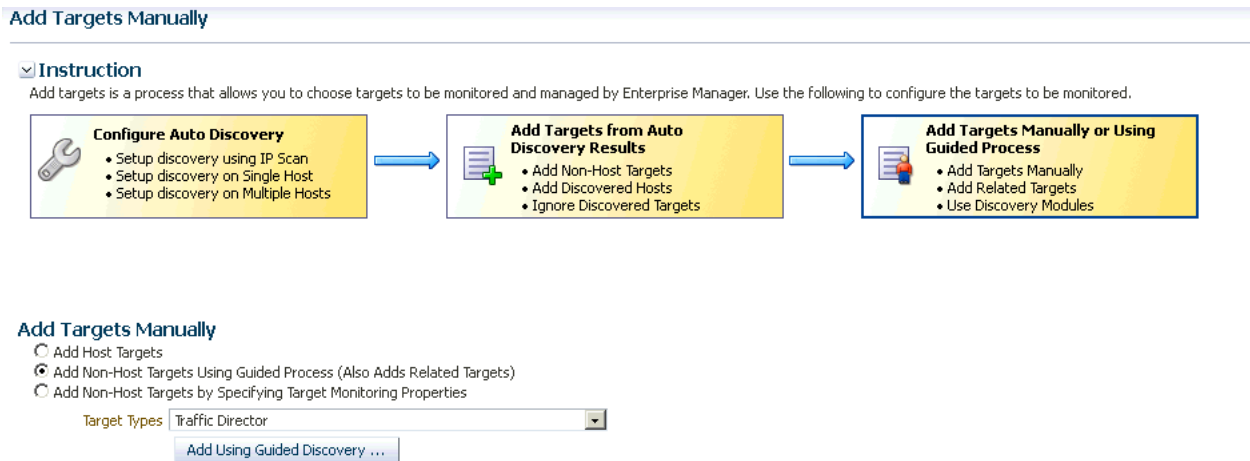
Note: You can view the Traffic Director Configuration details by selecting **Configuration** from the **Target** menu, then selecting **Last Collected**.

7.2.2 Adding Traffic Director Target Configuration

To manually add Traffic Director Configuration to Enterprise Manager, perform the following steps:

1. From Enterprise Manager, select **Add Target** from the **Setup** menu located at the top-right of the page, then select **Add Targets Manually**.
2. On the Add Targets Manually page, choose **Add Non-Host Targets Using Guided Process (Also Adds Related Targets)**. In the Target Types menu, select **Traffic Director**.
3. Click **Add Using Guided Discovery**.

Figure 7–1 shows how the resulting screen should look.

Figure 7–1 Filled in Add Targets Manually Screen

Fill out the pages as described in:

- [Finding Configurations and Instances](#)
- [Discovered Targets](#)
- [Viewing Results](#)

7.2.2.1 Finding Configurations and Instances

Use the Find Configurations and Instances page to supply the Administration Server Properties:

- **Administration Server Host**
The name of the host where the Traffic Director Administration Server is running. Search for a host by clicking the Search icon located at the end of the field.
Note: On selecting the host, the Agent URL field will be automatically populated.
- **Administration Server SSL Port**
SSL Port of the Administration Server.
- **User Name**
Name of the administrator allowed to access the Traffic Director Administration Server.
- **Password**
Password of the administrator allowed to access the Traffic Director Administration Server.
- **SNMP Port**
The port on which the SNMP agent is listening. All the SNMP agents on all Traffic Director instance hosts should be running on the same port.
- **Oracle Home**
Directory where the Traffic Director binaries have been installed.
- **Agent URL**
Enter the URL of the Management Agent running on the Administration Server host.

- Setup Prefix

Unique identifier for this setup. This ID will be prefixed to the names of the targets being discovered.

After you supply the information and click **Continue**, the Confirmation popup appears.

Click **Close** on the Confirmation popup and then click **Continue**.

Possible Error Messages

These are the typical messages you see when the entered information is incorrect.

Failed to find targets. Please check entered details OTD-70104 Unable to communicate with the administration server: Connection refused

Cause: This typically means the host or port entered is incorrect.

Action: Enter the correct host or port.

Failed to find targets. Please check entered details OTD-70104 Unable to communicate with the administration server: Invalid user or password

Cause: This typically means the user name/password entered is incorrect.

Action: Enter the correct user name and password.

Failed to find targets. Please check entered details OracleHome - xxxxxx not valid

Cause: This typically means the Oracle Home entered is incorrect.

Action: Check if the Oracle Home is correct and has no spelling errors.

7.2.2.2 Discovered Targets

The Add Traffic Director: Discovered Targets screen shows the discovered Traffic Director configurations and instances, along with the Management Agents that will be used for monitoring them.

Note: At this point, the targets have not yet been saved.

Click **Add Targets** to save them, or **Back** to go to the previous screen to review the details you entered.

The table on the page lists the Targets and Agent Assignments. The fields in the table are:

- Name—Name of the Traffic Director Configuration/Instance target.
- Type—Type of the target discovered. Discovered target types include Traffic Director Configuration and Traffic Director Instance.
- Agent URL—Management Agent that will be monitoring the target. All targets discovered are monitored by the Management Agent located on the Administration Server.

7.2.2.3 Viewing Results

The goal of the Add Traffic Director: Results page is to provide results. When you arrive at this page, the targets are already saved. Click **OK** to return to the Middleware page.

The Agent URL that you entered on the Find Configurations and Instances screen is used for monitoring all discovered targets. Once you assign the Management Agent, it cannot be changed.

For additional information refer to the *Oracle Enterprise Manager SNMP Support Reference Guide*.

7.3 About Traffic Director Instance

Once discovered, you can view the performance of entities like virtual-server, origin-servers, and instances, on the instance home page.

The performance visible in the instance home page (tables/charts showing metrics), is at the instance level and the metrics are calculated based on the data/traffic to/from that instance

You access this page by clicking an instance in the Instance region of the Traffic Director Configuration page.

The performance metrics, performance information, or performance summary visible on the instance home page is at the instance level.

The Traffic Director Instance page lists the following regions:

- **Summary**—Provides the general information regarding the instance including how long the instance has been up, its availability and the version of the instance. In addition, you can go directly to configure and manage the Traffic Director instances by using the link to the Traffic Director Admin Console.

Also, the Monitoring and Diagnostics section lists whether there are any incidents involving the Traffic Director instances. An incident is an event that represents an issue requiring resolution. Click the incident to determine what needs attention.

The Summary region also includes statistics for the following: Instance, Resource Usage, and Proxy Cache.

- **Response and Load**
- **CPU and Memory Usage**
- **Virtual Servers and Origin Servers**
- **Failover Groups**—Shows only the failover groups to which the instance belongs.
- **Exalogic**—Shows information about Traffic Director Instances associated to the Exalogic Elastic Cloud target on which the region has been added. The Exalogic region is a region you can add using the Exalogic Elastic Cloud target home page.

For seeing detailed information about the Traffic Director Instances/Configurations shown in this region, click the links to navigate to the respective target home pages.

Note: You can view the Traffic Director Configuration details by selecting **Configuration** from the **Target** menu, then selecting **Last Collected**.

7.4 About Traffic Director Refresh Flow

Oracle Traffic Director:adding to Exalogic targetAfter configurations are created, you can add new targets for newly added Configurations or Instances in the Traffic Director Administration Server, or delete the targets corresponding to Configurations or Instances that no longer exist in the Traffic Director Administration Server.

You can also modify target properties to reflect the addition or deletion of children. Only Configuration target properties are modified to reflect the addition or deletion of Instances.

You can access this flow by selecting **Refresh Configuration** from Traffic Director Configuration target menu.

Note: This flow adds and removes targets corresponding to all Traffic Director Administration Servers whose Configurations and Instances have been discovered under different setup prefixes.

7.4.1 Adding New Targets to Newly Added Configurations

To add new targets to newly added configurations, perform the following steps:

1. From the **Targets** menu, select **All Targets**, then select a Traffic Director target.
2. From the Traffic Director Configuration menu, select **Refresh Configuration**.
3. Click **Add Targets**.
4. On the resulting page, review that the newly added Configurations and their Instances are shown with refresh status *New* in the table. Click **Save Updates** to save the changes.

Now all new Configurations and their Instance targets are saved to the Repository and are being monitored.

7.4.2 Adding New Targets for Newly Added Instances of Configurations

To add new targets for newly added instances of configurations, perform the following steps:

1. From the **Targets** menu, select **All Targets**, then select a Traffic Director target.
2. From the Traffic Director Configuration menu, select **Refresh Configuration**.
3. Click **Add Targets**.
4. On the resulting page, review that the newly added Instances are shown as targets with refresh status *New* and their Configurations are shown with refresh status *Modified* in the table. Click **Save Updates** to save the changes.

Now all new Instance targets are saved to the Repository and are being monitored.

7.4.3 Deleting Targets of Configurations That Have Been Removed

To delete targets of configurations that have been removed, perform the following steps:

1. From the **Targets** menu, select **All Targets**, then select a Traffic Director target.
2. From the Traffic Director Configuration menu, select **Refresh Configuration**.
3. Click **Delete Targets**.
4. On the resulting page, review that the removed Configurations and their removed Instances are shown with refresh status *Deleted* in the table. Click **Save Updates** to save the changes.

Now all targets of removed Configurations and Instances have been removed from the Repository.

7.4.4 Deleting Targets of Instances That Have Been Removed

To delete targets of instances that have been removed, perform the following steps:

1. From the **Targets** menu, select **All Targets**, then select a Traffic Director target.

2. From the Traffic Director Configuration menu, select **Refresh Configuration**.
3. Click **Delete Targets**.
4. On the resulting page, review that the removed Instances are shown with refresh status *Deleted* in the table and their Configurations are shown with refresh status *Modified*. Click **Save Updates** to save the changes.

Now all targets corresponding to all removed Instances have been removed from the Repository.

Part III

Monitoring Oracle WebLogic Domain and Oracle GlassFish Server

The chapters in this part describe how you can monitor the Oracle WebLogic Domain and the Oracle GlassFish Server.

The chapters are:

- [Chapter 8, "Monitoring WebLogic Domains"](#)
- [Chapter 9, "Overview of Oracle GlassFish Server Management"](#)

Monitoring WebLogic Domains

When using Enterprise Manager version 12.1 and a Secure Socket Layer (SSL) protocol to discover and monitor WebLogic servers, the Intelligent Agent must be able to *trust* the server before it can establish a secure communication link. The Agent maintains a Java Keystore (JKS) truststore containing certificates of Certification Authorities (CAs) which it can trust when establishing a secure connection. The Agent comes with nine well-known CA certificates.

It is recommended that customers using WebLogic t3s in a production environment use certificates signed by a well-known Certification Authority (CA), such as VeriSign or Thawte, on their WebLogic servers. A few popular Root CA certificates are available out-of-box in the Agent's JKS-based truststore and does not require any action by the customer. However, if self-signed certificates or the default (out-of-box) demo certificate are being used on the Weblogic servers, then the following step is needed to explicitly import the Root CA certificate for these server certificates to the Agent's truststore.

The JKS Agent truststore is located at the following location:

```
$ORACLE_HOME/sysman/config/montrust/AgentTrust.jks
```

Note: ORACLE_HOME is the Management Agent's instance home.

Updating the Agent truststore is required on ALL Enterprise Manager Agents involved in the discovery and monitoring of the WebLogic domain using any secure protocol.

8.1 Updating the Agent Truststore

To update the Agent truststore (AgentTrust.jks), you use EMCTL. If the default demo certificate, or a self-signed certificate is being used on the WebLogic servers for t3s/iiops, then the Root CA certificate for this must be added to AgentTrust.jks in order for the Agent to be able to discover and monitor these WebLogic servers and J2EE applications using t3s. An EMCTL command is provided for this purpose.

```
emctl secure add_trust_cert_to_jks [-password <password> -trust_certs_loc <loc>
-alias <alias>]
```

Where:

- password = password to the AgentTrust.jks (if not specified, you will be prompted for the password at the command line)
- trust_certs_loc = location of the certificate file to import
- alias = alias for the certificate to import

8.1.1 Importing a Demo WebLogic Server Root CA Certificate.

To import the Root CA certificate for a Demo WebLogic server into the Agent's truststore, the EMCTL *secure* command needs to be executed from the host on which the Agent is located.

```
<ORACLE_HOME>/bin/emctl secure add_trust_cert_to_jks -password "welcome"
```

Note: ORACLE_HOME is the Management Agent's instance home.

The following example demonstrates a typical session using the secure command with the *add_trust_cert_to_jks* option.

Example 8-1 Sample Session

```
./emctl secure add_trust_cert_to_jks -password welcome
Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.2.0
Copyright (c) 1996, 2012 Oracle Corporation. All rights reserved.
```

```
Message      : Certificate was added to keystore
ExitStatus: SUCCESS
```

The default out-of-box password for the AgentTrust.jks is "welcome" and it is recommended that this be changed using the JDK keytool utility. If no password is specified along with the EMCTL command, the system will prompt you for the password.

8.1.2 Importing a Custom Root CA Certificate

If the WebLogic servers are secured with another certificate, such as a self-signed certificate, then that Root CA certificate must be imported into the Agent's truststore as follows:

```
<ORACLE_HOME>/bin/emctl secure add_trust_cert_to_jks -password "welcome"
trust_certs_loc <location of certificate> -alias <certificate-alias>
```

Note: ORACLE_HOME is the Management Agent's instance home.

8.2 Changing the Default AgentTrust.jks Password Using Keytool

The following JVM keytool utility command will let you change the default out-of-box password to the AgentTrust.jks.

```
<ORACLE_HOME>/jdk/bin/keytool -storepasswd -keystore AgentTrust.jks -storepass
welcome -new myNewPass
```

Note: ORACLE_HOME is the Management Agent's instance home.

8.3 Collecting JVM Performance Metrics for WebLogic Servers

In order to collect JVM performance metrics from platform MBeans, the Mbeans must be made accessible via the runtime MBeanServer. To do this, from the WebLogic console, set **PlatformMBeanServerEnabled=true**. *Domain->Advanced*

Note: This only applies to WebLogic server installations where Java Required Files (JRF) are not installed.

8.3.1 Setting the PlatformMBeanServerUsed Attribute

If you are using WebLogic server versions 9.2.0.40, 10.0.2.0, 10.3.1 and 10.3.2 and certain patch releases of 9.x, you must explicitly set the *PlatformMBeanServerUsed* attribute to *TRUE* in addition to setting the *PlatformMBeanServerEnabled* (shown in the previous section). You set the *PlatformMBeanServerUsed* attribute using the WebLogic Scripting Tool (WLST), as shown in the next section.

Note: From WebLogic server versions 10.3.3 onwards, the default out-of-box behavior enables platform MBeans to be accessible via runtime MBeanServers. Hence, this section can be skipped.

8.3.2 Activating Platform MBeans on WebLogicServer 9.x to 10.3.2 versions

The following WebLogic Scripting Tool session shown in [Example 8-2](#) demonstrates how to use check and set the PlatformMBeanServerUsed attribute.

User actions are shown in bold.

Example 8-2 Setting PlatformMBeanServerUsed

```
cd common/bin/
```

```
ade:[ adminsw_easvr ] [adminsw@mymachine bin]$ ./wlst.sh

CLASSPATH=/net/mymachine/scratch/shiphomes/wl/wl10/patch
wls1002/profiles/default/sys_manifest_classpath/weblogic
patch.jar:/net/mymachine/scratch/shiphomes/wl/wl10/patch
cie640/profiles/default/sys_manifest_classpath/weblogic
patch.jar:/net/mymachine/scratch/shiphomes/wl/wl10/jrockit_150
15/lib/tools.jar:/net/mymachine/scratch/shiphomes/wl/wl10/wlserver
10.0/server/lib/weblogic_sp.jar:/net/mymachine/scratch/shiphomes/wl/wl10/wlserver
10.0/server/lib/weblogic.jar:/net/mymachine/scratch/shiphomes/wl/wl10/modules/fea
ures/weblogic.server.modules
10.0.2.0.jar:/net/mymachine/scratch/shiphomes/wl/wl10/modules/features/com.bea.ci
.common-plugin.launch
2.1.2.0.jar:/net/mymachine/scratch/shiphomes/wl/wl10/wlserver
10.0/server/lib/webservices.jar:/net/mymachine/scratch/shiphomes/wl/wl10/modules/
rg.apache.ant
1.6.5/lib/ant-all.jar:/net/mymachine/scratch/shiphomes/wl/wl10/modules/net.sf.ant
ontrib_1.0b2.0/lib/ant-contrib.jar:
PATH=/net/mymachine/scratch/shiphomes/wl/wl10/wlserver
10.0/server/bin:/net/mymachine/scratch/shiphomes/wl/wl10/modules/org.apache.ant
1.6.5/bin:/net/mymachine/scratch/shiphomes/wl/wl10/jrockit_150
15/jre/bin:/net/mymachine/scratch/shiphomes/wl/wl10/jrockit_150
15/bin:/home/adminsw/products/valgrind/bin:/ade/adminsw
easvr/oracle/jdk/bin:/ade/adminsw
easvr/oracle/work/middleware/oms/perl/bin:/bin:/usr/local/bin:/usr/local/remote/p
ckages/firefox-1.5.0.3:/ade/adminsw_easvr/oratst/bin:/ade/adminsw
easvr/oracle/buildtools/bin:/ade/adminsw_easvr/oracle/emdev/merge:/ade/adminsw
easvr/oracle/emdev/utl:/ade/adminsw_easvr/oracle/utl:/pdp/pds/utl:/ade/adminsw
easvr/oracle/work/middleware/oms/bin:/ade/adminsw
easvr/oracle/nlsrtl3/bin:/opt/SUNWspro/bin:/usr/ccs/bin:/usr/bin:/usr/sbin:/ade/a
minsw
easvr/oracle/opmn/bin:/usr/X11R6/bin:/home/adminsw/products/valgrind/bin:/home/ad
minsw/products/valgrind/bin:/usr/kerberos/bin:/home/adminsw/products/valgrind/bin:
bin:/usr/bin:/usr/local/bin:/usr/X11R6/bin:/usr/local/ade/bin:/bin:/usr/local/bin
```

Your environment has been set.

```
CLASSPATH=/net/mymachine/scratch/shiphomes/wl/wl10/patch
wls1002/profiles/default/sys_manifest_classpath/weblogic
patch.jar:/net/mymachine/scratch/shiphomes/wl/wl10/patch
cie640/profiles/default/sys_manifest_classpath/weblogic
patch.jar:/net/mymachine/scratch/shiphomes/wl/wl10/jrocket_150
15/lib/tools.jar:/net/mymachine/scratch/shiphomes/wl/wl10/wlserver
10.0/server/lib/weblogic_sp.jar:/net/mymachine/scratch/shiphomes/wl/wl10/wlserver
10.0/server/lib/weblogic.jar:/net/mymachine/scratch/shiphomes/wl/wl10/modules/fea
tures/weblogic.server.modules
10.0.2.0.jar:/net/mymachine/scratch/shiphomes/wl/wl10/modules/features/com.bea.ci
.common-plugin.launch
2.1.2.0.jar:/net/mymachine/scratch/shiphomes/wl/wl10/wlserver
10.0/server/lib/webservices.jar:/net/mymachine/scratch/shiphomes/wl/wl10/modules/
rg.apache.ant
1.6.5/lib/ant-all.jar:/net/mymachine/scratch/shiphomes/wl/wl10/modules/net.sf.ant
ontrib
1.0b2.0/lib/ant-contrib.jar:/net/mymachine/scratch/shiphomes/wl/wl10/wlserver
10.0/common/eval/pointbase/lib/pbembedded51.jar:/net/mymachine/scratch/shiphomes/
l/wl10/wlserver
10.0/common/eval/pointbase/lib/pbtools51.jar:/net/mymachine/scratch/shiphomes/wl/
l10/wlserver_10.0/common/eval/pointbase/lib/pbclient51.jar
```

Initializing WebLogic Scripting Tool (WLST) ...

Welcome to WebLogic Server Administration Scripting Shell

Type help() for help on available commands

wls:/offline>

wls:/offline> connect('weblogic','welcome1','mymachine:7501')

Connecting to t3://mymachine:7501 with userid weblogic ...

Successfully connected to Admin Server 'AdminServer' that belongs to domain 'base domain'.

Warning: An insecure protocol was used to connect to the server. To ensure on-the-wire security, the SSL port or Admin port should be used instead.

wls:/base_domain/serverConfig> edit()

Location changed to edit tree. This is a writable tree with DomainMBean as the root. To make changes you will need to start an edit session via startEdit().

For more help, use help(edit)

wls:/base_domain/edit> startEdit()

Starting an edit session ...

Started edit session, please be sure to save and activate your changes once you are done.

wls:/base_domain/edit !> cd('JMX')

wls:/base_domain/edit/JMX !> ls()

drw- base_domain

wls:/base_domain/edit/JMX !> cd ('base_domain')

wls:/base_domain/edit/JMX/base_domain !> ls()

-rw-	CompatibilityMBeanServerEnabled	true
-rw-	DomainMBeanServerEnabled	true

```

-rw- EditMBeanServerEnabled true
-rw- InvocationTimeoutSeconds 0
-rw- ManagementEJBEnabled true
-rw- Name base_domain
-rw- Notes null
-rw- PlatformMBeanServerEnabled true
-rw- PlatformMBeanServerUsed false **
-rw- RuntimeMBeanServerEnabled true
-r-- Type JMX

-r-x freezeCurrentValue Void : String(attributeName)
-r-x isSet Boolean : String(propertyName)
)
-r-x restoreDefaultValue Void : String(attributeName)
-r-x unSet Void : String(propertyName)

```

```
wls:/base_domain/edit/JMX/base_domain !> set('PlatformMBeanServerUsed','true')
```

```
wls:/base_domain/edit/JMX/base_domain !> ls()
```

```

-rw- CompatibilityMBeanServerEnabled true
-rw- DomainMBeanServerEnabled true
-rw- EditMBeanServerEnabled true
-rw- InvocationTimeoutSeconds 0
-rw- ManagementEJBEnabled true
-rw- Name base_domain
-rw- Notes null
-rw- PlatformMBeanServerEnabled true
-rw- PlatformMBeanServerUsed true **
-rw- RuntimeMBeanServerEnabled true
-r-- Type JMX
-r-x freezeCurrentValue Void : String(attributeName)
-r-x isSet Boolean : String(propertyName)
)
-r-x restoreDefaultValue Void : String(attributeName)
-r-x unSet Void : String(propertyName)

```

```
wls:/base_domain/edit/JMX/base_domain !> activate()
```

Activating all your changes, this may take a while ...

The edit lock associated with this edit session is released once the activation is completed.

The following non-dynamic attribute(s) have been changed on MBeans

that require server re-start: **

MBean Changed : com.bea:Name=base_domain,Type=JMX

Attributes changed : PlatformMBeanServerUsed

Activation completed

```
wls:/base_domain/edit/JMX/base_domain> ade:[ adminsw_easvr ] [adminsw@mymachine bin]$
```

```
ade:[ adminsw_easvr ] [adminsw@mymachine bin]$
```

****NOTE:** *PlatformMBeanServerUsed* attribute is present in WebLogic releases 10.3.1.0 and 10.3.2.0 and also for certain patch releases of prior versions. If above *PlatformMBeanServerUsed* attribute is NOT present, or if it is present and already set to true, then running the commands are not necessary.

Overview of Oracle GlassFish Server Management

Oracle GlassFish Server (GlassFish Server) provides the server environment needed for the development and deployment of Java Platform, Enterprise Edition (Java EE platform) applications and web technologies based on Java technology.

Enterprise Manager Cloud Control provides the ability for you to monitor Oracle GlassFish Domains, Servers, and Clusters.

This chapter provides the following:

- [Before Getting Started](#)
- [Understanding the Oracle GlassFish Domain](#)
- [Understanding the Oracle GlassFish Server Home Page](#)
- [Understanding the Oracle GlassFish Cluster Home Page](#)
- [Viewing Collected Configuration Data for Oracle GlassFish Members](#)

9.1 Before Getting Started

Before you start using Oracle GlassFish, ensure you are familiar with the Oracle GlassFish concepts. Refer to the Oracle GlassFish Server documentation available at <http://www.oracle.com/technetwork/middleware/glassfish/documentation/index.html>.

Also, ensure you have been granted the roles and privileges needed to use Oracle GlassFish.

9.1.1 GlassFish Roles and Privileges

Before you start using Oracle GlassFish, ensure you have access to the following privileges:

Task	Privilege
Start and stop GlassFish targets	FMW_PROCESS_CONTROL_TARGET; CREATE_JOB
View start, stop, and restart menus	FMW_PROCESS_CONTROL_TARGET
Use start, stop, and restart menus	CREATE_JOB
Discovery and refresh	CREATE_PROPAGATING_GROUP

9.1.2 Adding Domain Certificate to Activate Start and Stop Operations

To start or stop Oracle GlassFish targets, you must manually export the domain certificate from the Oracle Glassfish domain for which you want to perform start and stop operations and add it to the AgentTrust.jks file of the Agent which is used to monitor that domain.

Perform the following steps to add the certificate. The first step extracts the Oracle GlassFish domain certificate to a temporary certificate file (server.cer). The second step adds this certificate to the Agent trust store.

1. This step should be run on the host where the Oracle GlassFish domain is installed, then copy this server.cer file to the Agent host.

```
keytool -keystore <GlassFish Domain>/config/cacerts.jks -export -alias  
<alias> -file <file> -noprompt -storepass <password>
```

Where:

<GlassFish Domain> - Domain where file is located
<alias> - Alias of Oracle Glassfish domain certificate in cacerts.jks (Default is slas)
<password> - Password for cacerts.jks (Default is changeit)
<file> - Temporary certificate file to which domain certificate is added, for example, server.cer

2. Import the Certificate to the AgentTrust.jks file using the keytool *or* emctl commands. This should be run on the Agent.

```
keytool -import -alias <alias> -file <file>  
-keystore $T_WORK/agentStateDir/sysman/config/montrust/AgentTrust.jks  
-noprompt -storepass <password>
```

Where:

<alias> - alias for the certificate to import
<file> - temporary file to which certificate was exported in step 1
<password> - password to the AgentTrust.jks file

OR

```
emctl secure add_trust_cert_to_jks [-password <password>  
-trust_certs_loc <loc> -alias <alias>]
```

Where:

<password> - password to the AgentTrust.jks file
<loc> - location of the certificate(certificate.cer) file to import
<alias> - alias for the certificate to import

9.2 Understanding the Oracle GlassFish Domain

As a GlassFish Server Administrator, you can use the GlassFish Domain home page to understand the overall health of the GlassFish Domain. This home page provides summary information about the domain, as well as specifics about clusters and servers.

Note: You will see errors if you do not have the server or cluster configured to enable the monitoring attributes and/or component monitoring levels required by Enterprise Manager. To collect and view metrics in Enterprise Manager, monitoring must be enabled on the servers. Go to the GlassFish Server Administration Console and click the **Configure Monitoring** link for each server. Ensure that both the Monitoring Service and Monitoring MBeans are enabled.

In addition, ensure the Monitoring Level is set to HIGH for the following components:

- Connector Connection Pool
 - Connector Service
 - EJB Container
 - JDBC Connection Pool
 - JMS Service
 - JVM
 - Transaction service
 - Web Container
-

Note: At the top of the page, the host that appears above the timestamp is the host of the agent monitoring the target. This may or may not be the Administration Server host.

Summary

The Summary section provides general information about the domain, link to the GlassFish Server Administration Console, and monitoring and diagnostics statistics.

- General

Shows the administration server for this domain, the host on which the administration server is deployed, the listener port and listener port for the secure sockets layer (SSL), and when the domain was last refreshed.

Note that the listener ports are associated with the administration server.

Click the name link to drill down to the Administration Server's Home page.

- Tools

Provides a link to the GlassFish Server Administration Console where you can configure and manage the GlassFish Server.

- Monitoring and Diagnostics

- Provides the number of incidents that occurred in the last 7 days. An incident is an event or a set of closely correlated events that represent an observed issue requiring resolution through immediate action or root-cause problem resolution.
- Provides the number and severity of incidents on any descendant target. Descendant targets are the members (children, grandchildren, and so on) within the domain. For example, a domain's descendant targets are any clusters or servers in that domain, as well as any other targets under them.

Note that the descendant target incident count does *not* include incidents on itself (only children).

- Alerts you to changes made to the configuration. Click the link next to Configuration Changes to learn what configurations changed and when.

Clusters

If the domain contains clusters, this region lists the clusters. If the domain does not contain clusters, the table appears with the message "No Clusters Found".

For each cluster, the name, status, servers, and incidents fields appear.

Servers

The domain home page lists all servers that are contained within the domain, including servers that are contained within any clusters in the domain.

For each server, the name, status, host, associated cluster, configuration data, as well as performance data appear. For additional information regarding a server, click the server name.

Domain Target Menu

The domain home page provides a menu of additional functions you can perform from this page. The menu is located under the name of the domain.

Menu options of particular interest are:

- **Control** - Allows you to start all servers, shut down all servers, create blackouts of all servers and end blackouts of all servers. Other than blackout operations, these operations do not impact the GlassFish Administration Server.
- **GlassFish Server Administration Console** - Opens a new window to the GlassFish Server Administration Console. This console allows for greater control and administration of the GlassFish Domains and its members such as servers and clusters.
- **Refresh GlassFish Domain** - Refreshes the domain.

This operation rediscovers the domain. When refresh is performed, the Management Agent connects to the Administration Server by way of the REpresentational State Transfer (REST) interface to obtain any changes in domain membership.

Changes in membership could include the addition or removal of new GlassFish Servers or the creation or removal of GlassFish Clusters. When a refresh is performed, the Administration Server must be up and running.

9.2.1 How to Access an Oracle GlassFish Domain

To access the Oracle GlassFish Domain home page, perform the following steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, select **Oracle GlassFish Full Hierarchy** in the Area field and click **Search**. In the table, click the GlassFish Domain in which you are interested. The GlassFish Domain page appears.

On this page, you can:

- Easily access the GlassFish Server Administration Console to make configuration changes to the domain and perform administration operations.

- Start up and shut down all servers and clusters except for the Domain Administration Server.
- Blackout the domain when the domain must go down for maintenance. By blacking out the domain, alert notifications are not sent. Click **Create Blackout...** at the top of the page to access this functionality. You can also use the **GlassFish Domain** menu: from the **Control** menu, select **Create Blackout...**
- Customize the layout and data displayed in target home pages. The changes you make are persisted for all targets' home pages of the particular target type you are customizing and are persisted for the user you are currently logged in as; this enables you to create customized consoles for monitoring various target types.
- Refresh the domain after creating a new server or deploying a new application. This enables you to see the most up-to-date information on this page.
- Perform configuration management operations such as viewing and analyzing collected configuration data, comparing configurations between servers, tracking configuration changes over time, and searching configuration data across the enterprise.

9.2.2 How to Add an Oracle GlassFish Domain To Be Monitored

There are two ways to access the functionality to add an Oracle GlassFish Domain to Enterprise Manager Cloud Control.

If you need to discover several domains, consider using the Enterprise Manager Command Line Interface (EMCLI). This allows you to discover multiple domains in one operation. See the *Oracle Enterprise Manager Command Line Interface* manual for additional information.

Method 1

1. From the **Targets** menu, select **Middleware**.
2. Click the **Add** button, then select **Oracle GlassFish Domain**.
3. On the **Add GlassFish Domain: Find Targets** page, provide the required information denoted by an asterisk. Click **Continue**.
4. Reassign the agents.

Method 2

1. From the Enterprise Manager Setup menu located at the top right of the screen, select **Add Target**, then select **Add Targets Manually**.
2. On the Add Targets Manually page, select **Add Targets Using Guided Process (Also Adds Related Targets)**.
3. In the Target Type menu, select **Oracle GlassFish Domain**.
4. Then click the **Add Using Guided Process** button.
5. On the **Add GlassFish Domain: Find Targets** page, provide the required information denoted by an asterisk. Click **Continue**.
6. Reassign the agents.

9.2.3 Adding an Oracle GlassFish Domain: Finding and Assigning Targets

After adding the domain to the Cloud Control console, consider performing the following tasks:

- Configure notification methods and your notification schedule to receive email and page notifications when potential problems occur, for example, GlassFish Server goes down unexpectedly, key performance metric threshold is reached, and so on.
- Create a monitoring template for GlassFish related components to set thresholds for key performance metrics and collection frequency of configuration and monitoring data. You can then apply this template to several GlassFish components to ensure that all components are monitored in a similar fashion.

Before you can add (discover) an Oracle Glassfish Domain as a managed target, you must provide the information Cloud Control requires to find the targets associated with the domain. Provide the required information as follows:

Field	Description
Administration Server Host	<p>Name of the Domain Administration Server Host. Select a host from the list provided. This is the host name for where the Administration Server is installed and running.</p> <p>Ensure that the Administration Server is up and accessible prior to initiating discovery. While the Administration Server must be up in order to perform discovery, it need not remain up for Cloud Control to monitor the availability and performance of the domain and its members. However, any time you want to refresh the domain (that is, rediscover the domain to begin monitoring newly created components or to remove recently removed components), you must ensure that the Administration Server is up.</p>
Port	<p>Administration Server Host port number. The default port number is 4848. This is the port on which the Administration Server is listening. The default port in the Discovery UI is 7001.</p> <p>If the port is configured for the HTTPS protocol only, then you must also specify 'HTTPS' as the protocol in the Advanced section of this page. For discovering secure domain, use the https protocol.</p> <p>If the agent needs to be secured, use emctl. If the default demo certificate, or a self-signed certificate is being used on the GlassFish Servers for t3s/iiops, then the Root CA certificate for this must be added to the AgentTrust.jks in order for the Agent to be able to discover and monitor these GlassFish Servers and Java EE applications using t3s. An emctl command is provided for this purpose.</p> <pre>emctl secure add_trust_cert_to_jks [-password <password> -trust_certs_loc <loc> -alias <alias>]</pre> <p>where</p> <ul style="list-style-type: none"> ■ alias - alias for the cert to import ■ password - password to the AgentTrust.jks (if not specified will be prompted for) ■ trust_certs_loc - location of the cert file to import
Username	User name of the Administration Server Domain. User needs CREATE_PROPAGATING_GROUP privilege for discovery/refresh.
Password	Password of the Administration Server Domain
Unique Domain Identifier	<p>Used as a prefix to ensure domain names are unique in environments with the same domain name. For example, if the Unique Domain Identifier is Domain01 and the domain name is production_domain then the domain name in Enterprise Manager would be Domain01_production_domain.</p> <p>The default Unique Domain Identifier is Domain01. This identifier is incremented each time an additional domain is discovered. For example, if you discover a domain with the name stage_domain, then the domain name in Enterprise Manager is Domain02_stage_domain.</p> <p>You can change the default identifier. However, the identifier must only contain alphanumeric characters and the special character '_'. No other special characters can be used in the identifier.</p>

Field	Description
Agent	<p>Agent used to discover the target; that is, the agent used to connect to the GlassFish Administration Server.</p> <p>The specified Management Agent can be local to the Administration Server (that is, installed on the same host machine as the Administration Server) or can be remote to the Administration Server (that is, installed on a different host machine as the Administration Server). The Management Agent uses the REST Interface to connect to the Administration Server (thus, requiring the Administration Server to be up) in order to discover the details of the domain and its members.</p> <p>This agent does not have to be the same agent that is used for monitoring. Typically, when you discover a target, you select the agent local to the GlassFish Administrator Server. When you choose the agent for monitoring, you choose the local agent to each managed server. If there is no local agent to each managed server, then by default, the agent used for discovery is used.</p> <p>You can choose the agent from the drop-down list. The agent must be a version 12.1 Management Agent.</p>

In the Advanced Section, provide additional information for discovering and assigning targets.

Field	Description
Protocol	Use either http or https to make the connection to the Administration Server. The default value for the Protocol field is http.
Service URL	Connection string used to make a JMX connection to the GlassFish domain.
External Parameters	<p>Optional field for passing parameters to the Java process used for connecting to the Administration Server. These parameters must begin with -D.</p> <p>You can name value pairs which are separated by a space (), such as -Dkerberoskey=a -Dparam2=b -Dparam3=c. For example: -Dname=foo -Dparam1=abcd</p>
Discovery Debug File Name	The agent side discovery messages for this session will be logged into this file. This file will be generated in the discovery agent's log directory <agent home>/sysman/log. If this file already exists, it will be updated.

Once you have provided all the information, click **Continue**.

A processing page appears indicating that Enterprise Manager is attempting to find targets. When processing is complete, the processing page displays the number of targets found. The Assign Agents page displays.

Note: If the process of discovering or refreshing a farm fails because the system has reached the maximum number of HTTP socket connections, you must edit the emd.properties file and increase the MaxIncomingConnections parameter to 500. After you make the change then bounce the agent.

Assigning Agents

The agents will be assigned automatically. If a local agent is found where a server is running, that agent is assigned. Otherwise the agent that you entered in the Find Targets page is assigned.

The Saving Target to Agent processing window appears indicating how many total targets have been added and successfully saved. It will also indicate the number of targets were unsuccessfully added.

If there are no warnings due to a failure to assign agents, the Show/Hide section is hidden by default. If there are any warnings, the Show/Hide section will automatically expand to display the Results table.

If the targets of the domain change in the future, you can use the Refresh Domain option to add or remove targets.

After Discovery

After adding the domain to the Cloud Control console, consider performing the following tasks:

- Configure notification methods and your notification schedule to receive email and page notifications when potential problems occur, for example, GlassFish Server goes down unexpectedly, key performance metric threshold is reached, and so on.
- Create a monitoring template for GlassFish related components to set thresholds for key performance metrics and collection frequency of configuration and monitoring data. You can then apply this template to several GlassFish components to ensure that all components are monitored in a similar fashion.

Notes

- If the process of discovering or refreshing a domain fails because the system has reached the maximum number of HTTP socket connections, you must edit the `emd.properties` file and increase the `MaxInComingConnections` parameter to 500. After you make the change, bounce the agent.
- When a user discovers a secure GlassFish Domain with a custom certificate, they also must import the certificate to the agent trust store of the following agents:
 - In the discovery agent
 - In all the agents used for monitoring the targets which belong to that GlassFish Domain
- If a local Management Agent is found on the same host machine as a GlassFish Server within the domain, then that agent is automatically assigned to monitor the GlassFish Server.
- If there are any validation errors in the discovery process, errors will display in a message box. Similarly, if agents to which targets are being pushed are down at the time of discovery, an error message displays stating that the operation for that target will fail.
- You can easily change the monitoring configuration of a target. For more information, see *Modifying the Monitoring Configuration of a Target*.
- You can refresh an existing domain by using the Refresh GlassFish Domain option on the GlassFish Domain menu.

9.2.4 Adding an Oracle GlassFish Domain: Displaying Results

After the Assign Agents page displays the number of targets discovered, you can use Cloud Control to display the targets that have been discovered in the domain. The Results page may indicate that all targets have been successfully added or that the process was partially successful. The Results page lists the number of successful targets along with the number of unsuccessful targets. You have the option of retrying targets with errors by using the Retry Targets with Errors feature.

If there were no warnings due to a failure to assign agents, the Show/Hide section is hidden by default. If there are any warnings, the Show/Hide section will automatically expand to display the Results table.

The Results table lists the Target, Type, Host, Configured Agent, and Status. The targets are displayed in an expanded hierarchy. If you had selected the option to manually assign an agent to a target, you can enter the agent in the Configured Agent field. Otherwise if a local agent is found where a server is running, that agent is assigned or the agent that you entered in the Find Targets page is assigned.

You can change the agent only for targets to which you can assign an agent. You cannot modify the monitoring agent of targets that use a parent agent such as J2EE applications, soa-infra, and so on.

The Status column displays the results of each target and displays the following values:

- Success (Green check mark) -- Target is pushed to the agent successfully
- Failure (Red X) -- Failed to push target to agent
- Already Saved -- Target already saved

If agents to which targets are being pushed are down at the time of discovery, an error message displays stating that the operation for that target will fail.

When all processing completes, you can choose to fix any issues caused by targets not being added because of errors related to agent assignments by pressing the Retry Targets with Errors button. The Retry Targets with Errors window appears. A scrollable table displays showing each target that was not added due to errors. Often targets are not added because agents were not started at the time of discovery. You can either change the agent in the Agent field or simply click Retry to initiate the process again.

Alternatively, you can leave this page and return later to address any issues. If you press Cancel, all targets that have been discovered will be monitored by Cloud Control and the Middleware page is then displayed.

Note: If the process of discovering or refreshing a domain fails because the system has reached the maximum number of HTTP socket connections, you must edit the `emd.properties` file and increase the `MaxInComingConnections` parameter to 500. After you make the change then bounce the agent.

9.2.5 Refreshing an Oracle GlassFish Domain

Enterprise Manager is not informed when changes are made to the domain configuration and membership. For example, if someone using the GlassFish Server Administration Console adds a new managed server, removes a server, adds a cluster, removes a cluster, and so on, Enterprise Manager does not know about the changes until the target is refreshed or rediscovered. It is only then that Enterprise Manager knows of the newly added targets and you can add them to Enterprise Manager for centralized management and monitoring.

By using Enterprise Manager, you are informed when target members are removed. You can remove these targets from Enterprise Manager. It is by design that Enterprise Manager does not automatically remove the targets.

When you refresh the domain, historical data is not lost unless you choose to delete a target that has been removed outside of Enterprise Manager. When new members are found, you must choose the agent to monitor them.

To manually refresh the GlassFish Domain to ensure you are monitoring the complete domain, follow these steps.

1. On the GlassFish Domain page, locate the GlassFish Domain menu.

2. From the menu, select **Refresh GlassFish Domain**.

9.3 Understanding the Oracle GlassFish Server Home Page

As a GlassFish Server Administrator, you can review the GlassFish Server home page to understand the overall health of the GlassFish Server. The GlassFish Server home page provides summary information about the server, and statistics regarding the most requested servlets and response and load.

The home page provides the following:

- Current status of GlassFish Server
- Key GlassFish Server performance metrics including: Java Message Service (JMS), Enterprise JavaBeans (EJB), Java Transaction API (JTA) usage
- Indication of any recent configuration changes
- Incidents

Summary

This section provides general information about the server, link to the GlassFish Server Administration Console, monitoring and diagnostics statistics, and so on.

- General

Shows the status and resource usage of the server. Click any link to get additional information on a metric.
- Servlets

Shows an overview of all the servlets present on the server. Some of the individual servlet metrics can be found in the Most Requested region.
- Connection Pool and JTA Usage

Shows metrics relating to connection pool and Java Transactions API (JTA) usage.
- Tools

Provides a link to the GlassFish Server Administration Console where you can configure and manage the GlassFish Server.
- Monitoring and Diagnostics
 - Provides the number of incidents that occurred in the last 7 days. An incident is an event or a set of closely correlated events that represent an observed issue requiring resolution through immediate action or root-cause problem resolution.
 - Alerts you to changes made to the configuration. Click the link next to Configuration Changes to learn what configurations changed and when.
- EJBs

Shows statistics pertaining to Enterprise JavaBeans (EJBs) accesses and the cache.
- JMS

Shows Java Message Service (JMS) statistics for this server.

Most Requested Servlets

This region lists the servlets that have had requests during the past 24 hours.

Response and Load

Shows the request processing time and requests per minute of all servlets over time.

GlassFish Server Menu

The server home page provides a menu of additional functions you can perform from this page. The menu is located under the name of the server.

Menu options of particular interest are:

- **Control** - Allows you to start, restart, and shut down the server, as well as create and end blackouts for the server. You cannot start and stop the GlassFish Administration Server.
- **GlassFish Server Administration Console** - Opens a new window to the GlassFish Server Administration Console. This console allows for greater control and administration of the GlassFish Servers.
- **Configuration** - Enables you to compare server configurations, track history, search, as well as study the last configuration data.
- **Monitoring** - You can customize the Performance Summary metrics, as well as the metric and collection settings where you can set thresholds and the frequency of collections.

9.3.1 How to Access an Oracle GlassFish Server

To access the Oracle GlassFish Server home page, perform the following steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, select **Oracle GlassFish Full Hierarchy** in the Area field and click **Search**. In the table, expand the Oracle GlassFish Domain node that contains the GlassFish Server of interest. Click the GlassFish Server and the GlassFish Server home page appears.

This page provides a summary of the server's health, as well as key performance indicators for the server, like Enterprise JavaBeans (EJBs) and Java Message Service (JMS).

In addition, you can:

- Easily access the GlassFish Server Administration Console (in the Summary region) to make configuration changes to the server and perform administration operations. Click the link, then log in with your user name and password.
- Create and end blackouts.
Blackouts allow Enterprise Manager users to suspend management data collection activity on one or more managed targets. For example, administrators use blackouts to prevent data collection during scheduled maintenance or emergency operations. By blacking out the server, alert notifications are not sent.
- Start up, shut down, and restart servers.
These operations call predefined jobs that perform the operations. Credentials are needed for each of these operations.
- Customize the layout in target home pages. For example, you can customize the page to show particular GlassFish Server metrics.

For more information, see *Personalizing a Cloud Control Page*.

- View historical performance metrics from the GlassFish Server menu by selecting **Monitoring**, then selecting **Performance Summary**. You can customize the Performance Summary page.
- Customize the layout and data displayed in target home pages. The changes you make are persisted for all targets' home pages of the particular target type you are customizing and are persisted for the user you are currently logged in as; this enables you to create customized consoles for monitoring various target types.
- Perform configuration management operations such as viewing and analyzing collected configuration data, comparing configurations between servers, tracking configuration changes over time, and searching configuration data across the enterprise.

9.4 Understanding the Oracle GlassFish Cluster Home Page

As a GlassFish Server Administrator, you can review the GlassFish Cluster home page to understand the overall health of the GlassFish Cluster by way of its status as well as its key performance and configuration data.

The home page provides the following:

- Current availability of GlassFish Cluster
- Key GlassFish Cluster resource usage metrics

These metrics are based on the servers within the cluster not on the cluster itself.
- Indication of any recent configuration changes
- Incidents

Summary

The Summary section provides general information about the cluster, link to the GlassFish Server Administration Console, and monitoring and diagnostics statistics.

- Availability

Shows the availability of the cluster. The cluster is considered up as long as one server in the cluster is up.
- Tools

Provides a link to the GlassFish Server Administration Console where you can configure and manage the GlassFish Cluster.
- Monitoring and Diagnostics
 - Provides the number of incidents that occurred in the last 7 days. An incident is an event or a set of closely correlated events that represent an observed issue requiring resolution through immediate action or root-cause problem resolution.
 - Provides the number and severity of incidents on any descendant target. Descendant targets are the members (children, grandchildren, and so on) within the domain. For example, a domain's descendant targets are any clusters or servers in that domain, as well as any other targets under them.

Note that the descendant target incident count does not include incidents on itself (only children).
 - Alerts you to changes made to the configuration. Click the link next to Configuration Changes to learn what configurations changed and when.

Servers

The domain lists all servers that are contained within the cluster.

For each server, the name, status, host, configuration data, as well as performance data appear. For additional information regarding a server, click the server name.

Resource Usage

Shows the CPU usage over time and the Heap Usage over time graphs for every server the cluster contains.

GlassFish Cluster Target Menu

The cluster home page provides a menu of additional functions you can perform from this page. The menu is located under the name of the cluster.

Menu options of particular interest are:

- **Control** - Allows you to start all servers, shut down all servers, create blackouts of all servers and end blackouts of all servers within the cluster.

Note: There is *no* rolling process control for the cluster, that is, all the servers are brought up (or down) together in parallel.
- **GlassFish Server Administration Console** - Opens a new window to the GlassFish Server Administration Console. This console allows for greater control and administration of the GlassFish Servers.

9.4.1 How to Access an Oracle GlassFish Cluster

To access the Oracle GlassFish Cluster home page, perform the following steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, select **Oracle GlassFish Full Hierarchy** in the Area field and click **Search**. In the table, click the GlassFish Cluster in which you are interested. The GlassFish Cluster home page appears.

On this page, you can:

- Easily access the GlassFish Server Administration Console (in the Summary region) to make configuration changes to the cluster and perform administration operations. Click the link, then log in with your user name and password.
- You can customize the layout and data displayed in target home pages to suit your specific needs. The changes you make are for a target type and user.

For more information, see *Personalizing a Cloud Control Page*.

9.5 Viewing Collected Configuration Data for Oracle GlassFish Members

Configuration data is available for Oracle GlassFish Domains, clusters, and servers. To view the configuration data:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, select **Oracle GlassFish** in the Area field and click **Search**. In the table, click the target of interest, that is, a GlassFish Domain, cluster, or server.
3. On the resulting page, click the menu located under the page title, for example, GlassFish Server, GlassFish Domain, or GlassFish Cluster.
4. Select **Configuration**, then select **Last Collected**.

In addition to viewing the Last Collected data, from the Configuration menu you can:

- Save current collected configuration data with which to compare future collections
- Create configuration searches against GlassFish related targets so you can search for specific configurations across a data center. Oracle provides the following predefined configuration searches: Oracle GlassFish Server: Ports and Oracle GlassFish Server: Datasources.
- Compare current configurations between two different GlassFish Servers, for example, production server versus QA server. You can also compare domains.
- Use predefined GlassFish Server configuration comparison template while comparing servers.

For example, use a template to ignore data in the comparison results (for instance configuration data that you EXPECT to be different) or to notify you of detected configuration differences that you deem critical.

Oracle also provides a template for comparing GlassFish domains.

9.6 Creating an Oracle GlassFish Server Configuration Comparison Template

Enterprise Manager provides a monitoring template for GlassFish which you can customize. For example, you can set metric thresholds and collection frequency (for both performance and configuration data), and then apply these settings in the template to several servers of the domain or across several domains.

This ensures a consistent way of monitoring across targets and eliminates the need to go to each server to specify thresholds and collection settings.

To create a configuration comparison template, perform these steps:

1. From Enterprise menu, select **Configuration**, then **Comparison Templates**.
2. Click the **Help** button located at the top right of the page for information on how to create a new comparison template.

Part IV

Managing Oracle SOA

The chapters in this part describe how you can discover and monitor Oracle BPEL Process Manager, Oracle Service Bus, and Oracle SOA Suite.

The chapters are:

- [Chapter 10, "Overview of Oracle SOA Management"](#)
- [Chapter 11, "Discovering and Monitoring Oracle BPEL Process Manager"](#)
- [Chapter 12, "Discovering and Monitoring Oracle Service Bus"](#)
- [Chapter 13, "Discovering and Monitoring the SOA Suite"](#)

Overview of Oracle SOA Management

The Oracle SOA Management Pack Enterprise Edition delivers comprehensive management capabilities for a Service-Oriented Architecture-based (SOA) environment. By combining SOA runtime governance, business-IT alignment, and SOA infrastructure management with Oracle's rich and comprehensive system management solution, Enterprise Manager Cloud Control significantly reduces the cost and complexity of managing SOA-based environments.

Table 10–1 Highlights of Oracle SOA Management Pack Enterprise Edition

Feature	Benefit
Centralized management console	Provides administrators managing SOA environments with a consolidated browser-based view of the entire enterprise, thereby enabling them to monitor and manage all of their components from a central location.
Discovery and service modeling	Provides discovery of the following: <ul style="list-style-type: none"> ■ Oracle SOA Infrastructure deployed to the WebLogic Server. ■ Oracle SOA Composite applications deployed to the SOA Infrastructure. ■ Oracle BPEL processes deployed to the Oracle BPEL Process Manager (BPEL Process Manager) server and the dependent partner links. ■ Oracle Service Bus-based business and proxy services. ■ Service modeling offers out-of-the-box automated system modeling capabilities for the SOA infrastructure.
Runtime governance	Defines SOAP tests to measure and record availability and performance of partner links (or any Web service) and business/proxy services for historical trending, troubleshooting, and root cause analysis purposes. Also provides an error hospital of process instances with drill-downs into instance details.
Infrastructure management	Monitors the availability and performance of the SOA infrastructure components. Both current and historic availability of targets (such as BPEL Process Manager or Oracle Service Bus) are recorded for troubleshooting and root cause analysis.
Configuration management	Collects configuration information for the BPEL Process Manager server/domains/processes and Oracle Service Bus. The parameters can be refreshed, saved, or compared with another target. Different versions of the same target can also be compared.
Deployment automation	Automates the deployment of the following: <ul style="list-style-type: none"> ■ SOA Artifacts Provisioning: This includes provisioning of SOA Composites, Oracle WebLogic Server Policies, Assertion Templates, and JPS Policy and Credential Stores. ■ BPEL processes on BPEL Process Managers ■ Oracle Service Bus resources from a source OSB domain to a target OSB domain. For detailed information on the provisioning procedures, see <i>Enterprise Manager Administrator's Guide for Software and Server Provisioning and Patching</i> .
Adapter metrics	Provides throughput and error metrics for different adapters in graphical format.
Business-IT alignment	Enables you to consolidate their IT and business management tools into a unified system. BAM-EM integration unites business KPIs and system metrics in one system for correlation and trending.

Table 10–1 (Cont.) Highlights of Oracle SOA Management Pack Enterprise Edition

Feature	Benefit
Service level management	Enables you to monitor services from the end-user's perspective using service tests or synthetic transactions, model relationships between services and underlying IT components, and report on achieved service levels.
Application Dependency and Performance	<p>Enables you to manage your SOA solutions by leveraging a model-driven top-down approach within your development, quality assurance (QA), staging, and production environments. Business application owners and operational staff can automatically discover your BPEL workflows and correlate them with the underlying Web services; Enterprise Service Buses (ESBs); and back-end Java 2 Platform, Enterprise Edition (Java EE) resources through detailed modeling and drill-down directly into the performance metrics at the component level.</p> <p>For more information, see Chapter 27, "Introduction to Application Dependency and Performance"</p>
Historical analysis and reporting	Store the collected metric and configuration data in a central repository, thereby enabling administrators to analyze metrics through various historical views and facilitate strategic trend analysis and reporting.
Instance Tracing	Allows you to trace the message flow across SOA Composites and SOA Infrastructure instances monitored by Enterprise Manager Cloud Control.
Business Transaction Management	Provides monitoring of business transactions as they flow across tiers and continuous discovery of components, transaction flow, service dependencies and relationships.
Dehydration Store	Shows the performance of the database that is used by the SOA Infrastructure. Using this data, the SOA administrator can identify problems that are causing the performance bottleneck.

Discovering and Monitoring Oracle BPEL Process Manager

This chapter describes how you can discover and monitor Oracle BPEL Process Manager (BPEL Process Manager) using Enterprise Manager Cloud Control.

In particular, this document covers the following:

- [Supported Versions](#)
- [Understanding the Discovery Mechanism](#)
- [Understanding the Discovery Process](#)
- [Setting Up Oracle Software Library](#)
- [Discovering BPEL Process Manager](#)
- [Configuring BPEL Process Manager](#)
- [Troubleshooting BPEL Process Managers](#)

11.1 Supported Versions

The following are the versions of BPEL Process Manager that are supported for monitoring in Enterprise Manager Cloud Control.

Table 11–1 Supported Versions

Supported BPEL Process Manager Version	Application Server Deployed To	Supported in Enterprise Manager
Oracle BPEL Process Manager 10.1.2	Oracle Application Server 10g Release 1 (10.1.2)	Enterprise Manager 10g Release 4 (10.2.0.4) or higher Enterprise Manager 11g Enterprise Manager 12c
Oracle BPEL Process Manager 10.1.3.1 and 10.1.3.3 (Part of Oracle SOA Suite 10.1.3.1 and 10.1.3.3)	Oracle Application Server 10g Release 1 (10.1.3.1) and (10.1.3.3)	Enterprise Manager 10g Release 3 (10.2.0.3) or higher Enterprise Manager 11g Enterprise Manager 12c
Oracle BPEL Process Manager 10.1.3.1 and 10.1.3.3 (Part of Oracle SOA Suite 10.1.3.1 and 10.1.3.3)	Oracle WebLogic Managed Server 9.2	Enterprise Manager 10g Release 5 (10.2.0.5) or higher Enterprise Manager 10 g Release 4 (10.2.0.4) with one-off patches applied. For details, see Section 11.3, "Understanding the Discovery Process" . Enterprise Manager 11g Enterprise Manager 12c

Table 11–1 (Cont.) Supported Versions

Supported BPEL Process Manager Version	Application Server Deployed To	Supported in Enterprise Manager
Oracle BPEL Process Manager 10.1.3.1 and 10.1.3.3 (Part of Oracle SOA Suite 10.1.3.1 and 10.1.3.3)	IBM WebSphere Application Server 6.1	Enterprise Manager 10g Release 5 (10.2.0.5) or higher Enterprise Manager 10g Release 4 (10.2.0.4) with one-off patches applied. For details, see Section 11.3, "Understanding the Discovery Process" . Enterprise Manager 11g Enterprise Manager 12c
Oracle BPEL Process Manager 10.1.3.4 (Part of Oracle SOA Suite 10.1.3.4)	Oracle Application Server 10g Release 1 (10.1.3.1) and (10.1.3.3)	Enterprise Manager 10g Release 5 (10.2.0.5) or higher Enterprise Manager 11g Release 1 (11.1.0.1) Enterprise Manager 12c
Oracle BPEL Process Manager 10.1.3.4 (Part of Oracle SOA Suite 10.1.3.4)	Oracle WebLogic Managed Server 9.2	Enterprise Manager 10g Release 5 (10.2.0.5) or higher Enterprise Manager 11g Enterprise Manager 12c
Oracle BPEL Process Manager 10.1.3.4 (Part of Oracle SOA Suite 10.1.3.4)	IBM WebSphere Application Server 6.1	Enterprise Manager 10g Release 5 (10.2.0.5) or higher Enterprise Manager 11g Enterprise Manager 12c

11.2 Understanding the Discovery Mechanism

The following describes the mechanism followed for discovering BPEL Process Managers in Enterprise Manager Cloud Control.

Table 11–2 Mechanism for Discovering BPEL Process Managers

BPEL Process Manager Version	Application Server Deployed To	Discovery Mechanism	Process
Oracle BPEL Process Manager 10.1.2	Oracle Application Server 10g Release 1 (10.1.2)	Manual/Automatic Discovery	<ul style="list-style-type: none"> If the Management Agent is installed before Oracle Application Server and BPEL Process Manager are installed, then you must manually discover that Oracle Application Server and BPEL Process Manager in Enterprise Manager Cloud Control. If the Management Agent is installed after Oracle Application Server and BPEL Process Manager are installed, then Enterprise Manager Cloud Control automatically discovers that Oracle Application Server and BPEL Process Manager <p>The Management Agent can be installed along with Enterprise Manager Cloud Control or separately as a standalone product.</p> <p>For discovery procedures, see Section 11.5.1, "Deployed to Oracle Application Server".</p>
Oracle BPEL Process Manager 10.1.3.1, 10.1.3.3, 10.1.3.4 (Part of Oracle SOA Suite 10.1.3.1, 10.1.3.3, 10.1.3.4)	Oracle Application Server 10g Release 1 (10.1.3.1) and (10.1.3.3)	Manual/Automatic Discovery	<ul style="list-style-type: none"> If the Management Agent is installed before Oracle Application Server and BPEL Process Manager are installed, then you must manually discover that Oracle Application Server and BPEL Process Manager in Enterprise Manager Cloud Control. If the Management Agent is installed after Oracle Application Server and BPEL Process Manager are installed, then Enterprise Manager Cloud Control automatically discovers that Oracle Application Server and BPEL Process Manager <p>The Management Agent can be installed along with Enterprise Manager Cloud Control or separately as a standalone product.</p> <p>For discovery procedures, see Section 11.5.1, "Deployed to Oracle Application Server".</p>

Table 11–2 (Cont.) Mechanism for Discovering BPEL Process Managers

BPEL Process Manager Version	Application Server Deployed To	Discovery Mechanism	Process
Oracle BPEL Process Manager 10.1.3.1, 10.1.3.3, 10.1.3.4 (Part of Oracle SOA Suite 10.1.3.1, 10.1.3.3, 10.1.3.4)	Oracle WebLogic Managed Server 9.2	Manual Discovery	First, manually discover Oracle WebLogic Managed Server. For procedures, see Section 11.5.2.1, "Discovering Oracle WebLogic Managed Server" . Then, manually discover BPEL Process Manager. For procedures, see Section 11.5.2.2, "Deployed to Oracle WebLogic Managed Server" .
Oracle BPEL Process Manager 10.1.3.1, 10.1.3.3, 10.1.3.4 (Part of Oracle SOA Suite 10.1.3.1, 10.1.3.3, 10.1.3.4)	IBM WebSphere Application Server 6.1	Manual Discovery	First, manually discover IBM WebSphere Application Server. For procedures, see Section 11.5.3.1, "Discovering IBM WebSphere Application Server" . Then, manually discover BPEL Process Manager. For procedures, see Section 11.5.3.2, "Deployed to IBM WebSphere Application Server" .
Oracle BPEL Process Manager 10.1.3.5	Oracle WebLogic Managed Server 10.x	Manual Discovery	First, manually discover Oracle WebLogic Managed Server. For procedures, see Section 11.5.2.1, "Discovering Oracle WebLogic Managed Server" . Then, manually discover BPEL Process Manager. For procedures, see Section 11.5.2.2, "Deployed to Oracle WebLogic Managed Server" .

11.3 Understanding the Discovery Process

The following describes the overall process involved in discovering and monitoring BPEL Process Manager in Enterprise Manager Cloud Control. Follow the instructions outlined against each step in this process to successfully discover and monitor your BPEL Process Manager.

Table 11–3 Discovery Process

Step	Requirement	Description
1	BPEL Process Manager	Install the BPEL Process Manager software in one of the following ways: <ul style="list-style-type: none"> For Oracle middleware, download and install the BPEL Process Manager using Oracle BPEL Process Manager 10.1.2, Oracle SOA Suite 10.1.3.1, 10.1.3.3, or 10.1.3.4 from the following URL: http://www.oracle.com/technology/software/tech/soa/index.html For non-Oracle middleware, download and install the BPEL Process Manager from the following URL: http://www.oracle.com/technology/software/products/ias/bpel/index.html
2	Enterprise Manager Cloud Control	To monitor BPEL Process Manager 10.x, install Enterprise Manager Cloud Control 12c. For information about installing the base release of Enterprise Manager Cloud Control, see the <i>Enterprise Manager Cloud Control Installation and Basic Configuration Guide</i> available at: http://www.oracle.com/technology/documentation/oem.html Oracle recommends that you install the Enterprise Manager Cloud Control components on a host that is different from the host where the BPEL Process Manager is installed. For example, if the BPEL Process Manager is installed on host1.xyz.com, then install and configure Oracle Management Service (OMS) and the Management Repository on host2.xyz.com.

Table 11–3 (Cont.) Discovery Process

Step	Requirement	Description
3	Oracle Management Agent (Management Agent)	<p>Install Oracle Management Agent 12c or higher on every host where BPEL Process Manager is installed.</p> <p>If Oracle Application Server/BPEL Process Manager and Enterprise Manager Cloud Control are all on the same host, then you do not have to install a separate Management Agent. The Management Agent that comes with Enterprise Manager Cloud Control is sufficient. However, if they are different hosts, then you must install a separate Management Agent on every host where BPEL Process Manager is installed.</p> <p>You can install the Management Agent in one of the following ways:</p> <ul style="list-style-type: none"> ■ Invoke the installer provided with Enterprise Manager 10g Cloud Control Release 3 (10.2.0.3) or higher, and select the installation type Additional Management Agent. ■ Use the Agent Deploy application within the Cloud Control console. ■ Use the full agent kit that is available at: http://www.oracle.com/technology/software/products/oem/htdocs/agentsoft.html <p>For information about installing the Management Agent, see the <i>Enterprise Manager Cloud Control Installation and Basic Configuration Guide</i> available at: http://www.oracle.com/technology/documentation/oem.html</p>
4	Discovery in Enterprise Manager Cloud Control	<p>BPEL Process Managers deployed to Oracle Application Servers are automatically discovered in Enterprise Manager Cloud Control.</p> <p>BPEL Process Managers deployed to Oracle WebLogic Managed Servers and IBM WebSphere Application Servers must be manually discovered in Enterprise Manager Cloud Control. For procedures to discover them, see Section 11.5, "Discovering BPEL Process Manager".</p>

11.4 Setting Up Oracle Software Library

If you are using Enterprise Manager 12c to discover and monitor the BPEL Process Manager deployed to Oracle WebLogic Managed Server 9.2 and IBM WebSphere Application Server 6.1, you must set up Oracle Software Library (Software Library) as described below:

To set up the Software Library:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. From the **Actions** menu, select **Administration**.
3. In the Software Library: Administration page, select the Storage Type and click **Add** from the Actions menu.
4. In the Add Software Library Location window, specify a valid directory path where you want to store the raw data for the components, and click **OK**.

Note: For more information about setting up the Software Library, see the *Enterprise Manager Advanced Installation and Configuration Guide* available at the following URL:

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

11.5 Discovering BPEL Process Manager

This section describes the procedures for discovering BPEL Process Managers. In particular, this section covers the following:

- [Deployed to Oracle Application Server](#)
- [Deployed to Oracle WebLogic Managed Server](#)

- [Deployed to IBM WebSphere Application Server](#)

11.5.1 Deployed to Oracle Application Server

A BPEL Process Manager deployed to Oracle Application Server is manually or automatically discovered in Enterprise Manager Cloud Control depending on when the Management Agent is installed.

- If the Management Agent is installed before Oracle Application Server and BPEL Process Manager are installed, then you must manually discover that Oracle Application Server and BPEL Process Manager in Enterprise Manager Cloud Control.
- If the Management Agent is installed after Oracle Application Server and BPEL Process Manager are installed, then Enterprise Manager Cloud Control automatically discovers that Oracle Application Server and BPEL Process Manager.

Note: You must install a Management Agent on every host where BPEL Process Manager is installed. If Oracle Application Server/BPEL Process Manager and Enterprise Manager Cloud Control are all on the same host, then you need not install a separate Management Agent. The Management Agent that comes with Enterprise Manager Cloud Control is sufficient. However, if they are different hosts, then you must install a separate Management Agent on every host where BPEL Process Manager is installed. The Management Agent can be installed along with Enterprise Manager Cloud Control or separately as a standalone product.

Also note that if you have added a new BPEL Process Manager to an Oracle Application Server that is already discovered and monitored in Enterprise Manager Cloud Control, then you must manually *rediscover* that Oracle Application Server.

To manually discover or *rediscover* Oracle Application Server:

1. From the **Targets** menu, select **Middleware**.
The Middleware page that lists all the middleware targets being monitored is displayed. In Enterprise Manager 10g Cloud Control Release 4 (10.2.0.4) or lower, the Middleware tab is Application Servers.
2. (Only for *Rediscovering*) In the Middleware page, select the Oracle Application Server that you want to rediscover and click **Remove**.
3. In the Middleware page, from the **Add** list, select **Oracle Application Server** and click **Go**. The Add Oracle Application Server Target: Specify Host page is displayed.
4. Enter the name of the host where that Oracle Application Server is running, and click **Continue**.

Enterprise Manager Cloud Control rediscovers that Oracle Application Server along with its core components and the newly added BPEL Process Manager.

11.5.2 Deployed to Oracle WebLogic Managed Server

To discover the BPEL Process Manager deployed to Oracle WebLogic Managed Server, you have to first discover and add Oracle WebLogic Managed Server to Enterprise Manager Cloud Control.

This section describes the procedures for the following:

- [Discovering Oracle WebLogic Managed Server](#)
- [Deployed to Oracle WebLogic Managed Server](#)

11.5.2.1 Discovering Oracle WebLogic Managed Server

To discover and add Oracle WebLogic Managed Server to Enterprise Manager Cloud Control:

1. From the **Targets** menu, select **Middleware**.

Enterprise Manager Cloud Control displays the Middleware page that lists all the middleware targets being monitored. In Enterprise Manager 10g Cloud Control Release 4 (10.2.0.4) or lower, the Middleware tab is Application Servers

2. In the Middleware page, from the **Add** list, select **Oracle Fusion Middleware / WebLogic Server Domain**, and click **Go**.

Enterprise Manager Cloud Control displays the Add Oracle Fusion Middleware / WebLogic Server Domain wizard that captures the details of the Oracle WebLogic Server Domain to be discovered and monitored.

3. In the wizard, specify the required details and click **Next** on each page to reach the end of the wizard.

For information about the details to be provided for each page of the wizard, click **Help** on each page.

4. In the last page of the wizard, click **Finish** to complete the discovery process and add the target to Cloud Control for monitoring purposes.

Enterprise Manager Cloud Control displays the Middleware page with a confirmation message that confirms that the Oracle WebLogic Manager Server has been successfully added to Cloud Control. In the Middleware page that shows all the middleware targets being monitored, you can see the Oracle WebLogic Managed Server you just added.

For additional information about Fusion Middleware discovery, see [Chapter 2, "Discovering Middleware Targets"](#).

11.5.2.2 Deployed to Oracle WebLogic Managed Server

To discover and add the BPEL Process Manager deployed to Oracle WebLogic Managed Server:

1. From the **Setup** menu, select **Add Targets**, then select **Add Targets Manually**. Now select the **Add Non-Host Targets by Specifying Target Monitoring Properties** option.
2. Select the target type from the drop-down list and click the torch icon to select a Monitoring Agent. Click **Add Manually**. The Select Application Server page of the Add BPEL Process Manager wizard is displayed.
 - a. In the Select Application Server page, provide the following details and click **Next**.

Table 11–4 *Select Application Server Page - Element Description*

UI Page Element	Description
Application Server Type	Select the type of application server where the BPEL Process Manager to be discovered is running.

Table 11–4 (Cont.) Select Application Server Page - Element Description

UI Page Element	Description
Application Server Name	Specify the name of the application server where the BPEL Process Manager to be discovered is running. If you are not sure about the name, click the each icon (torch icon) to view a list of application servers and select the appropriate one. The application server name must be suffixed with <i>oracleBPELServer</i> .

- b. In the Target Details page, provide the following details and click **Next**.

Table 11–5 Target Details Page - Element Description

UI Page Element	Description
Oracle Home	Specify the full path to the Oracle Application Server home directory where the BPEL Process Manager is installed. For example, /opt/app/orabpel/product/10.1.3.1/OracleAS.
Application Server Home	Specify the full path to the directory where Oracle WebLogic Managed Server (to which the BPEL target is deployed) is running. For example, /opt/wls9.2/weblogic9.2.

Note:

- Enterprise Manager Cloud Control checks the configuration settings of the associated application server and prefills the values for fields such as BPEL Process Manager Name, Display Name, Context Provider URL, and Oracle BPEL PM Console URL.
- At this point, if you encounter a discovery failure error, then follow the workaround steps given in [Table 11–9](#) to resolve the issue.

- c. In the Host Credentials page, specify the operating system credentials of the host where BPEL Process Manager is running. By default, the fields are prefilled with preferred credentials that are stored in the Management Repository for the selected host. You can either use these prefilled values or edit them to override the preferred credentials with your new credentials.
- d. In the Review page, review the details and click **Finish** to complete the discovery process and add the target to Enterprise Manager Cloud Control.
- Enterprise Manager Cloud Control displays the Agent home page with a confirmation message that confirms that the BPEL Process Manager has been successfully added for monitoring.
3. To verify whether the BPEL Process Manager has been added, click **Targets** and then **Middleware**.

Enterprise Manager Cloud Control displays the Middleware page that shows all the middleware targets being monitored, including the Oracle WebLogic Managed Server and the BPEL Process Manager you just added.

11.5.3 Deployed to IBM WebSphere Application Server

To discover the BPEL Process Manager deployed to IBM WebSphere Application Server, you have to first discover and add IBM WebSphere Application Server to Enterprise Manager Cloud Control.

This section describes the procedures for the following:

- [Discovering IBM WebSphere Application Server](#)
- [Deployed to IBM WebSphere Application Server](#)

11.5.3.1 Discovering IBM WebSphere Application Server

To discover and add IBM WebSphere Application Server to Enterprise Manager Cloud Control:

1. From the **Targets** menu, select **Middleware**.
Enterprise Manager Cloud Control displays the Middleware page that lists all the middleware targets being monitored.
2. In the Middleware page, select **IBM WebSphere Application Server** from the **Add** drop-down list and click **Go**.
Enterprise Manager Cloud Control displays the Add IBM WebSphere Application Server wizard that captures the details of the IBM WebSphere Application Server to be discovered and monitored.
3. In the Add IBM WebSphere Application Server wizard, specify the required details and click **Next** on each page to reach the end of the wizard.
For information about the details to be provided for each page of the wizard, click **Help** on each page.
4. In the last page of the Add IBM WebSphere Application Server wizard, click **Finish** to complete the discovery process and add the target to Enterprise Manager Cloud Control for monitoring purposes.
Enterprise Manager Cloud Control displays the Middleware page with a confirmation message that confirms that the IBM WebSphere Application Server has been successfully added for monitoring. In the Middleware page that shows all the application server being monitored, you can see the IBM WebSphere Application Server you just added.

11.5.3.2 Deployed to IBM WebSphere Application Server

To discover and add the BPEL Process Manager deployed to IBM WebSphere Application Server:

1. From the **Setup** menu, select **Add Targets**, then select **Add Targets Manually**. Now select the **Add Non-Host Targets** by Specifying Target Monitoring Properties option.
2. Select the target type from the drop-down list and click the torch icon to select a Monitoring Agent. Click **Add Manually**. The Select Application Server page of the Add BPEL Process Manager wizard is displayed.
 - a. In the Select Application Server page, provide the following details and click **Next**.

Table 11–6 Select Application Server Page - Element Description

UI Page Element	Description
Application Server Type	Select IBM WebSphere Application Server from the list.
Application Server Name	Specify the name of IBM WebSphere Application Server where the BPEL Process Manager to be discovered is running. If you are not sure about the name, click the search icon (torch icon) to view a list of application servers and select the appropriate one. The application server name must be suffixed with <i>oracleBPELServer</i> .

- b. In the Target Details page, provide the following details and click **Next**.

Table 11–7 Target Details Page - Element Description

UI Page Element	Description
Oracle Home	Specify the full path to the Oracle Application Server home directory where the BPEL Process Manager is installed. For example, /opt/app/orabpel/product/10.1.3.1/OracleAS.
Application Server Home	Specify the full path to the directory where IBM WebSphere Application Server (to which the BPEL target is deployed) is running.
BPEL Application Installation Location	Specify the full path to the installation directory where the BPEL application is installed. For example, if the BPEL application is installed in <\$WEBSPPHERE_HOME>/profiles/AppSrv01/installedApps/sta00114Cel101/CollaxaWebApplications-sta00114Node01.ear, then specify the path as <\$WEBSPPHERE_HOME>/profiles/AppSrv01/installedApps. Here, replace \$WEBSPPHERE_HOME with the full path of the application home location.

Note: Enterprise Manager Cloud Control checks the configuration settings of the associated application server and prefills the values for fields such as BPEL Process Manager Name, Display Name, Context Provider URL, and Oracle BPEL PM Console URL.

- c. In the Host Credentials page, specify the operating system credentials of the host where BPEL Process Manager is running. By default, the fields are prefilled with preferred credentials that are stored in the Management Repository for the selected host. You can either use these prefilled values or edit them to override the preferred credentials with your new credentials.
- d. In the Review page, review the details and click **Finish** to complete the discovery process and add the target to Enterprise Manager Cloud Control. Enterprise Manager Cloud Control displays the Agent home page with a confirmation message that confirms that the BPEL Process Manager has been successfully added for monitoring.

Note: At this point, if you encounter a discovery failure error, then follow the workaround steps given in [Table 11–10](#) and resolve the issue.

3. To verify whether the BPEL Process Manager has been added, select **Middleware** from the **Targets** menu. /

Enterprise Manager Cloud Control displays the **Middleware** page that shows all the middleware targets being monitored, including the IBM WebSphere Application Server and the BPEL Process Manager you just added.

11.6 Configuring BPEL Process Manager

After discovering BPEL Process Manager, you must perform the following configuration steps:

- [Specifying Details for Monitoring BPEL Process Manager](#)
- [Registering BPEL Process Manager Credentials and Host Credentials](#)

11.6.1 Specifying Details for Monitoring BPEL Process Manager

Follow these steps to specify the details required for monitoring BPEL Process Managers. If the values are prefilled, then validate them.

1. In the BPEL Process Manager Home page, select **Target Setup**, then select **Monitoring Configuration** from the **BPEL Process Manager** menu.
2. In the Monitoring Configuration page, specify the following details. If these values are prefilled, then validate them.

- **BPEL Admin Username** - Specify the BPEL administrator user ID.
- **BPEL Password** - Specify the BPEL admin password.

When adding the credentials, validate the following two criteria:

- BPEL Admin User ID and password should have BPEL Admin role
- The same credentials should succeed for the BPEL console login operation

- **Initial Context Factory** - Specify the initial context factor. You can copy the following string value:

```
com.evermind.server.rmi.RMIInitialContextFactory
```

- **Context Provider URL** - Specify the context provider URL. You can copy the following string value:

```
opmn:ormi://<host>:<opmn_port>:home/orabpel
```

Note: Replace the <host>,<opmn port> with the correct host address and opmn port number details for the Oracle Application Server where the BPEL Process Manager is deployed.

To retrieve SOA Applications Server OPMN PORT details, follow these steps:

1. Open the configuration file `$SOA_ORACLE_HOME/opmn/conf/opmn.xml`. `$SOA_ORACLE_HOME` corresponds to SOA Application server home location.
2. Identify the value of the request port attribute in the configuration file.

-
- **BPEL Repository Host Name** - Specify the BPEL Dehydration store (database) host name.

- **BPEL Repository Port** - Specify the BPEL Dehydration store (database) port.
 - **BPEL Repository SID** - Specify the BPEL Dehydration store (database) SID.
 - **BPEL Repository User Name** - Specify the BPEL Dehydration store (database) user name. By default, the user name is `orabpel`.
 - **BPEL Repository Password** - Specify the BPEL Dehydration store (database) password. By default, the password is `welcome1`.
 - **Recoverable Instances Time Threshold (Days)** - Specify the number of days for which the retryable instances must be shown.
 - **Process Aggregate State** - Specify 5, a numeric value that signifies the **constant** state of the BPEL target.
3. Click **OK** to save the settings.

11.6.2 Registering BPEL Process Manager Credentials and Host Credentials

Follow these steps to register the credentials of the BPEL Process Manager, and the credentials of the host where BPEL Process Manager is running.

1. From the **Setup** menu, select **Security**, then select **Preferred Credentials**.
2. Select the Host target type and click **Manage Preferred Credentials**.
3. Select **Normal Host Credentials** in the Credential Set column in the Default Preferred Credentials section and click **Set**.
4. In the Select Named Credential window, enter the user name and password and click **Save** to return to the Preferred Credentials page.
5. Select Oracle BPEL Process Manager target type and click **Manage Preferred Credentials**.
6. Select **Monitoring Administrator Credentials** in the Credential Set column in the Default Preferred Credentials section and click **Set**.
7. In the Select Named Credential window, enter the user name and password and click **Save**.

11.7 Troubleshooting BPEL Process Managers

This section describes the errors you might encounter while discovering BPEL Process Managers, and the workaround steps you can follow to resolve each of them.

This section covers the following:

- [Discovery Errors on Target Details Page](#)
- [Discovery Errors on Review Page](#)
- [Discovery Errors on Review Page](#)

11.7.1 Discovery Errors on Target Details Page

The following error occurs in the Target Details page of the Add BPEL Process Manager wizard where you provide details about the BPEL Process Manager installed on Oracle WebLogic Managed Server.

Table 11–8 Errors on Target Details Page While Adding BPEL Process Manager Deployed to Oracle WebLogic Managed Server

Error Message	Workaround Steps
Oracle BPEL Process Manager not found in the selected Application Server. Select another Application Server.	<p>This error may occur if BPEL is not deployed on the selected Application Server or if the configuration data has not been collected.</p> <p>To resolve this issue:</p> <ol style="list-style-type: none"> 1. Select another Application Server. 2. Navigate to the Application Server Home page and select Configuration, then select Last Collected from the Application Server target menu.

11.7.2 Discovery Errors on Review Page

The following errors occur in the Review page of the Add BPEL Process Manager wizard when you are about to add a BPEL Process Manager installed on Oracle WebLogic Managed Server, to Enterprise Manager Cloud Control for monitoring purposes.

Table 11–9 Errors on Review Page While Adding BPEL Process Manager Deployed to Oracle WebLogic Managed Server

Error Message	Workaround Steps
Discovery Failure - Oracle BPEL Process Manager target discovery failed due to incorrect host credentials.	<ol style="list-style-type: none"> 1. In the last page of the Add BPEL Process Manager wizard where you see this error message, click Previous to reach the Host Credentials page. 2. In the Host Credentials page, specify the correct host credentials or set the preferred credentials for the specific host. Ensure that these are Agent user credentials.
Oracle BPEL Process Manager Discovery Failed - Unable to connect to Oracle BPEL Process Manager. The possible reasons can be incorrect path or insufficient permission to access Oracle BPEL Process Manager home location or inaccessible Oracle BPEL Process Manager home location. Review the specified value.	<ol style="list-style-type: none"> 1. In the last page of the Add BPEL Process Manager wizard, click Previous repeatedly to reach the Target Details page. 2. In the Target Details page, verify the Oracle home location of the BPEL Process Manager. 3. In the Target Details page, verify the installation location of the associated application server.
Oracle BPEL Process Manager Discovery Failed - Unable to connect to Oracle BPEL Process Manager. The possible reasons can be incorrect path or insufficient permission to access Oracle BPEL Process Manager home location or inaccessible Oracle BPEL Process Manager home location. Review the specified value.	Ensure that the BPEL directories have read permission for the Agent user.

11.7.3 Discovery Errors on Review Page

The following errors occur in the Review page of the Add BPEL Process Manager wizard when you are about to add a BPEL Process Manager installed on IBM WebSphere Application Server, to Enterprise Manager Cloud Control for monitoring purposes.

Table 11–10 Error on Review Page While Adding BPEL Process Manager Deployed to IBM WebSphere Application Server

Error Message	Workaround Steps
Discovery Failure - Oracle BPEL Process Manager target discovery failed due to incorrect host credentials.	<ol style="list-style-type: none"> 1. In the last page of the Add BPEL Process Manager wizard where you see this error message, click Previous to reach the Host Credentials page. 2. In the Host Credentials page, specify the correct host credentials or set the preferred credentials for the specific host. Ensure that these are Agent user credentials.
Oracle BPEL Process Manager Discovery Failed - Unable to connect to Oracle BPEL Process Manager. The possible reasons can be incorrect path or insufficient permission to access Oracle BPEL Process Manager home location or inaccessible Oracle BPEL Process Manager home location. Review the specified value.	<ol style="list-style-type: none"> 1. In the last page of the Add BPEL Process Manager wizard where you see this error message, click Previous repeatedly to reach the Target Details page. 2. In the Target Details page, verify the BPEL application installation location. For example, the BPEL application may be installed at the following location: <code><\$WEBSPPHERE_HOME>/profiles/AppSrv01/installedApps/sta00114Cell01/CollaxaWebApplications-sta00114Node01.ear</code> In this case, the path you specify must look like this: <code><\$WEBSPPHERE_HOME>/profiles/AppSrv01/installedApps</code> Note: Replace \$WEBSPPHERE_HOME with the absolute application home location. 3. In the Target Details page, verify the application server home location of the associated application server. 4. In the Target Details page, verify the Oracle home location of the BPEL Process Manager.

11.7.4 Display Errors on Processes Page

Sometimes, after the discovery of a BPEL Process Manager, the BPEL process may occasionally not be listed in the BPEL Process Manager Processes page in Enterprise Manager Cloud Control.

There are two causes for this and two ways to ensure they display on the Processes page. The sections below discuss these causes and workaround steps to fix them.

11.7.4.1 No Credentials Specified for Monitoring BPEL Process Manager

You may not have specified the credentials required for monitoring BPEL Process Managers. To address this, do the following:

1. In the BPEL Process Manager Home page, select **Target Setup**, then select **Monitoring Configuration** from the **BPEL Process Manager** menu.
2. In the Monitoring Configuration page, check the following fields:

- **BPEL Admin Username** - Provide the BPEL administrator user ID.
- **BPEL Password** - Provide the BPEL admin password.

When adding the credentials validate the following two criteria:

- BPEL Admin User ID and password should have BPEL Admin role
- The same credentials should succeed for the BPEL console login operation

- **Initial Context Factory** - In case this field is empty, copy the following string value:

com.evermind.server.rmi.RMIInitialContextFactory

- **Context Provider URL** - In case this field is empty, copy the following highlighted string value:

opmn:ormi://<host>:<opmn_port>:home/orabpel

Note: Replace the <host>,<opmn port> with the correct host address and opmn port number details for the Oracle Application Server where the BPEL Process Manager is deployed.

3. Click **OK** to save the settings.

11.7.5 Retrieving the OPMN Port

To retrieve SOA Applications Server OPMN PORT details, follow these steps.

1. Open the configuration file \$SOA_ORACLE_HOME/opmn/conf/opmn.xml. \$SOA_ORACLE_HOME corresponds to SOA Application server home location.
2. Identify the value of the request port attribute in the configuration file.

11.7.6 javax.naming.NameNotFoundException Error

The following error occurs in the error details page when incorrect provider URL is specified.

Table 11–11 javax.naming.NameNotFoundException Error - Workaround Steps

Error Message	Workaround Steps
oracle.sysman.emSDK.emd.fetchlet.FetchletException: java.lang.Exception: Failed to create "ejb/collaxa/system/ServerBean" bean; exception reported is: "javax.naming.NameNotFoundException: ..."	<ol style="list-style-type: none">1. Validate the format of the string.2. Verify if the OPMN port is correct.3. Verify if the <oc4j_instance> name is properly substituted with the correct value, that is, the OC4J name value. The format must be like this: opmn:ormi://<host>:<opmn_port>:home/orabpel

Figure 11–1 javax.naming.NameNotFoundException Error



11.7.7 javax.naming.NamingException Error

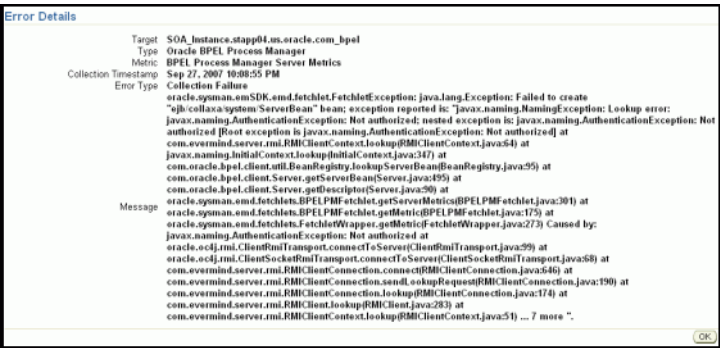
The following error occurs in the error details page when incorrect password is specified.

Table 11–12 *javax.naming.NamingException Error - Workaround Steps*

Error Message	Workaround Steps
oracle.sysman.emSDK.emd.fetchlet.FetchletException: java.lang.Exception: Failed to create "ejb/collaxa/system/ServerBean" bean; exception reported is: "javax.naming.NamingException: Lookup error:..."	1. Validate the values specified for BPEL Admin username and BPEL Password fields in the Monitoring Configuration page. (Confirm the validity of credentials by using the same credentials to log in to the BPELConsole).

(See [Figure 11–2](#))

Figure 11–2 *javax.naming.NamingException Error*



11.7.8 javax.naming.NoInitialContextException Error

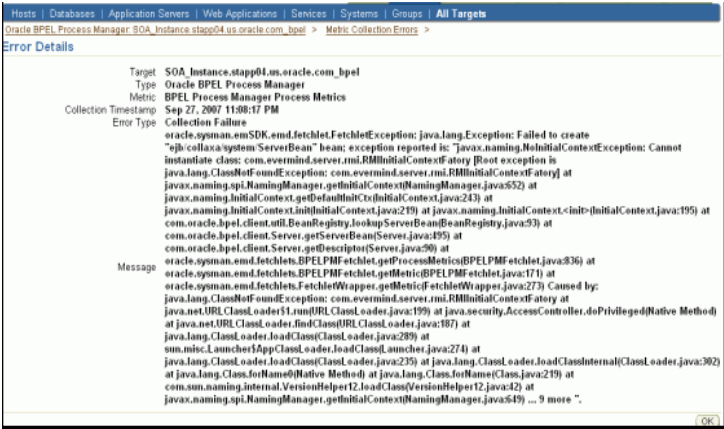
The following error occurs in the error details page when incorrect *Initial Context Factory* value is specified.

Table 11–13 *javax.naming.NoInitialContextException Error - Workaround Steps*

Error Message	Workaround Steps
oracle.sysman.emSDK.emd.fetchlet.FetchletException: java.lang.Exception: Failed to create "ejb/collaxa/system/ServerBean" bean; exception reported is: "javax.naming.NoInitialContextException: Cannot instantiate class:..."	1. Provide the following value for the Initial Context Factory field in the Monitoring Configuration page: com.evermind.server.rmi.RMIInitialContextFactory

(See [Figure 11–3](#))

Figure 11–3 *javax.naming.NoInitialContextException Error*



11.7.9 Error While Creating BPEL Infrastructure Services

The following error occurs when you are creating a new BPEL infrastructure service.

Table 11–14 *javax.naming.NoInitialContextException Error - Workaround Steps*

Error Message	Workaround Steps
An error encountered while discovering the dependencies. Please try again.	<div>1. Apply patch 10849036 on the OMS and try creating the BPEL infrastructure service again: com.evermind.server.rmi.RMIInitialContextFactory</div>

11.7.10 Metric Collection Errors for BPEL Process Manager Partner Link Metrics

The following metric collection error appears on the home page when you monitor BPEL 10.1.3.3 or 10.1.3.4 using Oracle Management Agent 12c:

Table 11–15 *Metric Collection Errors for BPEL Process Manager Partner Link Metrics - Workaround Steps*

Error Message	Workaround Steps
java.rmi.UnmarshalException: Error deserializing return-value: java.io.InvalidClassException: javax.xml.namespace.QName; local class incompatible: stream classdesc serialVersionUID = -916876369326528164, local class serialVersionUID = -9120448754896609940 at com.oracle.bpel.client.util.ExceptionUtils.handleServerException(ExceptionUtils.java:82) at com.oracle.bpel.client.BPELProcessHandle.getDescriptor(BPELProcessHandle.java:207) at oracle.sysman.emd.fetchlets.BPELPMFFetchlet.getPartnerLinkMetrics(BPELPMFFetchlet.java:873) at oracle.sysman.emd.fetchlets.BPELPMFFetchlet.getMetric(BELPMFFetchlet.java:235) at oracle.sysman.emd.fetchlets.FetchletWrapper.getMetric(FetchletWrapper.java:382)	<div>Follow the workaround described in the My Oracle Support Note 735128.1. You can access My Oracle Support at the following URL: https://support.oracle.com/CSP/ui/flash.html</div>

11.7.11 Agent Monitoring Metric Errors

This error occurs if the same Agent is used to monitor the BPEL 10g and OSB, and BPEL 10g and SOA 11g targets.

Table 11–16 Metric Errors During Agent Monitoring

Error Message	Workaround
The following exception has occurred: Exception at getPartnerLinkMetrics: java.lang.NoClassDefFoundError: Could not initialize class javax.rmi.PortableRemoteObject	The same Management must not be used to monitor BPEL 10g and OSB, and BPEL 10g and SOA 11g targets.

Discovering and Monitoring Oracle Service Bus

This chapter describes how you can discover and monitor Oracle Service Bus (OSB) using Enterprise Manager Cloud Control.

In particular, this document covers the following:

- [Supported Versions](#)
- [Understanding the Discovery Mechanism](#)
- [Understanding the Discovery Process](#)
- [Downloading One-Off Patches](#)
- [Discovering Oracle Service Bus](#)
- [Enabling Management Packs](#)
- [Monitoring Oracle Service Bus in Cloud Control](#)
- [Generating Oracle Service Bus Reports Using BI Publisher](#)
- [Troubleshooting Oracle Service Bus](#)

12.1 Supported Versions

The following are the versions of OSB that are supported for monitoring in Enterprise Manager Cloud Control Release 12c.

- Aqualogic Service Bus 2.6
- Aqualogic Service Bus 3.0
- Oracle Service Bus 10gR3
- Oracle Service Bus 11.1.1.2.0
- Oracle Service Bus 11.1.1.3.0
- Oracle Service Bus 11.1.1.4.0
- Oracle Service Bus PS4
- Oracle Service Bus PS5
- Oracle Service Bus PS6 (11.1.1.7.0)

12.2 Understanding the Discovery Mechanism

The OSB deployed to Oracle WebLogic Managed Server is automatically discovered in Enterprise Manager Cloud Control when that Oracle WebLogic Managed Server is discovered and added to Enterprise Manager Cloud Control.

The discovery of OSB depends on whether the Oracle WebLogic Managed Server is already being monitored in Enterprise Manager Cloud Control.

- If Oracle WebLogic Managed Server is not being monitored in Cloud Control, then first discover and add it to Cloud Control; this will automatically discover the OSB that is deployed to it.
- If Oracle WebLogic Managed Server is already being monitored in Cloud Control, then refresh the membership of the Oracle WebLogic Server Domain to which the Oracle WebLogic Managed Server belongs. This will automatically discover the OSB that is deployed to it.

For instructions to discover OSB, see [Section 12.5, "Discovering Oracle Service Bus"](#).

12.3 Understanding the Discovery Process

The following table describes the overall process involved in discovering and monitoring OSB in Enterprise Manager Cloud Control. Follow the instructions outlined against each step in this process to successfully discover and monitor your OSB.

Table 12–1 *Discovery Process*

Step	Requirement	Description
1	Oracle Service Bus	Install the OSB software. Note: Before you launch the OSB Deployment Procedure, ensure that Sun JDK has been installed.
2	Enterprise Manager Cloud Control	Install Enterprise Manager 12c. For information about installing the base release of Enterprise Manager Cloud Control, see the <i>Enterprise Manager Cloud Control Basic Installation and Configuration Guide</i> available at: http://www.oracle.com/technology/documentation/oem.html Oracle recommends that you install the Enterprise Manager Cloud Control components on a host that is different from the host where OSB is installed. For example, if OSB is installed on host1.xyz.com, then install and configure Oracle Management Service (OMS) and the Management Repository on host2.xyz.com.

Table 12–1 (Cont.) Discovery Process

Step	Requirement	Description
3	Oracle Management Agent (Management Agent)	<p>Install Oracle Management Agent 12c on the host where OSB is installed.</p> <p>If OSB and Enterprise Manager Cloud Control are on the same host, then you do not have to install a separate Management Agent. The Management Agent that comes with Enterprise Manager Cloud Control is sufficient. However, if they are different hosts, then you must install a separate Management Agent on the host where OSB is installed. Alternatively, the Management Agent can also be installed on a different host and made to remotely monitor the OSB target on another host.</p> <p>You can install the Management Agent in one of the following ways:</p> <ul style="list-style-type: none"> ■ Invoke the installer provided with Enterprise Manager 12c, and select the installation type Additional Management Agent. Then apply the 10.2.0.5 Agent patch on it. ■ Use the Agent Deploy application within the Enterprise Manager 12c. ■ Use the full agent kit that is available at: http://www.oracle.com/technology/software/products/oem/htdocs/agentsoft.html <p>For information about installing the Management Agent, see the <i>Enterprise Manager Cloud Control Basic Installation and Configuration Guide</i> available at: http://www.oracle.com/technology/documentation/oem.html</p>
4	One-Off Patches	<p>The support for discovering and monitoring of OSB is enabled only when the one-off patches as described in Section 12.4, "Downloading One-Off Patches" are applied to the WebLogic Server Home where OSB is running.</p>
5	Discovery in Enterprise Manager Cloud Control	<p>OSB is automatically discovered when the Oracle WebLogic Server Domain to which it is deployed is discovered and added to Enterprise Manager Cloud Control.</p>

12.4 Downloading One-Off Patches

To view OSB services in Enterprise Manager Cloud Control, you must apply the following patches to your OSB servers.

Table 12–2 One-Off Patches

Oracle Service Bus Version	ID	Password
Oracle Service Bus 2.6	EMMU	83XNT2D4
Oracle Service Bus 2.6.1	9NAF	TLZE4IPI
Oracle Service Bus 3.0	RPCD	JJEC2EY2
Oracle Service Bus 10.3.0	9HPA	FFLQHDHP
Oracle Service Bus 10.3.1	No Patch Required	
Oracle Service Bus 11.1.1.3.0 and 11.1.1.4.0, and Oracle Service Bus PS4, PS5 and PS6	No Patch Required	

You can apply the patches in one of the following ways:

- **Online mode** - Using the SmartUpdate tool available with Oracle WebLogic Managed Server
- **Offline mode** - Manually copying the JAR files and classes to the OSB directories

For information about downloading these patches and applying them in either offline or online mode, see My Oracle Support Note 804148.1. You can access My Oracle Support at:

<https://support.oracle.com/CSP/ui/flash.html>

Note: After applying the patches, restart the WebLogic domain and all of the management agents monitoring the domain.

12.5 Discovering Oracle Service Bus

The OSB deployed to Oracle WebLogic Managed Server is automatically discovered in Enterprise Manager Cloud Control when that Oracle WebLogic Managed Server is discovered and added to Enterprise Manager.

Before discovering OSB, identify whether the Oracle WebLogic Managed Server is already being monitored in Enterprise Manager.

- If Oracle WebLogic Managed Server is not being monitored in Enterprise Manager, then first discover and add it to Enterprise Manager Cloud Control; this will automatically discover the OSB that is deployed to it.
- If Oracle WebLogic Managed Server is already being monitored in Enterprise Manager, then refresh the membership of the Oracle WebLogic Server Domain to which the Oracle WebLogic Managed Server belongs. This will automatically discover the OSB that is deployed to it.

This section outlines the instructions for discovering OSB for the cases described above. In particular, this section covers the following:

- [Discovering OSB Deployed to WLS Not Monitored by Enterprise Manager](#)
- [Discovering OSB Deployed to WLS Monitored by Enterprise Manager](#)

12.5.1 Discovering OSB Deployed to WLS Not Monitored by Enterprise Manager

To discover OSB deployed to Oracle WebLogic Manager Server that is not monitored in Cloud Control, first discover that Oracle WebLogic Manager Server in Enterprise Manager Cloud Control; this will automatically discover the OSB that is deployed to it. To discover Oracle WebLogic Manager Server, follow these steps:

1. From the **Targets** menu, select **Middleware**.

Enterprise Manager Cloud Control displays the Middleware page that lists all the middleware targets being monitored.

2. In the Middleware page, select **Oracle Fusion Middleware/WebLogic Server Domain** from the **Add** drop-down list and click **Go**.

Enterprise Manager Cloud Control displays the Add Oracle Fusion Middleware / WebLogic Server Domain wizard that captures the details of the Oracle WebLogic Server Domain to be discovered and monitored.

3. In the Add Oracle Fusion Middleware / WebLogic Server Domain wizard, specify the required details and click **Next** on each page to reach the end of the wizard.

For information about the details to be provided for each page of the wizard, click **Help** on each page.

4. In the last page of the Add Oracle Fusion Middleware / WebLogic Server Domain wizard, click **Finish** to complete the discovery process and add the target to Cloud Control for monitoring purposes.

Enterprise Manager displays the Middleware page with a confirmation message that confirms that the Oracle WebLogic Manager Server has been successfully added to Cloud Control.

In the Middleware page that shows all the middleware targets being monitored, you can see the Oracle WebLogic Managed Server and the OSB you just added. Note that, at this point, OSB will be the last target listed in the table. To see it nested under its Oracle WebLogic Managed Server, click **Refresh** on this page. Alternatively, navigate to another tab or page, and then return to the Middleware page.

Note:

- After discovering and adding OSB to Enterprise Manager Cloud Control, you can monitor its status from the OSB Home page. You can use the Services page to view a list of services.

For the first collection that happens, you will see the value "0" for all the metrics that are enabled in Oracle Enterprise Manager Release 12c. This is an expected behavior. From the second collection onwards, you should see the actual metric values. However, if you still see the value "0", then perhaps the service monitoring is turned off. To resolve this issue, on the Services page, click Launch Console to access the OSB Console, and turn on the service monitoring and set the level to "pipeline" or "action"

- In the case of clustered OSB domain, the Management Agent installed on Admin Server host should be used to discover the entire domain. This constraint is not applicable for version 12.1.0.2 of Cloud Control. This is only valid up to version 12.1.0.1 of Cloud Control.
-

For additional information about Fusion Middleware discovery, see [Chapter 2, "Discovering Middleware Targets"](#).

12.5.2 Discovering OSB Deployed to WLS Monitored by Enterprise Manager

To discover OSB deployed to Oracle WebLogic Managed Server that is already being monitored in Cloud Control, refresh the membership of the Oracle WebLogic Server Domain to which the Oracle WebLogic Managed Server belongs. This will automatically discover the OSB that is deployed to it.

To refresh the membership of the Oracle WebLogic Server Domain to which the Oracle WebLogic Managed Server belongs, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the **Middleware** page, select the **Oracle WebLogic Server Domain** target from the list of Middleware targets being monitored.
3. On the Oracle WebLogic Server Domain Home page, in the General section, click **Refresh Domain**. Enterprise Manager Cloud Control displays the membership page that lists the OSB that is currently not being monitored. Click **OK**.

Enterprise Manager Cloud Control refreshes the membership and returns to the Oracle WebLogic Server Domain Home page.

Note: On the Oracle WebLogic Server Domain Home page, in the Status section, the legend of the status pie chart may not show an increased count to indicate the newly added OSB target. This is an expected behavior because Enterprise Manager Cloud Control takes a few seconds to reflect the membership details in this section.

4. Click the **Members** tab and verify whether the OSB has been added.

12.6 Enabling Management Packs

Besides monitoring the status of OSB, if you want to gain access to additional value-added features, then you must enable the Management Pack for SOA.

To enable the Management Pack for SOA:

1. From the **Setup** menu, select **Management Packs**, then select **Management Pack Access**.

Enterprise Manager Cloud Control displays the Management Pack Access page.

2. In the Management Pack Access page, from the Search list, select **Oracle Service Bus**.

Enterprise Manager Cloud Control lists all the Oracle Service Bus targets being monitored.

3. From the table, for the Oracle Service Bus target you are interested in, enable the SOA Management Pack Enterprise Edition and click **Apply**.

12.7 Monitoring Oracle Service Bus in Cloud Control

Enterprise Manager Cloud Control helps you monitor the health of Oracle Service Bus targets deployed to Oracle WebLogic Managed Servers. When you discover Oracle WebLogic Manager Servers, Cloud Control automatically discovers the Oracle Service Bus targets deployed to them and adds them for central monitoring and management.

For each Oracle Service Bus target being monitored, Cloud Control provides information about its status, availability, performance, services, alerts, business services, and proxy services. It also allows you to view the latest configuration details, save them at a particular time, and compare them with another Oracle Service Bus instances. Oracle Service Bus also provides a graphical view representation for the dependencies between proxy services and business services.

In addition to monitoring capabilities, Cloud Control also allows you to black out an Oracle Service Bus target and create infrastructure services. While blackout helps you suspend the monitoring of the target for a temporary period (for example, during maintenance), infrastructure services are dependency services that are created to identify the infrastructure components on which the Oracle Service Bus target depends.

12.7.1 Enabling Monitoring for OSB Services

If you are not able to view OSB data on Enterprise Manager pages, it may be because monitoring is disabled for OSB Services. Before you can view OSB data in Enterprise Manager, check to see if monitoring is enabled for OSB Services. You can do that by following these steps:

1. Click the **OSB Console** link in the OSB Home Summary region.
2. Log into the OSB target.
3. Click **Create** (to create the new session).
4. On the left navigation panel, click **Operations** and then choose **Global Settings**.
5. Determine whether monitoring is enabled. To enable monitoring, select the check box to update it.
6. Activate the session.

12.8 Generating Oracle Service Bus Reports Using BI Publisher

You can use Enterprise Manager to print Oracle Service Bus reports using BI Publisher Enterprise Reports. Oracle Business Intelligence (BI) Publisher is an enterprise reporting solution for authoring, managing, and delivering highly formatted documents. Oracle BI Publisher also allows you to build custom reporting applications that leverage existing infrastructure. Reports can be designed using familiar desktop products and viewed online or scheduled for delivery to a wide range of destinations.

For example, you can generate an OSB Services Report that describes the way OSB services have been performing over a period of time. The report provides charts to list the top 5 OSB Services and a table with critical metric details for all the services.

The following table describes the OSB-related reports you can choose.

Table 12–3 OSB Reports

OSB Report	Description
OSB Service Summary Report	The OSB Service Summary Report provides information about the Average Response Time, Open Instances Count, Fault Instances Count, and Web Service Security Violation Count for the selected service. The OSB Service Summary Report displays a chart with the top 5 OSB services based on Average Response Time or Throughput across the selected OSB services for the specified time period. The report can be sorted based on a performance metric (for example, Average Response Time) or a usage metric (for example, Instance Count). As part of the report parameters setting, you can use options that allow you to select the OSB Service by Projects or by selecting individual services.
OSB Service Operations Summary Report	The OSB Service Operations Summary Report provides internal operation level details for the selected service. The details in the report cover the Average Response Time, Open Instances Count, Fault Instances Count, and Web Service Security Violation Count. The report can be sorted based on a performance metric or a usage metric. As part of the report parameters setting, you can use options that allow you to select the OSB Service by Projects or by selecting individual services.
OSB Proxy Service Flow Component Performance Summary Report	A message flow is composed of components that define the logic for routing and manipulating messages as they flow through a proxy service. The OSB Proxy Service Flow Component Performance Summary Report provides internal flow component details for the selected proxy service. The details in the report display the Average Response Time, Open Instances Count, Fault Instances Count, and Web Service Security Violation Count. The report can be sorted based on a performance metric or a usage metric. As part of the report parameters setting, you can use options that allow you to select the OSB Service by Projects or by selecting individual services.

To print OSB reports using BI Publisher Enterprise reports, follow these steps:

1. From the Enterprise menu, click **Reports**, and then click **BI Publisher Enterprise Reports**.

Enterprise Manager Cloud Control displays the login page for BI Publisher Enterprise Reports.

2. Enter your credentials to log into BI Publisher.

The BI Publisher Enterprise page displays, showing you Recent reports, Others, and Favorites. You can use this page to create a new report, submit a report job, and perform other tasks.

3. Click the Report you want to display.

4. On the Report page, use the parameter filters to tailor the report structure that displays, then click **Refresh**.

You can view the OSB Services Report using the filters based on the various search parameters available at the top of the page, such as Target Name, Date Range, and so on. Similarly, you can view the report based on the Sort By option as well, allowing you to sort the report by Service Name or Average Response Time, for example.

You can refresh the report anytime by clicking the Refresh icon on the upper right side of the OSB Service Report tab. You can hide or display the search parameters by clicking the Parameters icon. You can choose to view the report in various formats such as HTML, PDF, RTF, Excel, and PowerPoint by clicking the View Report icon. Likewise you can display more available actions by clicking the Actions icon. For more help about using BI Publisher, click the help icon.

12.9 Troubleshooting Oracle Service Bus

This section describes the errors you might encounter while discovering OSB, and the workaround steps you can follow to resolve each of them.

12.9.1 Required Patches Missing

The following error occurs when you try to discover OSB from an Oracle WebLogic Admin Server that has not been patched with the required one-off patches.

Table 12–4 *oracle.sysman.emSDK.emd.fetchlet.FetchletException Error - Workaround Steps*

Error Message	Workaround Steps
oracle.sysman.emSDK.emd.fetchlet.FetchletException: The MBean is not available on the OSB instance. The required EM plug in patch should be missing on OSB instance.	Apply the one-off patches as described in Section 12.4 , " Downloading One-Off Patches ".

12.9.2 System and Service

The following error occurs if configuration information has not been collected for the selected Application Server.

Table 12–5 Create System and Service Error - Workaround Steps

Error Message	Workaround Steps
An error encountered while discovering the dependencies. This may occur if some configuration information is missing. Check whether the configuration information was collected for the dependent targets and then try again.	Collect the latest configuration data by navigating to the Application Server Home page and clicking Configuration and then select Last Collected from the Application Server menu.

12.9.3 SOAP Test

The following error occurs when the Management Agent is upgraded to Enterprise Manager 12c with OMS 10.2.0.5.

Table 12–6 SOAP Test Error - Workaround Steps

Error Message	Workaround Steps
Add SOAP Test failed. The selected service has an invalid or incorrect WSDL URL. Check whether the Oracle Service Bus Target URL value is valid in the Monitoring Configuration page of the selected target. To access the Monitoring Configuration page, go to the Oracle Service Bus Homepage and from the Related Links section, select Monitoring Configuration.	<p>If the Management Agent has been upgraded to 12c with OMS 10.2.0.5, the following workaround must be applied to support the SOAP test.</p> <p>In the Monitoring Configuration page for the OSB target, set the Server URL to Access Proxy Services property to the URL for the specific WebLogic Server target. The URL must be in the format: <code>http://<host>:<port>/</code>. For example, <code>http://stade61.us.oracle.com:7001/</code></p>

Discovering and Monitoring the SOA Suite

This chapter describes how you can discover and configure the components of the SOA Suite 11g using Enterprise Manager Cloud Control.

In particular, this document covers the following:

- [New Features in This Release](#)
- [List of Supported Versions](#)
- [Monitoring Templates](#)
- [Overview of the Discovery Process](#)
- [Discovering the SOA Suite](#)
- [Metric and Collection Settings](#)
- [Configuring Instance Tracing](#)
- [Setting Up and Using SOA Instance Tracing](#)
- [Monitoring Dehydration Store](#)
- [Viewing the Service Topology](#)
- [Publishing a Server to UDDI](#)
- [Generating SOA Reports Using BI Publisher](#)
- [Generating SOA Reports Using Information Publisher](#)
- [Provisioning SOA Artifacts and Composites](#)
- [Diagnosing Issues and Incidents](#)
- [Verifying Target Monitoring Setup](#)
- [Searching Faults in the SOA Infrastructure](#)
- [Creating and Tracking Bulk Recovery Jobs](#)
- [Recovering BPEL/BPMN Messages](#)
- [Troubleshooting](#)

13.1 New Features in This Release

The new features that have been introduced in the 12c version of the SOA Suite are:

- **Target Setup Verification** allows you to perform a series of tests that can help diagnose and repair setup problems required for target monitoring.

- **Fault Management** allows you to perform a real-time search and recovery of faults and BPEL/BPMN messages. In addition to this, You can perform bulk recovery of faults, and track and monitor the status of submitted bulk recoveries.
- **Oracle RAC Database Support for Dehydration Store Performance** allows you to monitor the health and performance of RAC-based Dehydration Store for Multi Data Source or GridLink Data Source type.
- **Instance Tracing Enhancements** allows you to perform the following operations:
 - Tracing instances using the Instance Id.
 - Viewing payloads for BPEL component instances.
 - Tracing instances at SOA Infrastructure level.
- **Support Work Bench Enhancements** is enhanced to generate BPEL and Mediator dumps. The enhanced capability is supported for SOA PS6 (11.1.1.7.0).

13.2 List of Supported Versions

The following are the versions of the SOA Suite 11g that are supported in Enterprise Manager Cloud Control 12c:

- 11.1.1.2.0 (PS1)
- 11.1.1.3.0 (PS2)
- 11.1.1.4.0 (PS3)
- 11.1.1.5.0 (PS4)
- 11.1.1.6.0 (PS5)
- 11.1.1.7.0 (PS6)

13.3 Monitoring Templates

The following Oracle-certified default templates are being shipped for PS1 and Ps2 agents. [Table 13–1](#) describes the available templates, and the agents they are applicable to:

Table 13–1 *Monitoring Templates*

Target Type	Agent Name	Template Name
SOA Infrastructure	PS1 Agent	Oracle Certified Fusion Apps Template for SOA Infrastructure for FMW Plugin 12.1.0.3.0
SOA Infrastructure	PS1 Agent	Oracle Certified FMW Template for SOA Infrastructure for FMW Plugin 12.1.0.3.0
SOA Infrastructure	PS2 Agent	Oracle Certified Fusion Apps Template for SOA Infrastructure
SOA Infrastructure	PS2 Agent	Oracle Certified FMW Template for SOA Infrastructure
SOA Composite	PS1 Agent	Oracle Certified Fusion Apps Template for SOA Composite for FMW Plugin 12.1.0.3.0
SOA Composite	PS1 Agent	Oracle Certified FMW Template for SOA Composite for FMW Plugin 12.1.0.3.0
SOA Composite	PS2 Agent	Oracle Certified Fusion Apps Template for SOA Composite

Table 13–1 (Cont.) Monitoring Templates

Target Type	Agent Name	Template Name
SOA Composite	PS2 Agent	Oracle Certified FMW Template for SOA Composite

Note: The templates created using older versions of OMS (12c PS1, 12c BP1, and so on) should not be used in 12c PS2.

13.4 Overview of the Discovery Process

The following describes the overall process involved in discovering and monitoring SOA Suite 11g in Enterprise Manager Cloud Control. Follow the instructions outlined against each step in this process to successfully discover and monitor the SOA Suite.

Table 13–2 Understanding the Discovery Process

Oracle SOA Suite Version	Application Server Deployed To	Discovery Mechanism	Process
Oracle SOA Suite 11.1.x	Oracle WebLogic Managed Server	Manual Discovery	<ol style="list-style-type: none"> 1. First, manually discover Oracle WebLogic Managed Server. For procedures, see Section 11.5.2.1, "Discovering Oracle WebLogic Managed Server". 2. To monitor the SOA Suite, you can use an agent running locally on the Administration Server of the WebLogic domain or a remote management agent running on another host that is not part of the WebLogic domain. <p>Note: If you use a remote agent to monitor the SOA Suite, then certain the following operations like:</p> <ul style="list-style-type: none"> - Instance Tracing, Viewing Recoverable Instances, Fault Recovery, and Provisioning SOA Artifacts is not supported. - Host Metrics cannot be captured by the remote agent. 3. To ensure the all the metric data is collected, add the <code>soa-infra-mgmt.jar</code> and the <code>oracle-soa-client-api.jar</code> files to the <code>\$AGENT_HOME/plugins/oracle.sysman.emas.agent.plugin_12.1.0.0.0/archives/jlib/</code> (the Agent Home directory). If the <code>extjlib</code> directory does not exist, it can be created under <code>\$FMW_PLUGIN_HOME/archives/jlib</code>. This step is required only if you are using a remote agent to monitor the SOA Suite. <p>Note: For SOA PS3 and higher, the <code>jrf-api.jar</code> file must also be present in the Agent Home directory.</p>

13.5 Discovering the SOA Suite

This section describes the procedure for discovering the SOA Suite 11g.

- [Discovering the SOA Suite](#)
- [Configuring the SOA Suite](#)

13.5.1 Discovering the SOA Suite

You can use a local or a remote Management Agent to perform the discovery process. as follows:

- [Discovering the SOA Suite Using a Local Agent](#)
- [Discovering the SOA Suite Using a Remote Agent](#)

Discovering the SOA Suite Using a Local Agent

In the case of discovery using a local agent, you need to use a Management Agent that is running on the same host as the Administration Server.

1. From the **Targets** menu, select **Middleware**.
Oracle Enterprise Manager Cloud Control displays the Middleware page that lists all the middleware targets being monitored.
2. In the Middleware page, from the **Add** list, select Oracle Fusion Middleware / WebLogic Domain and click **Go**. Specify the Administration Server Host, Port, User Name, Password, Agent (local or remote) and the JMX Protocol and click **Continue**.
3. You will return to the Middleware page. You will see the SOA instances under the WebLogic Domain.

Note: SOA Composites that are created after the discovery of SOA Suite Domain are not displayed automatically. To view all the SOA Composites, navigate to the Home page of the WebLogic Server target and select the **Refresh Domain** option from the menu.

Discovering the SOA Suite Using a Remote Agent

You can discover the SOA Suite 11g using a remote agent which may be running on a host that is different from the host on which the Administration Server is running. In this case, you may not be able to perform certain operations like instance tracing, fault recovery, viewing recoverable instances, provisioning SOA Artifacts, and capturing the host metrics.

To ensure the all the metric data is collected, add the `soa-infra-mgmt.jar` and the `oracle-soa-client-api.jar` files to the `$AGENT_HOME/plugins/oracle.sysman.emas.agent.plugin_12.1.0.0.0/archives/jlib/` (the Agent Home directory). If the `extjlib` directory does not exist, it can be created. This step is required only if you are using a remote agent to monitor the SOA Suite.

Note: For SOA PS3 and higher, the `jrf-api.jar` file must also be present in the Agent Home directory.

13.5.2 Configuring the SOA Suite

After discovering the SOA Suite 11g, you must perform the following additional configuration steps:

1. Set the instance state in the Common Properties page:

- a. Click **Fusion Middleware Control** on the SOA Infrastructure Home Page.
 - b. Navigate to the Home page of the SOA Infrastructure target.
 - c. Select **Common Properties** from the SOA-Infra drop-down menu.
 - d. On the Common Properties page, select the **Capture Composite Instance State** check box.
2. Set the SOA database details like the host name, port, and credentials.
 - a. Navigate to the Middleware page in Enterprise Manager Cloud Control.
 - b. Select a SOA Infrastructure home from the list and click **Configure**.
 - c. The SOA Infrastructure Home page is displayed. Click **Monitoring Configuration** from the SOA-Infra drop-down menu.
 - d. Set the SOA database details in the Monitoring Configuration page.
 3. Set preferred credentials for the WebLogic Domain.
 - a. From the **Setup** menu, select **Security**, then select **Preferred Credentials**.
 - b. Select the Oracle WebLogic Domain target and click **Managed Preferred Credentials**.
 - c. Select WebLogic Administrator Credentials in the Target Preferred Credentials and click **Set**.
 - d. Enter the user name and password in the Select Named Credentials window and click **Save**.

13.6 Metric and Collection Settings

For the following metrics the collection schedule is not available on the Metric and Collection Settings page. Detailed steps to update the collection intervals are listed in the following table:

Table 13–3 Metric and Collection Settings

Target Type	Metric Name	Collection Interval Update Steps
SOA Infrastructure	Response	<p>Navigate to the associated weblogic server where SOA is deployed, to do so, follow these steps:</p> <ol style="list-style-type: none"> 1. From the Targets menu, select Middleware. 2. On the Middleware page, click the SOA Infrastructure target. 3. On the SOA Infrastructure target page, from WebLogic Server menu, select Monitoring, and click Metric and Collection Settings. 4. Click Other Collected Items tab. 5. Click Collection Schedule corresponding to Application Metrics to update the collection interval. <p>Note: This change is applicable to all the applications deployed in that WebLogic server.</p>

Table 13–3 (Cont.) Metric and Collection Settings

Target Type	Metric Name	Collection Interval Update Steps
SOA Composite	Response	<p>For SOA PS5 (11.1.1.6.0) or earlier, follow these steps:</p> <ol style="list-style-type: none"> 1. From the Targets menu, select Middleware. 2. On the Middleware page, click the SOA Composite target. 3. On the SOA Composite target page, from SOA Composite menu, select Monitoring, and click Metric and Collection Settings. 4. Click Other Collected Items tab. 5. Update the collection interval for the metric SOA Composite Status (11.1.1.6.0 and earlier) <p>For SOA PS6 (11.1.1.7.0) onwards, navigate to the associated SOA Infrastructure where SOA composite is deployed. To do so, follow these steps:</p> <ol style="list-style-type: none"> 1. From the Targets menu, select Middleware. 2. On the Middleware page, click the SOA Infrastructure where SOA composite is deployed. 3. On the SOA Infrastructure target page, from SOA Infrastructure menu, select Monitoring, and click Metric and Collection Settings. 4. Click Other Collected Items tab. <p>Update the collection interval for the metric SOA Composite Status.</p> <p>Note: This change is applicable to all the SOA composites which are deployed in that SOA Infrastructure</p>
		<p>Navigate to the associated SOA Infrastructure where soa composite is deployed. To do so, follow these steps:</p> <ol style="list-style-type: none"> 1. From the Targets menu, select Middleware. 2. On the Middleware page, click the SOA Infrastructure where SOA composite is deployed. 3. On the SOA Infrastructure target page, from SOA Infrastructure menu, select Monitoring, and click Metric and Collection Settings. 4. Click Other Collected Items tab. <p>Update the collection interval for the metric SOA Infrastructure - Recoverable Faults.</p>
SOA Composite	SOA Composite - Component Detail Metrics	<p>Navigate to the associated SOA Infrastructure where soa composite is deployed. To do so, follow these steps:</p> <ol style="list-style-type: none"> 1. From the Targets menu, select Middleware. 2. On the Middleware page, click the SOA Infrastructure where SOA composite is deployed. 3. On the SOA Infrastructure target page, from SOA Infrastructure menu, select Monitoring, and click Metric and Collection Settings. 4. Click Other Collected Items tab. <p>Update the collection interval for the metric SOA Infrastructure - Recoverable Faults.</p>

Table 13–3 (Cont.) Metric and Collection Settings

Target Type	Metric Name	Collection Interval Update Steps
SOA Composite	SOA Composite - Recoverable And Rejected Messages	<p>Navigate to the associated SOA Infrastructure where soa composite is deployed. To do so, follow these steps:</p> <ol style="list-style-type: none"> 1. From the Targets menu, select Middleware. 2. On the Middleware page, click the SOA Infrastructure where SOA composite is deployed. 3. On the SOA Infrastructure target page, from SOA Infrastructure menu, select Monitoring, and click Metric and Collection Settings. 4. Click Other Collected Items tab. <p>Update the collection interval for the metric SOA Infrastructure - Recoverable And Rejected Messages.</p>

13.6.1 Viewing Application Dependency and Performance (ADP) Metrics

If the SOA instance is being monitored by the ADP Manager, additional metrics such as Arrival Rate, Minimum, Maximum, and Average Response Time will be collected.

Tip: The ADP Manager must be registered before it can collect the metric data. For details on registering the ADP Manager, see *Oracle Enterprise Manager Cloud Control Advanced Installation and Configuration Guide*.

13.7 Configuring Instance Tracing

Before enabling Instance Tracing, ensure that you meet the following prerequisites:

- Ensure that the SOA infrastructure is monitored by an Oracle Management Agent (Management Agent) that is running on the same host as the SOA infrastructure.
- Set the following as preferred credentials:
 - Credentials of the host on which the SOA server is running.
 - Administrator credentials of the Oracle WebLogic Domain.

To enable Instance Tracing for any SOA Infrastructure instances involved in executing composite instances:

1. Follow the sequence listed under Step 3 of the [Section 13.5.2, "Configuring the SOA Suite"](#) section.
2. To view the state of the listed instances, enable the Capture Composite State flag by following the sequence listed Step 1 of the [Section 13.5.2, "Configuring the SOA Suite"](#) section.

13.8 Setting Up and Using SOA Instance Tracing

Starting with Enterprise Manager 12.1.0.3 PS2 release, Instance Tracing functionality is also available at the SOA Infrastructure level, in addition to the existing SOA Composite level.

The section contains the following topics:

- [SOA Instance Tracing at SOA Infrastructure Level](#)

- [SOA Instance Tracing at SOA Composite Level](#)

13.8.1 SOA Instance Tracing at SOA Infrastructure Level

At SOA Infra level, an instance search is performed across SOA Composites, and a particular SOA composite instance is traced. To do so, follow these steps:

1. In Cloud Control, from **Targets** menu, select **Middleware**.
2. On the Middleware page, click the **SOA Infrastructure** target.
3. On the SOA Infrastructure home page, from **SOA Infrastructure** menu, select **Instance Tracing**.
4. On the Trace Instance page, search for the instance you want to trace by specifying search criteria such as Instance ID, Instance Creation Time Window, Instance Count, Sensor Values and so on. Click **Search**.

Note: You must set the Preferred Credentials for the WebLogic Domain and the Monitoring Agent Host to retrieve search data. These settings for Instance tracing can be done from the Target Setup Verification page. To do so, from the **Targets** menu, select **Middleware**. On the Middleware page, click the target you are interested in. For example, SOA Infrastructure. On the Home page of the target, from the target-specific menu, select **Target Setup**, and click **Verification**.

- **WebLogic Domain:** To set the Administrator Credentials for the WebLogic Domain (SOA Domain Target), select **Security**, then select **Preferred Credentials** from the **Setup** menu. Select Oracle WebLogic Domain as the Target Type and click **Manage Preferred Credentials**. Specify the WebLogic Administrator credentials and click **OK**.
 - **Monitoring Agent Host:** To set the Preferred Credentials for the host on which the Management Agent has been installed, select **Security**, then select **Preferred Credentials** from the **Setup** menu. Select Host as the Target Type and click **Manage Preferred Credentials**. Specify the Normal Host Credentials for the host and click **OK**.
-

All the instances across all the SOA Composites are displayed. A new QBE feature has been introduced into the instance table which allows you to filter the instances based on different fields like: Instance id, Partition name, Composite name, Instance State, Tracing job status, Name, ECID, and Conversation Id.

5. Select the **Instance Id** of a particular SOA Composite that you want to trace, click **Trace**.
6. In the Trace Instance dialog box, click **Add** to add one or more SOA Infrastructure targets for which the trace information is to be gathered. Specify the WebLogic Administration and Host Credentials for these targets and click **OK**.

A trace job is submitted to collect the instance trace data from the selected targets and the flow trace for the instance is generated.

7. When the job has been completed, click **Tracing Job Status** to view the flow trace data.

8. Click the **Component Instance** to drill down to the component instance audit trail which shows the activity level execution details.

To view the payloads for BPEL component instance, click the component instance to drill down to the component audit trail, and then click **View Payload** to view the payloads.

13.8.2 SOA Instance Tracing at SOA Composite Level

To configure and trace a SOA Composite instance, follow these steps:

1. In Cloud Control, from **Targets** menu, select **Middleware**.
2. On the Middleware page, click the **SOA Composite** target.
3. On the SOA Composite home page, from the **SOA Composite** menu, select **Trace Instance**.
4. On the Trace Instance page, search for the instance you want to trace by specifying search criteria such as Instance ID, Instance Creation Time Window, Instance Count, Sensor Values and so on. Click **Search**.

Note: You must set the Preferred Credentials for the WebLogic Domain and the Monitoring Agent Host to retrieve search data. These settings for Instance tracing can be done from the Target Setup Verification page. To do so, from the **Targets** menu, select **Middleware**. On the Middleware page, click the target you are interested in. For example, SOA Infrastructure. On the Home page of the target, from the target-specific menu, select **Target Setup**, and click **Verification**.

- **WebLogic Domain:** To set the Administrator Credentials for the WebLogic Domain (SOA Domain Target), select **Security**, then select **Preferred Credentials** from the **Setup** menu. Select Oracle WebLogic Domain as the Target Type and click **Manage Preferred Credentials**. Specify the WebLogic Administrator credentials and click **OK**.
 - **Monitoring Agent Host:** To set the Preferred Credentials for the host on which the Management Agent has been installed, select **Security**, then select **Preferred Credentials** from the **Setup** menu. Select Host as the Target Type and click **Manage Preferred Credentials**. Specify the Normal Host Credentials for the host and click **OK**.
-

The list of instances that match the search criteria are displayed. A new QBE feature has been introduced into the instance table which allows you to filter the instances based on different fields like: Instance id, Partition name, Composite name, Instance State, Tracing job status, Name, ECID, and Conversation Id.

5. Select the **Instance Id** for tracing instances present within a SOA Domain, to trace instances that span across various SOA domains in your infrastructure, click **Trace**.
6. In the Trace Instance dialog box, click **Add** to add one or more SOA Infrastructure targets for which the trace information is to be gathered. Specify the WebLogic Administration and Host Credentials for these targets and click **OK**.

A trace job is submitted to collect the instance trace data from the selected targets and the flow trace for the instance is generated.

7. When the job has been completed, click **Tracing Job Status** to view the flow trace data. Click the **Log Viewer** link to access the log messages related to the instance.
8. Click the **Component Instance** to drill down to the component instance audit trail which shows the activity level execution details.

To view the payloads for BPEL component instance, click the component instance to drill down to the component audit trail, and then click **View Payload** to view the payloads.

13.9 Monitoring Dehydration Store

The Dehydration Store Diagnostics feature provides a dedicated view that allows you to analyze the behavior of the SOA Dehydration database. You can monitor SQL performance metrics and table growth specifically in the context of the SOA Suite's use of the database. The view displays both throughput and wait bottleneck data which allows you to monitor the general health of the target database instance. Using Active Session History, you can track usage data and display it as a table space chart, a growth rate chart, or an execution chart.

Note: In addition to monitoring Oracle standalone database, the Dehydration Store now supports reviewing the general health of the RAC database engine, and identifying problems that are causing performance bottlenecks.

You can also monitor Real Application Cluster (RAC) databases. For RAC, you can monitor Multi Data Source and GridLink Data Sources. In RAC scenario, the dehydration store tab lists all the associated database nodes in the form of a drop down menu. You can select any particular instance from the **Show Database Instance** menu, and view the associated metric data.

13.9.1 Enabling Monitoring of the SOA Dehydration Store

To configure and enable monitoring of the SOA Dehydration Store, follow these steps:

1. From the **Targets** menu, select **Databases** to check if the database target representing the SOA Dehydration Store has been discovered in Enterprise Manager.
2. Check if at least one configuration for the SOA Infrastructure and WebLogic Server targets is available.
3. On the monitoring configuration for the SOA Infrastructure target, the following fields related to SOA Repository must be configured:
 - SOA Repository Connection Descriptor: The connection URL string specified for the JDBC data source on the WebLogic server. This configuration is collected as part of the configuration collection mechanism for the SOA Server instance. For example,

On Single Instance Database

host:port/sid (or service_name)

On RAC Database

- Multi Data Source:

(DESCRIPTION= (ADDRESS_LIST= (ADDRESS= (PROTOCOL=TCP) (HOST=<Host

```
1>) (PORT=<Port 1>)) (CONNECT_DATA=(SERVICE_NAME=<Service
Name>) (INSTANCE_NAME=<Instance 1>)); (DESCRIPTION=(ADDRESS_LIST
=(ADDRESS=(PROTOCOL=TCP) (HOST=<Host 2>) (PORT=<Port 2>))) (CONNECT_DATA
=(SERVICE_NAME=<Service Name>) (INSTANCE_NAME=<Instance 2>)))
```

– GridLink:

```
(DESCRIPTION=(ADDRESS_LIST=
ADDRESS=(PROTOCOL=TCP) (HOST=<Host>) (PORT=<Port>)))
(CONNECT_DATA=(SERVICE_NAME=<Port>))
```

- SOA Repository Host Name: The database listener host for the SOA database instance. This is optional if the connection string has already been configured.
- SOA Repository Port: - The database listener port for the SOA database instance. This is optional if the connection string has already been configured.
- SOA Repository Schema Name: The schema name configured for SOA Dehydration Store.
- SOA Repository User Name: The schema name configured for SOA Dehydration Store.
- SOA Repository Password: The password for the SOA schema user.
- SOA Repository SID: The SID for the SOA database instance.

If you do not see data after these configuration details have been specified, you must wait for the next collection interval.

13.9.2 Viewing the SOA Dehydration Store Data

To view the dehydration diagnostics data, follow these steps:

1. From the **Targets** menu, select **Middleware**, then click on a SOA Infrastructure target.
2. In the SOA Infrastructure Home page, click the **Dehydration Store** tab.
3. The following details area displayed:
 - Throughput indicators that provide details of the general health of the database instance.
 - Wait bottleneck issues related to the CPU, I/O, and Wait events.
 - Tablespace utilization for the SOA schema.
 - Performance data recorded by the ASH.
 - Key SOA tables and tablespace details related to the SOA schema.

13.10 Viewing the Service Topology

The Service Topology provides a graphical end to end view of the composite applications. It depicts the various application components and their interactions happening at runtime. It allows you to view the service level dependencies among the components and provides key performance statistics and incidents information for them. Composite applications are distributed in nature and this view helps you quickly visualize the structure, status, availability, dependencies, configuration changes, and performance of business-critical distributed applications from one place and easily identify any availability or performance issues.

The Service Topology Viewer shows the following:

- **Service to Service Calls:** It allows you to view the Service to Service calls between any two SOA entities (Composites/J2EE Applications/OSB/BPEL 10g instances) by clicking the link between the entities.
- **Dependency Highlighting:** It allows you to view the dependencies for any service. If you click on a service, all the services that it is dependent on and vice versa are highlighted.
- **Database Associations:** Shows all the databases used by the SOA Composites, BPEL 10g instance and J2EE applications.
- **External Services:** Shows the services that are used by a SOA Composite application but are external to it or not managed by Enterprise Manager Cloud Control.

For more details, refer to the Enterprise Manager Online Help.

13.11 Publishing a Server to UDDI

To publish a service to UDDI, navigate to the Services and References Home page, select a service from the table and click **Publish to UDDI** from the menu. The Publish Service to UDDI window is displayed with the following fields:

- **Service Name:** The name of Web Service to be published to the UDDI Registry. This is a Read Only field.
- **Service Description:** The description of the selected Web Service.
- **Service Definition Location:** The URL location of the Service Definition. This is a Read Only field.
- **UDDI Source:** A logical name for an external UDDI registry source. Select the UDDI Source from the drop-down list.
- **Business Name:** The name of the data structure in the UDDI registry. Select a Business Name that has been registered with the UDDI from the list.

Click **OK** to start the process that publishes the web service to UDDI or click **Cancel** to cancel publishing the service.

13.12 Generating SOA Reports Using BI Publisher

You can use Enterprise Manager to print SOA reports using BI Publisher Enterprise Reports. Oracle Business Intelligence (BI) Publisher is an enterprise reporting solution for authoring, managing, and delivering highly formatted documents. Oracle BI Publisher also allows you to build custom reporting applications that leverage existing infrastructure. Reports can be designed using familiar desktop products and viewed online or scheduled for delivery to a wide range of destinations.

The following table describes the SOA reports that can be generated using BI Publisher:

Table 13–4 SOA Reports

SOA Report	Description
SOA Infrastructure Performance Report	<p>The SOA Infrastructure Performance Summary Report provides information about the average response time, error rate, throughput, system faults, business faults, web service policy violation faults for selected SOA Composite. It displays a chart with the top 5 SOA Composites based on average response time or throughput across the selected SOA composites for specified time period.</p> <p>The report can be sorted based on performance metric (average response time) or the usage metric (instance count). As part of the report parameters setting, you can use options that allow you to select the SOA Composite by Partitions or by selecting individual composites.</p>
SOA Composite Detailed Performance Report	<p>The SOA Composite Detailed Performance Summary Report provides information about the average response time, error rate, throughput, system faults, business faults, web service policy violation faults for each selected composite assembly part such as service, reference, and service component. This is an in-depth report that provides complete details about the each assembly part in the SOA Composite.</p> <p>It displays a chart with the top 5 SOA Composites based on average response time or throughput across the selected SOA Composites for a specified time period. The report can be sorted based on performance metric (average response time) or the usage metric (instance count).</p> <p>As part of the report parameters setting, you can use options that allow you to select the SOA Composite by Partitions or by selecting individual composites.</p>
Top 5 SOA Composites (From Dehydration Store)	<p>This report shows how the SOA Composites have been performing over a period of time. Charts listing the top 5 SOA composites are displayed and critical metric data for all the SOA composites are displayed in a table.</p>

To print SOA reports using BI Publisher, follow these steps:

1. From the **Enterprise** menu, select **Reports**, then select **BI Publisher Enterprise Reports**.
Enterprise Manager Cloud Control displays the login page for BI Publisher Enterprise Reports.
2. Enter your credentials to log into BI Publisher.
3. The BI Publisher Enterprise page displays, showing you Recent reports, Others, and Favorites. You can use this page to create a new report, submit a report job, and perform other tasks.
4. Click the Report you want to view.
5. You can select different filters such as SOA Composite Name, Partition Name, Date Range, and so on to view the report. You can also select a Sort By option to sort the report on Composite Name, Sorted Instances, and so on.
6. You can refresh the report anytime by clicking the **Refresh** icon on the upper right side of the SOA Report tab. You can hide or display the search parameters by clicking the Parameters icon. You can choose to view the report in various formats such as HTML, PDF, RTF, Excel, and PowerPoint by clicking the **View Report** icon. Likewise you can display more available actions by clicking the **Actions** icon. For more help about using BI Publisher, click the help icon.

13.13 Generating SOA Reports Using Information Publisher

This section describes the procedure to create SOA Reports.

1. From the Targets menu, select **Middleware**, and click on a SOA Infrastructure target. The SOA Infrastructure Home page appears.
2. From the SOA Infrastructure menu, select the **Information Publisher Reports**.
The out-of-box SOA reports are displayed under the SOA Performance Reports section.
3. Select a report from the section (for example, you can select **Pending Instance Statistics**) and click **Create Like**. The Create Report Definition page is displayed.
4. In the General page, enter the following details:
 - a. Enter the BPEL Process Name as the title.
 - b. Click the Set Time Period to set the time interval for the report.
 - c. Click the **Run report using target privileges of the Report Owner (SYSMAN)** check box in the Privileges section.
5. Click the **Elements** tab and click the **Set Parameters** icon for the Pending Instance Statistics Element in the table.
6. In the Set Parameters page, click the torch icon to select a Composite Name. The Result Set Size with default values for the Pending Instance Statistics report is displayed.
7. Select a Component Name from the list, enter the Result Set Size and click **Continue** to return to the Elements page.
8. The selected target name is displayed in the Elements table.
9. To schedule periodic report generation, click the **Schedule** tab.
10. Specify the schedule type and other details and click **OK**.
11. You will return to the Report Home page where the newly scheduled report is displayed in the table. Click the report name to view the details.

13.14 Provisioning SOA Artifacts and Composites

The SOA Artifacts Deployment Procedure allows you to:

- Provision SOA Artifacts from a reference installation or from a gold image
- Create a gold image of the SOA Artifacts
- Provision SOA Composites either from the Software Library or from another accessible location.

For more details on the SOA Artifacts Deployment Procedure, see the *Enterprise Manager Administrator's Guide for Software and Server Provisioning and Patching*.

13.15 Diagnosing Issues and Incidents

To diagnostic data for problems and incidents, access the Support Workbench page. To do so, navigate to the SOA Infrastructure Home page, and from the **SOA Infrastructure** menu, select **Diagnostics**, then select **Support Workbench**.

Enter the credentials for the host on which the WebLogic server is running and the WebLogic credentials for the WebLogic server. Click **Continue** to log into the Support Workbench page. On this page, you can do the following:

- View problem or incident details.
- View, create, or modify incident packages.
- View health checker findings
- Close resolved problems.

13.16 Verifying Target Monitoring Setup

As a prerequisite, verify the target monitoring setup before you perform any operations on the SOA infrastructure. Use the Target Setup Verification page to run a series of diagnostic scans and verify if you have met all functional as well as system-level prerequisites required for monitoring targets in Enterprise Manager. This helps you discover and repair all target monitoring setup-related issues beforehand.

This section describes the following:

- [Running Functionality-Level Diagnostic Checks](#)
- [Running System-Level Diagnostic Checks](#)
- [Repairing Target Monitoring Setup Issues](#)

Note: If the following error is displayed while trying to access the Target Setup Verification page, then you must upgrade the Management Agent version to 12.1.0.3 or higher:

Current agent version is not supported for Functionality Check scan. Upgrade to 12.1.0.3.0 agent version or higher for performing this scan.

Note: You will not be able to click on torch icon available next to database system field in Dehydration Store repair pop-up if association exists between Database System and SOA Infrastructure. It is enabled only when association is missing. When it is missing, you can select appropriate database system target from target selector popup. Pop-up can be launched by clicking on the torch icon.

13.16.1 Running Functionality-Level Diagnostic Checks

To run diagnostic scans on the functionalities associated with an Enterprise Manager target and identify any setup issues, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the target you are interested in. For example, SOA Infrastructure.
3. On the Home page of the target, from the target-specific menu, select **Target Setup**, and click **Verification**.
4. On the Target Setup Verification page, in the Functionality Check section, click **Scan**.
5. If setup problems are detected, repair them. See [Section 13.16.3](#).

13.16.2 Running System-Level Diagnostic Checks

To run diagnostic scans on the system components that monitor an Enterprise Manager target and check their availability rate, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the target you are interested in. For example, SOA Infrastructure.
3. On the Home page of the target, from the target-specific menu, select **Target Setup**, and click **Verification**.
4. On the Target Setup Verification page, in the System Check section, click **Scan**.

Note:

- The **Availability(%)** column shows the availability rate of the target types. Click the rate to drill down and view more details.
 - The **Time since last collection (min)** column shows the total time elapsed since the last metric collection for the target.
-

13.16.3 Repairing Target Monitoring Setup Issues

To repair any target monitoring setup-related issues, follow these steps:

1. Run functionality-level diagnostic scan to identify setup-related issues. See [Section 13.16.1](#).
2. If setup issues are found, in the Functionality Check section, click the repair icon against the functionality that requires to be fixed.

In the dialog that appears, enter the required details, and click **Re-Scan and Save** to validate the credentials, run the functionality check again, and save the details in Enterprise Manager. If you are sure the credentials are correct, then click **Save** to save the details without running the check again.

Note:

- The host credentials you are expected to provide are credentials of the host where the Management Agent, which monitors the SOA Infrastructure, is running.
- While repairing dehydration store issues, you are required to provide SOA dehydration store configuration details such as the database system and the SOA repository credentials.

If the configuration information has been collected, and if the association between the database system and the SOA infrastructure already exist, then the database system is prepopulated by default, and you need to enter only the credentials of the SOA repository. Otherwise, click the torch icon and manually select the database system with which the SOA infrastructure communicates, and enter the credentials of the SOA repository.

The connection descriptor is prepopulated by default, is an editable field, and appears in multiple rows if its an Oracle RAC database. Do not modify the descriptor unless you want to correct it.

The data source type is displayed only if the database system is an Oracle RAC database. The data source type can be either Multi Data Source or GridLink Data Source. Also note, that the data source type appears as 'NA' if details of the database system do not match with the connection descriptor.

13.17 Searching Faults in the SOA Infrastructure

This section describes how you can search faults in the SOA infrastructure. In particular, you can perform the following tasks:

- [Overview of Faults and Fault Types in SOA Infrastructure](#)
- [Overview of the Recovery Actions for Resolving Faults](#)
- [Prerequisites for Searching, Viewing, and Recovering Faults](#)
- [Searching and Viewing Faults](#)
- [Recovering a Few Faults Quickly \(Simple Recovery\)](#)

13.17.1 Overview of Faults and Fault Types in SOA Infrastructure

The following are the types of SOA composite application faults displayed in Enterprise Manager Cloud Control:

- **Business:** Application-specific faults that are generated when there is a problem with the information being processed (for example, a social security number is not found in the database).
- **System:** Network errors or other types of errors such as a database server or a web service being unreachable.
- **Oracle Web Service Manager (OWSM):** Errors on policies attached to SOA composite applications, service components, or binding components. Policies apply security to the delivery of messages.

The following are the categories of SOA composite application faults in Enterprise Manager Cloud Control:

- **Recoverable**
 - Business faults and some specific system faults
 - Oracle Mediator input file path and output directory mismatch
 - An Oracle BPM Worklist user is not authorized to perform relevant (expected) actions
- **Nonrecoverable**
 - Rejected messages
 - Most system faults
 - Non-existent references
 - Service invocation failures
 - Policy faults
- **Rejected Messages**

A fault is classified as a rejected message based on where it occurs. If a fault occurs before entering a SOA composite, without generating a composite instance, it is classified as a rejected message. A system or a policy fault can be identified as a rejected message.

13.17.2 Overview of the Recovery Actions for Resolving Faults

Recovery actions enable you to recover or resolve the SOA composite application faults. The following describes the recovery actions supported for different SOA engines.

Table 13–5 Overview of the Recovery Actions for Resolving Faults

Recovery Action	Description	Applicable To SOA Engine Type
Retry	Retries the instance directly. An example of a scenario in which to use this recovery action is when the fault occurred because the service provider was not reachable due to a network error. The network error is now resolved.	<ul style="list-style-type: none"> ■ BPEL ■ BPMN ■ Mediator
Abort	Terminates the entire instance.	<ul style="list-style-type: none"> ■ BPEL ■ BPMN ■ Mediator
Continue	Ignores the fault and continues processing (marks the faulting activity as a success).	<ul style="list-style-type: none"> ■ BPEL ■ BPMN
Rethrow	Rethrows the current fault. BPEL fault handlers (catch branches) are used to handle the fault. By default, all exceptions are caught by the fault management framework unless an explicit rethrow fault policy is provided.	<ul style="list-style-type: none"> ■ BPEL ■ BPMN
Replay	Replays the entire scope again in which the fault occurred.	<ul style="list-style-type: none"> ■ BPEL ■ BPMN

13.17.3 Prerequisites for Searching, Viewing, and Recovering Faults

Meet the following prerequisites before searching, viewing, and recovering SOA composite application faults:

- Ensure that the SOA infrastructure is monitored by an Oracle Management Agent (Management Agent) that is running on the same host as the SOA infrastructure. At the moment, there is no support for searching, viewing, and recovering faults from a SOA infrastructure that is monitored by a remote Management Agent.
- Set the following as preferred credentials. These credentials can be set from the Target Setup Verification page. To do so, from the **Targets** menu, select **Middleware**. On the Middleware page, click the target you are interested in. For example, SOA Infrastructure. On the Home page of the target, from the target-specific menu, select **Target Setup**, and click **Verification**:
 - Credentials of the host on which the SOA server is running.
 - Administrator credentials of the Oracle WebLogic Domain.

13.17.4 Searching and Viewing Faults

To search and view SOA composite application faults, follow these steps:

1. Meet the prerequisites. See [Section 13.17.3](#).
2. From the **Targets** menu, select **Middleware**.
3. On the Middleware page, click the SOA Infrastructure target.
4. On the SOA Infrastructure target page, click **Faults and Rejected Messages**.
5. In the Faults and Rejected Messages tab, set the search criteria. See [Section 13.17.4.1](#).
6. Click **Search**.
7. View the faults:
 - To know the total faults in the SOA infrastructure, see **Total Faults in SOA Infrastructure**, which is placed in the footer of the results table.
 - To know the number of faults displayed in the table (out of the total number of faults in the SOA infrastructure), see **Displayed Faults**, which is placed in the footer of the results table.
 - To view details of each fault, see the results table.
 - To hide or unhide columns in the table, from the **View** menu, select **Columns**, then select the column name you want to hide or unhide.
 - To filter or perform a fine search for a particular column, enter a search keyword in the textbox placed above the column header. See [Section 13.17.4.7](#).
For example, to filter and list all faults related to the BPEL engine type, in the **Engine Type** column, type `bpel`.
 - To sort the fault details alphabetically, click the column header based on which you want to sort the details.
 - To find out the number of rows to which the search results have been restricted, see the note below the table.

For example, the following note appears if the rows were restricted to 20.

This table of search results is limited to 20 fault instances. Narrow the

results by using the search parameters.

13.17.4.1 Setting Search Criteria

To search for faults and messages, enter details as described in the following table, and click **Search**.

Table 13–6 *Setting Search Criteria*

Field	Description
Composite	Select the composite for which you want to search faults.
Error Message Contains	Enter the keywords you are looking for in the error messages of the faults so that faults with only such keywords can be searched. The value can be an alphanumeric value, and can have special characters, including wildcard characters. The value must not exceed 100 characters.
Fault Id	Enter the ID of the fault you want to search. The value must be a numeric value.
ECID	Enter the Execution Context ID so that only faults occurring within that particular context ID can be searched. The value must be a numeric value.
Fault Time From	Select the date and time from when you want to search for faults. By default, the date and time when you access the page appears. For example, if you select 02/25/2013 02:45:32 AM, then all faults that occurred from this date and time will be searched for in the selected composite.
Fault Time To	Select the date and time till when you want to search for faults. For example, if you select 02/28/2013 02:45:32 AM, then all faults that occurred till this date and time will be searched for.
Fault Type	To search for all types of faults, select All Faults . To search for a particular type of fault, select one of the following: <ul style="list-style-type: none"> ■ System Faults, indicate all network errors or other types of errors such as a database server or a web service being unreachable. ■ Business Faults, indicate all application-specific faults that were generated when there was a problem with the information processed (for example, a social security number is not found in the database). ■ OWSM Faults, indicate Oracle Web Service Manager errors on policies attached to SOA composite applications, service components, or binding components. Policies apply security to the delivery of messages.
Component Type	To search for faults in all the component types, select All . To search for faults in a particular component type, select an appropriate component type: BPEL, BPMN, Mediator, Business Rules, or Human Workflow.

Table 13–6 (Cont.) Setting Search Criteria

Field	Description
Fault State	<p>To search for faults in any state, select All. To search for faults in a particular state, select one of the following:</p> <ul style="list-style-type: none"> ▪ Recoverable, indicates business faults and some specific system faults, Oracle Mediator input file path and output directory mismatch faults, and other faults related to Oracle BPM Worklist, where the user is not authorized to perform any relevant (expected) actions. ▪ Nonrecoverable, indicates rejected messages, most system faults, non-existent references, service invocation failures, and policy faults.
Instance ID	Enter the composite instance ID from which you want to search faults. The value must be a numeric value.
Restrict Search Result (rows)	<p>By default, the search results are restricted to 500 rows in the table. If you want to modify this limit or restriction, enter a suitable value.</p> <p>The highest value you can enter as the limit depends on the limit set on the OMS. When no limit is set on the OMS, the limit that is honored by default is 2000, so the default range you can enter in the Restrict Search Result (rows) field is 1 to 2000.</p> <p>To modify this maximum limit set on the OMS, run the following command: <code>emctl set property -name oracle.sysman.core.uifwk.maxRows -value <max_limit_value></code></p> <p>Note: The higher the value you set as the limit, the longer the time it takes to retrieve the faults, and that entering a higher value than the default in Restrict Search Result (rows) can lead to longer time to get the faults, and hence a longer load time.</p>

13.17.4.2 Finding Total Faults in the SOA Infrastructure

To find the total faults in the SOA infrastructure, follow these steps:

1. Search for faults in the SOA infrastructure. See [Section 13.17.4](#).
2. Once the search results appear, see **Total Faults in SOA Infrastructure**, which is placed at the bottom-right corner, below the table.

Note: While retrieving the total faults in the SOA infrastructure, the **Restrict Search Result (rows)** field in the search criteria is not considered. For example, if there are a total of 700 faults, and if you enter 500 for this field, then the search is performed to list only 500 faults in the table, but the **Total Faults in SOA Infrastructure** field displays 700.

13.17.4.3 Limiting Faults Searched and Retrieved from the SOA Infrastructure

When you search for faults in the SOA infrastructure, the search might result in numerous faults. By default, the search results are restricted to 500 rows in the table. However, you can choose to modify this limit if you want.

To modify the limit, set the **Restrict Search Result (rows)** field to a suitable value while setting the search criteria (see [Section 13.17.4.1](#)). Then search.

The highest value you can enter as the limit depends on the limit set on the OMS. When no limit is set on the OMS, the limit that is honored by default is 2000, so the default range you can enter in the **Restrict Search Result (rows)** field is 1 to 2000.

To modify the maximum value set on the OMS, run the following command:

```
emctl set property -name oracle.sysman.core.uifwk.maxRows -value <max_limit_value>
```

Caution: The higher the value you set as the limit, the longer the time it takes to retrieve the faults, and that entering a higher value than the default in Restrict Search Result (rows) can lead to longer time to get the faults, and hence a longer load time

13.17.4.4 Searching Only Recoverable Faults

There might be numerous faults in the SOA infrastructure, but you can search and view only the recoverable faults. For example, there might be 700 faults in total, but there maybe only 550 recoverable faults; you can search and list only those 550 faults if you want.

To search only for recoverable faults, while searching for faults (see [Section 13.17.4](#)), set the search criteria with the **Fault State** list set to **Recoverable**. If you set it to **All**, then faults that are recoverable and not recoverable are searched and listed.

13.17.4.5 Searching Faults in a Particular Service Engine

There might be faults across various service engines such as BPEL, BPMN, Mediator, Business Rules, and Human Workflow. You can search and view only faults occurred in a particular service engine.

To search for faults in a particular service engine, set the search criteria with the **Component Type** list set to a particular service engine of interest (see [Section 13.17.4.1](#)). Then search.

13.17.4.6 Searching Faults by Error Message

There might be numerous errors in the SOA infrastructure, but you might be interested only in those errors that contain some keywords of your interest. For example, you might be interested only in errors that contain the word `ORAMED`. You can search and view faults with such keywords.

To search faults by error messages, set the search criteria with the **Error Message Contains** field set to some keywords of your interest (see [Section 13.17.4.1](#)). Then search.

Note:

- By default, the entered keywords are searched anywhere in the error message.
 - The keywords you enter are case sensitive.
 - The only wildcard character permitted is `%`, which signifies all or anything after, before, or between two keywords. For example, `BPEL%fault` will result in faults with the error message `BPEL is a fault`.
-

13.17.4.7 Filtering Displayed Search Results

When you set the search criteria and search for faults in the SOA infrastructure, and when the search results appear in the results table, you can filter the search results further to show only those rows or fault instances that interest you, based on a keyword entered in the column header.

For example, from the displayed fault instances, to filter and view only the *bpel* service engine's results, enter the keyword *bpel* in the textbox placed above the **Component Type** column header. This is essentially the value shown in the *bpel* fault instance row for the **Component Type** column.

To filter the displayed search results, follow these steps:

1. Search for faults in the SOA infrastructure. See [Section 13.17.4](#).
2. Once the results appear in the table, in the textbox placed above the header of the column you want to filter, enter a search keyword.

For example, to filter and list all faults related to the BPEL engine type, in the textbox placed above the **Engine Type** column header, type *bpel*.

13.17.5 Recovering a Few Faults Quickly (Simple Recovery)

To recover only a few SOA composite application faults quickly, follow these steps:

1. Meet the prerequisites. See [Section 13.17.3](#).
2. From the **Targets** menu, select **Middleware**.
3. On the Middleware page, click the SOA Infrastructure target.
4. On the SOA Infrastructure target page, click **Faults and Rejected Messages**.
5. In the Faults and Rejected Messages tab, set the search criteria. See [Section 13.17.4.1](#).
6. Click **Search**.
7. In the table, select one or up to 5 faults at a time, and from the **Recovery Options** menu, select an appropriate recovery action that matches your requirement. For information on the recovery actions, see [Section 13.17.2](#).
8. Enterprise Manager displays an informational message with one of the following mentioned to confirm whether or not it can submit the recovery job successfully. Click **OK**, and take the necessary action if required.
 - If you have selected more than 5 faults, then the recovery job is not submitted. Select 5 or fewer faults, and try again. Alternatively, select 5 or more, and try a bulk recovery. See [Section 13.18.1](#).
 - If there are no recoverable faults, then the recovery job is not submitted.
 - If there are faults that are recoverable and not recoverable, then the recovery job is submitted only for recoverable jobs. You can track the recover job. See [Section 13.18.2](#).
9. Perform Step (1) to Step (5) again to verify if the faults you selected for recovery still appear in the search results. If they do not appear, then the recovery operation for those faults has been successfully submitted.

13.18 Creating and Tracking Bulk Recovery Jobs

This section describes the following:

- [Recovering Faults in Bulk \(Bulk Recovery\)](#)
- [Tracking Bulk Recovery Jobs, and Viewing Their Results and Errors](#)
- [Creating Bulk Recovery Jobs Using EMCLI and Web Services](#)

13.18.1 Recovering Faults in Bulk (Bulk Recovery)

To directly recover a large number of faults from the SOA database, follow these steps to perform a bulk recovery:

1. Meet the prerequisites. See [Section 13.17.3](#).
2. Access the Create Bulk Recovery or Bulk Recovery Jobs page in one of the following ways. *(If you are already on one of these pages, ignore this step.)*
 - Directly to the Bulk Recovery Jobs page. See [Section 13.18.1.1](#).
 - From the SOA Infrastructure Faults and Rejected Messages tab. See [Section 13.18.1.2](#).
3. On the Create Bulk Recovery page, in the Composites section, click **Add** to add additional composites for which faults must be searched. You can add up to 10 composites.

Note: You can add only up to 10 composites.

4. In the Fault Details section, set the details of the faults you want to recover. See [Section 13.18.1.3](#).
5. In the Recovery Options section, set the recovery and batch parameters. See [Section 13.18.1.4](#).
6. In the Job Parameters section, schedule the bulk recovery job. See [Section 13.18.1.5](#).
7. To verify the number of faults that will be recovered for the given criteria, click **Estimate Faults**.

A pop-up appears informing you of the total number of faults in the SOA Infrastructure based on the criteria you have set. Based on the count, you can decide whether or not you want to proceed. If required, you can adjust the settings. For example, you can modify the fault time period.

8. If you are ready to perform the bulk recovery, click **Submit**.
9. Track the status of the bulk recovery job. See [Section 13.18.2](#).
10. Search for faults again to verify if the faults you selected for recovery still appear in the search results. If they do not appear, then the recovery operation for those faults has been successfully submitted. See [Section 13.17.4](#).

13.18.1.1 Navigating to the Bulk Recovery Jobs Page Directly

To navigate directly to the Bulk Recovery Jobs page, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the SOA Infrastructure target.
3. On the SOA Infrastructure target page, from the **SOA Infrastructure** menu, select **Fault Management**, then select **Bulk Recovery**.
4. On the Bulk Recovery Jobs page, click **Create Job**.

13.18.1.2 Navigating to the Create Bulk Recovery Page From SOA Infrastructure

To navigate to the Create Bulk Recovery page from the SOA Infrastructure Faults and Rejected Messages tab, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the SOA Infrastructure target.
3. On the SOA Infrastructure target page, click **Faults and Rejected Messages**.
4. In the Faults and Rejected Messages tab, set the search criteria. See [Section 13.17.4.1](#).
5. Click **Search**.
6. In the table, select one or more faults, and click **Bulk Recover**.
7. If you want to automatically include other faults, which you have not explicitly selected in the table, but belong to the same composite as the faults you have selected in the table, then in the Populate Options dialog, select **Same composites as selected faults**.

Enterprise Manager automatically identifies all other faults belonging to the same composites as the selected faults, and lists them in the next page for bulk recovery.

13.18.1.3 Setting Fault Details for Recovering Faults in Bulk

To set the fault details while recovering faults in bulk, follow these steps:

1. On the Create Bulk Recovery page, in the Fault Details section, select an engine type and enter a keyword you are looking for in the error messages so that only faults with such error messages are recovered.
2. From the **Time Period** list, select a time period for which the faults must be recovered. The default is **Last 1 Day**, but if you want to customize the time period, you can select **Custom** from the list, and enter the start and end time stamp.

Note: If you schedule a repeating job (see [Section 13.18.1.5](#)) with a custom fault time period, then all executions of the job use the same fault time period to recover faults. The time zone of the fault time you set is the time zone of the host where the Oracle Management Agent, which monitors the target, is installed. This information is displayed on the top-right corner of the Home page for any target.

13.18.1.4 Setting Recovery and Batch Details for Recovering Faults in Bulk

To set the recovery and batch details for recovering faults in bulk, follow these steps:

1. On the Create Bulk Recovery page, in the Recovery Options section, from the **Recovery Action** list, select a recovery action. For information on recovery actions, see [Section 13.17.2](#).
2. By default, **Batch by Fault Time** is enabled so that faults can be grouped into multiple, smaller units or batches based on the time they were created, and run sequentially. Oracle recommends that you keep the option enabled to simplify the fault recovery process. However, if you do not want to create batches for some reason, then deselect this option.
3. If you keep the **Batch by Fault Time** option enabled, then do the following:
 - a. By default, the batches are created with faults that occurred within every 10 minutes. If you want to change this time period, then enter a value in minutes

in the **Batch Time Period** field. The minimum time period is 5 minutes and the maximum time period is 360 minutes.

- b. By default, the delay time between two batches is set to 300 seconds. If you want to change this delay time, then enter a value in seconds in the **Delay between batches (secs)** field. The minimum delay time is 5 seconds and the maximum delay time is 900 seconds.

Example

You might want to recover faults that occur from 02/08/2013 2:00 am to 02/09/2013 2:00 am, in batches of 1 hour, with a delay of 300 seconds between each batch.

To do so, while setting the fault details (see [Section 13.18.1.3](#)), select **Custom** from the **Time Period** list, and set **Fault Start Time** to 02/08/2013 2:00 am and **Fault End Time** to 02/09/2013 2:00 am. Then, while setting the recovery and batch details (as described in this topic), set **Batch Time Period (mins)** to 60 and **Delay between batches (secs)** to 300.

Based on these settings, the first batch starts recovering faults that occurred on 02/08/2013 between 2:00 am and 3:00 am. Then, there is a delay of 300 seconds. And then, the next batch starts recovering faults that occurred on 02/08/2013 between 3:00 am and 4:00 am. This pattern of recovery continues until the final batch recovers faults that occurred on 02/09/2013 between 1:00 and 2:00 am.

13.18.1.5 Scheduling Bulk Recovery Jobs to Run Once or Repeatedly

To schedule bulk recovery jobs, on the Create Bulk Recovery page, in the Job Parameters section, select one of the following options:

- To run the jobs only once, select one of these options:
 - **Immediately**, if you want to run the job immediately.
 - **Later**, if you want to run the job just once, at a schedule date and time, and not immediately.
- To run the jobs repeatedly at a set frequency, from the **Repeat** list, select an appropriate option, and set the frequency.

Note: For a repeating job, if you set a custom fault time period (see [Section 13.18.1.5](#)), then all executions of the job use the same fault time period to recover faults. The time zone of the fault time you set is the time zone of the host where the Oracle Management Agent, which monitors the target, is installed. This information is displayed on the top-right corner of the Home page for any target.

- To set a grace period, select **Do not run if it cannot start within**, and set an appropriate grace period.

A grace period is a period of time that defines the maximum permissible delay when attempting to run a scheduled job. If the job system cannot start the execution within a time period equal to the scheduled time + grace period you set, then it skips the job. By default, all jobs are scheduled with indefinite grace periods.

Example

You might want to recover faults that occur from 02/08/2013 2:00 am to 02/09/2013 2:00 am, in batches of 1 hour, with a delay of 300 seconds between each batch, repeatedly every 2 days, starting from 02/10/2013 10:00 am to 02/14/2013 10:00 am.

To do so, while setting the fault details (see [Section 13.18.1.3](#)), select **Custom** from the **Time Period** list, and set **Fault Start Time** to 02/08/2013 2:00 am and **Fault End Time** to 02/09/2013 2:00 am. Then, while setting the recovery and batch details (as described in this topic), set **Batch Time Period (mins)** to 60 and **Delay between batches (secs)** to 300.

Then, while scheduling the recovery job, select **Repeating**, set **Every** to **2** and select **Days**, then set **Starting** to 02/10/2013 10:00 am and **Ending** to 02/14/2013 10:00 am.

Based on these settings, on 02/10/2013 10:00 am, the first batch starts recovering faults that occurred on 02/08/2013 between 2:00 am and 3:00 am. Then, there is a delay of 300 seconds. And then, the next batch starts recovering faults that occurred on 02/08/2013 between 3:00 am and 4:00 am. This pattern of recovery continues until the final batch recovers faults that occurred on 02/09/2013 between 1:00 and 2:00 am. This recovery job starts all over again on 02/12/2013 10:00 am, and repeats itself on 02/12/2013 10:00 am.

13.18.2 Tracking Bulk Recovery Jobs, and Viewing Their Results and Errors

To track bulk recovery jobs and view their results and errors, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the SOA Infrastructure target.
3. On the SOA Infrastructure target page, from the **SOA Infrastructure** menu, select **Fault Management**, then select **Bulk Recovery**.
4. On the Bulk Recovery Jobs page, you can view details such as the job name, the date and time when it ran or is scheduled to run, the user who scheduled the job, the current status of the job, the faults that were recovered and not recovered.

Note:

- This page lists only those jobs that ran in the last three days, and only for the current user. The list also include jobs submitted via EM Command Line Interface (EM CLI).
 - The Recovered Faults column displays the number of faults for which the recovery has been attempted by the SOA Suite so far for the job. The Not Recovered Faults column displays the number of faults for which the recovery could not be attempted by the SOA Suite due to some errors.
-

5. Click the name of the job for which you want to view more details such as its actual results, the job failure errors, and the recovery errors.

Enterprise Manager displays the Bulk Recovery Job Details page that provides the following information.

Section Name	Description
Results	<p>Provides the status of the recovery job, essentially details such as the composite selected for recovery, the ID of the faults selected for recovery, and the recovery attempt status of the faults, which can be either Recovered or Not Recovered.</p> <p>Note: The recovery status indicates only the recovery attempt status, and not the actual recovery status of the fault. To know the actual recovery status, search for the fault ID. How?</p>
Job Failure Error	<p>Provides details of the failed recovery jobs. The errors are shown from the last point of failure of the job.</p> <p>The details include:</p> <ul style="list-style-type: none"> ■ Failed Step Name, the name of the step of the recovery job that failed. The job has two steps, mainly <i>pre-check</i> and <i>recover_faults</i>. The <i>recover_faults</i> step runs once for every batch of the job, or only once if batching is not enabled. ■ Failed Step Output, the output of the failed job step. ■ Parsed Error Message, any known error message that is as part of the step output. ■ Composite, the name of the composite selected for the recovery job. ■ Fault Time From, the start date and time from when faults occurred and for which the recovery job was submitted ■ Fault Time To, the end date and time till when faults occurred and for which the recovery job was submitted ■ Error Details, the error indicating that the status of the recovery job could not be retrieved. This means that the recovery was attempted for the given composite and time period, but there was a time-out while retrieving the status. <p>To verify if the faults were recovered, search for the faults for the given composite and time period again. How?</p> <p>To run the recovery job again, reduce the batch time period to a value lower than the value you entered earlier (see Section 13.18.1.4), and submit the bulk recovery job again (see Section 13.18.1).</p>

13.18.3 Creating Bulk Recovery Jobs Using EMCLI and Web Services

You can create Bulk Recovery Jobs for a SOA Infrastructure Target from the command line using EM CLI, from EM Job Systems using Web-Service Interface, or from Cloud Control UI.

This section contains the following sections:

- [Creating Bulk Recovery Jobs Using EMCLI](#)
- [Viewing the Submitted Jobs and Outputs Using EMCLI](#)
- [Creating Bulk Recovery Jobs through Web-Service](#)

13.18.3.1 Creating Bulk Recovery Jobs Using EMCLI

In EM CLI, use EM job system's `get_jobs` command to submit bulk recovery jobs. The inputs to the job are supplied using a properties file.

To create the bulk recovery job using EMCLI, follow these steps:

1. Log in to EMCLI. For example:

```
emcli login -username=sysman
```

2. Find out the input parameters to be entered in the property file to run the bulk recovery job. To do so, run the following command:

```
emcli describe_job_type -type=SOABulkRecovery
```

3. Use any editor to open the properties file, and provide your inputs. You can then save and close the properties file.

Using any editor, create a new text file. For example, temp.properties

Here's a sample Property File:

```
target_list=<soa-infra target name>:oracle_soainfra
variable.CompositeList=<composite1 target name>, <composite 2 target name>
variable.BatchDelay=300
variable.BatchSize=10
variable.EnableBatching=1
variable.EngineType=BPEL
variable.ErrorMsg=xxxx
variable.FaultStartTime=01-01-2013 00:00:00 PST
variable.FaultEndTime=01-02-2013 00:00:00 PST
variable.FaultTimePeriod=Custom
variable.RecoveryAction=Continue
```

Note: Currently, Oracle supports only one SOA-Infrastructure target to be entered in the *target_list* property.

4. Run the following command to submit a bulk recovery job with the updated property file as an input:

```
emcli create_job -name=bulk522 -job_type=SOABulkRecovery -input_file=property
file:/tmp/temp.properties
```

5. Set the preferred credentials or named credentials for the WebLogic Domain and SOA Server Host. By default, the job uses the preferred credentials, that is, WebLogic Administrator Credentials for the WebLogic domain and Normal Host Credentials for the SOA Server hosts.

To set the preferred credentials, run the following commands:

Setting WebLogic Domain Credentials:

```
emcli set_preferred_credential -target_type=weblogic_domain -target
name=<weblogic domain target name> -set_name=WLCreditsNormal -credential
name=<existing named credential name> -credential_owner=<user>
```

Setting SOA Host Credentials:

```
emcli set_preferred_credential -target_type=host -target_name=<host target
name> -set_name=HostCreditsNormal -credential_name=<existing named credential
name> -credential_owner=<user>
```

Alternately, you can override the preferred credentials by supplying the named credentials as an input to the property file for the current submission.

Following example describes how to set the named credentials for the WebLogic Domain and SOA Server host:

```
target_list=<SOA-Infra TargetName>:oracle_soainfra
cred.SOAAgentHostCred.<slc01nbo.us.oracle.com>:<host>=NAMED:xxxx
cred.SOADomainCreds.<target_name>:<target_type>=NAMED:xxxx
```

13.18.3.2 Viewing the Submitted Jobs and Outputs Using EMCLI

The following table describes certain other operations that can be performed using EMCLI commands.

Table 13–7 EMCLI Commands For Bulk Recovery

EMCLI Command	Description	Example
get_jobs	This EMCLI command to view all the Bulk Recovery Jobs that have been submitted.	emcli get_jobs -targets=<SOA-Infra target name>:oracle_soainfra -format=name:csv grep BULK521
get_job_execution_detail	<p>This EMCLI command to view the output of the Bulk Recovery job execution. To view the details of the job steps, you need to supply the Execution ID of the job.</p> <p>Note: The output of the job that is displayed using the EMCLI command is unstructured. For a complete and structured report of the output, log in to Enterprise Manager Cloud Control. From Enterprise menu, select Job, and then click Activity. On the Job Activity Page, in the Advanced Search region, enter the name of the job, and then click Go. Select the job, and drill down to the steps by click Expand All.</p>	<p>Run the following command to get the Execution ID of the job:</p> <pre>emcli get_jobs -targets=<SOA-Infra target name>:oracle_soainfra -format=name:csv grep BULK521</pre> <p>Use the Execution ID in the following command to view the details of the job submitted:</p> <pre>emcli get_job_execution_detail -execution=D4081BAB8942F246E040F00A5AA93E04 -xml -showOutput</pre>

13.18.3.3 Creating Bulk Recovery Jobs through Web-Service

In addition to EM User Interface and EMCLI, you can also use Web-Service Interface provided by EM job system to create Bulk Recovery Jobs. The web-service interface of the Job System is available by default in an EM installation, and the URL for the WSDL is as follows:

```
<protocol>://<machine>:<port>/em/websvcs/extws/JobControlService?wsdl
```

The EM job system web services are implemented as Simple Object Access Protocol (SOAP) end-points. Client programs can access these end-points using a variety of languages like Java, C++, and Ruby. The web service is used by sending a SOAP request message to one of the end-points, and retrieving the corresponding response message.

Typically, the operations exposed by Job system in the Web-Service Interface is very similar to the EMCLI operations such as `create_job`, `describe_job_type`, and so on.

13.19 Recovering BPEL/BPMN Messages

To find recoverable instances of the BPEL or BPMN Service Engine, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the SOA Infrastructure target.
3. On the SOA Infrastructure target page, from the **SOA Infrastructure** menu, select **Service Engine**, then select **BPEL/BPMN**. Based on the selection, the home page of the service engine is displayed.

4. On the home page, select the **Recovery** tab.
5. To recover messages in which faults occurred, select one or more messages in the table, up to a maximum of 5 messages at a time, and click **Recover**.

Search again to verify if the faults you selected for recovery still appear in the search results. If they do not appear, then the recovery operation for those faults has been successfully submitted.

Note: To mark messages so that they are never delivered, select one or more message in the table, and click **Cancel**.

13.20 Troubleshooting

This section describes the errors you might encounter while discovering the SOA Suite 11g and the workaround steps you can follow to resolve each of them.

This section covers the following:

- [Discovery](#)
- [Monitoring](#)
- [Instance Tracing](#)
- [Recent Faults](#)
- [Fault Management](#)
- [Application Dependency and Performance Integration](#)
- [Information Publisher Reports](#)
- [BI Publisher Reports](#)
- [Systems and Services](#)
- [BPEL Recovery](#)
- [SOA License Issue](#)
- [Dehydration Store Issue](#)

13.20.1 Discovery

The following error occurs when the SOA instances are being discovered.

Table 13–8 Error Message

Error Message	Workaround Steps
New SOA Composite deployed on the SOA Server from JDeveloper are not displayed automatically in Enterprise Manager Cloud Control.	To discover the newly deployed SOA Composites in Enterprise Manager Cloud Control, you must run the Refresh Farm menu option for the associated WebLogic Domain.

13.20.2 Monitoring

The following error occurs when the collection frequency causes a delay in the collection of configuration data.

Table 13–9 Error Message

Error Message	Workaround Steps
All metrics are not displayed.	Enterprise Manager Cloud Control uses the Management Agent to collect metric data. For the first collection, the agent may need 15 minutes to upload the metric data.

13.20.3 Instance Tracing

The following error occurs when the instance is traced.

Instance Search Fails - Same reason as BPEL first column. If Management Agent is down or unreachable.

Table 13–10 Error Message

Error Message	Workaround Steps
Instance Tracing Job Fails	<ol style="list-style-type: none"> 1. Navigate to the Jobs page, and locate the Instance Tracing job (TRACE SOA INSTANCE ID + Instance ID + Submitted time) and view the output to identify the step that has failed. 2. Resolve the issue and run the job again by clicking Retry on the Jobs page. 3. Navigate to the Instance Tracing page to view the trace results. You can also submit a new job by running the Trace Instance option on the Instance Tracing page.

13.20.4 Recent Faults

The following errors occur when:

- All instances with faults are not displayed as only the last 10 values are collected.
- The most recently collected fault instances do not appear in the Faults and Messages page.

Table 13–11 Error Message

Error Message	Workaround Steps
All instances with faults are not populated in Enterprise Manager Cloud Control.	By default, you can only view the latest 10 faults collected during the last 15 minutes. To view additional faults, navigate to Fusion Middleware by clicking the link in the General section on the target Home page.

13.20.5 Fault Management

This section contains the troubleshooting information for fault management:

- [Bulk Recovery](#)
- [Fault Search and Recovery](#)
- [Fault Management and Instance Tracing Errors](#)

13.20.5.1 Bulk Recovery

In general when there is a Bulk Recovery Error, follow these steps to navigate to the page that describes the errors:

1. In Cloud Control, from the **Targets** menu, select **Middleware**.

2. On the Middleware page, click the SOA Infrastructure target.
3. On the SOA Infrastructure target page, from the **SOA Infrastructure** menu, select **Fault Management**, then select **Bulk Recovery**.
4. On the Bulk Recovery Jobs page, select the job that has failed.
5. On the Bulk Recovery Job Details page, in the Job Failure Error section, check the **Parsed Error Message** and **Error Details** fields to understand about the error because of which the job failed.

The following are some of the error messages that you may see in the Parsed Error Message field along with their suggested fixes:

Table 13–12 Error Message

Error Message	Workaround Steps
<pre>java.lang.IllegalArgumentException: Invalid Job Identifier! The specified identifier does not match any valid fault recovery jobs.</pre>	<ol style="list-style-type: none"> 1. Ensure the SOA Infrastructure is up and running. 2. Choose a smaller value for Batch Time Period parameter that is entered while Creating a Bulk Recovery job. This will ensure lesser number of faults are recovered in each batch.
<pre>java.lang.IllegalStateException The results job xxxx are not available because the processing has not yet completed.</pre>	<ol style="list-style-type: none"> 3. Choose Fault time period appropriately excluding the fault time period for which faults have already been recovered by current job. To do so, follow these steps: <ol style="list-style-type: none"> a. The failed job details gives the Composite, Fault Time From and Fault Time To for which the recovery failed. b. Choose new Fault Time From of the new job as the Fault Time To of the failure point of the failed job. 4. Submit another bulk recovery job with same parameters but with the reduced Batch Time Period value, and the new Fault Time From and Fault Time To.
<pre>t3://slc03dms.us.oracle.com:8001 javax.naming.CommunicationException [Root exception is java.net.ConnectException: t3://slc03dms.us.oracle.com:8001 /soa-infra: Destination unreachable; nested exception is: java.net.ConnectException: Connection refused; No available router to destination]</pre>	<p>Ensure that the SOA Infrastructure is up and running, and submit another bulk recovery job with the same parameters.</p>

13.20.5.2 Fault Search and Recovery

The following error occurs when you are unable to connect to the SOA Infrastructure target:

Table 13–13 Error Message

Error Message	Workaround Steps
Error connecting to SOA Infra t3://slc03dms.us.oracle.com:80 01.	Ensure that the SOA Infrastructure is up and running.

13.20.5.3 Fault Management and Instance Tracing Errors

The following errors occur when the SOA database is not functional:

Table 13–14 Error Message

Error Message	Workaround Steps
Error occurred when getting faults Java.rmi.RemoteException: EJB Exception: ; nested exception is: java.lang.RuntimeException: java.lang.RuntimeException: weblogic.jdbc.extensions.PoolD isabledSQLException: weblogic.common.resourcepool.R esourceDisabledException: Pool SOALocalTxDataSource is Suspended, cannot allocate resources to applications.	Ensure the SOA Database is up and running.
t3://slc03dms.us.oracle.com:80 01 javax.naming.CommunicationExce ption [Root exception is java.net.ConnectException: t3://slc03dms.us.oracle.com:80 01/soa-infra: Destination unreachable; nested exception is: java.net.ConnectException: Connection refused; No available router to destination]	Ensure the SOA Database is up and running.
Error occurred when getting faults oracle.sysman.emSDK.agent.comm .exception.ConnectException: Unable to connect to the agent at https://slc03dms.us.oracle.com :3872/emd/main/ [Connection refused]	Ensure the SOA Database is up and running.

13.20.6 Application Dependency and Performance Integration

When you click on the Application Dependency and Performance link in the SOA Instance Home page, you may see a blank page. This error may occur if:

- Application Dependency and Performance is not being used to monitor the SOA instance.
- Application Dependency and Performance has not been registered in Enterprise Manager Cloud Control.

Table 13–15 Error Message

Error Message	Workaround Steps
Missing ADP Data - Add the metrics - and add one for blank page.	To monitor data collected using ADP, the ADP Manager must be registered and configured.

13.20.7 Information Publisher Reports

This section lists report related errors.

Table 13–16 Error Message

Error Message	Workaround Steps
Report generation fails due to invalid database details.	<ol style="list-style-type: none"> 1. Navigate to the All Targets page. 2. Select the SOA Infrastructure target on which the specific SOA Composite has been deployed and click Configure. 3. In the Monitoring Configuration page, specify the database connection details and the credentials and click OK.
No targets found message for Oracle SOA Composite Reports.	You cannot use the out-of-box reports directly. You must use the Create Like option to generate custom reports based on the SOA Composite Target type.
Report generation fails due to invalid host details.	Set valid credentials for the host target on which the SOA Infrastructure instance is running.

13.20.8 BI Publisher Reports

This section lists BI Publisher report related errors.

Table 13–17 Error Message

Error Message	Workaround Steps
Exception Encountered For One of SOA BIP Report If SOA Dehydration Is Not Configured	<p>If the SOA Dehydration store details are not configured in BI Publisher, the SOA Composite Report (from Dehydration Store) is not generated, and the following exception message is displayed:</p> <p>The report cannot be rendered because of an error, please contact the administrator. Parameter name: P_PARTITION_NAME Can not establish database connection(EMSOA)</p> <p>To work around this issue, you must manually create the SOA database connection by choosing JDBC Connection from the Administration menu after the BI Publisher setup has been configured. The name of the data source name should be EMSOA. Use the following steps to create the EMSOA data source:</p> <ol style="list-style-type: none"> 1. From the Enterprise menu, select Reports, and then select BI Publisher Reports. The BI Publisher Enterprise login page appears. 2. Enter your credentials to log in to BI Publisher. 3. Click the Administration link available at the top right corner. 4. Navigate to the Data Sources page by clicking the JDBC Connection link in the Data Sources section. Click Add Data Source. 5. Enter EMSOA in the Data Source field, specify the driver type, driver class, connection string, user name, and password. Click Test Connection to ensure that the connection can be established successfully. 6. Click Apply. The newly created EMSOA jdbc data source appears on the Data Sources page. <p>Once you have created the EMSOA data source, the issue should be resolved.</p>

13.20.9 Systems and Services

The following error occurs when you try to refresh a service that has not been created.

Table 13–18 Error Message

Error Message	Workaround Steps
Create Service option does not work.	System and service creation depends on the configuration collection of the SOA Infrastructure and related targets. Check the log file for details.
Refresh Service option does not work.	The Refresh Service function works for an existing Infrastructure service. In case the service does not exist, it should be created using the Create Service menu option.

13.20.10 BPEL Recovery

The following error occurs when invalid credentials are provided.

Table 13–19 Error Message

Error Message	Workaround Steps
Invalid Host and WebLogic Domain Credentials	For the BPEL Recovery functionality to work, the host credentials and WebLogic Domain credentials must to be available in the preferred credential store. Set the valid credentials and try again.

13.20.11 SOA License Issue

The following error occurs if the SOA Management Pack EE has not been enabled.

Table 13–20 Error Message

Error Message	Workaround Steps
The page requested is part of the SOA Management Pack EE.	<p>The SOA Management Pack EE must be enabled for the specific SOA Infrastructure target. To enable the license, follow these steps:</p> <ol style="list-style-type: none"> 1. From the Setup menu, select Management Packs, then select Management Pack Access. 2. Select SOA Infrastructure in the Target Type drop-down box. 3. Uncheck and check the SOA Management Pack EE. 4. Click Apply and navigate to the SOA Composite page.

13.20.12 Dehydration Store Issue

Data is not displayed on the Dehydration Store page.

Table 13–21 Error Message

Error Message	Workaround Steps
Data is not displayed in the Dehydration Store page.	<p>This error may occur if there is a data mismatch between the values specified for the database target and the WebLogic Server Datasource. To resolve this issue, follow these steps:</p> <ol style="list-style-type: none"> 1. Compare the Database Host and SID value of the database target with the value collected for the WebLogic Server JDBC Datasource configuration. 2. If the values are different, select Services from the Targets menu. Select DataSources, then select SOALocalTxtSource, then click Connection Pool to update the Datasource Connection URL

Part V

Managing Oracle Business Intelligence

The chapter in this part describes how you can discover, monitor, and administer Oracle Business Intelligence instance and Oracle Essbase targets in Enterprise Manager Cloud Control 12c.

This part contains the following chapter:

- [Chapter 14, "Discovering and Monitoring Oracle Business Intelligence Instance and Oracle Essbase"](#)

Discovering and Monitoring Oracle Business Intelligence Instance and Oracle Essbase

Oracle Business Intelligence (Oracle BI), a part of Oracle Business Analytics, is a combination of technology and applications that provide a range of business intelligence capabilities, such as enterprise performance management, financial performance management, data integration, data warehousing, as well as a number of query, reporting, analysis, and alerting tools.

You can use Enterprise Manager Cloud Control 12c to monitor certain Oracle Business Intelligence targets. Monitoring the status, performance, and health of Oracle Business Intelligence targets enables you to set up a more efficient business intelligence system.

By monitoring a target using Enterprise Manager, you obtain a complete and up to date overview of the status, availability, performance, and health of the target. Enterprise Manager displays complex target performance data in a simple form, using graphs and pie charts. It also keeps you informed about target metrics crossing their threshold levels, target alerts, and target incidents that require user action.

This chapter explains how to monitor Oracle BI Instance and Oracle Essbase targets in Enterprise Manager Cloud Control 12c. It consists of the following sections:

- [Overview of Oracle Business Intelligence Targets You Can Monitor](#)
- [Understanding the Monitoring Process](#)
- [Discovering Oracle Business Intelligence Instance and Oracle Essbase Targets](#)
- [Monitoring Oracle Business Intelligence Instance and Essbase Targets](#)
- [Administering Oracle Business Intelligence Instance and Essbase Targets](#)

14.1 Overview of Oracle Business Intelligence Targets You Can Monitor

This section gives an overview of the Oracle Business Intelligence targets you can monitor using Enterprise Manager Cloud Control 12c Release 2 (12.1.0.2.0) or higher. It contains the following:

- [Oracle Business Intelligence Instance](#)
- [Oracle Essbase](#)

14.1.1 Oracle Business Intelligence Instance

Oracle Business Intelligence Instance (BI Instance) is a logical grouping of Business Intelligence components that can be configured as a unit to deliver a single integrated business intelligence capability. Every BI Instance target is part of a WebLogic domain.

For information on WebLogic domains, refer to *Oracle Fusion Middleware Creating Domains Using the Configuration Wizard*.

A BI Instance target consists of a number of components, which can be monitored individually using Enterprise Manager. [Table 14–1](#) describes these components.

Table 14–1 Oracle Business Intelligence Instance Components

Component	Description
BI Server	This component provides query and data access capabilities for Oracle Business Intelligence, and provides services for accessing and managing the enterprise semantic model.
BI Presentation Server	This component provides the framework and interface for the presentation of Oracle Business Intelligence data to web clients. It maintains an Oracle BI Presentation Catalog service on the file system for customizing this presentation framework.
BI Cluster Controller	This component manages Oracle Business Intelligence Server (BI Server) clusters. It also manages the active-passive clustering of the Oracle Business Intelligence Scheduler (BI Scheduler) components.
BI Scheduler	This component provides extensible scheduling for analyses to be delivered to users at specified times.
BI Java Host	This component provides component services that enable Oracle BI Presentation Services to support various components such as Java tasks for Oracle BI Scheduler, Oracle BI Publisher, and graph generation. It also enables Oracle BI Server query access to Hyperion Financial Management and Oracle Online Analytical Processing (OLAP) data sources.

14.1.2 Oracle Essbase

Oracle Essbase is a multidimensional database management system that provides business performance management solutions for meeting the complex calculation requirements of analysts across an enterprise.

Oracle Essbase consists of an Online Analytical Processing (OLAP) server that provides an environment for deploying pre-packaged applications and developing custom analytic and performance management applications. Every Essbase target is part of a WebLogic domain. For information on WebLogic domains, refer to *Oracle Fusion Middleware Creating Domains Using the Configuration Wizard*.

Using Enterprise Manager, you can monitor the Essbase server and every deployed Essbase application individually.

14.2 Understanding the Monitoring Process

To monitor Oracle Business Intelligence Instance (BI Instance) and Oracle Essbase targets, follow these steps:

1. Install Oracle Business Intelligence.

For information on how to install Oracle Business Intelligence, refer to *Oracle Fusion Middleware Installation Guide for Oracle Business Intelligence*.

2. Install the Enterprise Manager Cloud Control 12c Release 2 (12.1.0.2.0) system, or higher. If you are using an earlier version of Enterprise Manager Cloud Control, upgrade it to 12c Release 2 (12.1.0.2.0) or higher.

For information on how to install the Enterprise Manager Cloud Control 12c system, refer to *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.

For information on how to upgrade to Enterprise Manager Cloud Control 12c Release 2 (12.1.0.2.0) or higher, refer to *Oracle Enterprise Manager Cloud Control Upgrade Guide*.

Note: Oracle recommends that you install the Enterprise Manager Cloud Control system on a different host, other than the one on which you have installed Oracle Business Intelligence.

3. If the host on which you installed Oracle Business Intelligence does not have Oracle Management Agent (Management Agent) installed, install a Management Agent of version 12.1.0.2.0 or higher. If the host has a Management Agent of version earlier than 12.1.0.2.0 installed, upgrade the Management Agent to 12.1.0.2.0 or higher.

For information on how to install a Management Agent, refer to *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.

For information on how to upgrade a Management Agent, refer to *Oracle Enterprise Manager Cloud Control Upgrade Guide*.

4. If the Management Agent on the host on which you installed Oracle Business Intelligence does not have the Enterprise Manager for Oracle Fusion Middleware plug-in installed, deploy the 12.1.0.3.0 version, or a higher version of this plug-in on the host. If an earlier version of this plug-in is deployed to the Management Agent on the host, upgrade it to 12.1.0.3.0 or higher.

The 12.1.0.3.0 Enterprise Manager for Oracle Fusion Middleware plug-in is downloaded by default to the OMS host when you install a 12.1.0.2.0 OMS. The 12.1.0.4.0 Enterprise Manager for Oracle Fusion Middleware plug-in is downloaded by default to the OMS host when you install a 12.1.0.3.0 OMS.

For information on how to deploy a plug-in and upgrade an existing plug-in, refer to the Using Plug-Ins chapter of the *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

5. Discover the required BI Instance and Essbase targets.

BI Instance and Essbase targets are automatically discovered when you discover the WebLogic domain that they are part of.

The BI Instance and Essbase targets you want to monitor may be part of an undiscovered WebLogic domain, or a previously discovered WebLogic domain.

For information on how to discover BI Instance and Essbase targets part of an undiscovered WebLogic domain, see [Section 14.3.1](#).

For information on how to discover BI Instance and Essbase targets part of a previously discovered WebLogic domain, see [Section 14.3.2](#).

6. Monitor the BI Instance and Essbase targets.

For information on how to monitor BI Instance and Essbase targets, see [Section 14.4](#).

14.3 Discovering Oracle Business Intelligence Instance and Oracle Essbase Targets

Oracle Business Intelligence Instance (BI Instance) and Oracle Essbase targets you want to discover may be part of an undiscovered WebLogic domain, or a discovered WebLogic domain.

This section contains the following:

- [Discovering Targets of an Undiscovered WebLogic Domain](#)
- [Discovering New or Modified Targets of a Discovered WebLogic Domain](#)

Note: This section is applicable only for Oracle Business Intelligence Enterprise Edition 11g targets.

14.3.1 Discovering Targets of an Undiscovered WebLogic Domain

To discover BI Instance and Essbase targets part of an undiscovered WebLogic domain, first discover the WebLogic domain that the targets are part of. To do so, either enable the automatic discovery of WebLogic domains, or discover the required WebLogic domains manually. After discovering the WebLogic domains, you must promote the targets and assign Management Agents to monitor them.

The following sections explain how to perform these actions.

Enabling Automatic Discovery of Targets

In this method, you enable the automatic discovery of Fusion Middleware targets to automatically discover the various WebLogic domains in the enterprise. Also, you promote the BI Instance and Essbase targets part of the WebLogic domains, and assign Management Agents to monitor these targets. For information on how to use this method, see [Section 2.1](#).

Discovering Targets Manually

In this method, you manually discover WebLogic domains. Also, you promote the BI Instance and Essbase targets part of the WebLogic domains, and assign Management Agents to monitor these targets. For information on how to use this method, see [Section 2.2](#).

14.3.2 Discovering New or Modified Targets of a Discovered WebLogic Domain

In a typical enterprise, WebLogic domains are not static. New or modified domain members, such as BI Instance and Essbase targets, may be added to a discovered WebLogic domain at any point of time. Either enable the automatic discovery of these added targets, or discover them manually. After discovering these targets, you must promote the targets and assign Management Agents to monitor them.

The following sections explain how to perform these actions.

Enabling Automatic Discovery of Targets

In this method, you enable the automatic discovery of new or modified WebLogic domain member targets, such as BI Instance and Essbase targets. Also, you promote the new or modified domain member targets, and assign Management Agents to monitor them. For information on how to use this method, see [Section 2.3.1](#).

Discovering Targets Manually

In this method, you manually check a WebLogic domain for new members, such as BI Instance and Essbase targets, and discover them. Also, you promote the new or modified domain member targets, and assign Management Agents to monitor them. For information on how to use this method, see [Section 2.3.2](#).

14.4 Monitoring Oracle Business Intelligence Instance and Essbase Targets

To monitor Oracle Business Intelligence Instance (BI Instance) and Essbase targets, navigate to the home page of the required target.

To navigate to the home page of a BI Instance or Essbase target, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.

Using the target home page, you can perform a number of monitoring tasks. These tasks are described in this section, which contains the following:

- [Performing General Monitoring Tasks](#)
- [Performing Target-Specific Monitoring Tasks](#)

Note: This section is applicable only for Oracle Business Intelligence Enterprise Edition 11g targets.

14.4.1 Performing General Monitoring Tasks

This section explains how to perform general BI Instance and Essbase target monitoring tasks, such as viewing target status and availability, performance, health, alerts, incidents, and so on.

This section contains the following elements:

General

- [Viewing Target General and Availability Summary](#)
- [Viewing Target Status and Availability History](#)

Performance

- [Viewing Target Performance or Resource Usage](#)
- [Viewing Target Metrics](#)
- [Viewing or Editing Target Metric and Collection Settings](#)
- [Viewing Target Metric Collection Errors](#)

Health

- [Viewing Target Health](#)
- [Viewing Target Alert History](#)

- [Viewing Target Incidents](#)
- [Viewing Target Logs](#)

Configuration, Jobs, and Compliance

- [Viewing Target Configuration and Configuration File](#)
- [Viewing Target Job Activity](#)
- [Viewing Target Compliance](#)

14.4.1.1 Viewing Target General and Availability Summary

To view a general summary of the target details, navigate to the **Summary** section, by following these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.

The **Summary** section provides background information about the target, which helps you locate the target binaries, log files, metadata and configuration files, for viewing or editing purposes.

[Table 14–2](#) describes the elements of the **Summary** section.

Table 14–2 Target General and Availability Summary

Element	Description
Up Since	<i>(Displayed only when the target is up)</i> Time the target was last started successfully.
Down Since	<i>(Displayed only when the target is down)</i> Time the target was last stopped.
Availability	Percentage availability of the target.
Version	Version of the target software.
Oracle Home	Location of the target binaries.
Oracle Instance	Location of the target content files, metadata, configuration files and log files.
Port	Port used by the target for communication.
Running Applications (Only for Essbase Server targets)	Number of Essbase applications currently up and running.
Unexposed Applications (Only for Essbase Server targets)	Number of Essbase applications currently not being accessed by any user.
Connected Users (Only for Essbase Server targets)	Number of users currently connected through one or more of the applications.
Storage Type (Only for Essbase application targets)	Type of data storage used by the application.

Table 14–2 (Cont.) Target General and Availability Summary

Element	Description
Cubes (Only for Essbase application targets)	Number of cubes contained in the application.
Query Tracking (Only for Essbase application targets)	Whether or not query tracking, that is, tracking data combinations having a large number of data values that require aggregation, is enabled.
Memory Usage (MB) (Only for Essbase application targets)	Memory used by the application in MB.
Threads (Only for Essbase application targets)	Number of application threads.

14.4.1.2 Viewing Target Status and Availability History

To view the status and availability history of a target, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.
4. From the BI Instance component menu or Essbase target menu displayed on the target home page, select **Monitoring**, then select **Status History**.

Sometimes, due to network problems and system errors, the target might be down, or the Oracle Management Service (OMS) might not be able to reach the Management Agent that monitors the target. The Availability (Status History) page provides information about when, and for how long these situations occurred for a particular target. This information is essential for troubleshooting target related incidents.

The Availability (Status History) page consists of the **Overall Availability**, **Downtime History**, and **General** sections. The **Overall Availability** section consists of a pie chart depicting the availability of the target, from the time it was discovered. The **Downtime History** section provides detailed information about the periods when the target was down.

[Table 14–3](#) describes the elements of the **General** section.

Table 14–3 Target Status and Availability History

Element	Description
Current Status	Current status of the target, whether it is up and running, or down.
Up Since	<i>(Displayed only when the target is up)</i> Time the target was last started successfully.
Down Since	<i>(Displayed only when the target is down)</i> Time the target was last stopped.
Availability (%)	Percentage availability of the target.
Down Time (minutes)	Duration for which the target was down.

Table 14–3 (Cont.) Target Status and Availability History

Element	Description
Blackout Time (minutes)	Total duration of blackouts set on the target.
Agent Down Time (minutes)	Duration for which the Oracle Management Agent monitoring the target was down.
System Error Time (minutes)	Duration for which the target could not be monitored, due to a system error.
Status Pending Time (minutes)	Duration for which the status of the target could not be determined.

14.4.1.3 Viewing Target Performance or Resource Usage

To view the performance or resource usage of a target, navigate to the **Response or CPU and Memory Usage** section, by following these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target. Graphs depicting the target performance or target resource usage are displayed.
4. (Optional) To view the performance or resource usage data in a tabular format, click **Table View**.

Note: For the BI Instance, BI Server, and BI Presentation Server targets, you can view only performance data, and not resource usage data, on the target home page. For other BI Instance component targets, and Essbase targets, you can view only resource usage data and not performance data on the target home page.

Target Performance

The **Response and Load** section displays the performance of the BI Instance, BI Server, or BI Presentation Server target. For these targets, the **Response and Load** section can consist of the following graphs:

- The variation of Average Query Time with time
Average Query Time is the average time the BI Server or BI Presentation Server takes to execute a query. The Average Query Time is collected and uploaded to the Oracle Management Repository every fifteen minutes, by default.
- The variation of Server Queries (per second) with time
Server Queries (per second) is the number of queries processed by the BI Server or BI Presentation Server in one second. Server Queries (per second) is collected and uploaded to the Oracle Management Repository every fifteen minutes, by default.
- The variation of Completed Requests (per second) with time
Completed Requests (per second) is the number of requests completed by the BI Presentation Server in one second. Completed Requests (per second) is collected and uploaded to the Oracle Management Repository every fifteen minutes, by default.

Carefully observing these graphs can sometimes provide early warnings about server overloading, reduced server access, and so on. Analyzing graphical data collected over a long period of time can help you set up a more efficient BI Server or BI Presentation Server.

For detailed information on target performance, access the Performance Summary page. To access this page, from the **Business Intelligence Instance, BI Server** or **BI Presentation Services** menu, select **Monitoring**, then select **Performance Summary**.

Target Resource Usage

The **CPU and Memory Usage** section displays the resource usage of the target. It consists of two graphs:

- The variation of CPU Usage (%) with time
CPU Usage specifies the percentage of CPU time used by the target. A large value of CPU Usage can cause the Business Intelligence components and applications to slow down, reducing their performance. The CPU Usage is collected and uploaded to the Oracle Management Repository every fifteen minutes, by default.
- The variation of Memory Usage (MB) with time
Memory Usage specifies the amount of memory used by the target. A large value of Memory Usage can cause the Business Intelligence components and applications to slow down. The Memory Usage is collected and uploaded to the Oracle Management Repository every fifteen minutes, by default.

Carefully observing these graphs can sometimes provide early warnings about application overloading, component downtime, and so on.

14.4.1.4 Viewing Target Metrics

To view all the metrics collected for a particular target, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.
4. From the BI Instance component menu or Essbase target menu displayed on the target home page, select **Monitoring**, then select **All Metrics**.

The All Metrics page displays details about all the metrics collected for a particular target. The average value, threshold values, collection schedule, and metric value history is displayed for each collected metric.

14.4.1.5 Viewing or Editing Target Metric and Collection Settings

To view and edit the metric and collection settings for a particular target, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.

4. From the BI Instance component menu or Essbase target menu displayed on the target home page, select **Monitoring**, then select **Metric and Collection Settings**.
5. To edit the collection schedule or thresholds of a metric, or any other collected item, click the corresponding icon present in the **Edit** column.

The Metric and Collection Settings page provides details about target metric collection thresholds and target metric collection schedules. Using this page, administrators can edit the warning threshold and critical threshold values of target metrics and other collected items, as well as the time intervals at which these are collected.

14.4.1.6 Viewing Target Metric Collection Errors

To view the metric collection errors for a particular target, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.
4. From the BI Instance component menu or Essbase target menu displayed on the target home page, select **Monitoring**, then select **Metric Collection Errors**.

The Metric Collection Errors page provides details about the errors encountered while obtaining target metrics. These details give you an idea of the metrics that may not represent the performance of the target accurately, as errors were encountered while collecting them.

14.4.1.7 Viewing Target Health

To view a summary of the health of the target, navigate to the **Monitoring and Diagnostics** section, by following these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.

The **Monitoring and Diagnostics** section specifies the number of abnormal occurrences related to the target that require user action, and the number of changes made to the target configuration, within a particular time interval. This information is useful to administrators who want to quickly get an idea of the overall health of the target, and know the number of issues that need to be resolved. For more details on target configuration, access **Configuration** from the BI Instance component menu or Essbase target menu.

[Table 14–4](#) describes the elements of the **Monitoring and Diagnostics** section.

Table 14–4 Target Health

Element	Description
Incidents	The number of unresolved situations or issues that impact the target negatively, and hence require user action. The displayed integer is also a link to the Incident Manager page.

Table 14–4 (Cont.) Target Health

Element	Description
Descendant Target Incidents (Only for Essbase Server Targets)	The number of incidents related to Essbase applications. The displayed integer is also a link to the Incident Manager page.
Configuration Changes	The number of changes made to the target configuration in the last seven days. The displayed integer is also a link to the Configuration History page.

14.4.1.8 Viewing Target Alert History

To view the alert history of a particular target, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.
4. From the BI Instance component menu or Essbase target menu displayed on the target home page, select **Monitoring**, then select **Alert History**.

The Alert History page provides details about target metrics, such as the periods when a particular metric was beyond its critical threshold value, the periods when the metric could not be calculated, and so on. These details help you plan corrective measures for metric-related problems, before any severe damage or prolonged downtime can occur.

Table 14–5 describes the elements of the Alert History page.

Table 14–5 Target Alert History

Element	Description
Metric	Parameter related to the performance of the target.
History	Condition of the metric at various times. The condition can have the values Critical, Warning, Clear, and No Data.

14.4.1.9 Viewing Target Incidents

To view the incidents related to the target, navigate to the **Incidents** section, by following these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.

The **Incidents** section provides details about the various events, related to the target, that negatively impact the business intelligence system. These events require user action. The details provided by this section, such as the incident summary, severity, target, target type, and so on, are essential for troubleshooting.

For detailed reports on target incidents, access the Incident Manager page. To access this page, from the BI Instance component menu or Essbase target menu, select **Monitoring**, then select **Incident Manager**.

For details on the elements of the **Incidents** section, refer to *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

14.4.1.10 Viewing Target Logs

To view the log messages related to a particular target, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the BI Instance component or Essbase target menu displayed on the target home page, select **Logs**, then select **View Log Messages**.
4. (Optional) To view or download the target log files, click **Target Log Files**, select the required log file, then click **View Log File** or **Download**, respectively.
5. (Optional) To export log messages to a file, from the Log Messages page, select the required messages. From the **Export Messages to File** menu, click the file format you want to export the selected messages to. Choose a location, and download the file.

The target logs are a repository of target error messages, warnings, and notifications. They can be used for tracing the intermediate steps of an operation, and are essential for troubleshooting incidents and problems.

You can use the Log Messages page to view all log messages, search for a particular message, view messages related to a message, export messages to a file, view the target log files, and download the log files. For more information about log files, refer to *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

For the BI Instance target, this page displays log messages related to all system components and Java EE components. For the BI Instance component targets and Essbase targets, this page displays only those log messages that are related to the target.

[Table 14–6](#) describes the elements of the Log Messages page.

Table 14–6 Target Log Messages

Element	Description
Time	Date and time when the log message was created.
Message Type	Type of the log message. Message Type can be Incident Error, Error, Warning, Notification, Trace, or Unknown. These types represent the decreasing severity of messages, with Trace representing the least severe message and Incident Error representing the most severe message. Unknown indicates that Message Type is not known.
Message ID	9-digit string that uniquely identifies the message within the framework.
Message	Text of the log message.
Execution Context ID (ECID)	Global unique identifier of the execution of a particular request, in which a target component participates. You can use the ECID to correlate error messages from different target components.
Relationship ID	Identifier which distinguishes the work done by a particular thread on a particular process, from the work done by any other thread on the same, or any other process, on behalf of the same request.

Table 14–6 (Cont.) Target Log Messages

Element	Description
Component	Target component that generated the message.
Module	Identifier of the module that generated the message.
Incident ID	Identifier of the incident to which the message corresponds.
Instance	Oracle Instance containing the target component that generated the message.
Message Group	Group containing the message.
Message Level	An integer value representing the severity of the message. Ranges from 1 (most severe) to 32 (least severe).
Hosting Client	Identifier of the client or security group related to the message.
Organization	Organization ID for the target component that generated the message. This ID is <code>oracle</code> for all Oracle components.
Host	Name of the host where the message was generated.
Host IP Address	Network address of the host where the message was generated.
User	User whose execution context generated the message.
Process ID	Identifier of the process or execution unit that generated the message.
Thread ID	Identifier of the thread that generated the message.
Upstream Component	Component that the message generating component works with, on the client side.
Downstream Component	Component that the message generating component works with, on the server side.
Detail Location	URL linking to additional information about the message.
Supplemental Detail	Detailed information about the message, more detailed than the message text.
Target Log Files	Link to the target log files.
Log File	Log file containing the message.

14.4.1.11 Viewing Target Configuration and Configuration File

To view the configuration data of a target, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.
4. From the BI Instance component menu or Essbase target menu displayed on the target home page, select **Configuration**, then select **Last Collected** to access the Target Configuration browser.
5. (Optional) To export target configuration data to a configuration file, click **Export**. The exported target configuration data is stored in a `.xls` file.

Use the Target Configuration browser to view the latest configuration data of the target. Using the browser, you can also search for configuration data, view saved

target configurations, compare target configurations, and view the target configuration history.

14.4.1.12 Viewing Target Job Activity

To view the past, currently running, and scheduled jobs related to a target, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.
4. From the BI Instance component menu or Essbase target menu displayed on the target home page, select **Job Activity**.

The Job Activity page displays target jobs related to target administrative tasks, such as starting the target, stopping the target, target blackouts, and so on.

Use the Job Activity page to search for a particular job and retrieve job details such as the owner, status, scheduled start time, and so on. You can also use the Job Activity page to perform target job administration tasks, such as creating, editing, suspending, and resuming a job.

14.4.1.13 Viewing Target Compliance

To view the compliance of a target to compliance standards or compliance frameworks, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.
4. From the BI Instance component menu or Essbase target menu displayed on the target home page, select **Compliance**, then select **Results**.
5. To view the compliance results of a target with respect to a particular compliance standard, select **Compliance Standards**. To view the compliance results of a target with respect to a particular compliance framework, select **Compliance Frameworks**.

Use the Compliance Results page to view the compliance of a target to compliance standards and compliance frameworks. This page also lists the number of violations made to compliance standards and compliance frameworks, hence giving you an idea of whether the targets in your enterprise adhere to established standards or not.

For more information on target compliance, refer to *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

14.4.2 Performing Target-Specific Monitoring Tasks

This section explains how you can perform target-specific BI Instance and Essbase target monitoring tasks, such as viewing BI Instance dashboard reports, BI Instance scheduler reports, Essbase application data storage details, and so on.

This section contains the following:

BI Instance

- [Viewing Oracle Business Intelligence Dashboard Reports](#)
- [Viewing Oracle Business Intelligence Scheduler Reports](#)
- [Viewing Oracle Business Intelligence Instance Key Metrics](#)

Essbase

- [Viewing Oracle Essbase Applications Summary](#)
- [Viewing Oracle Essbase Application Data Storage Details](#)

14.4.2.1 Viewing Oracle Business Intelligence Dashboard Reports

To view Oracle Business Intelligence dashboard reports, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance. Click the BI Instance name.
3. From the **Business Intelligence Instance** menu, select **Dashboard Reports**.
4. From the **View** list, select the set of dashboard reports you want to view.

Note: To view Oracle Business Intelligence dashboard reports in Enterprise Manager Cloud Control, you must enable usage tracking. For information on how to enable usage tracking, refer to the Managing Usage Tracking chapter of the *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Using this page, you can view the dashboard usage in the past 7 days, the dashboards that failed in the past 24 hours, the top dashboards by resource usage in the past 7 days, and the top users by resource usage in the past 7 days. These details tell you which dashboards are the most popular, which dashboards failed recently, which dashboards use the maximum resources, and which user is the most active. An in-depth analysis of these details can provide important insights into the functioning of an enterprise.

Note: Without specifying the correct credentials on the Monitoring Credentials page, you cannot access certain dashboard reports. Hence, ensure that you specify the appropriate credentials on the Monitoring Credentials page, before accessing the Dashboard Reports page.

To access the Monitoring Credentials page, from the **Business Intelligence Instance** menu, select **Target Setup**, then select **Monitoring Credentials**.

[Table 14–7](#) describes the elements of the Dashboard Reports page.

Table 14–7 Oracle Business Intelligence Dashboard Reports

Element	Description
User	User who accessed the dashboard.
Total Sessions	Total number of user sessions which accessed the dashboard.

Table 14–7 (Cont.) Oracle Business Intelligence Dashboard Reports

Element	Description
Last Accessed On	Time when the dashboard was last accessed.
Dashboard	Dashboard name.
Error Code	Dashboard error code.
Error Message	Dashboard error message.
Repository	Name of the repository accessed by the dashboard.
Subject Area	Information about business areas, or the groups of users in an organization.
Start Time	Time when the server received the logical request for the dashboard.
End Time	Time when the server completed servicing the logical request for the dashboard.
View Log Messages	View log messages related to the dashboard.
Total Time	Total time taken to service all logical requests made for a particular dashboard. Note: In the Top Users by Resource Usage in Last 7 Days reports, this element represents the total time taken to service all logical requests made by a particular user.
Database Time	Time taken by the database to complete all physical requests made for a particular dashboard. Note: In the Top Users by Resource Usage in Last 7 Days reports, this element represents the time taken by the database to complete all physical requests made by a particular user.
Compile Time	Time taken to convert all logical requests made for a particular dashboard. Note: In the Top Users by Resource Usage in Last 7 Days reports, this element represents the time taken to convert all logical requests made by a particular user, to physical requests.
Failed Logical Requests	Number of logical requests made for the dashboard that failed. Note: In the Top Users by Resource Usage in Last 7 Days reports, this element represents the number of logical requests made by a particular user that failed.
Total Logical Requests	Total number of logical requests made for the dashboard. Note: In the Top Users by Resource Usage in Last 7 Days reports, this element represents the total number of logical requests made by a particular user.

14.4.2.2 Viewing Oracle Business Intelligence Scheduler Reports

To view Oracle Business Intelligence scheduler reports, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance. Click the BI Instance name.
3. From the **Business Intelligence Instance** menu, select **Scheduler Reports**.
4. From the **View** list, select the set of scheduler reports you want to view.

Using this page, you can view the BI Instance target jobs that failed in the past 24 hours, and the BI Instance target jobs that have been scheduled to begin later. These

details inform you about the jobs that failed recently and the jobs scheduled to take place in the future, giving you a summary of the BI Instance past and future job activity.

[Table 14–8](#) describes the elements of the Scheduler Reports page.

Table 14–8 Oracle Business Intelligence Instance Scheduler Reports

Element	Description
Job Name	Name of the job, as specified by the user who created it.
Instance ID	ID of the job instance.
Job ID	ID of the job.
Start Time	Time the job started.
End Time	Time the job ended or failed.
Error Message	Error message of the failed job.
User	User who created the job.
Scheduled Time	Time the job is scheduled to begin.
Script Type	Type of script to be executed.

14.4.2.3 Viewing Oracle Business Intelligence Instance Key Metrics

To view the key metrics related to the BI Instance target, navigate to the **Metrics** section by following these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance. Click the BI Instance name.

The **Metrics** section displays the key metrics used to monitor the performance of the BI Instance. Analyzing these metrics provides early warnings of errors and incidents, and helps you identify problem areas quickly.

To view all BI Instance metrics, access the All Metrics page. To access this page, from the **Business Intelligence Instance** menu, select **Monitoring**, then select **All Metrics**. For more information on this page, see [Section 14.4.1.4](#).

[Table 14–9](#) describes the elements of the **Metrics** section.

Table 14–9 Oracle Business Intelligence Instance Key Metrics

Metric	Description
Request Processing Time (ms)	Average time, in milliseconds, taken by the BI Servers to process a request. This metric is collected from the time the BI Analytics application was last started.
SOA Request Processing Time (ms)	Average time, in milliseconds, taken by the Oracle WebLogic Server cluster to process a web services request. This metric is collected from the time the BI SOA application was last started.
Average Query Time (seconds)	Average time, in seconds, taken by the BI Servers to process a query. This metric is collected from the time the BI Server was last started.
Active Sessions	Total number of active sessions for the BI Instance. This metric is collected from the time the BI Analytics application was last started.

Table 14–9 (Cont.) Oracle Business Intelligence Instance Key Metrics

Metric	Description
Requests (per minute)	Average number of requests, per minute, received by the BI Servers. This metric is collected from the time the BI Analytics application was last started.
SOA Requests (per minute)	Average number of servlet and/or JavaServer Pages (JSP) invocations, per minute, for web services requests across the Oracle WebLogic Server cluster. This metric is collected from the time the BI SOA application was last started.
Presentation Services Requests (per second)	Average number of requests, per second, received by the BI Presentation Servers. This metric is collected from the time the BI Presentation Server was last started.
Server Queries (per second)	Average number of queries, per second, completed by the BI Servers. This metric is collected from the time the BI Server was last started.
Failed Queries	Number of failed BI Server queries. This metric is collected from the time the BI Presentation Server was last started.

14.4.2.4 Viewing Oracle Essbase Applications Summary

To view a summary of Oracle Business Intelligence Essbase applications, navigate to the **Applications** section, by following these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having an Essbase Server target. Click the Essbase Server name.

The **Applications** section provides details about the status, resource usage, and data storage type of the various Essbase applications under the Essbase server. This section is useful to administrators who want to quickly obtain an overview of the availability and storage details of the Essbase applications being monitored.

Note: If the applications displayed in the **Applications** section are different from the ones displayed in the **Target Navigation** window, refresh the Oracle Fusion Middleware farm. To do this, from the **Target Navigation** window, click the Oracle Fusion Middleware farm name. From the **Farm** menu, click **Refresh WebLogic Domain**. Click **Add/Update Targets**.

Table 14–10 describes the elements of this section.

Table 14–10 Oracle Essbase Applications Summary

Element	Description
Name	Name of the application.
Status	Application status, whether the application is up or down.
Storage Type	Type of application data storage.
Memory Usage (MB)	Memory, in MB, used by the application.
Cubes	Number of cubes contained in the application.

14.4.2.5 Viewing Oracle Essbase Application Data Storage Details

To view details about how data for an Oracle Business Intelligence Essbase application is stored, navigate to the **Cubes** section, by following these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having an Essbase Server target. Click the Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required Essbase application.

The **Cubes** section provides structural and usage information about the cubes contained in the Essbase application. These details tell you about how data storage is designed for the application, and how accessible the application data is at the moment.

Table 14–11 describes the elements of this section.

Table 14–11 Oracle Essbase Application Data Storage Details

Element	Description
Name	Name of the cube.
Dimensions	Number of dimensions the cube has.
Connected Users	Number of users currently connected to the cube data.
Locks	Number of data block locks currently held on the cube.
Data Cache Size (KB)	Size, in KB, of the buffer in memory that holds uncompressed data blocks.

14.5 Administering Oracle Business Intelligence Instance and Essbase Targets

To administer Oracle Business Intelligence Instance (BI Instance) and Essbase targets using Enterprise Manager Cloud Control, navigate to the home page of the required target. For information on how to do this, see [Section 14.4](#).

Using Enterprise Manager Cloud Control, you can perform general, as well as target specific administration tasks.

This section contains the following:

- [Performing General Administration Tasks](#)
- [Performing Target-Specific Administration Tasks](#)

Note: This section is applicable only for Oracle Business Intelligence Enterprise Edition 11g targets.

14.5.1 Performing General Administration Tasks

This section explains how to perform general BI Instance and Essbase target administration tasks, such as starting, stopping, or restarting the target, administering target access privileges, administering target blackouts, and so on.

This section contains the following:

- [Starting, Stopping, or Restarting the Target](#)
- [Administering Target Access Privileges](#)

- [Administering Target Blackouts](#)
- [Viewing Target Monitoring Configuration](#)

14.5.1.1 Starting, Stopping, or Restarting the Target

To start, stop, or restart a target, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.
4. Click **Start Up**, **Shut Down**, or **Restart** to start, stop, or restart the target, respectively. Alternatively, from the BI Instance component menu or Essbase target menu, select **Control**, then select **Start Up**, **Shut Down**, or **Restart**.

To run certain patching and maintenance tasks, you may need to stop the target, perform the task, and restart it once the operation is complete.

14.5.1.2 Administering Target Access Privileges

To manage the access privileges for a target, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.
4. From the BI Instance component menu or Essbase target menu displayed on the target home page, select **Target Setup**, then select **Administrator Access**.
5. Click **Add** to grant target access privileges to a role or an administrator.

Use the Access page to set target privileges for roles and administrators. The available privileges are View, Operator, and Full.

View only allows you to view the target in the console, whereas Operator allows you to view targets, and perform all administrative actions except deleting targets. Full allows you to view targets, and perform all administrative actions.

14.5.1.3 Administering Target Blackouts

To administer the blackouts for a target, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.
4. From the BI Instance component menu or Essbase target menu displayed on the target home page, select **Monitoring**, then select **Blackouts**.

Blackouts suspend data collection on a monitored target. Blackouts are useful when you want to perform scheduled maintenance tasks on monitored targets.

Use the Blackouts page to search for existing target blackouts, edit existing blackouts, define new blackouts, and stop blackouts. You can also create and stop blackouts using the BI Instance component menu, or the Essbase target menu. To create or stop a blackout, from the BI Instance component menu, or the Essbase target menu, select **Control**, then select **Create Blackout** or **End Blackout**, respectively.

14.5.1.4 Viewing Target Monitoring Configuration

To view the monitoring configuration details for a particular target, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance or Essbase Server target. Click the BI Instance or Essbase Server name.
3. From the navigation tree in the **Target Navigation** window, click the name of the required target.
4. From the BI Instance component menu or Essbase target menu displayed on the target home page, select **Target Setup**, then select **Monitoring Configuration**.

The Monitoring Configuration page provides information about instance properties of the target, which provide internal details about target monitoring.

[Table 14–12](#) describes the elements of the Monitoring Configuration page.

Table 14–12 Target Monitoring Configuration

Element	Description
Canonical Path	Component path of the form <code>instance_name/component_name</code> .
Oracle Instance Home	Location of the target content files, metadata, configuration files and log files.
DB Class String	String needed to form a JDBC connection with a target repository.
DB Connection String	String that specifies information about the target repository, and the means to connect to it.
DB Password	Repository database password.
DB User Name	Repository database user name.
Domain Home	Domain home directory of the WebLogic domain that the target is a part of.
Is JRF Enabled	Whether Oracle Java Required Files (JRF) is applied to the target instance or not.
Monitoring Mode	Indicates whether the Enterprise Manager instance uses a repository while monitoring the target or not. Repo indicates that a repository is used, whereas Repo-less indicates that a repository is not used.
Version	Version of the target software.

14.5.2 Performing Target-Specific Administration Tasks

This section explains how to perform target-specific BI Instance and Essbase target administration tasks, such as viewing BI Instance component failovers, and editing BI Instance monitoring credentials.

This section contains the following:

- [Viewing Oracle Business Intelligence Component Failovers](#)

- [Editing Oracle Business Intelligence Monitoring Credentials](#)

14.5.2.1 Viewing Oracle Business Intelligence Component Failovers

To view the BI Instance component failovers, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having a BI Instance target. Click the BI Instance name.
3. Select the **Availability** tab, then select **Failover**.

This page displays the risk levels of BI Instance component failure, the recommended backup actions to prevent component failures, and the backup or secondary hosts for components that have failovers configured. Administrators can use this information to plan failovers for BI Instance components that have a high risk of failure.

For more information on the recommended backup actions to avoid BI Instance component failures, refer to *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

14.5.2.2 Editing Oracle Business Intelligence Monitoring Credentials

To edit the BI Instance monitoring credentials, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. From the navigation tree, select the Oracle Fusion Middleware farm having an BI Instance. Click the BI Instance name.
3. From the **Business Intelligence Instance** menu, select **Target Setup**, then select **Monitoring Credentials**.
4. Edit the required fields, then click **Save**.

This page enables you to specify and edit the credentials required to connect to the database which stores scheduling and usage tracking information. Without specifying the correct credentials on this page, you cannot access certain dashboard reports. Hence, ensure that you specify the appropriate credentials on this page, before accessing the Dashboard Reports page.

[Table 14–13](#) describes the elements of the Monitoring Credentials page.

Table 14–13 Oracle Business Intelligence Instance Monitoring Credentials

Element	Description
Database Type	Type of the database.
Hostname	Name of the host on which the database is installed.
Port	Port used for communicating with the database.
Service Name	Name of the database service.
Username	User name used for database login.
Password	Password used for database login.

Part VI

Using JVM Diagnostics and MDA Advisor

The chapters in this part provide information regarding JVM Diagnostics and MDA Advisor.

The chapters are:

- [Chapter 15, "Introduction to JVM Diagnostics"](#)
- [Chapter 16, "Using JVM Diagnostics"](#)
- [Chapter 17, "Troubleshooting JVM Diagnostics"](#)
- [Chapter 18, "Using Middleware Diagnostics Advisor"](#)

Introduction to JVM Diagnostics

This chapter provides an overview of JVM Diagnostics. It contains the following sections:

- [Overview](#)
- [New Features in this Release](#)
- [Supported Platforms and JVMs](#)
- [User Roles](#)

15.1 Overview

Mission critical Java applications often suffer from availability and performance problems. Developers and IT administrators spend a lot of time diagnosing the root cause of these problems. Many times, the problems occurring in production environments either cannot be reproduced or may take too long to reproduce in other environments. This can cause severe impact on the business.

Oracle Enterprise Manager Cloud Control 12c's JVM Diagnostics enables administrators to diagnose performance problems in Java application in the production environment. By eliminating the need to reproduce problems, it reduces the time required to resolve these problems. This improves application availability and performance. Using JVM Diagnostics, administrators will be able identify the root cause of performance problems in the production environment without having to reproduce them in the test or development environment. It does not require complex instrumentation or restarting of the application to get in-depth application details. Application administrators will be able to identify Java problems or Database issues that are causing application downtime without any detailed application knowledge. The key features of JVM Diagnostics are:

- [Java Activity Monitoring and Diagnostics with Low Overhead](#)
- [In-depth Visibility of JVM Activity](#)
- [Real Time Transaction Tracing](#)
- [Cross-Tier Correlation with Oracle Databases](#)
- [Memory Leak Detection and Analysis](#)
- [JVM Pooling](#)
- [Real-time and Historical Diagnostics](#)

15.1.1 Java Activity Monitoring and Diagnostics with Low Overhead

JVM Diagnostics provides in-depth monitoring of Java applications without slowing them down. It helps you to identify the slowest requests, slowest methods, requests waiting on I/O, requests using a lot of CPU cycles, and requests waiting on database calls. It also identifies the end-user requests that have been impacted by resource bottlenecks. Application resources that are causing the performance bottleneck are also visible.

15.1.2 In-depth Visibility of JVM Activity

JVM Diagnostics provides immediate visibility into the Java stack. You can monitor thread states and Java method/line numbers in real time and you can proactively identify issues rather than diagnosing issues like application crashes, memory leaks, and application hangs after they occur.

15.1.3 Real Time Transaction Tracing

If a particular request is hanging or if the entire application is slow, administrators can perform a real-time transaction trace to view current Java application activity. You can see the offending threads and their execution call stacks. You can also analyze various bottleneck resources such as how much time a thread spent in waiting for a database lock. Complex problems such as activity in one thread (or request) affecting the activity in the other thread or rest of the JVM can be found very quickly.

Sometimes the monitoring interval (default 2 seconds) that is in use is too coarse grained. The Java thread of interest may be too short lived or the amount of monitoring data collected may be insufficient. In such cases, you can run a JVM Trace to get fine-grained details of the JVM activity. This feature allows you to monitor your Java application at a very high frequency (default of once every 50ms) for a short period of time. This allows you to identify interdependency of threads, bottleneck resources (DB, I/O, CPU, Locks, Network) & top methods.

15.1.4 Cross-Tier Correlation with Oracle Databases

JVM Diagnostics facilitates tracing of Java requests to the associated database sessions and vice-versa enabling rapid resolution of problems that span different tiers. Administrators can drill down from a JVM Thread in a DB Wait State to the associated Oracle database session. Additionally, they can now drill up from the SQL query to the associated JVM and related WebLogic Server targets (this is applicable only if the database and JVM are being monitored by Enterprise Manager).

This feature highlights the slowest SQL queries and helps administrators to tune SQL and the database to improve application performance. This facilitates smooth communication between the database administrators and application administrators by isolating the problems to the database or the application tier.

15.1.5 Memory Leak Detection and Analysis

Memory leaks lead to application slowdowns and eventually cause applications to crash. JVM Diagnostics alerts administrators on abnormalities in Java memory consumption. Administrators can use JVM Diagnostics and take heap dumps in production applications without disturbing the application. They can take multiple heap dumps over a period of time, analyze the differences between the heap dumps and identify the object causing the memory leak. Heap analysis can be performed even across different application versions. Differential Heap Analysis with multiple heap dumps makes it easy to identify memory leaks.

15.1.6 JVM Pooling

JVM Diagnostics allows administrators to group sets of JVMs together into JVM pools. This provides the console user with a single view across all related JVMs. Hence all JVM's that make up a single application or a single cluster may be grouped together in an application. This allows administrators to visualize problems naturally and intuitively.

15.1.7 Real-time and Historical Diagnostics

With JVM Diagnostics, you can perform real-time and historical diagnostics on your Java applications. This provides you with detailed insight on the root causes of production problems without having to reproduce the same problem in a Test or QA environment. You can play back transactions interactively from the browser and view the time spent in the network and the server.

Apart from the real-time data, you can also analyze historical data to diagnose problems that occurred in the past. You can view historical data that shows the time taken by end-user requests and the breakdown by Servlet, JSP, EJB, JDBC, and SQL layers.

15.2 New Features in this Release

This section lists some of the new JVM Diagnostics features in Oracle Enterprise Manager Cloud Control 12c:

- Automated JVM Diagnostics Engine and Agent Deployment Procedure.
- Customer provisioning scripts for deploying the JVM Diagnostics Agent in production environments.
- New home pages available for JVM Pool and JVM targets.
- Top databases chart added to the JVM and JVM Pool target pages.
- Metric Palette for JVMs and JVM Pools is now available and associated with WebLogic Server targets.
- JVM Diagnostics integrated with the Middleware Diagnostics Advisor.
- Automatic correlation between JVM and WebLogic Server targets.
- JVM Diagnostics region is now included in Composite and Fusion Applications.
- JVM threshold violations are now integrated with the Event subsystem.
- Heap snapshots can be taken in HPROF format for external analysis. Heap snapshots can now be imported.
- Performance Diagnostics and Live Thread Analysis pages can be directly accessed from the JVM target page.
- Bi-directional integration between JVM threads and database sessions through Live Thread Analysis.
- Thread State Transition charts and class histograms are now available.
- RMI Thread State is now supported.
- Wait Time on SQL (DB Wait), Thread Stack Local Objects Browser, Depth and number of locks are now displayed in the Live Thread Analysis page.
- Sample analyzer can be accessed from the Thread State Transition Chart.

- Windows 64, WLS Virtual Edition, Solaris x86-64 and AIX 64 platforms are now supported.
- JVM Diagnostics Engine is now supported for all OMS platforms.
- The DB Agent is not required for cross tier correlation if you are using Oracle Database 11gR2 or later versions.
- You can drill down to JVM Diagnostics from RUEI and ESS.
- Integration between JVM Diagnostics and Coherence, JVM Diagnostics and ECM is now available.
- The performance of the JVMTI engine has improved.
- Integration with Enterprise Manager Offline Diagnostics feature. You can now take snapshots of diagnostic data, export this data, or import diagnostic data for a particular collection.
- JVM targets can now be integrated with JBoss targets.
- You can save, compare, edit, and delete class histograms.
- The `emctl jvmd verb` can now be used to manage the JVM Diagnostics Engine. See [Section 16.11, "Using emctl to Manage the JVM Diagnostics Engine"](#) for details.

15.3 Supported Platforms and JVMs

The JVM Diagnostics Engine is supported on all platforms on which the Oracle Management Service has been certified.

For the latest certification information, refer to My Oracle Support Note 1415144.1. You can access My Oracle Support at the following URL:

<https://support.oracle.com/CSP/ui/flash.html>

15.4 User Roles

To use JVM Diagnostics, you must have either of the following Enterprise Manager system privileges:

- JVM Diagnostics User: Allows you to view JVM Diagnostics data.
- JVM Diagnostics Administrator: Allows you to manage JVM Diagnostics operations such as creating and analyzing heap and thread snapshots, tracing threads, and so on.

You can define these privileges in the Setup pages. For more details on defining these privileges, see the *Enterprise Manager Security* chapter in the Enterprise Manager Cloud Control Administrator's Guide.

Using JVM Diagnostics

This chapter describes the various tasks you can perform using JVM Diagnostics. In particular, it contains the following:

- [Installing JVM Diagnostics](#)
- [Setting Up JVM Diagnostics](#)
- [Accessing the JVM Diagnostics Pages](#)
- [Managing JVM Pools](#)
- [Managing JVMs](#)
- [Managing Thread Snapshots](#)
- [Analyzing Trace Diagnostic Images](#)
- [Uploading Trace Diagnostics Images](#)
- [JVM Offline Diagnostics](#)
- [Viewing JVM Diagnostics Threshold Violations](#)
- [Using emctl to Manage the JVM Diagnostics Engine](#)

16.1 Installing JVM Diagnostics

The JVM Diagnostics Engine runs as an Enterprise JavaBeans (EJB) Technology on a WebLogic Server. The JVM Diagnostics Agent is deployed on the targeted JVM (the one running a production WebLogic Server). It collects real-time data and transmits it to the JVM Diagnostics Engine. This data is stored in the Management Repository, and the collected information is displayed on Enterprise Manager Cloud Control console for monitoring purposes. The communication between the JVM Diagnostics Engine and the JVM Diagnostics Agent can be a secure (SSL) or non-secure connection.

The Application Performance Management Page is a GUI based screen that enables you to deploy and monitor the health of the JVM Diagnostics Engine in a reliable and an efficient manner. Using the Application Performance Management Page, you can achieve the following:

- Deploy JVM Diagnostics Engine
- Monitor the availability of all the JVM Diagnostics Engines
- Access information about the JVM Diagnostics Engines like hosts to which the engines are deployed, the current status, the port on which they are running, version, and so on.

For more details on installing JVM Diagnostics, see Enterprise Manager Cloud Control Basic Installation Guide.

16.1.1 Monitoring a Standalone JVM

If you need to monitor a standalone JVM, you can manually deploy the JVM Diagnostics Agent by following these steps:

1. From the **Setup** menu, select **Middleware Management**, then select **Application Performance Management**. Check if the JVM Diagnostics Engine has been enabled. The JVM Diagnostics Engine is independent and will run in a WLS target in the same domain as the Cloud Control OMS server that is already running. This is essentially another WLS container in the same domain and is very lightweight, so it can run well on the same server as the Cloud Control OMS instances.
2. Select the JVM Diagnostics Engine row in the Application Performance Management Engines table and click **Configure**. In the JVM Diagnostics Setup page, click the **JVMs and Pools** tab and then click **Downloads** to download the `jamagent.war` file.

3. Add the `jamagent.war` to the CLASSPATH as follows:

```
CLASSPATH=$CLASSPATH:/scratch/ssmith/jvmd/jamagent.war
export CLASSPATH
```

4. Start JVM Diagnostics with the JVM Diagnostics Agent as follows:

```
$JAVA_HOME/bin/java -cp $CLASSPATH $JVM_OPT $SYS_OPT jamagent.jamrun
[$JAMAGENT_PARAMS_LIST] $TARGET_CLASS $TARGET_CLASS_PARAMS
```

where `[$JAMAGENT_PARAMS_LIST]` refers to the JVM Diagnostics Agent parameters such as `jamloglevel`, `jamconshost`, `jamconsport`, `jamconsretr`, `jamtimeout`, and so on.

If these parameters are not specified, they are picked up from the `jamagent.war` file.

An example is given below:

```
CLASSPATH="$CLASSPATH:/scratch/ssmith/jamagent.war"
$JAVA_HOME/bin/java -cp $CLASSPATH $JVM_OPT $SYS_OPT jamagent.jamrun
jamconshost=10.229.187.109 jamconsport=3800
oracle.ad4j.groupidprop= MyJVMPool1/JVM50
```

16.2 Setting Up JVM Diagnostics

Follow these steps to set up and configure JVM Diagnostics:

1. Select the JVM Diagnostics Engine row in the Application Performance Management Engines table and click **Configure**.
 - **Name:** The name assigned to the JVM Diagnostics Engine. This ID identifies the JVM Diagnostics Engine in all the processes.
 - **Host:** The machine on which the JVM Diagnostics Engine has been deployed.
 - **Port:** The port of the machine on which the JVM Diagnostics Engine has been deployed.
 - **SSL Port:** The SSL Port of the machine on which the JVM Diagnostics Engine has been deployed.
 - **Build Version:** The build version of this JVM Diagnostics Engine.

- **Status:** The status of the JVM Diagnostics Engine (Active/Inactive).

Select the JVM Diagnostics Engines row and click **Configure** to configure the JVM Diagnostics Engine parameters, JVMs and Pools, databases, and heap loader. The following tabs are displayed:

- [Configuring the JVM Diagnostics Engine](#)
- [Configuring JVMs and Pools](#)
- [Register Databases](#)
- [Configuring the Heap Analysis Hosts](#)

16.2.1 Configuring the JVM Diagnostics Engine

You can configure the JVM Diagnostics Engine and create new idle thread rules or delete existing ones.

1. Click the **JVMD Configuration** tab. The following page appears.

Figure 16–1 JVMD Configuration Page

The screenshot shows the 'JVMD Configuration' page with the following sections:

- JVMD Engine Parameters:**
 - JVMD Engine Log Level: Info-3
 - Cross Tier Log Level: Info-3
 - Agent Request Timeout (secs): 30
 - Enable Monitoring: ☒
- Advanced Settings:**
 - Monitoring Aggregation Interval (secs): 90
 - Purge Detail Data older than (hours): 24
 - Thread Stack Repository Insertion Rate(%): 100
 - System Sample Interval (secs): 30
 - Purge Aggregated Data Older than (days): 30
 - Retry changing threads: ☐
- Idle Thread Rules:**
 - Buttons: New Rule, Remove
 - Message: Adding a rule will make threads idle if they meet the rule
 - Table with 3 columns: Rule Id, Rule Type, Rule Value.

Rule Id	Rule Type	Rule Value
1	Current Call	weblogic.socket.PostSocketMuxer->processSockets
2	Current Call	weblogic.socket.PostSocketMuxer->poll
3	Current Call	weblogic.socket.EPollSocketMuxer->processSockets
4	Current Call	weblogic.socket.NTSocketMuxer->getIoCompletionResult
5	Current Call	weblogic.socket.DevPollSocketMuxer->processSockets
6	Current Call	weblogic.socket.DevPollSocketMuxer->doPoll
7	Current Call	weblogic.socket.DevPollSocketMuxer->processSockets

2. You can modify the following details in the JVMD Engine Parameters region.
 - **JVMD Manager Log Level:** The log level for console diagnostics messages. Log levels 1 to 5 are supported where:
 - 1 = Error
 - 2 = Warning
 - 3 = Info
 - 4 = Debug
 - 5 = Trace

The default log level is 3.
 - **Cross Tier Log Level:** The log level for cross-tier diagnostic messages. Log levels 1 to 5 are supported where:
 - 1 = Error

- 2 = Warning
- 3 = Info
- 4 = Debug
- 5 = Trace

The default log level is 3.

- **Agent Request Timeout:** The number of seconds that the JVM Diagnostics Engine waits for the JVM Diagnostics Agent to respond. You can increase this value if the monitored JVMs are extremely busy and the console times out and disconnects while waiting for a response.
- **Monitoring Aggregation Interval:** The frequency at which the detailed monitoring samples should be aggregated into summary data.
- **System Sample Interval:** The frequency at which system details (cumulative CPU counters, heap size, and number of GCs) should be collected in monitoring.
- **Purge Data Older Than (Days):** The period for which the detailed monitoring samples should be retained.
- **Enable Monitoring:** Select this check box to start or stop monitoring.
- **Thread Stack Repository Insertion Rate (%):** Enter a number between 1 and 100. The thread stacks will be stored in the repository at the specified rate.
- **Retry Changing Threads:** If a thread stack changes during a sample (this can happen when a thread is using CPU), JVM Diagnostics will skip that thread for that sample. If you find missing samples, use this feature to retrace the changed stacks. This will retry (up to 5 times) threads with changing stacks. It will also make system calls to get the stack if possible.

Note: This field is not applicable to the JVMTI (level 0) optimization.

Click **Save** to save the parameters.

3. In the Idle Thread Rules region, you can mark a thread as idle by adding it to an Idle Thread Rule. All threads that have been marked as idle will not be monitored. Click **New Rule** to create a new Idle Thread Rule.

- **Rule Type:** The Rule Type can be:
 - **Monitor (Waiting on Lock):** Select this type if you want to ignore threads that are locked with a lock of the specified name.
 - **Current Call:** Select this type if you want to ignore all threads that are making a call to the selected function. You can specify a wild card here, for example, if you specify `weblogic.servlet.*`, all threads that meet this criteria will be ignored.
- **Rule Value:** The Rule Value should contain the class name, method, followed by class+method. An example of a Current Call is `weblogic.socket.PosixSocketMuxer->processSockets`. An example of a Monitor (Waiting on Lock) is `weblogic.socket.PosixSocketMuxer$1`.

All threads that meet the criteria specified in the Idle Thread Rule will not appear in the View Active Threads screen.

16.2.2 Configuring JVMs and Pools

You can group sets of JVMs into JVM pools that provide monitoring information across all related JVMs in a single view. You can view all the JVM pools in the WebLogic Domain, create a new JVM pool, and edit existing JVM pools.

1. From the **Setup** menu, select **Middleware Management**, then select **Application Performance Management**. Select the JVM Diagnostics Engine row in the Application Performance Management Engines table and click **Configure**.
2. Click the **JVMs and Pools** tab. The list of all available JVM pools is displayed. For each pool, you can set the Poll Enabled flag and specify the Poll Interval. If the **Polling Enabled** flag is set to **Y**, JVMs belonging to this pool will be polled for active requests periodically based on the Poll Interval.

Figure 16–2 Configure JVMs and Pools

JVMD Diagnostics Setup Page Refreshed May 6, 2013 11:27:49 AM PDT

JVMD Configuration **JVMs and Pools** Register Databases Heap Analysis Hosts

JVMs and Pools

View

Name	Target Name	Type	Status	Member Status Summary	Ver
Default	Default	Java Virtual Machine Pool	n/a	0 1 0 0	
EMGC_DOMAIN_jvmpool	/EMGC_EMGC_DOMAIN/EMGC_DOMAIN_jvmpool	Java Virtual Machine Pool	n/a	0 1 0 0	
EMGC_OMS1_jvm	/EMGC_EMGC_DOMAIN/EMGC_DOMAIN/EMGC_OMS1_jvm	Java Virtual Machine	↑	0 0 0 0	1.6

Columns Hidden: 5

3. Click **Create Pool** to create a new pool.
 - a. In the Add JVM Pool Information page, enter the name and description of the JVM pool.
 - b. In the **Poll Interval** field, specify the sample interval for JVMs belonging to this pool when monitoring (polling) is enabled.
 - c. Check the **Poll Enabled** check box to poll the JVMs belonging to this pool.
 - d. Click **Save** to save the JVM Pool information.
4. To delete a pool, check the select check box and click the **Remove** icon. You cannot delete the **Default** and **Other** pools.
5. Click **Downloads** to download JVM Diagnostics components. You can download the following components:
 - **JVMD Agent:** Contains JVM Diagnostics Agent binaries for all supported platforms.
 - **LoadHeap:** Contains scripts to upload heap snapshots to the repository.
 - **JVMD Agent with MDA (WebLogic only):** Contains JVM Diagnostics and MDA Agents for the Weblogic server on supported platforms.
 - **DB Agent:** Contains supported platforms executables for establishing cross tier analysis for legacy database versions.
 - **JVMD Engine:** Collects JVM metrics from JVM Diagnostics Agents for real time view and historical data.

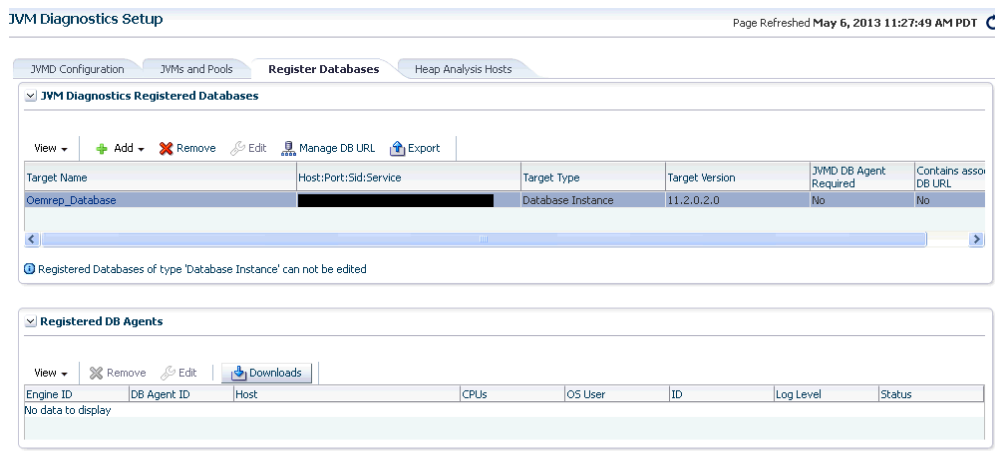
6. Select a JVM Pool and click **Configure**. See [Section 16.4.4, "Configuring a JVM Pool"](#) for details.
7. Select a JVM and click **Configure**. See [Section 16.5.9, "Configuring a JVM"](#) for details.

16.2.3 Register Databases

Follow these steps to register the database:

1. From the **Setup** menu, select **Middleware Management**, then select **Application Performance Management**. Select the JVM Diagnostics Engine row in the Application Performance Management Engines table and click **Configure**.
2. Click the **Register Databases** tab. The JVM Diagnostics Registered Databases page appears.

Figure 16–3 Register Databases



3. The list of registered databases is displayed. The database name, host, Oracle SID for the monitored database, and listener port number is displayed. You will also see a flag indicating whether the database agent is required.
4. You can do the following:
 - **Add a Database Instance:** From the Add menu, select **Database Instance** to register an Enterprise Manager database target.
 - **Add a Custom Database:** From the Add menu, select **Custom Database** to register an external database target. Specify the Name, Host, Port, SID, Instance ID, Username, and Password and click **Test Connection** to validate the database details. After the validation, click **OK** to register the database.
 - **Remove:** Select a database from list and click **Remove** to remove a registered database.
 - **Manage DB URL:** Use this option to establish cross tier correlation between JVM Diagnostics and Database Diagnostics. Select a database and click **Manage DB URL**. In the **Associate / Disassociate a Registered Database**, select a Database URL and click **Add** and specify the URL of the database to be associated.

- **Edit:** You cannot edit a Database Instance. Only custom databases can be edited. Select a custom database from the list and click **Edit**.

Note:

- The DB User must have select privileges on the GV_\$SESSION, GV_\$SESSION_WAIT, GV_\$PROCESS, GV_\$SQLTEXT, GV_\$SQLAREA, GV_\$LOCK, and GV_\$LATCHNAME fixed views in the target database.

To grant select privileges to an user such as jvmdadmin, enter a command as follows:

```
SQL> grant select on SYS.GV_$LATCHNAME to jvmdadmin
```

- Multiple registrations may be necessary for a single database agent if different database users are running multiple applications.
 - For Oracle RAC databases, each node must be registered.
-
5. After the database has been registered, the JVM Diagnostics Engine will start monitoring the cross-tier JVM calls between applications being monitored for a particular JVM and the underlying database.
 6. In the Registered DB Agents region, the following details are displayed:
 - **Manager ID:** The ID of the Manager to which this DB Agent is connected.
 - **DB ID:** This is a unique ID given to the DB Agent and is used to identify the DB Agent in all the processes.
 - **Host:** The host on which the DB Agent is running.
 - **CPUs:** The number of CPUs.
 - **OS User:** The user who started the DB Agent.
 - **ID:** The port on which the DB Agent is running (Default is 5555).
 - **Log Level:** The log level of the DB Agent.
 - **Status:** The status of the DB Agent (active or inactive)
 7. Check the **Select** check box and click the **Edit** icon to edit the DB Agent.
 8. Click **Downloads** to manually download the various binaries such as JVM Diagnostics Agent, JVM Diagnostics Engine, Database Agent, Load Heap, and deploy them. You can download the following:
 - **JVM Diagnostics Agent WAR File:** The JVM Diagnostics Agent Parameters web.xml parameters window is displayed. From the Available Managers drop-down, you can select entries that are in the format <host>:<port> - for normal communication, <host>:<port>(secure communication) for communication over the SSL Port or you can select Other. If you select Other, you need to specify the Manager IP Address and the Manager Port to which the JVM Diagnostics Agent is connecting to. While downloading the JVM Diagnostics Agent, you can modify the following parameters:
 - **Tuning Timeouts Parameters:** You can modify the Connection Retry Time, GC Wait Timeout, Long Request Timeout, and Idle Agent Timeout.

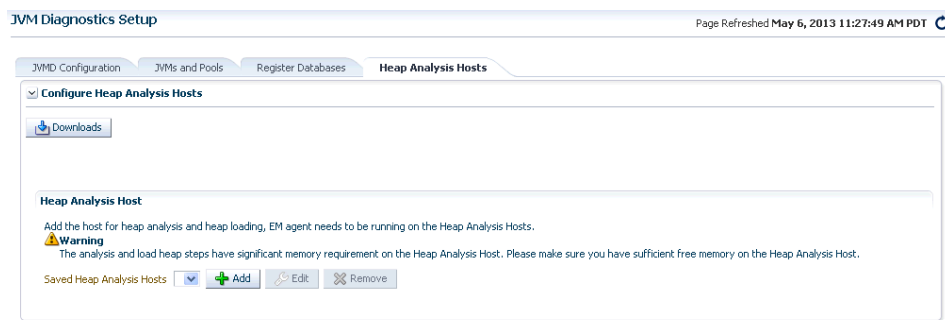
- **Target Association Parameters:** If you select WebLogic Server, you can specify the Target Name, and Pool Name. If you select Other Server, you can specify the Group ID Property and Pool Name.
- **Logging Parameters:** You can modify the Agent Log Level.
- **Optimization Level:** You can modify the Optimization Level.
- **JVM Diagnostics Engine EAR File:** You can open the jammanager.ear file or save it to a specified location.
- **Load Heap:** The loadheap.zip is saved to a specified location.
- **DB Agent:** Specify the DB Agent parameters. In the Available Engines drop-down, you can select entries that are in the format <host>:<port> - for normal communication, <host>:<port>(secure communication) for communication over the SSL Port, or you can select Other. If you select Other, you need to specify the Manager IP Address and the Manager Port to which the JVM Diagnostics Agent is connecting to. Specify the Agent Log Level in the Logging Parameters region and click **Download**.

16.2.4 Configuring the Heap Analysis Hosts

To configure the heap analysis hosts, follow these steps:

1. From the **Setup** menu, select **Middleware Management**, then select **Application Performance Management**. Select the JVM Diagnostics Engine row in the Application Performance Management Engines table and click **Configure**.
2. Click the **Heap Analysis Hosts** tab. The Configure Heap Analysis Hosts page appears.

Figure 16–4 Configuring the Heap Analysis Hosts



3. Click **Downloads** and select **Loadheap** to download the loadheap.zip file. This file contains scripts that can be used to upload heap snapshots to the JVM Diagnostics repository.
4. To configure a heap analysis host, click **Add** and enter the following details:
 - **Alias:** Enter an alias for the host.
 - **Heap Analysis Host:** The heap analysis host on which the Management Agent has been deployed.
5. Click **Save** to save the configuration.

16.2.5 Viewing Registered JVMs and Managers

Follow these steps to view a list of registered JVMs and JVM Managers:

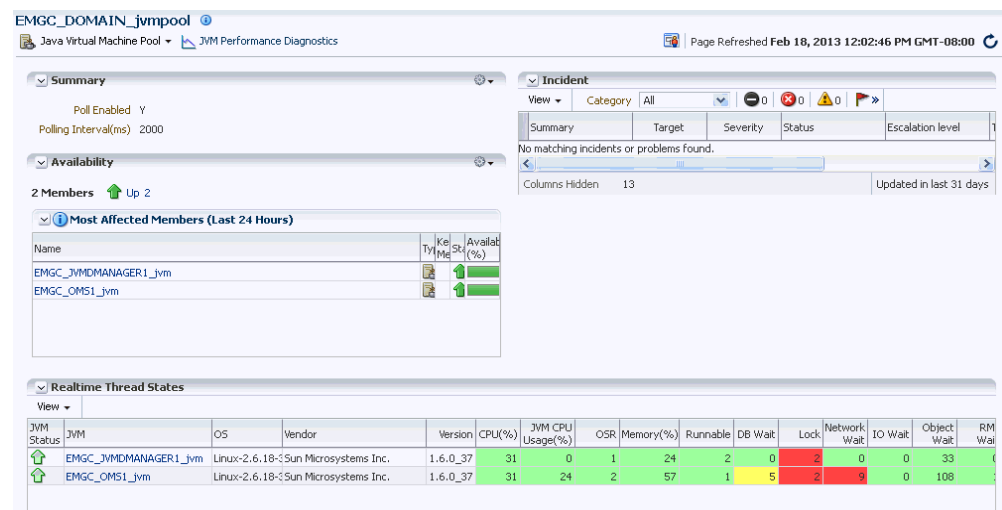
1. From the **Setup** menu, select **Middleware Management**, then select **Application Performance Management**.
2. The list of registered JVM Diagnostics Engines are displayed.
 - Name: The name assigned to the JVM Diagnostics Engine. This ID identifies the JVM Diagnostics Engine in all the processes.
 - Host: The machine on which the JVM Diagnostics Engine has been deployed.
 - Port: The port of the machine on which the JVM Diagnostics Engine has been deployed.
 - SSL Port: The SSL Port of the machine on which the JVM Diagnostics Engine has been deployed.
 - Build Version: The build version of this JVM Diagnostics Engine.
 - Status: The status of the JVM Diagnostics Engine (Active/Inactive)

Select the JVM Diagnostics Engines row and click **Configure** to configure the JVM Diagnostics Engine parameters, JVMs and Pools, databases, and heap loader.

16.3 Accessing the JVM Diagnostics Pages

After you have deployed the JVM Diagnostics Engine and configured JVM Diagnostics, you can start using the features. Navigate to the **Middleware** tab in Enterprise Manager and click on a JVM Pool or Java Virtual Machine target. The Home page for the target is displayed.

Figure 16–5 JVM Pool Home Page



To start using JVM Diagnostics, select the JVM Diagnostics option from the JVM Pool drop-down menu and choose the appropriate option.

You can also access the JVM Diagnostics pages from the WebLogic Server, WebLogic Domain, JBoss Server, or Cluster target Home pages. To do so, click on a target to

navigate to the Home page. From the **Target** menu, select **Diagnostics**, then select the appropriate JVM Diagnostics menu option.

16.4 Managing JVM Pools

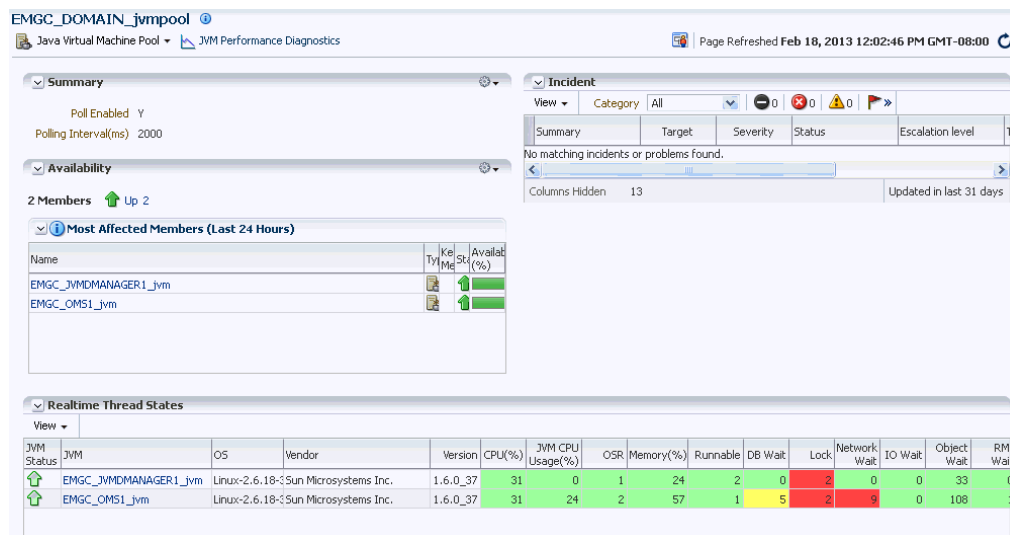
You can group sets of JVMs into JVM pools that provide monitoring information across all related JVMs in a single view. You can monitor all the JVMs in a pool, view historical and real time data for the JVM pool, manage threads and heap snapshots, create a new pool, and edit an existing JVM pool. JVMs and JVM Pools are now targets in Enterprise Manager. You can do the following:

- [Viewing the JVM Pool Home Page](#)
- [Viewing the JVM Pool Performance Diagnostics Page](#)
- [Viewing the JVM Pool Live Thread Analysis Page](#)
- [Managing Thread Snapshots](#)
- [Analyzing Heap Snapshots](#)
- [Configuring a JVM Pool](#)
- [Removing a JVM Pool](#)
- [Add to Group](#)

16.4.1 Viewing the JVM Pool Home Page

The JVM Pool Home page shows the details of all JVMs in the pool.

Figure 16–6 JVM Pool Home Page



It shows the following details:

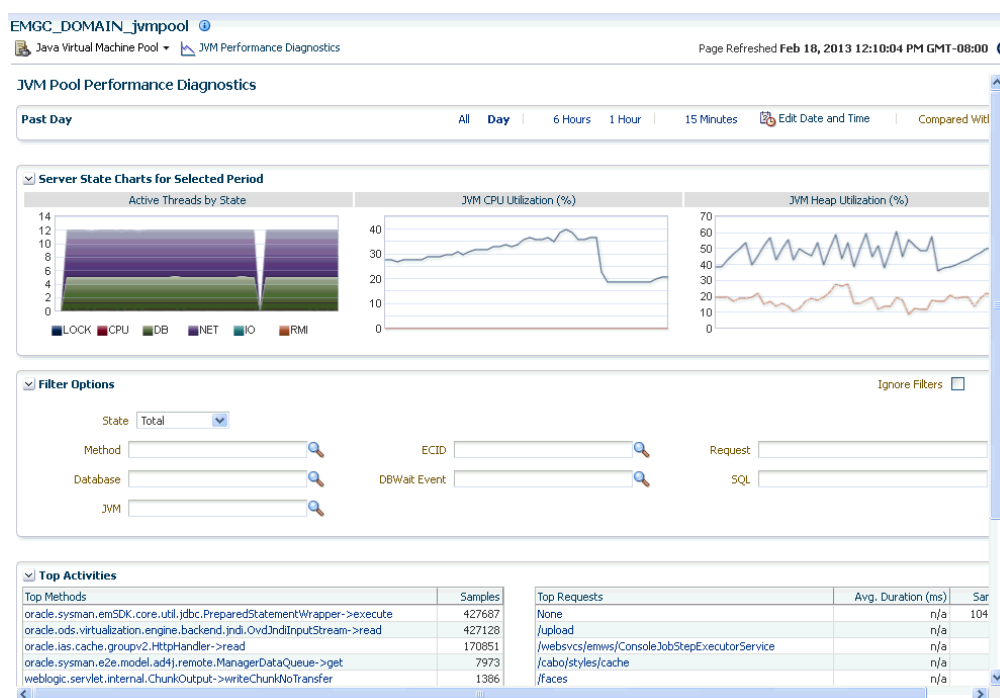
- **Summary:** Shows whether polling is enabled and the Polling Interval.
- **Availability:** This region shows the availability status of the members in the JVM Pool. Click on a Member link to drill down JVM Home Page.

- **Incident:** This region shows any open incidents that have occurred, the type, and category of the incident. Click the Summary link to drill down to the Incident Details page.
- **Realtime Thread States:** This region shows the realtime thread status for each JVM in the pool. The current activity of the JVM including CPU usage, memory, number of threads waiting for a database response, number of threads waiting for synchronization lock, and other details are displayed. Click on the JVM link to drill down to the JVM Performance Diagnostics page.

16.4.2 Viewing the JVM Pool Performance Diagnostics Page

You can view the summary and detailed information of the selected JVM Pool on this page. You can also compare the JVM pool data across two specific time periods. To view this page, select **Middleware** from the **Targets** menu and click on a JVM Pool target. Select the **JVM Performance Diagnostics** option from the **Java Virtual Machine Pool** menu or click on the **JVM Performance Diagnostics** link next to the Java Virtual Machine Pool menu.

Figure 16–7 JVM Pool Performance Diagnostics Page



This page shows the summary details of the JVM pools which include the Server State Charts, and a list of Top Methods and Top Requests. You can view the Server State Charts, list of Top Methods, Top Requests, Top DBStates, and Top SQLs.

You can filter the data that is displayed by specifying various criteria such as Method Name, JVM Name, Thread State, DBState, and so on. You can also view server state charts for each JVM in the pool. Check the **Ignore Filters** checkbox if you want to ignore the specified filters.

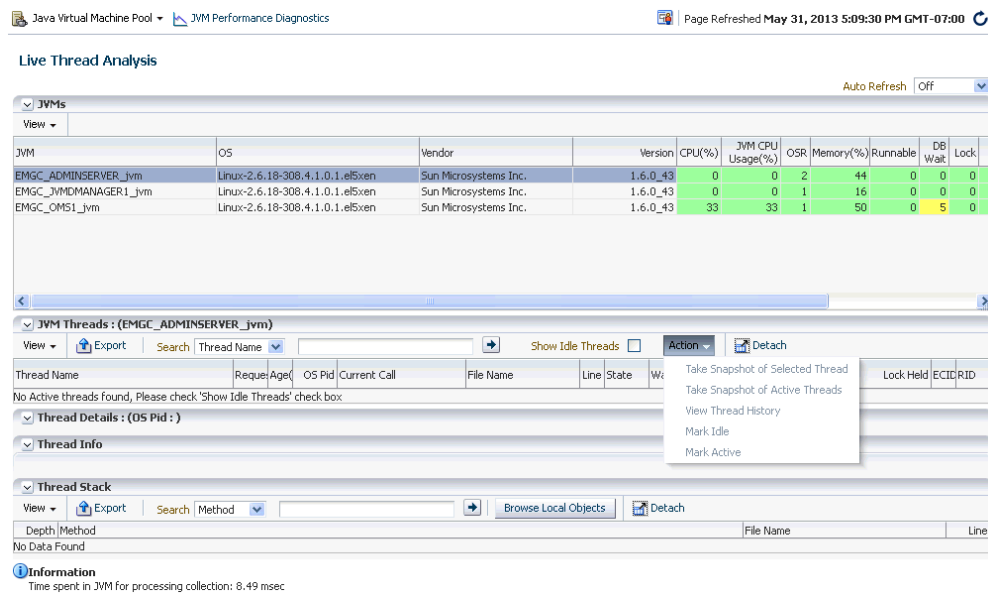
In the Top Activities region, you can see the top methods, top requests, top DBWait events, top SQLs, and top databases in the JVM Pool. You can click on the link to drill down to the Details page.

You can compare data between two periods. Select the **Compared With** checkbox and specify the comparison period to view the comparison data.

16.4.3 Viewing the JVM Pool Live Thread Analysis Page

This page shows the real-time data for all the JVMs in the selected pool. This data is useful in analyzing the various active and idle threads on the JVM. During analysis you can drill down from the thread level to methods used in the thread to local variables that are part of the method. From the **Targets** menu, select **Middleware**, and then click on a Java Virtual Machine Pool target. Select the **Live Thread Analysis** option from the **Java Virtual Machine Pool** menu. The JVM Pool Live Thread Analysis Page appears.

Figure 16–8 JVM Pool Live Thread Analysis Page



This page shows the following:

- **JVMs:** This table shows the list of JVMs and the current status of each JVM. The current activity of the JVM including CPU usage, memory, number of threads waiting for a database response, number of threads waiting for synchronization lock, and other details are displayed
- **JVM Threads:** This table shows a list of all the threads running in the JVM. For each thread, the following details are displayed:
 - Thread Name: Name of the thread. Click on the link to view the JVM Stack.
 - Request: Application request being processed by the thread.
 - OS PID: Operating system and Thread IDs for this thread.
 - Current Call: Lowest user method being executed by the thread.
 - File Name: Name of the file that contains the class and method for the current call.
 - Line: Line number in the method currently being executed.

- **State:** The current state of the thread. This can be DB Wait, RMI Wait, or Network Wait. If the ADP Agent is running, the Request Name and Request Age are displayed. If a thread is in the **DB Wait State**, the Waiting On column displays the name of the database the thread is waiting on, and the time the thread has to wait is displayed in the Wait Time column.

Click on the link to view the database diagnostics details. See [Section 16.5.4.1, "Cross Tier Analysis"](#) for more details. You can track database issues and determine the application request responsible for the database activity. You can also view the complete call stack including the method and line number information.

Note: To view the database diagnostics details, ensure that:

- The JVM Diagnostics Agent is running on the JVM that initiated the request.
 - The monitored database must be registered by the JVM Diagnostics Engine.
-

You can do the following:

- **Export:** Select a thread and click Export to export the thread details along with thread stacks information to an Excel file.
- **Show Idle Threads:** Select this check box to list all the Idle Threads in the JVM Threads table.
- **Take Snapshot of Selected Thread:** Select a thread in the list, and from the **Action** menu, select **Take Snapshot of Selected Thread**. The Thread Snapshot page is displayed. You can configure the snapshot settings and click **Take Thread Snapshot**. A snapshot file with details of the selected thread is generated. From the **Java Virtual Machine Pool** menu, select **Thread Snapshots** to view additional details.
- **Take Snapshot of Active Threads:** This option allows you to take a snapshot of all the active threads. From the **Action** menu, select **Take Snapshot of Active Threads**, the Thread Snapshot page is displayed. You can configure the snapshot settings and click **Take Thread Snapshot**. A snapshot file with details of all the active threads is generated. From the **Java Virtual Machine Pool** menu, select **Thread Snapshots** to view additional details.
- **View Thread History:** Select a thread and from the **Action** menu, select **View Thread History**. The historical data for the thread for the last 30 minutes is displayed.
- **Mark Idle:** Select a thread and from the **Action** menu, select **Mark Idle**. The selected thread will be marked as Idle based on the Idle Thread Rule.
- **Mark Active:** Select an Idle thread and from the **Action** menu, select **Mark Active** to change the status to Active.
- **Thread Info:** This region shows the detailed information for a selected thread. Details of thread including Current Call, Request, ECID, State, Waiting On, and Wait Request are displayed. If the thread is in the DB Wait State, click on the link to drill down to the Database Home page
- **Thread Stack:** The Thread Stack table shows the details of the selected thread such as the depth of the thread, methods used in the thread, file where the method is used, and the line number. You can drill down from the method level to a lower

level. Select a method from the table and click **Browse Local Objects**. A popup window is displayed which shows the local variables, objects, their classes, and values. You can export these details to a file by clicking **Export**. You are prompted to specify the directory in which the file is to be stored. Enter the path and click **Save** to save file in .csv format.

You can refresh the data that is displayed by specifying the **Auto Refresh** period.

16.4.4 Configuring a JVM Pool

From the **Targets** menu, select **Middleware**, and then click on a Java Virtual Machine Pool target. Select the **Configure JVM Pool** option from the **Java Virtual Machine Pool** menu. You can do the following:

- Modify the JVM pool details. You can enable or disable monitoring of pools or change their polling intervals by updating the pool properties. Click **Save** to save the changes.
- Configure the JVM pool thresholds. See [Section 16.4.4.1, "Updating Pool Thresholds"](#).

16.4.4.1 Updating Pool Thresholds

Follow these steps to edit the pool thresholds:

1. From the **Setup** menu, select **Application Performance Management**. Select the JVM Diagnostics Engine row in the Application Performance Management Engines table and click **Configure**.
2. Click the **JVMs and Pools** tab to view and configure the JVMs and JVM Pools.
3. Select a JVM Pool and click **Configure**. The Edit JVM Pool Information page appears.

Figure 16–9 Edit JVM Pool Threshold Values

The screenshot shows two panels from the Oracle Enterprise Manager Cloud Control interface. The top panel, titled 'Edit JVM Pool Information : /EMGC_EMGC_DOMAIN/EMGC_DOMAIN_jvmpool', contains fields for 'Pool' (set to /EMGC_EMGC_DOMAIN/EMGC_DOMAIN_jvmpool), 'Description' (set to /EMGC_EMGC_DOMAIN/EMGC_DOMAIN_jvmpool), 'Poll Interval (ms)' (set to 2000), and a checked 'Poll Enabled' checkbox. The bottom panel, titled 'Edit JVM Pool threshold Values', displays a table of thresholds for the selected JVM pool.

Metric	Comparison Operator	Threshold	Trigger Samples	Corrective Actions
EMGC_EMGC_DOMAIN/EMGC_DOMAIN_jvmpool				
Host CPU Usage (%)				
Critical	>=	90	50	<None> + Add
Warning	>=	70	50	<None> + Add
DB Wait (thread count)				
Critical	>=	6	50	<None> + Add
Warning	>=	3	50	<None> + Add
GC Overhead(%)				
Critical	>=	90	50	<None> + Add

4. In the Edit JVM Pool Threshold Values region, the following details are displayed:
 - Level: Thresholds violations can have a level of R (red) or Y (yellow).
 - Metric: The attribute or metric that is being monitored.

- **Threshold:** The Critical and Warning threshold for the metric. A violation occurs when the threshold is exceeded after a minimum number of samples have been monitored.
5. Click **Add** to add a corrective action. Select a corrective action from the list and click **OK**. You can select:
- **No Action:** No corrective is defined.
 - **Trace Dump:** Select this option to trace a particular thread, or all active threads in response to a threshold violation. You can define the following parameters:
 - * **Poll Interval:** Interval after which snapshot should be repeated.
 - * **Poll Duration:** Duration for which the snapshot should be taken.
 - * **Thread Details:** You can specify if the thread details need be included in the snapshot.
 - * **Try Changing Threads:** Sometimes the stack associated with the thread may change rapidly which makes taking the snapshot difficult. If you select this parameter, you can suspend the thread and take the snapshot.
 - * **Include Network Waits:** Specify if network wait threads need to be included in the snapshot.
 - * **All Threads:** Specify if all threads (active and idle) must be included in the snapshot.
 - * **Allow Trace Interrupt:** Indicate whether the trace process can be interrupted.
 - **Heap Dump:** Select this option to generate a heap dump in response to a threshold violation. The Heap Snapshot Type can be:
 - **Txt:** Text (txt) for analysis in JVM Diagnostics.
 - **HPROF:** Binary format for analysis with external tools.

If a corrective action (trace dump or heap dump) is generated due to a threshold violation, the trace or heap dump files are displayed in the Event Details page. See [Section 16.10, "Viewing JVM Diagnostics Threshold Violations"](#).
6. Click **Save** to save the threshold values.

16.4.5 Removing a JVM Pool

You will see a warning message if you select the **Remove Target** option from the JVM Pool menu. The message displays the name of the target being deleted and that when a pool is deleted, all the JVM targets in the pool are also displayed. Click **Yes** to delete the JVM Pool or **No** to return to the JVM Pool Home page.

16.4.6 Add to Group

Select this option to add the JVM Pool to one or more groups. A pop-up window appears with a list of groups on which you have Operator privileges. Select one or more groups and click **Add** to add the target to the group.

16.5 Managing JVMs

You can monitor a specific JVM in a pool, view historical and real time data, and so on. You can do the following:

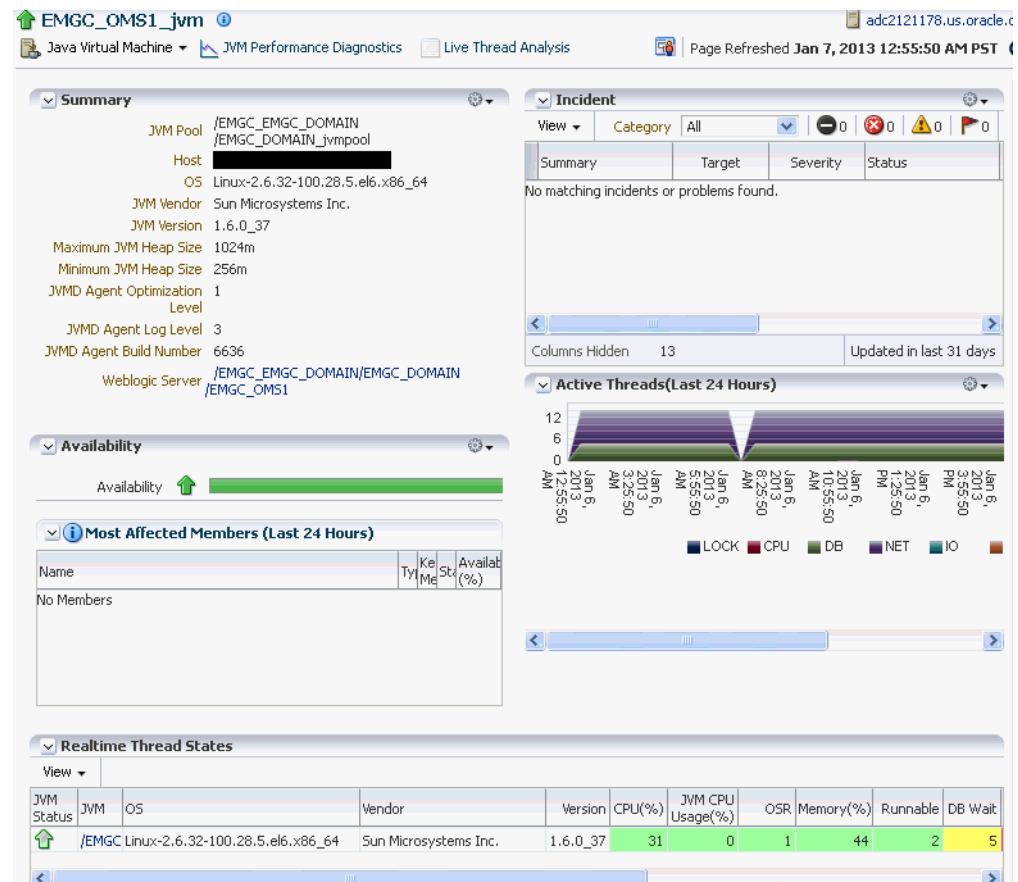
- [Viewing the JVM Home Page](#)
- [Viewing the JVM Performance Diagnostics Page](#)
- [Viewing the JVM Diagnostics Performance Summary](#)
- [Viewing the JVM Live Thread Analysis Page](#)
- [Viewing the JVM Live Heap Analysis Page](#)
- [Managing Thread Snapshots](#)
- [Analyzing Heap Snapshots](#)
- [Configuring a JVM](#)
- [Removing a JVM](#)
- [Add JVM to Group](#)

16.5.1 Viewing the JVM Home Page

The JVM Home page shows the summary and configuration information of all the JVMs in the JVM pool. Follow these steps to view the JVM Home page:

1. From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine target.

Figure 16–10 JVM Home Page



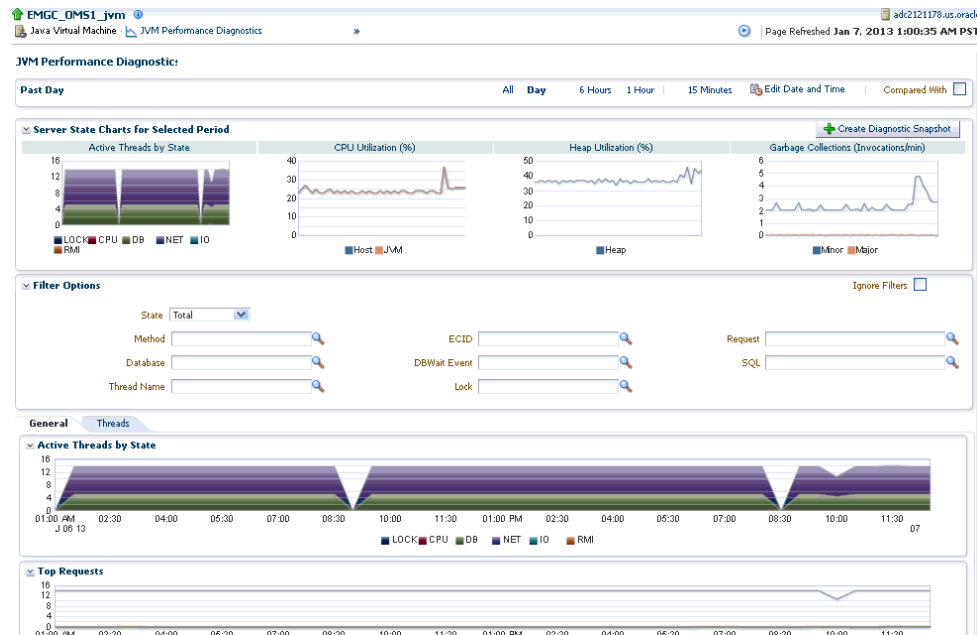
2. The JVM Home page with the following details is displayed.

- Summary: Shows details of the JVM such as the JVM Pool it belongs to, the host, Agent Optimization Level, Agent Log Level, JVM Version, and vendor details.
- Incident: This region shows any open incidents that have occurred, the type, and category of the incident. Click on the **Summary** link to drill down to the Incident Details page.
- Availability: The availability status of the JVM. Click on a **Member** link to drill down JVM Home Page.
- Active Threads: The chart shows the number of active threads in the JVM in the last 24 hours.
- Realtime Thread States: Shows the state of the various threads in the JVM in the color coded columns. The current activity of the JVM including CPU usage, memory, number of threads waiting for a database response, network response, or number of threads waiting for synchronization lock, idle threads, and configuration changes are displayed. Click on a JVM to view the list of threads in the JVM and the details of each thread.

16.5.2 Viewing the JVM Performance Diagnostics Page

This page shows the summary and detailed information for a specific JVM. To view this page, select **Middleware** from the **Targets** menu and click on a Java Virtual Machine target. Select the **JVM Performance Diagnostics** option from the **Java Virtual Machine** menu.

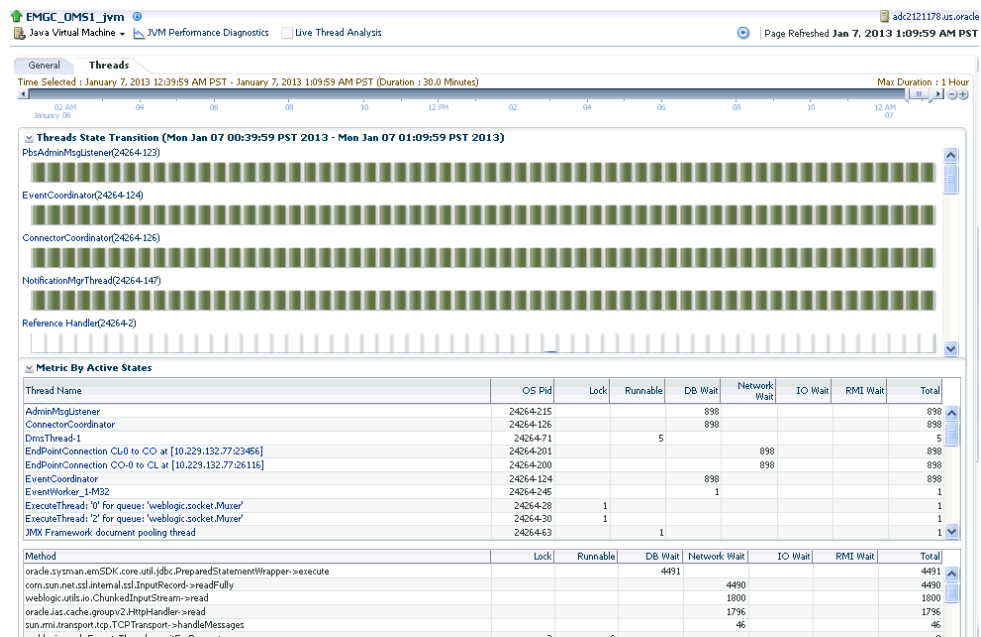
Figure 16–11 JVM Performance Diagnostics Page



This page shows the summary details of the JVM which include the Server State Charts, Active Threads by State, Top Methods, Top Requests, Top DBWait Events, Top SQLs, and Top Databases. You can filter the data that is displayed by specifying various criteria such as Method Name, JVM Name, Thread State, DBState, Lock, and so on.

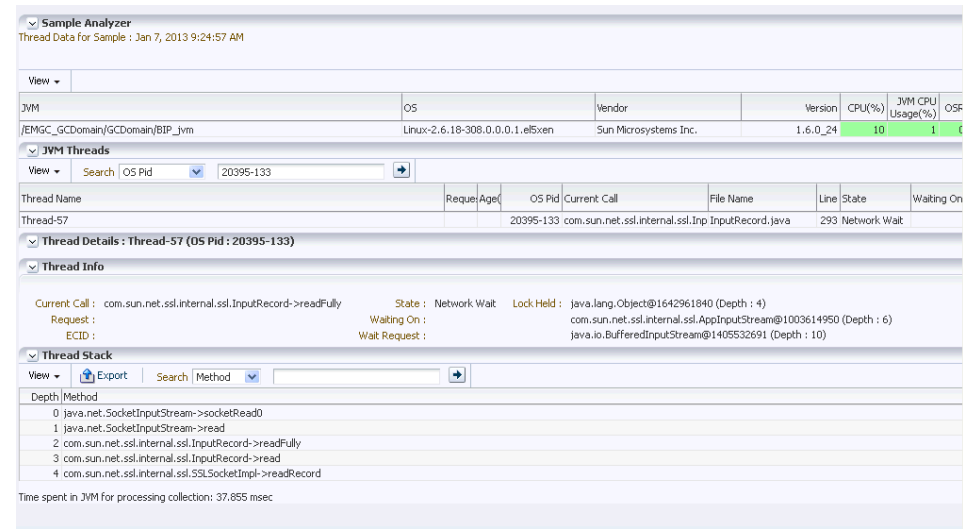
Click on the **Threads** tab to view the Thread State Transition chart. This chart shows how the threads have transitioned from one state to the other in the selected period. You can change the time interval and move it to a different time period by using the quick time selection control at the top of the page. You can hover over the colored bars to see the transition changes from one state to the other, for example from **Runnable** to **Not Active** or to **Runnable**.

Figure 16–12 Thread State Transition



Click on a bar graph in the **Thread State Transition** chart to view the Sample Analyzer which provides a detailed analysis on the state of the thread. This feature allows you to analyze each sample (JVM snapshot at a specific time) in the monitored data.

Figure 16–13 Sample Analyzer



In the JVM Performance Diagnostics page (Figure 16–11), you can do the following:

- You can compare data between two periods. Select the **Compared With** checkbox and specify the comparison period. The data for the selected period and the current period is displayed.


- Click the **Create Diagnostics Snapshot** link to collect diagnostic data for the JVM target for a specific period and analyze this data in offline. For more details, see [Section 16.9, "JVM Offline Diagnostics"](#)

16.5.3 Viewing the JVM Diagnostics Performance Summary

You can view the performance metrics (system and active threads) for a JVM target on the Performance Summary page. A set of charts are displayed on this page for the JVM target. To view the JVM performance metrics, follow these steps:

- From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine target.
- Select the **Performance Summary** option from the **Java Virtual Machine** menu. The following page appears:

Figure 16–14 Performance Summary Page



- A set of default charts that show the values of the JVM performance metrics over a period of time are displayed. Review the metrics for any periods of time where the Warning or Critical Thresholds were reached.

If any of the metrics exceed the Warning Thresholds or Critical Thresholds, it could indicate memory is a factor in the JVM performance and availability. It could mean there is a memory leak or that the JVM heap configuration is too small for the application load. If the heap configuration is correct, you must review the real time heap data. See [Section 16.5.5, "Viewing the JVM Live Heap Analysis Page"](#) for details. You can then create a snapshot that can be examined for leaks. See [Section 16.5.7, "Taking a Heap Snapshot"](#) for details.

- Click the **Show Metric Palette** button. The metric palette has a folder for the current target (JVM) and the related targets. You can add or remove metric charts. Leaf nodes act as check boxes. Clicking a leaf node on causes a chart to be added.

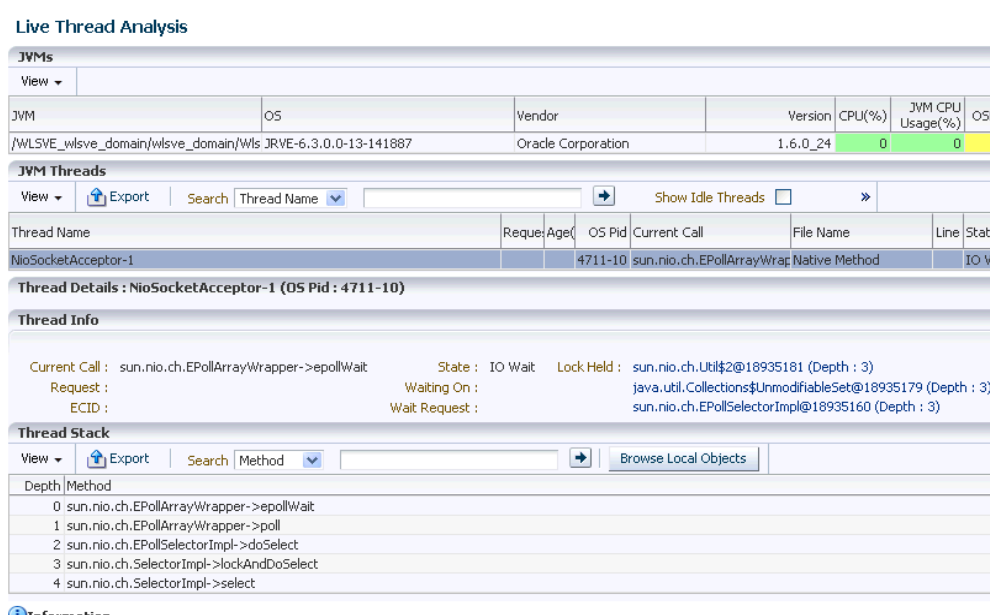
Clicking it off removes the metric. Dragging a leaf node from the palette to a chart or legend adds the metric to that chart.

16.5.4 Viewing the JVM Live Thread Analysis Page

This page shows the real-time data for a selected JVM. This data is useful in analyzing the various active and idle threads on the JVM. During analysis you can drill down from the thread level to methods used in the thread, to local variables that are part of the method. To view this page, select **Middleware** from the **Targets** menu and click on a Java Virtual Machine target. Select the **Live Thread Analysis** option from the **Java Virtual Machine** menu.

You can also access this page by clicking the **Live Thread Analysis** link on top left corner of the page.

Figure 16–15 JVM Live Thread Analysis Page



This page shows the following:

- **JVM Threads:** This table shows a list of all the threads running in the JVM. Click on a thread to view the thread details in the Thread Info table. For each thread, the Thread Name, Request, Age, OS PID, Current Call, File Name, Line, State, Waiting On, Wait Time, Lock Held, ECID.

If the ADP or BTM Agent is running, the Request Name and Request Age is displayed. If a thread is in the **DB Wait State**, the Waiting On column displays the name of the database the thread is waiting on, and the time the thread has to wait is displayed in the Wait Time column. You can click on the link in the DB Wait column to view the database diagnostics details. This is helpful in tracking database issues and determining the application request responsible for the database activity.

Note: You can view the database diagnostics details if:

- The JVM Diagnostics Agent is running on the JVM that initiated the request.
 - The monitored database must be registered by the JVM Diagnostics Engine.
-

You can perform the following actions:

- **Take a Snapshot of a Selected Thread or Active Threads:** Select a thread from the list and choose the **Take Snapshot of a Selected Thread** option from the Actions menu. The Thread Snapshot page is displayed where you take a snapshot. If you select the **Take Snapshot of Active Threads** option, you can take a snapshot of all active threads running on this JVM. You can specify the following parameters for each snapshot:
 - * Poll Interval: Interval after which snapshot should be repeated.
 - * Poll Duration: Duration for which the snapshot should be taken.
 - * Thread Details: You can specify if the thread details need be included in the snapshot.
 - * Try Changing Threads: Sometimes the stack associated with the thread may change rapidly which makes taking the snapshot difficult. If you select this parameter, you can suspend the thread and take the snapshot.
 - * Include Network Waits: Specify if network wait threads need to be included in the snapshot.
 - * All Threads: Specify if all threads (active and idle) must be included in the snapshot.
 - * Allow Trace Interrupt: Indicate whether the trace process can be interrupted.
- **Mark Idle:** Select a thread from the list and click **Actions > Mark Idle** to mark a thread as idle.
- **Mark Active:** If you selected the Show Idle Threads check box, a list of idle threads is displayed. Select a thread and click **Actions > Mark Active** to mark it as an active thread.
- **Show Idle Threads:** Select this check box to list only the idle threads in the JVM Threads table.
- **Export:** Click **Export** to export the thread details along with thread stacks information to an Excel file.
- **View Thread History:** Select a thread and from the Action menu, select **View Thread History**. The historical data for the thread for the last 30 minutes is displayed.
- **Thread Info:** This section shows the detailed information for a selected thread. Details of thread including Current Call, Request, ECID, State, Waiting On, and Wait Request are displayed. If the thread is in the DB Wait State, click on the link to drill down to the Database Home page. See [Section 16.5.4.1, "Cross Tier Analysis"](#) for more details.
- **Thread Stack:** The Thread Stack table shows the details of the selected thread such as the depth of the thread, methods used in the thread, file where the method is

used, and the line number. You can drill down from the method level to a lower level. You can do the following:

- **Browse Local Objects:** Select a method from the table and click **Browse Local Objects**. A popup window is displayed which shows the local variables, objects, their classes, and values.
- **Export:** You can export the details of a selected thread to a file by clicking **Export**. You are prompted to specify the directory in which the file is to be stored. Enter the path and click **Save** to save file in .csv format.
- **Auto Refresh:** You can refresh the data that is displayed by specifying the Auto Refresh period.

16.5.4.1 Cross Tier Analysis

You can trace any JVM activity from the JVM thread to the database. You can view cross tier correlation for live threads and historical monitored data.

Before you establish cross tier correlation, ensure that the database is an Enterprise Manager target and has been registered with JVM Diagnostics. To register the database, select the **Application Performance Management** option from the **Setup** menu, then click **Setup JVM Diagnostics** in the Application Performance Management page. Click on the **Register Databases** tab. The list of registered database targets is displayed.

If the JVMD Agent Required field has a **No** value, you can proceed with the cross tier analysis. If the field has a **Yes** value, you must ensure that the root user or the same OS user who started the database must be running the JVM Diagnostics Database Agent on the target database machine. If the JVMD Agent Required field has a **Status Unavailable** value, you cannot perform cross tier analysis as the JDBC connection to the database cannot be established.

To view the cross tier correlation for live threads, follow these steps:

1. From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine target. Select the **Live Thread Analysis** option from the **Java Virtual Machine** menu. See [Figure 16–8](#).
2. In the JVM Threads column, select a thread with a DB Wait State.
3. The thread details are displayed in the Thread Info section. If cross tier correlation has been established, you can see `SID=<value>"SERIALNUM=<value>` when you hover over the State field. Click the **DB Wait** link.
4. The Database Details popup is displayed which shows the host, port, SID, user, and JDBC URL for the target database. Click the **Register All DB Targets** link to register all database targets with JVM Diagnostics and refresh the Live Thread Analysis page.

Oracle Database 11g Release 2 supports special cross tier requirements for JVM Diagnostics and cross tier correlation is automatically established when you click the **Register All DB Targets** link. If you are using an earlier Oracle Database version, the registration process prompts you to run the JVM Diagnostics Agent on the host machine.

5. Click **View Register Databases** to navigate to the Registered Databases page where you can manually register the database target with JVM Diagnostics and check the value in the DB Agent Required column. Click **Add**, to add a database instance or a custom database. If you register a database as a custom database, the DB Name is displayed in the Waiting field in the Threads Info section but the cross tier correlation cannot be established.

To view the cross tier correlation for historical monitored data, follow these steps:

1. From the Targets menu, select **Middleware**, then click on a Java Virtual Machine target. Select the **JVM Performance Diagnostics** option from the **Java Virtual Machine** menu.
2. The three tables Top Databases, Top SQLS, and Top DBWait Events related to the cross tier are displayed. The Top Databases table shows the top databases in which JVM or JVMs in the pool have activities. The Top DBWait Events table shows the top DB Wait events caused by the JVM threads in the database. The Top SQLs table shows the list of SQL calls sorted by the number of samples. Click a SQL call to view the charts for that call. Click the top field of a table to launch a popup which shows the names of the database on which activity was performed.
3. Click the Database Name link to drill down to the Database Diagnostics page which shows the corresponding database activity.
4. Click the Top SQLs and Top DBWait Events links to navigate to the SQL Details page and the ASH Viewer page of database diagnostics.

If cross tier correlation has been established, you can view JVM Diagnostics activities for a database (drilling up from Database Diagnostics to JVM Diagnostics). Click the JVM Diagnostics link in the Performance page to drill up to the JVM Performance Diagnostics page. Data relevant to the time interval, database and other filters is displayed.

16.5.4.2 JVM Diagnostics - Oracle Real Application Cluster Drill-Down

Oracle Real Application Cluster (Oracle RAC) databases have a complex configuration of database instances and listeners. User applications use Oracle RAC services to connect to the database instead of SIDs that are used for single instance databases. User applications can connect to Oracle RAC listeners that are listening on different machines than the actual database instances. For cross tier correlation to be established, all the listener and database instances must be discovered targets in Enterprise Manager. Cross tier correlation can be established by using either of the following options:

1. If all the database instances in the Oracle RAC have been discovered in Enterprise Manager, follow these steps to establish cross tier correlation:

- For every database instance in the Oracle RAC, add details of all the listeners servicing that database instance by specifying the `jvmd_db_listeners_additional_info` property whose value must be in the format "`<listener name:listener host:listener IP:listener port,...`".

For instance, if the Oracle RAC has two database instances, l1 and l2 which are served by listeners L1 and L2 where L1 is listening on hostname H1, ip IP1 and port PORT1 and L2 is listening on hostname H2, ip IP2, and port PORT2, the value for the `jvmd_db_listeners_additional_info` property must be specified as '`L1:H1:IP1:PORT1,L2:H2:IP2:PORT2`'. The `jvmd_db_listeners_additional_info` property is used to ensure that host name and port number used in the JDBC URL in the application running on the server monitored by JVM Diagnostics matches the right database instance.

For example, if the JDBC URL is

`"jdbc:oracle:thin:@xxxx.oracle.com:1521:sid"`, the `"jvmd_db_listeners_additional_info"` property for all the database instances in the Oracle RAC, to which this URL is connecting must be `"LISTENER:xxxx.oracle.com:IP:1521"`

- To add the `jvmd_db_listeners_additional_info` target property to the database instance, follow these steps for each database instance:
 - Enter the following command:


```
insert into mgmt_target_properties (TARGET_GUID, PROPERTY_NAME,
PROPERTY_VALUE) values ((select target_guid from mgmt_targets where
target_name='<DB Instance Name>' and target_type='<oracle_data-
base>'), 'jvmd_db_listeners_additional_info', LISTENER_
<host1>:<host2>:<IP_add1>:<Service1>, LISTENER_
<host3>:<host4>:<IP_add2>:<Service2>')
```
 - Restart the JVM Diagnostics Engine.
 - Remove (if already registered) and re-register the database instance target (member of RAC) with JVM Diagnostics.
 - Ensure that this database instance is not registered as a custom database with JVM Diagnostics.
- 2. If all the database instances in the Oracle RAC have not been discovered in Enterprise Manager:
 - Register the database instances with JVM Diagnostics as custom databases. To register the custom databases, use the same hostname, IP address, and port number, as specified in the JDBC URL (in the application running on the server monitored by JVM Diagnostics).

16.5.5 Viewing the JVM Live Heap Analysis Page

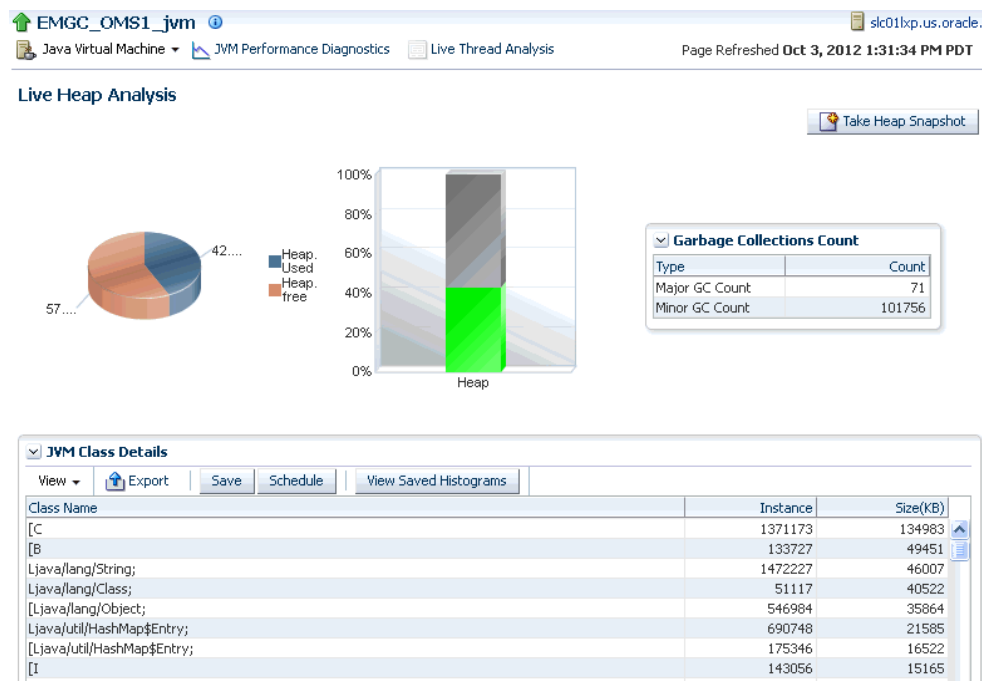
This page shows the real time organization of all objects in the JVM Heap. To view this page, select **Middleware** from the **Targets** menu and click on a JVM Pool target. Select the **Live Heap Analysis** option from the **Java Virtual Machine** menu.

If the JVM is running at Optimization Level 0, the following details are displayed:

- Pie Charts and Bar Charts that show the percentage utilization of the entire heap (free versus used)
- Major and Minor Garbage Collections Count
- Options to take snapshots of active threads or a specific thread.
- JVM Class Details table that displays all the classes in the JVM Heap in decreasing order of their size (in KB). You can export this data to an .xls file using the **Export** option.

Note: For JVMs running at optimization level, the following details are displayed:

- JVM Heap Memory Usage table where the usage (in KB) in various heap spaces.
 - JVM Heap Number of Objects table which displays the number of objects in various heap spaces.
-

Figure 16–16 JVM Live Heap Analysis Page

The following details are displayed:

- **Garbage Collections:** The number of objects that have been added to the garbage collection. The type of garbage collection i.e. minor or major, and the number of garbage collections of a particular type is displayed.
- **JVM Class Details:** This table provides a summary of the heap usage by different types of objects in the heap.
 - **Class Name:** The name of the space within the JVM heap.
 - **Instance:** The number of heap objects for number of instances of classes in a heap space.
 - **Size:** The size of the JVM heap.

You can do the following:

- **Take Heap Snapshot:** Click **Take Heap Snapshot** to take a heap snapshot. See [Section 16.5.7, "Taking a Heap Snapshot"](#) for details.
- **Manage Class Histograms:** A class histogram is displayed in the form of a table when the optimization level of the jamagent is 0. The histogram displays the top 300 data rows sorted by the size. You can perform several operations with histograms. For more details, see [Section 16.5.6, "Working with Class Histograms"](#).

16.5.6 Working with Class Histograms

A class histogram is displayed in the form of a table when the optimization level of the jamagent is 0. The histogram displays the top 300 data rows sorted by the size. You can perform various operations on class histograms. This section describes the following:

- [Saving a Class Histogram](#)
- [Viewing Saved Histograms](#)

- [Scheduling a Histogram Job](#)
- [Comparing Class Histograms](#)
- [Deleting Class Histograms](#)

16.5.6.1 Saving a Class Histogram

To save a class histogram, follow these steps:

1. From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine target. Select the **Live Heap Analysis** option from the Java Virtual Machine menu.
2. In the Live Heap Analysis page, scroll down to the JVM Class Details table. Click **Save**.
3. In the Save Class Histogram window, enter a name for the snapshot and a description and click **OK**.

16.5.6.2 Viewing Saved Histograms

To view saved histograms, follow these steps:

1. From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine target. Select the **Live Heap Analysis** option from the Java Virtual Machine menu.
2. In the Live Heap Analysis page, scroll down to the JVM Class Details table. Click **View Saved Histograms**. The Available Heap Snapshots page appears.
3. Scroll down to the Available Class Histograms table to view a list of saved class histograms.

16.5.6.3 Scheduling a Histogram Job

Scheduling will allow you to insert JVM Class Histogram data into the repository by running the job at the defined time. To schedule a class histogram job, follow these steps:

1. From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine target. Select the **Live Heap Analysis** option from the Java Virtual Machine menu.
2. In the Live Heap Analysis page, scroll down to the JVM Class Details table. Click **Schedule**. The Schedule Settings job appears.
3. Enter a name and description for the job to be scheduled.
4. Specify the schedule as **Immediate** or **Later**. If you select **Later**, you can specify if the job needs to be run only once or repeated at specified intervals.
5. Select the frequency at which you want to repeat the job from the **Repeat** drop-down list.
6. Select the option for the Grace Period. If you select the grace period, the job will remain active and run within the specified grace period.
7. Click **OK** to schedule the histogram job. A confirmation window appears indicating that the job has successfully submitted.

To view the job status, from the Enterprise menu, select **Job**, then select **Activity**. Select the Job Type as **All**, and Target Type as **Targetless** to see the histogram job.

16.5.6.4 Comparing Class Histograms

The compare functionality allows you to compare any two class histogram snapshots listed in the table. To compare class histograms, follow these steps:

1. From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine target. Select the **Live Heap Analysis** option from the Java Virtual Machine menu.
2. In the Live Heap Analysis page, scroll down to the JVM Class Details table. Click **View Saved Histograms**. The Available Heap Snapshots page appears.
3. Scroll down to the Available Class Histograms table to view a list of saved class histograms. Select any two class histograms and click **Compare**. The Compare Class Histograms page appears. The Class Name, Instance Size (size of each snapshot), and Number of Instances (for each snapshot) are displayed.

16.5.6.5 Deleting Class Histograms

To delete class histograms, follow these steps:

1. From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine target. Select the **Live Heap Analysis** option from the Java Virtual Machine menu.
2. In the Live Heap Analysis page, scroll down to the JVM Class Details table. Click **View Saved Histograms**. The Available Heap Snapshots page appears.
3. Scroll down to the Available Class Histograms table to view a list of saved class histograms. Select the class histogram you want to delete and click **Remove**. A confirmation message is displayed. Click **OK** to delete the class histogram.

16.5.7 Taking a Heap Snapshot

A heap snapshot is a snapshot of JVM memory. It shows a view of all objects in the JVM along with the references between those objects. It can be used to study memory usage patterns and detect possible memory leaks. To take a heap snapshot, follow these steps:

1. From the **Targets** menu, select **Middleware**, then click on a JVM target. The JVM Home page is displayed.
2. Select **Live Heap Analysis** from the **Java Virtual Machine** menu.
3. Click **Take Heap Snapshot**. The Load Heap Snapshot page appears.

Figure 16–17 Heap Snapshot Page

Java Virtual Machine JVM Performance Diagnostics Live Thread Analysis Page Refreshed May 6, 2013 7:23:55 PM PDT

Load Heap Snapshot [Submit](#) [Cancel](#)

Heap Snapshot Formats and Analysis Options

Heap Snapshot Type ☒ JVMD Format (txt) ☐ HPROF Format (binary)

Heap Analysis Options [Load Heap Data to Repository](#)

Heap Snapshot Time

☒ Now ☐ Schedule [5/6/2013 7:23 PM](#)

JVM Diagnostics Agent Host Credentials

Specify credentials for host where JVM Diagnostics Agent is running

Select credential from one of the following options.

Credential ☒ Preferred ☐ Named ☐ New

Preferred Credential Name [Normal Host Credentials](#)

Credential Details Default preferred credentials are not set.

Heap Analysis Host

Add the host for heap analysis and heap loading. EM agent needs to be running on the Heap Analysis Hosts.

Warning
The analysis and load heap steps have significant memory requirement on the Heap Analysis Host. Please make sure you have sufficient free memory on the Heap Analysis Host.

Saved Heap Analysis Hosts [test](#) [Add](#) [Edit](#) [Remove](#)

Heap Analysis Host Details

Attribute	Value
Alias	test
Host Name	slc03rdo.us.oracle.com

4. Specify the Heap Snapshot Type. This can be:
 - JVMD Format (txt) for analysis in JVM Diagnostics.
 - HPROF Format (binary) for analysis with external tools.
5. In the Heap Analysis Options field, select the required option from the drop down menu.
 - **Take a Heap Snapshot Only:** If you want to take a heap snapshot and load it to the repository at a later date, you must leave this field blank. Specify the schedule and click **Submit**. The heap snapshot is generated and the file name in which it is stored is displayed. You can upload the heap snapshot and analyze it using appropriate options from the Heap Snapshots menu.
 - **Load Heap Data to Repository:** Select this option to take a heap snapshot and automatically load it to the repository. If you select this option, you must ensure that the following prerequisites are met:
 - The Management Agent must be deployed on the host machine on which the JVM target is running.
 - The Heap Loader Host is a standalone machine (with high CPU and Memory) on which the Management Agent has been deployed.
 - DB Client Home which is the location of ORACLE_HOME where sqllldr & sqllplus are present.
 - There should be sufficient disk space in the system temp directory.
 - A JVM Diagnostics DB User must have been created using the create_jvm_diagnostic_db_user script.
 - **Generate Memory Leak Report:** Select this option to generate a memory leak report. The memory leak report tab shows the potential memory leak sources by identifying frequent patterns in the heap graph.

- **Generate Anti-pattern Report:** Select this option to generate an anti-pattern report. This report shows the summary or one kind of anti-pattern issue. This option is available only if the Heap Snapshot Type is HPROF (binary).
6. In the Heap Snapshot Time field, specify whether the heap snapshot is taken immediately or at a later date.
 7. Specify the credentials for the host on which the JVM Diagnostics Agent is running.
 8. If the Heap Loader Host has not been configured, click Add. Provide an Alias for the host and select the host (Heap Analysis Host) target on which the Management Agent is running. Click Save.
 9. If the Heap Loader Host has already been configured, the Available Heap Loaders are displayed. Select a heap loader from the list and enter the credentials for the Heap Loader host.

Note:

- If preferred credentials for JVM Target & Heap Loader host are set, then the Enter Credentials region will not be displayed.
 - If the Named Credentials for the JVM Diagnostics DB User is set, the Enter Credentials region will not be displayed.
-

10. Click **Submit** to submit the heap snapshot job. A confirmation message is displayed. Click **Yes** to continue. The job details are displayed in the Heap Analysis Job page. Click on the link to view the job status.

16.5.8 Analyzing Heap Snapshots

The JVM Diagnostics memory analysis feature allows you to not only find the objects responsible for the growth but also track their reachability from the root-set. With this feature, you can find the dangling reference responsible for memory leaks. To find a memory leak, you take snapshots of the JVM heap at different points in time. Between the snapshots, your JVM and Java applications continue running at full speed with zero overhead.

A heap snapshot is a snapshot of JVM memory. Each snapshot stores information about the objects in the heap, their relationships and root-set reachability. You can load the snapshots into the repository, and compare them to see where the memory growth has occurred. Click **Heap Snapshots and Class Histograms** from the menu in the JVM Pool or JVM Home page. The following page appears:

Figure 16–18 Available Heap Snapshots

The screenshot displays the 'Available Heap Snapshots' section of a web application. At the top, there's a search bar labeled 'Search' with a 'Snap Name' input field. Below this is a toolbar with buttons for 'View', 'Detach', 'Remove', 'Detail', and 'Load'. The main table has columns: Date, JVM, Vendor, Version, Heap Snap Name, Description, Size(MB), and Used(MB). A single row is visible with the following data: Date: Feb 8, 2013 12:41:33 AM, JVM: /EMGC_EMGC_DOMAIN/EMGC, Vendor: (empty), Version: (empty), Heap Snap Name: heapdump_EMGC_EMGC, Description: This heap is not loaded, Size(MB): (empty), Used(MB): (empty). Below the table, it says 'Rows Selected 1'. The bottom section, 'Available Class Histograms', has a toolbar with 'View', 'Remove', 'Detail', 'Compare', and 'Detach' buttons, and a table with columns: Snap Date, Snap Id, JVM Name, JVM Version, Snap Name, and Snap Description. The table is currently empty, showing 'No Data to display'.

This page contains the following regions:

- **Available Heap Snapshots:** You can specify the Target Name and Target Type to filter the heap snapshots that are displayed. You can also specify the Heap ID in the Snap Name field to search for specific heap snapshots and display them. The following details is displayed:
 - **Heap ID:** The identification number for the heap snapshot.
 - **Date:** The date on which the heap snapshot was taken.
 - **JVM Name:** The server on which the JVM is running.
 - **Vendor:** The name of the JVM Vendor.
 - **Size:** The total size of the Java heap. An adequate heap size helps improve the performance of the application.
 - **Used:** The amount of heap that has already been used.

Note: If the heap snapshot was taken in HPROF format, the value in the Size and Used fields will be 0.

- **Used(%):** The percentage of heap used.

You can do the following:

- Click **Create** to take a heap snapshot. See Taking a Heap Snapshot.
- Select a heap snapshot and click the **Detail** link to drill down to the Roots page. See Viewing Heap Usage by Roots.
- Select a heap snapshot and click **Load** to load the heap snapshot to the repository. See Uploading Heap Snapshots.

- Select a heap snapshot and click **Reports** to download heap reports to the local host. These reports must have been generated and loaded to the repository for the selected heap snapshot. You can download the Memory Leak Report and the Antipattern Report.
- **Available Class Histograms:** The list of saved histograms with details such as date on which the snapshot was taken, Snap ID, Timestamp, JVM Name and Version, Description are displayed. See Section 15.5.6, "Working with Class Histograms" for more details.

16.5.8.1 Viewing Heap Usage by Roots

To view the heap usage by each class of root, follow these steps:

1. From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine or a Java Virtual Machine Pool target.
2. Select **Heap Snapshots and Class Histograms** from the Java Virtual Machine or Java Virtual Machine Pool target.
3. The list of available heaps is displayed. Select a heap snapshot and click **Details** to view the number of objects and memory reachable from each root. Click on the **Root** tab to view the objects directly reachable from the root. The following details are displayed:

Figure 16–19 Heap Roots

The screenshot shows the 'Heap Roots' page with the following data:

Root	Objects	Total Memory(KB)	Adjusted Memory(KB)
System Classes	7387482	390,946	390,946
JNI	5558906	284,959	1,938
ThreadOther	5557931	283,175	11
System Other	5553568	282,978	1
ObjectMonitor	5553551	282,978	0
Thread Handles	5553553	282,978	0

Click on the **Root** link to view the objects directly reachable from the root. The following details are displayed:

- **Root:** The name of the root is displayed here. Click on the name to drill down to the Top 40 Objects page.
- **Object:** The total number of objects reachable from this root.
- **KB:** The total amount of memory reachable from this root.
- **Adj:** The adjusted memory reachable from this root. This parameter is useful in tracking the memory leak hot-spots.
- **Retained Memory:** The total size of all objects that would be removed when garbage collection is performed on this node.

Click **Compare with** to compare the heap snapshot with another one. See [Comparing Heap Snapshots](#) for details.

16.5.8.1.1 Top 40 Objects This page shows the top 40 objects reachable from a root. The objects are sorted in descending order by the ascending memory reachable from the object (or the difference of the adjusted memory reachable when comparing two heaps). This view provides a lot of rich detailed information like the amount of memory used by an object, amount of memory reachable by an object (total memory used by all the children), and number of objects reachable from a given object.

Figure 16–20 Top 40 Objects

Top 40 Objects in System Mirrors Page Refreshed Jul 31, 2011 9:10:17 PM PDT

Heaps > Roots > Heap Top

Signature	Root	Type	Field	Space	Bytes	Len	Children
java/lang/Class	C	Klass		Perm	424		13
NULL		Constant Pool	IFLD	Perm	3,856		2
S		Array	ITIF	Perm	408	196	1
NULL		Symbol	IIMP	Perm	176		1
java/lang/Class		Instance	JMR	Perm	96		2
S		Array	IINC	Perm	80	32	1
[I		Array	IMOR	Perm	32	4	1
[I		Array	ILIR	Perm	32	4	1
NULL		Symbol	NAM	Perm	32		1
[I		Array	IAKS	Perm	32	4	1
NULL		Symbol	ISFN	Perm	24		1
I		Array	IMET	Perm	16	0	1
sun/reflect/ReflectionFactory		Instance	quickCheckMemberAc Old		8		1
java/lang/reflect/AnnotatedElement		Klass	SKS	Perm	344		10
[I	C	Other		Perm	328		1
java/lang/String	C	Instance		Perm	24		1

The following details are displayed:

- **Signature:** The signature of the object. Click on the link to drill down to the Heap Object Information page.
- **Root:** This is the internal root identifier.
- **Type:** The type of the object which can be Klass, Instance, Method, and so on.
- **Space:** The heap space in which the object is present.
- **Bytes:** The amount of space used by the object.
- **Len:** If the object is an array, the length of the array is displayed here.
- **Children:** The number of descendants reachable from the object.
- **Adj:** Adjusted memory reachable from this object.
- **Retained Memory:** The total size of all objects that would be removed when garbage collection is performed on this node.
- **Depth:** Indicates how far this object is from the root.

16.5.8.1.2 Heap Object Information This page shows information about a specific object in the heap snapshot. The following details are displayed:

Figure 16–21 Heap Object Information

Heap Object Page Refreshed May 6, 2013 7:57:41 PM PDT

Heaps > Heap 22 Detail > Heap Object

May 6, 2013 6:05:26 AM : 718/1015 MB

▼ Heap Object Information

Gar	Space	Type	Signature	Bytes	Len	Children	Adj	Depth	Retained Objects	Retained Memory(Bytes)
N	Midd	Instance	Loracke/jsp/runtimev2/JspTimeoutThread;	128	0	267,684	13,664,384	5	1,312	33

▼ Roots

Type	Field
ThreadOther	Thread-206_
JNI	
Thread Handles	

▼ Object Children

Gar	Space	Type	Signature	Field	Bytes	Len	Children	Adj	Depth
N	Midd	Instance	Loracke/jsp/runtimev2/JspPageTable;	PRECOMPILE_OF	144	0	267,651	13,663,072	6
N	Midd	Instance	Ljava/lang/ThreadLocal\$ThreadLocalMap;	NULL	24	0	29	1,064	6
N	Midd	Instance	Ljava/security/AccessControlContext;	NULL	32	0	2	72	6
N	Midd	Array	C	[C	48	16	1	48	6
N	Midd	Instance	Ljava/lang/ThreadGroup;	threads	48	0	318,213	19,151,256	3
N	Midd	Instance	Lweblogic/utils/classloaders/ChangeAwareClassLoader;	NULL	104	0	185,231	12,351,344	3
N	Midd	Instance	Lweblogic/servlet/internal/WebAppServletContext;	versionId	208	0	771,704	39,872,044	2

▼ Object Parents

Gar	Space	Type	Signature	Field	Bytes	Len	Children	Adj	Depth
N	Midd	Array	Ljava/lang/Thread;		1,040	138	318,207	19,150,984	4
N	Midd	Instance	Loracke/jsp/runtimev2/JspPageTable;	NULL	144	0	267,651	13,663,072	6

- **Heap Object Information**
 - Gar: Indicates whether this object is garbage or reachable from the root.
 - Space: The heap space in which the object is present.
 - Type: The type of the object which can be Klass, Instance, Method, and so on.
 - Signature: The signature of the object.
 - Bytes: The amount of space used by the object.
 - Len: If the object is an array, the length of the array is displayed here.
 - Children: The number of descendants reachable from the object.
 - Adj: Adjusted memory reachable from this object.
 - Depth: Indicates how far this object is from the root.
- **Roots**
 - Type: The type of root which can be Klass, Instance, Method, and so on.
 - Field: If the root is a local thread, this field contains information about the thread and method.
- **Object Children**
 - Gar: Indicates whether this child is garbage or reachable from the root.
 - Space: The heap space in which the child is present.
 - Type: The type of the child which can be Klass, Instance, Method, and so on.
 - Signature: The signature of the child. Click on the link to drill down to the Details page.
 - Bytes: The amount of space used by the child.
 - Len: If the child is an array, the length of the array is displayed here.
 - Children: The number of descendants reachable from the child.
 - Adj: Adjusted memory reachable from this child.

- Depth: Indicates how far this child is from the root.
- Object Parents
 - Gar: Indicates whether this parent is garbage or reachable from the root.
 - Space: The heap space in which the parent is present.
 - Type: The type of the parent which can be Klass, Instance, Method, and so on.
 - Signature: The signature of the parent. Click on the link to drill down to the Details page.
 - Bytes: The amount of space used by the parent.
 - Len: If the parent is an array, the length of the array is displayed here.
 - Children: The number of descendants reachable from the parent.
 - Adj: Adjusted memory reachable from this parent.
 - Depth: Indicates how far this parent is from the root.

16.5.8.1.3 Comparing Heap Snapshots To find a memory leak, you can take snapshots of the JVM Heap at different points in time. Each snapshot stores information about the objects in the heap, their relationships and root-set reachability. You can compare two heap snapshots to see where the memory growth has occurred.

1. From the **Targets** menu, select **Middleware**, then click on a JVM or JVM Pool target.
2. Select **Heap Snapshots** option from the **Java Virtual Machine** or **Java Virtual Machine Pool** menu.
3. The list of available heaps is displayed. Click the **Compare Heaps** tab. The first heap in the list is selected for comparison and you are prompted to select the second heap.
4. The two heaps are compared and a comparison table is displayed in the Diff Heaps page. The details of each heap with the following details are displayed:
 - Objects: The total number of objects reachable from the root.
 - KB: The total amount of memory reachable from the root.
 - Adj: The adjusted memory reachable from this root. This parameter is useful in tracking the memory leak hot-spots. It provides a better representation of the memory used by an object by ignoring backwards pointing references from child objects to their respective parent object.
 - Delta: The difference in the total memory and adjusted reachable memory of the two heaps that are being compared.
5. Click on the root-set with the most growth to diagnose the memory leak.
6. Click the **View Summary** button to see a bottom up view of memory reachable by class of objects.

16.5.8.2 Viewing Heap Usage by Objects

Click the **Usage** tab to view the heap usage by objects. The following details are displayed:

- Object Type: The type of object, Instance, Array, Klass, and so on.
- Garbage: Indicates if this is garbage or reachable from root.

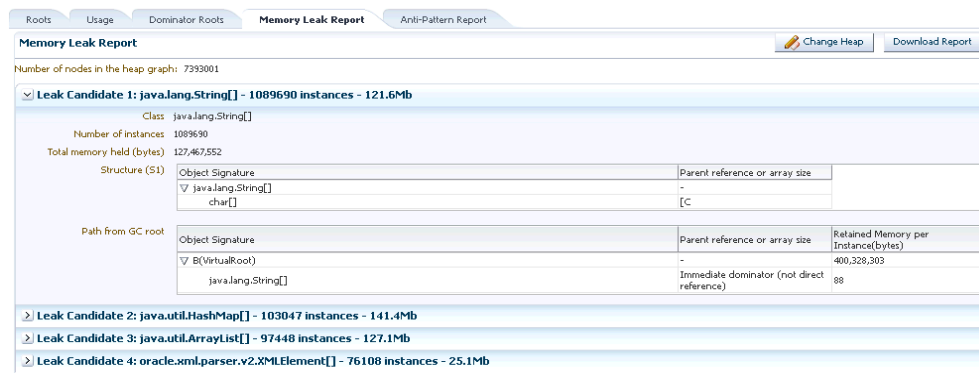
- **Objects:** The total number of objects.
- **Total Memory:** The total amount of memory reachable by root.

Click **Compare with** to compare the heap snapshot with another one. See [Comparing Heap Snapshots](#) for details.

16.5.8.3 Memory Leak Report

Click the **Memory Leak Report** tab to view the memory leak report.

Figure 16–22 Memory Leak Report



The memory leak report shows the potential memory leak sources by finding frequent patterns in the heap graph. This tab shows a list of memory leak candidates which contain the most frequent patterns in a heap and could represent potential memory leak sources. Click **Download Report** to download the memory leak report in .txt format.

16.5.8.4 Anti-Pattern Report

Click the **Anti-Pattern Report** tab. The Anti-Pattern report is divided to different sections. Each section either shows the summary or one kind of anti-pattern issue.

The first section contains a summary of the most acute problems detected by JOverflow. The second section contains the total number of Java classes and Java objects. It also contains a histogram for top memory usage objects grouped by the Class. The third section shows the reference chains for high memory consumers. Each anti-pattern section calculates the overhead that shows the amount of memory that could be saved if the problem is eliminated.

16.5.9 Configuring a JVM

From the **Targets** menu, select **Middleware**, and then click on a Java Virtual Machine target. Select the **Configure JVM Target** option from the **Java Virtual Machine** menu. The Edit JVM Information page is displayed. You can change the JVM Pool, location of the Heap Dump Directory, and the Log Level. Click **Save** to save the changes.

16.5.10 Removing a JVM

You will see a warning message if you select the **Remove Target** option from the JVM menu. The message displays the name of the target being deleted. Click **Yes** to delete the JVM or **No** to return to the JVM Home page.

16.5.11 Add JVM to Group

Select this option to add the JVM to one or more groups. A pop-up window appears with a list of groups on which you have Operator privileges. Select one or more groups and click **Add** to add the target to the group.


16.6 Managing Thread Snapshots

If a particular request is slow or hanging or if the entire application is slow, you can run the real-time transaction trace to view current Java application activity. You can look at the offending threads and their execution stack and analyze how much time a thread spent in waiting for DB wait or wait on a lock. Complex problems such as activity in one thread (or request) affecting the activity in the other thread or rest of the JVM can be found very quickly.

You can trace all active threads and generate a trace file which contains details such as resource usage, thread states, call stack information, and so on. During tracing, the state and stack of the target thread is sampled at set intervals for the desired duration. Follow these steps to trace active threads:

1. From the **Targets** menu, select **Middleware**, then select a Java Virtual Machine target.
2. Select the **Thread Snapshots** option from the Java Virtual Machine menu. The Thread Snapshots page appears.

Figure 16–23 Thread Snapshots Page



All the traces that have been loaded into the repository using the **Trace Active Threads** option are displayed here. For each thread, the Thread Snapshot ID, the date, JVM Name, Thread Name, Duration, and the number of samples taken during the trace is displayed. The Thread column indicates if all threads or only active threads have been traced.

3. Click **Create** to take a thread snapshot of all active threads in the JVM. The Thread Snapshot of All Active Threads page appears where you can trace the active threads. See [Section 16.6.1, "Tracing Active Threads"](#) for details.
4. Select a thread and click the **Details** link to drill down to the Diagnostic Image Analysis page.
5. Select a thread snapshot and click **Import** to upload a thread snapshot from your local machine. The Import Thread Snapshot dialog box is displayed. Click **Browse** and select the thread snapshot to be imported and click **OK**.

16.6.1 Tracing Active Threads

To trace active threads, follow these steps:

1. Click **Create** in the Thread Snapshots page. The Thread Snapshot of All Active Threads page appears.

Figure 16–24 Tracking Active Threads

Java Virtual Machine JVM Performance Diagnostics Live Thread Analysis Page Refreshed Jun 1, 2013 12:03:39 AM PDT

Thread Snapshot of All Active Threads

Poll Interval (ms)

Poll Duration (sec)

Description

Thread Details ☒

Try Changing Threads ☐

Include Network Waits ☐

All Threads ☐

Allow Trace Interrupt ☐

Take Thread Snapshot

2. Specify the following details:
 - **Poll Interval (ms):** The time interval between successive samples. The default value is 200 ms but can be changed.
 - **Poll Duration (sec):** The duration for which the thread snapshot should be taken.
 - **Trace Thread Details:** If this box is unchecked, only the last user call for the active thread will be stored. If the box is checked, all calls for the active thread will be stored, so you can view the call stack. Checking the box increases the overhead and space requirements
 - **Try Changing Threads:** If a thread stack changes during a sample (this can happen when a thread is using CPU), JVM Diagnostics will skip that thread for that sample. If you find missing samples, use this feature to retrace the changed stacks. This will retry (up to 5 times) threads with changing stacks. It will also make system calls to get the stack if possible.
 - **Include Network Waits:** Most JVMs have large number of idle threads waiting for network events. If you leave this check box unchecked, idle threads will not be included in the trace. Checking this box increases the overhead and space requirements.
 - **Trace All Threads:** Check this box if both idle and active threads will be included in the trace.
 - **Allow Trace Interrupt:** Allows you to interrupt the trace process.
3. Click **Take Thread Snapshot** and click **OK** to generate the trace file. When the trace has been completed successfully, click [here](#) link to view the trace data in the Diagnostics Image Analysis page.

16.7 Analyzing Trace Diagnostic Images

A trace diagnostic image contains details such as resource usage, thread states, call stack information etc. The trace diagnostic image captures thread data at short intervals. If an application is hanging or is slow, you can analyze these threads and find out the application tier that causing the delay.

On the Diagnostic Image Analysis page, you can:

- Click **Description** to view details of the thread snapshot being analyzed. The following Server State charts are displayed:
 - Active Threads by State: This chart shows the status of all threads in the JVM. The threads can be in different states like RMI, IO, NET, DB, CPU, and LOCK.
 - CPU Utilization by JVM: This chart shows the CPU utilization in the JVM.
 - Heap Utilization by JVM: This chart shows the heap utilization in the JVM.
- You can filter the data that is displayed by specifying various criteria such as Method Name, JVM Name, Thread State, DBState, and so on. Check the Ignore Filters check box if you want to ignore the specified filters. The Active Threads by State, Top Requests, Top Methods, Top SQLs, Top DBWait Events, and Top Databases charts are displayed.
- Click on the **Threads** tab to view the Thread State Transition, Metric By Active States, and Method data.
- Click **Change / Manage Traces**. A list of available trace files for the JVM Pool is displayed in the List of Traces Loaded page. The following details are displayed:
 - Trace ID: The port number for the JVM.
 - Date: The date and time the trace was run.
 - JVM Name: The name of the host on which the JVM is running.
 - Thread: Indicates whether the trace was run for a specific thread or all active threads.
 - Seconds: The trace duration in seconds.
 - Samples: The number of samples in the trace.
- Select a trace file and click **Analyze**. The trace details are displayed in the Diagnostic Image Analysis page.
- Click **Offline Diagnostics** to create a diagnostic snapshot.

16.8 Uploading Trace Diagnostics Images

You can upload trace data and heap data to a file and view it at a later date.

Uploading Trace Data

This file can be uploaded and stored on the OMS or on the database. To upload trace data (trace diagnostic image), follow these steps:

1. From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine Pool or a Java Virtual Machine target.
2. Select the **Thread Snapshots** option from the Java Virtual Machine Pool or Java Virtual Machine menu. The list of available thread snapshots is displayed.
3. Click **Upload**. The Upload Trace Diagnostic Images page is displayed.
4. When you run the Trace Threads option, the trace file is saved on your local machine or on the OMS. To upload the trace file:
 - **Thread Snapshot Path:** Click **Modify** and specify the location of the diagnostic snapshots on the OMS.
 - **Thread Snapshot File:** Click **Browse** to select the Trace Diagnostic Image File and click **Start Upload**.

5. The list of trace files present on the specified location on the OMS is displayed. Select a file and click the **Load** icon to load the trace file to the database.
6. A message indicating that the trace file has been loaded successfully is displayed. Click the **Loaded Successfully, Goto Saved Trace** link to view the Diagnostic Image Analysis page or click **Ignore** to remain on the same page.

Uploading Heap Data

A heap snapshot is a snapshot of JVM memory. It shows a view of all objects in the JVM along with the references between those objects. You can upload these heap snapshots into the repository and analyze them to detect memory leaks. To upload a heap snapshot, follow these steps:

1. From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine target.
2. Select the **Heap Snapshots** option from the Java Virtual Machine menu. The list of available heap snapshots is displayed.
3. Click **Upload**. The Upload Trace Diagnostic Image page is displayed.
4. When you run the **Take Heap Snapshot** option, the heap snapshot is saved on your local machine or on the OMS. To upload the heap snapshot:
 - **Thread Snapshot Path:** Click **Modify** and specify the location of the diagnostic snapshots on the OMS.
 - **Thread Snapshot File:** Click **Browse** to select the Trace Diagnostic Image File and click **Start Upload**.
5. The list of heap snapshots present on the specified location on the OMS is displayed. Select a file and click the **Load** icon to load the heap snapshot to the database.
6. A message indicating that the heap snapshot has been loaded successfully is displayed. Click the **Loaded Successfully, Goto Saved Snapshot** link or click **Ignore** to remain on the same page.
7. The Diagnostic Image Analysis page is displayed. Click **Description** to view details of the heap snapshot being analyzed. The following Server State charts are displayed:

The following details are displayed:

- **Garbage Collections:** The number of objects that have been added to the garbage collection.
 - **JVM Heap Memory Usage:** This table provides a summary of the heap usage by different types of objects in the heap.
 - **JVM Heap No of Objects:** This table provides details of all the objects in the heap.
8. Click **Change / Manage Heap Snapshots**. A list of available heap snapshots for the JVM is displayed.
 9. Select a heap snapshot file and click **Analyze** to view the Diagnostic Image Analysis page.

16.9 JVM Offline Diagnostics

Diagnostic data for one or more JVM targets can be collected for a specific period and analyzed in an offline mode. This section describes the various options that are

available to collect live JVM data and analyze it in offline mode. It contains the following sections:

- Creating a Diagnostic Snapshot
- Using the Diagnostics Snapshot Page
- Analyzing a Diagnostics Snapshot
- Viewing a Diagnostics Snapshot

16.9.1 Creating a Diagnostic Snapshot

You can create diagnostic snapshots for one or more JVM targets for a specified period. To create a diagnostic snapshot, specify the following:

1. From the **Targets** menu, select **Middleware**.
2. Select the **Manage Diagnostic Snapshots** option from the **Middleware Features** menu.

The Create Diagnostic Snapshot option is also available in the JVM Performance Diagnostics page. Navigate the Performance Diagnostics page for a JVM, specify the time range for which you want to create the collection and click **Create Diagnostic Snapshot**.

3. Click **Create** in the Diagnostic Snapshots page. You can navigate to this page by clicking **Offline Diagnostics** on the Diagnostic Image Analysis page.
4. Enter a name and description for the diagnostic snapshot.
5. Specify the duration for the diagnostic snapshot.
6. Click **Add**. Select one or more JVM targets for which the diagnostic data is to be collected.

Note: The JVM targets that you select must belong to the same JVM Pool.

7. Select the diagnostic types for the selected target and click **OK**. You will see a pop-up window that indicates that the diagnostic snapshot is being created. Click **Close** after the diagnostic snapshot has been created. You will return to the Diagnostic Snapshots page.

16.9.2 Using the Diagnostic Snapshots Page

You can collect diagnostic data for one or more JVM targets and analyze them in an offline mode. This page shows the list of diagnostic snapshots that have been created. You can specify search criteria to retrieve a specific snapshot. You can do the following:

- **Create:** Click **Create** to create diagnostic snapshots for one or more JVMs. The Create Diagnostic Snapshot page is displayed.
- **Export:** Select a file and click **Export** to export the diagnostic data to a file. Enter the location in which the file is to be stored. You can review and analyze the saved file in an offline mode on the same or a different host machine.
- **Import:** Click **Import** to import an exported file with diagnostic data for a particular collection object. Specify the name of the file and upload the file from

your system. You can analyze the exported file and view a summary of the diagnostic snapshot.

- **Analyze:** Select a file and click **Analyze**. The Analyze Diagnostic Snapshot page is displayed.
- **Delete:** Select a diagnostic snapshot from the list and click **Delete**. A confirmation message is displayed. Click **OK** to delete the diagnostic snapshot.
- **View:** Select a file and click **View**. The View Diagnostic Snapshot page is displayed.

16.9.3 Analyzing a Diagnostic Snapshot

This page displays the summary details of the diagnostic snapshot and a summary of all the diagnostic types of the diagnostic snapshot. You can view the thread stack, thread states, CPU Utilization, Heap Utilization, Active Threads Graphs, and Garbage Collections.

To analyze a diagnostic snapshot, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. Select the **Manage Diagnostic Snapshots** option from the **Middleware Features** menu.
3. In the Diagnostic Snapshots page, select a snapshot from the list and click **Analyze**.
4. You can analyze details for each JVM for the specified time interval. Click **More Details** to view detailed diagnostics information for the JVM. The Diagnostic Image Analysis page is displayed.

16.9.4 Viewing a Diagnostic Snapshot

This page displays the summary of the targets, target types and the diagnostic information collected.

1. From the **Targets** menu, select **Middleware**.
2. Select the **Manage Diagnostic Snapshots** option from the **Middleware Features** menu.
3. In the Diagnostic Snapshots page, select a snapshot from the list and click **View**.
4. The summary details for the selected JVM target, target types, and the diagnostic information collected for the JVM is displayed.

16.10 Viewing JVM Diagnostics Threshold Violations

An event is a discrete occurrence detected by Enterprise Manager related to one or more managed entities at a particular point in time which may indicate normal or problematic behavior. Examples of events include: a database target going down, performance threshold violation, change in application configuration files, successful completion of job execution, or job failure.

JVM Diagnostics threshold violations are now integrated with the Enterprise Manager Event subsystem. When a threshold violation occurs, an Enterprise Manager event is generated. To view the event, follow these steps:

1. From the **Enterprise** menu, select **Monitoring**, then select **Incident Manager**.

- In the View panel, click **Events without Incidents**. The JVM Diagnostics events are displayed if there are any outstanding JVMD threshold violations.

Figure 16–25 Incident Manager: Events without Incidents

Incident Manager: Events without incidents

Page Refreshed Sep 8, 2011 11:28:06 AM PDT

Severity	Message	Target Name	Target Type	Event Reported	Category	Last Updated	Incident ID
Warning	Memory Utilization is 80.141%, crossed w		Host	Sep 8, 2011 12:44:34	Capacity	Sep 8, 2011 12:44:37	
Warning	Host CPU Usage (%) exceeds warning lev		Java Virtual Mach	Sep 6, 2011 2:28:48	Performance	Sep 6, 2011 2:28:48	
Warning	JVM Heap Usage (%) exceeds warning lev		Java Virtual Mach	Sep 6, 2011 2:18:15	Performance	Sep 6, 2011 2:18:15	
Warning	Agent Unreachable cleared. The current st		Oracle HTTP Serv	Sep 5, 2011 4:34:32	Availability	Sep 5, 2011 4:52:10	
Error	Host CPU Usage (%) exceeds critical level		Java Virtual Mach	Sep 5, 2011 4:05:04	Performance	Sep 5, 2011 4:05:04	
Error	Host CPU Usage (%) exceeds critical level		Java Virtual Mach	Sep 5, 2011 4:03:57	Performance	Sep 5, 2011 4:03:58	
Error	Threads in DB Wait exceeds critical level		Java Virtual Mach	Sep 5, 2011 4:00:24	Performance	Sep 5, 2011 4:00:24	
Error	Threads in Network Wait exceeds critical le		Java Virtual Mach	Sep 2, 2011 7:14:57	Performance	Sep 2, 2011 7:14:57	
Warning	There have been 288 failed login attempts	Oemrep_Databas	Database Instanc	Sep 2, 2011 2:24:38	Security	Sep 2, 2011 2:24:39	
Warning	Path does not exist: /ade/aime_a1013/ora:webtierhome_0_	Oracle Home		Sep 2, 2011 2:03:27	Error	Sep 2, 2011 2:04:22	
Warning	Path does not exist: /ade/aime_a1013/ora:webtierhome_0_	Oracle Home		Sep 2, 2011 2:03:27	Error	Sep 2, 2011 2:04:22	
Warning	Path does not exist: /ade/aime_a1013/ora:webtierhome_0_	Oracle Home		Sep 2, 2011 2:03:27	Error	Sep 2, 2011 2:04:02	

Columns Hidden: 5 Row count: 13

- Click on the link in the Target Name column of a JVM Diagnostics Event.

The JVMD threshold violations will show up in the Incidents table of the JVM or JVM Pool Home page only if the events have been promoted to incidents. For more information on promoting events to incidents, see the Enterprise Manager Cloud Control Administrator's Guide.

16.11 Using emctl to Manage the JVM Diagnostics Engine

You can use emctl commands to start and stop the JVM Diagnostics Engine. This is useful when a running JVM Diagnostics Engine may cause an upgrade process to fail.

Table 16–1 Extended JVMD emctl Commands

Command	Description
emctl extended oms jvmd list	Queries and lists all the JVM Diagnostics Managed servers from the Repository.
emctl extended oms jvmd start -server=<server_name1>,<server_name2>... For example: emctl extended oms jvmd start -server=EMJVMDMANAGER,MYJVMDMGR	Starts the JVM Diagnostics Managed servers mentioned in command line arguments. The servers could be running on the same local host on which the OMS is running or can be running on a remote host.
emctl extended oms jvmd start -all	Starts all JVM Diagnostics Managed servers on the same local host on which the OMS is running.
emctl extended oms jvmd start -global	Starts all JVM Diagnostics Managed servers, even if they are running on remote hosts (remote to this OMS host).
emctl extended oms jvmd stop -server=<server_name1>,<server_name2>... For example: emctl extended oms jvmd stop -server=EMJVMDMANAGER,MYJVMDMGR	Stops the JVM Diagnostics Managed servers mentioned in command line arguments. The servers could be running on the same local host on which the OMS is running or can be running on a remote host.
emctl extended oms jvmd stop -all	Stops all JVM Diagnostics Managed servers that are running on the same local host on which the OMS is running.
emctl extended jvmd stop -global	Stop all JVM Diagnostics Managed servers, even if they are running on remote hosts (remote to this OMS host).

Table 16–1 (Cont.) Extended JVMD emctl Commands

Command	Description
emctl extended oms jvmd status -server=<server_name1>,<server_name2>... For example: emctl extended oms jvmd stop -server=EMJVMDMANAGER,MYJVMDMGR	Shows the status of the JVM Diagnostics Managed servers mentioned in command line arguments. The servers could be running on the same local host on which the OMS is running or can be running on a remote host.
emctl extended oms jvmd status -all	Status of all the JVM Diagnostics Engines in this domain
emctl extended oms jvmd -help	Shows the online help for the JVM Diagnostics commands.

Troubleshooting JVM Diagnostics

This chapter describes the errors you may encounter while deploying and using JVM Diagnostics and the workaround steps you can follow to resolve each of them. It contains the following sections:

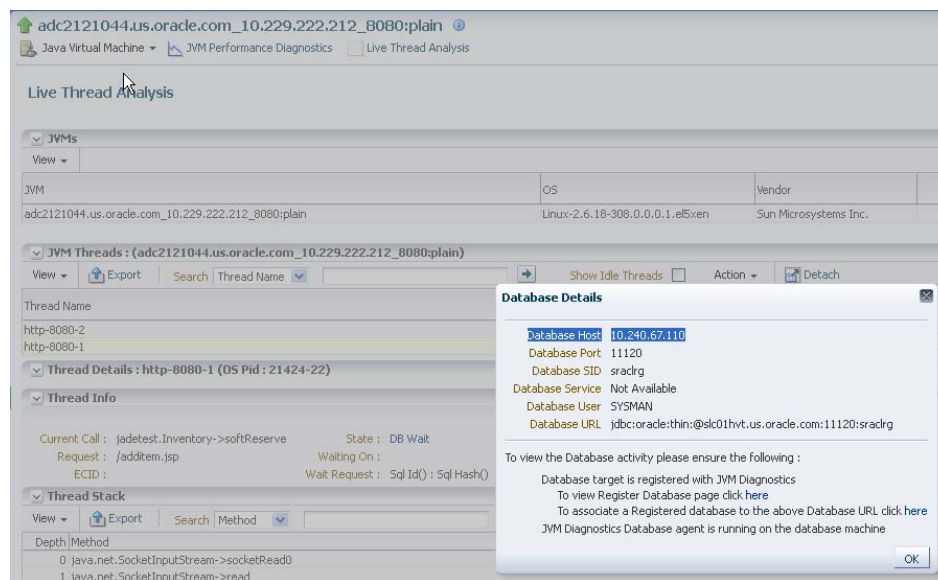
- [Cross Tier Functionality Errors](#)
- [Trace Errors](#)
- [Deployment Script Execution Errors](#)
- [LoadHeap Errors](#)
- [Heap Dump Errors on AIX 64 and AIX 32 bit for IBM JDK 1.6](#)
- [Errors on JVM Diagnostics UI Pages](#)
- [Frequently Asked Questions](#)

17.1 Cross Tier Functionality Errors

This section lists the errors that show the status of the JVM Diagnostics Engine. Cross tier functionality errors may occur due the following:

- Mismatched database connection information
- Insufficient user privileges
- Database Agent Errors

In the Performance Diagnostics page, if the Top SQLs / Top DBWait Events graph contains **Unknown** entries and the Top Databases graph contains **Non-Defined** entries, and the Database Details popup window appears when you click the **DB Wait** link in the Live Thread Analysis page, cross tier correlation cannot be established.

Figure 17–1 Live Thread Analysis (Cross Tier)

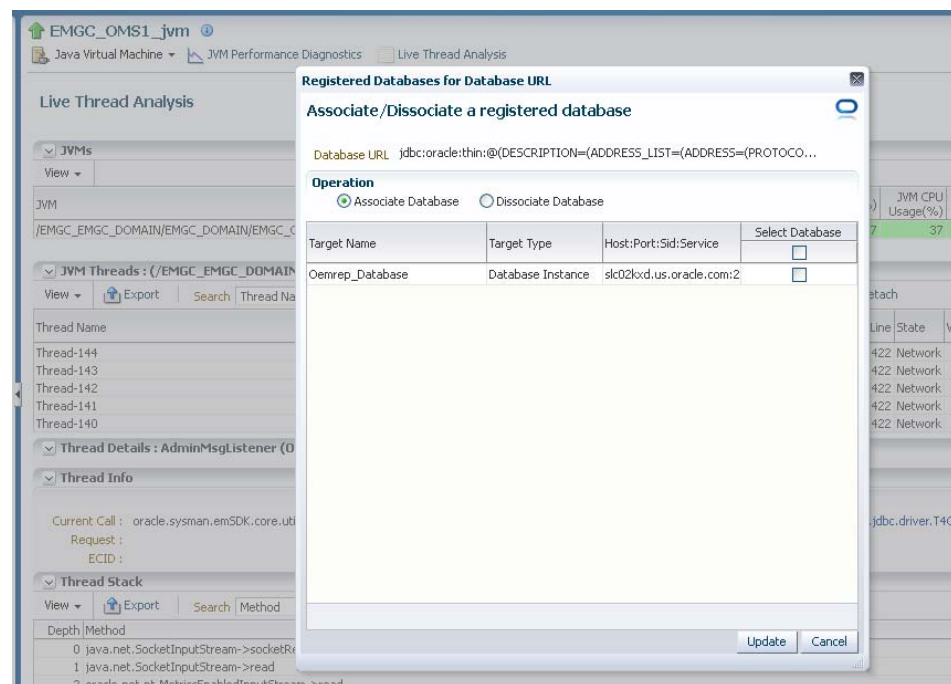
Note: If cross tier correlation is successful, when you click on the **DB Wait** link in the Live Thread Analysis page, the Database Diagnostics page for the database instance is displayed. In this case, the Top SQLs / Top DBWait Events and Top Databases graphs in the JVM Performance Diagnostics page will not contain **Unknown** and **Not Defined** entries respectively. For custom databases, the DB Wait link is not enabled.

Solution:

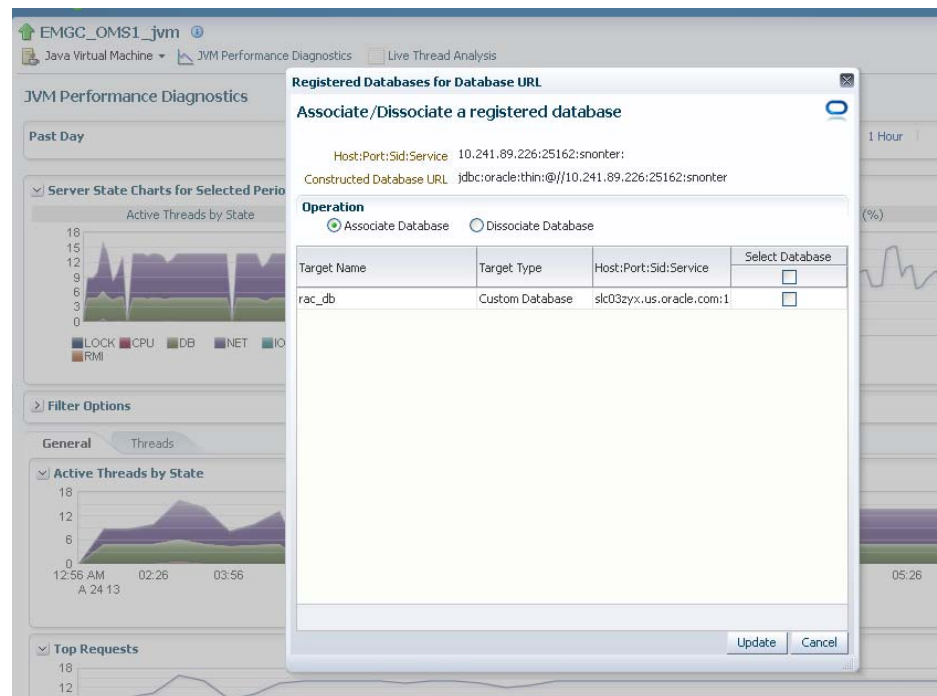
- If cross tier correlation cannot be established due to database mismatch, check if the database has been registered. From the **Setup** menu, select **Middleware Management**, then **Application Performance Management**. Select a JVM Diagnostics Engine and click **Configure**. Click the **Register Databases** tab and check whether database has been registered. If the database has not been registered, click the **DBWait** link to examine the JDBC connection string and verify if it matches the database registered with JVM Diagnostics. For example, if the JDBC connection string contains SID, the database registered needs to have SID. Similarly, the service name, and the hostname of the database in the JDBC connection string must match that of the registered database. Another example of such information that requires matching is the hostname of the database
- If it is a custom database, the user may have insufficient privileges. In this case, check whether the user has permissions on the `v$active_services`, `v$instance`, `v$session`, `v$sqltext`, `v$sqlprocess`, and `v$session_wait` tables.
- From the **Setup** menu, select **Middleware Management**, then select **Application Performance Management**. Select a JVM Diagnostics Engine and click **Configure**. Click the **Register Databases** tab and check whether a JVM Diagnostics DB Agent is required for the registered database. If yes, then ensure that the database agent is running on a machine on which the database is installed with the correct IP address and port number.

- If JDBC URL returned by JVM Diagnostics Agent is for one of registered databases, but cross tier correlation cannot be established due to database mismatch, wrong host name, and so on, the JDBC URL must be associated with a registered database(s). You can associate a JDBC URL with a database from the following pages:
 - **Live Thread Analysis Page:** From the **Java Virtual Machine** menu, select **Live Thread Analysis**. In the JVM Threads table, select a thread that is in the DB Wait state and click **Manage DB URL**. In the **Associate / Disassociate a Registered Database**, select a JDBC URL and click **Add** and specify the URL of the registered database with which is to be associated.

Figure 17–2 Live Thread Analysis: Associate / Disassociate a Registered Database

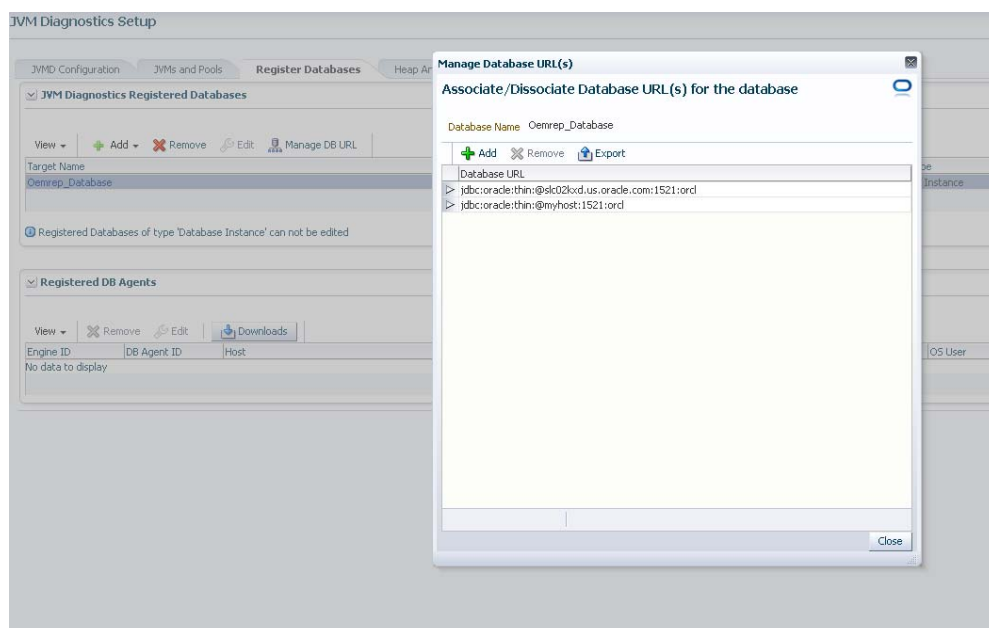


- **JVM Performance Diagnostics Page:** From the **Java Virtual Machine** menu, select **Performance Diagnostics**. In the JVM Threads table, select a thread that is in the DB Wait state and click **Manage DB URL**. In the **Associate / Disassociate a Registered Database**, select a JDBC URL and click **Add** and specify the URL of the registered database with which is to be associated.

Figure 17–3 Performance Diagnostics: Associate / Disassociate a Registered Database

- **Registered Databases Page:** From the **Setup** menu, select **Middleware Management**, then select **Application Performance Management**. Select the JVM Diagnostics Engine row in the Application Performance Management Engines table and click **Configure**.

Click the **Register Databases** tab. The JVM Diagnostics Registered Databases page appears. The list of registered databases is displayed. Select a database and click **Manage DB URL**. In the Associate / Disassociate a Registered Database, select a Database URL and click **Add** and specify the URL of the database to be associated.

Figure 17–4 Setup: Associate / Disassociate a Registered Database

- If cross tier correlation cannot be established due to mismatch of the JVM Diagnostics Agent host name with the machine name stored in `V$SESSION` table of the database (for instance, inconsistent logical naming of machine), do the following:
 - Update the `v$SESS_MACHINE` column of the `jam_jvm` table in the Enterprise Manager repository (for example, update `jam_jvm` set `V$SESS_MACHINE = 'JVMD Agent Machine name'` where `jam_jvm_id = 'jam_jvm_id'`) with the right value as specified in the `V$SESSION` of the database)
- If cross tier correlation cannot be established as the database is inaccessible to the JVM Diagnostics Manager, check the database name in the log file and check if the database is down or inactive, the Listener is down. If this is the case, the JVM Diagnostics Manager cannot connect to the database to establish the cross tier correlation.

If, after following all the above steps, cross tier correlation still cannot be established, you need to purge the JVMD Manager log file (*.out). From the **Setup** menu, select **Middleware Diagnostics** and then select **Application Performance Management**. Select a JVM Diagnostics Engine and click **Configure** and temporarily set the JVMD Engine Log Level and Cross Tier Log Level to Trace.

Turn the monitoring off temporarily (if possible) and navigate to the Live Thread Analysis page when the application is making DB calls (There should be at least on Thread in Db wait) and send the JVMD Manager logs to report the issue. Return to the previous log level and turn monitoring on again.

17.2 Trace Errors

This section lists errors that occur during tracing. The following error occurs if the Poll Duration has a large value and causes a timeout.

Error: weblogic.transaction.internal.TimedOutException: Transaction timed out after 30 seconds.

Solution: This error does not affect the Trace functionality and can be ignored.

17.3 Deployment Script Execution Errors

This section lists the errors that occur when you run the deployment script.

- **Error:** Script Exception: Error occurred while performing deploy: The action you performed timed out after 600,000 milliseconds.

Solution: To resolve this issue, check if the lock for the target WebLogic domain Administration Console has already been acquired. If it has been acquired, release it and run the script again by following these steps:
 - Login to the WebLogic Administration Console: *http://<machine address>:<weblogic port>/console*.
 - If you see the Activate Changes and Undo All Changes buttons in the left pane, click on these buttons to clear them. If the buttons are not cleared, click Undo All Changes and run the script again.
- **Error:** If the user name and password for the WebLogic Administration Server are incorrect, you may see the following error:

Caused by: java.lang.SecurityException: User: <username>, failed to be authenticated.

This message is typically embedded in a long error message trail.

You may also see the following exception:

```
javax.naming.AuthenticationException [Root exception is
java.lang.SecurityException: User: weblogic, failed to be authenticated.]
at
weblogic.jndi.internal.ExceptionTranslator.toNamingException(ExceptionTranslator.java:42)
at
weblogic.jndi.WLInitialContextFactoryDelegate.toNamingException(WLInitialContextFactoryDelegate.java:788)
at
weblogic.jndi.WLInitialContextFactoryDelegate.pushSubject(WLInitialContextFactoryDelegate.java:682)
at
weblogic.jndi.WLInitialContextFactoryDelegate.newContext(WLInitialContextFactoryDelegate.java:469)
at
weblogic.jndi.WLInitialContextFactoryDelegate.getInitialContext(WLInitialContextFactoryDelegate.java:376)
at weblogic.jndi.Environment.getContext(Environment.java:315)
at weblogic.jndi.Environment.getContext(Environment.java:285)
at
weblogic.jndi.WLInitialContextFactory.getInitialContext(WLInitialContextFactory.java:117)
at javax.naming.spi.NamingManager.getInitialContext(NamingManager.java:235)
at
javax.naming.InitialContext.initializeDefaultInitCtx(InitialContext.java:318)
at javax.naming.InitialContext.getDefaultInitCtx(InitialContext.java:348)
at javax.naming.InitialContext.internalInit(InitialContext.java:286)
at javax.naming.InitialContext.<init>(InitialContext.java:211)
```

Solution: Enter the correct user name and password for the WebLogic Administration Server and run the script again.

- **Error:** If the path to the `weblogic.jar` (on the Management Agent) as specified in the agent auto deployment page is incorrect, or if the user does not have read permissions on the `weblogic.jar` (typically seen on hosted machines with the name `aime`), the following exception may occur:

```
Exception in thread "main" java.lang.NoClassDefFoundError:
javax/enterprise/deploy/spi/exceptions/TargetException
Caused by: java.lang.ClassNotFoundException:
javax.enterprise.deploy.spi.exceptions.TargetException
at java.net.URLClassLoader$1.run(URLClassLoader.java:200)
at java.security.AccessController.doPrivileged(Native Method)
at java.net.URLClassLoader.findClass(URLClassLoader.java:188)
at java.lang.ClassLoader.loadClass(ClassLoader.java:307)
at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:301)
at java.lang.ClassLoader.loadClass(ClassLoader.java:252)
at java.lang.ClassLoader.loadClassInternal(ClassLoader.java:320)
```

- **Error:** If the WebLogic Administration Console is locked, the agent deployment job may not work as expected. You will see a message that the `agent.log` files cannot be deployment since the WebLogic Domain is locked.

Solution: JVM Diagnostics Agents are deployed by using `t3/t3s` protocols. Make sure that `t3/t3s` ports are open.

- **Error:** If you are deploying to an SSL enabled WebLogic Domain using the demo certificate, you may see an error if the WebLogic Server demo certificate has not been imported to the keystore.

Solution: You must import the WebLogic Server demo certificate to the keystore of the Management Agent.

- **Error:** While copying the `deployer.zip` or `javadiagnosticagent.ear` files, errors like broken pipe appear.

Solution: The Oracle Management Service and the Management Agent must be installed by the same user or users belonging to the same group.

- **Error:** Agent deployment fails on WebLogic 10.0.

Solution: To deploy the JVM Diagnostics Agent, `weblogic.jar` and `MANIFEST.MF` (of the `weblogic.jar`), are required. If, in some WebLogic 10.0 installations, some of the jar files are missing, agent deployment will fail.

- **Error:** JVM D AGENT DEPLOYMENT FAILED FOR WEBLOGIC 9.2 TARGET.

The following exception occurs:

```
EM Agent home : /scratch/aime/agsh_0819/core/12.1.0.2.0
MIDDLEWARE_HOME : /scratch/aime/mw923
IS_WEBLOGIC9 : true
em agent state dir : /scratch/aime/agsh_0819/agent_inst
acsera home : /tmp/ad4j_1345730608009/4910760210525348050
wls admin url : t3://adc6160336.us.oracle.com:7001
wls username : weblogic
target : AdminServer?
weblogic jar path :
/scratch/aime/mw923/weblogic92/server/lib/weblogic.jar&&ls
/scratch/aime/mw923/weblogic92/server/lib/wljmxclient.jar&&ls
/scratch/aime/mw923/weblogic92/server/lib/wlcipher.jar
application name : HttpDeployer?
agent keystore location :
/scratch/aime/agsh_0819/agent_inst/sysman/config/montrust/AgentTrust.jks
Command used for deployment:
```

```

/scratch/aime/agsh_0819/core/12.1.0.2.0/jdk/bin/java -cp
/tmp/ad4j_1345730608009/4910760210525348050/ADPAgent/lib/mips.jar:/scratch/aime
e/mw923/weblogic92/server/lib/weblogic.jar&&ls
/scratch/aime/mw923/weblogic92/server/lib/wljsxclient.jar&&ls
/scratch/aime/mw923/weblogic92/server/lib/wlscipher.jar
-Dweblogic.security.SSL.ignoreHostnameVerify=true
-Djava.security.egd=file:/dev/./urandom
-Dweblogic.security.SSL.trustedCAKeyStore=/scratch/aime/agsh_0819/agent_inst/
sysman/config/montrust/AgentTrust.jks-Dsun.lang.ClassLoader.allowArraySyntax=
true -Dbea.home=/scratch/aime/mw923
com.acsera.ejb.Deployer.RemoteHttpDeployerShell -deploy -adminurl
t3://adc6160336.us.oracle.com:7001 -upload -source
/tmp/ad4j_1345730608009/4910760210525348050/ADPAgent/deploy/HttpDeployer.ear
-targets AdminServer? -username weblogic -name HttpDeployer?
-usenonexclusivelock

```

The application will be first undeployed on the targeted server

Usage: java [-options] class [args...]

(to execute a class)

or java [-options] -jar jarfile

(to execute a jar file)

where options include:

d32 use a 32-bit data model if available

```

-d64 use a 64-bit data model if available
-client to select the "client" VM
-server to select the "server" VM
-hotspot is a synonym for the "client" VM [deprecated]
The default VM is server,
because you are running on a server-class machine.
-cp <class search path of directories and zip/jar files>
-classpath <class search path of directories and zip/jar files>
A : separated list of directories, JAR archives,
and ZIP archives to search for class files.
-D<name>=<value>
set a system property
-verbose[:class|gc|jni]
enable verbose output
-version print product version and exit
-version:<value>
require the specified version to run
-showversion print product version and continue
-jre-restrict-search | -jre-no-restrict-search
include/exclude user private JREs in the version search
-? -help print this help message
-X print help on non-standard options
-ea[:<packagename>...|:<classname>]
-enableassertions[:<packagename>...|:<classname>]
enable assertions
-da[:<packagename>...|:<classname>]
-disableassertions[:<packagename>...|:<classname>]
disable assertions
-esa | -enablesystemassertions
enable system assertions
-dsa | -disablesystemassertions
disable system assertions
-agentlib:<libname>[=<options>]

```

```

load native agent library <libname>, e.g. -agentlib:hprof
see also, -agentlib:jdwp=help and -agentlib:hprof=help
-agentpath:<pathname>[=<options>]
load native agent library by full pathname
-javaagent:<jarpath>[=<options>]
load Java programming language agent, see
java.lang.instrument
-splash:<imagepath>
show splash screen with specified image
/scratch/aime/mw923/weblogic92/server/lib/wljmxclient.jar
ls: invalid line width: eblogic.security.SSL.ignoreHostnameVerify=true
Status returned from the java process is 512

```

17.4 LoadHeap Errors

This section lists loadheap errors.

- **Error:** The following error occurs during the heapdump operation.

```

glibc detected * free(): invalid next size (fast): 0x0965d090" ./loadheap.sh:
line 237: 32357 Aborted ./bin/${bindir}/processlog in=$infile hdr=${sumdata}
obj=${objdata} rel=${reldata} root=${rootdata} osum=${objsumdata}
rrel=${rootrel} heap=${heap_id} skip=$skipgarbage db=$dbtype $* Error
processing file /tmp/heapdump6.txt

```

Solution: Check if the heapdump operation has been successfully completed. Open the heapdump6.txt file and check if there is a heapdump finished string at the end of the file. If you see this string, load the finished dump file.

- **Error:** Heapdump already in progress, cannot take another heapdump.

Solution: Check if the heapdump operation has been successfully completed. Open the heapdump6.txt file and check if there is a heapdump finished string at the end of the file.

- **Error:** loadheap.sh created unusable unique indexes.

Solution: Run the loadheap/sql/cleanup.sql shipped with loadheap.zip to fix the unique indexes.

17.5 Heap Dump Errors on AIX 64 and AIX 32 bit for IBM JDK 1.6

The following error occurs when you try to deploy the JVM Diagnostics Agent on IBM JDK 1.6:

Error: The following can occur when the JVM Diagnostics Agent is deployed on JDK 1.6.

```

Jam Agent : can_tag_objects capability is not set. Copy /tmp/libjamcapability.so
to another directory and restart Java with argument -agentpath: <Absolute path of
libjamcapability.so>

```

Solution: Deploy the latest jamagent.war and add -agentpath:<Absolute path of libjamcapability.so after copying to another directory> to the java arguments.

- This message appears only after the JVM Diagnostics Agent has connected to JVM Diagnostics Engine. Secondly, this argument should be a JVM argument (and not a program argument)
- If the server is started using the WebLogic Administration Console (through nodemanager). these arguments can be specified in the Administration Console

under **Server Start**. If the server is started from the command line (startWeblogic.sh or startManagedServer.sh), these arguments have to be specified in the startWeblogic.sh. If there are multiple servers, make sure a check for the server name is present in the startWeblogic.sh to ensure that the path for the libjamcapability.so is separate for each server.

- A sample entry to be made in startWeblogic.sh is below:

```
if [ "${SERVER_NAME}" = "AdminServer" ] ; then
echo "***** MODIFIED ADMIN SERVER"
JAVA_OPTIONS="${JAVA_OPTIONS} -agentpath:<Absolute path of
libjamcapability.so.X after copying to another directory>
export JAVA_OPTIONS
fi
```

- The message "Capabilities Added by libjamcapability.so" during server startup (before the jamagent logs appear) confirms that libjamcapability.so was loaded fine.

17.6 Errors on JVM Diagnostics UI Pages

This section lists the user interface errors.

- **Error:** This is an Agent timeout error:

```
JAM Console:Socket timed out after recv -- client adc2100083.us.oracle.com:7001
is not Active [0] secs
JAM Console jamlooptimeout=[3]
JAM CONSOLE: JVM 1 is not active
JAM Cons ErrProcessing Request:128 JVM 1 is not active jamDAL: jamreq returned
128 return status < 0 from jamDalInst.processRequest
```

Solution: To resolve this error, increase the Agent Request Timeout (secs) and Agent Loop Request Timeout (secs).

- **Error:** The JVM Diagnostics Agent is up and running but is not displayed in the real time pages.

Solution: If the log file shows JAMMANAGER: OLD AGENT or NULL POOL or wrong optimization level, this indicates that the old JVM Diagnostics Agent or Dbagent is being used. To resolve this issue, select **Application Performance Management** from the **Setup** menu. The list of Application Performance Management Engines is displayed. Select the JVM Diagnostics Engine row, click **Configure** then click the **Register Databases** tab. Click the **Downloads** button in the Registered DB Agents region, and select **JVMD Agent** from the JVMD Component list. Specify the JVM Diagnostics Agent web.xml Parameters, click **Download**, then click **OK** to download the jamagent.war.

- **Error:** You do not have the necessary privileges to view this page.

Solution: Ensure that you have the required JVM Diagnostics Administrator or User privileges to view the JVM Diagnostics data.

17.7 Frequently Asked Questions

This section lists some of the questions you may have while using JVM Diagnostics. It includes the following:

- [Location of the JVM Diagnostics Logs](#)
- [JVM Diagnostics Engine Status](#)

- [JVM Diagnostics Agent Status](#)
- [Monitoring Status](#)
- [Running the create_jvm_diagnostic_db_user.sh Script](#)
- [Usage of the Try Changing Threads Parameter](#)
- [Significance of Optimization Levels](#)
- [Custom Provisioning Agent Deployment](#)
- [Log Manager Level](#)
- [Repository Space Requirements](#)

17.7.1 Location of the JVM Diagnostics Logs

You can find the JVM Diagnostics logs in the following locations:

- The JVM Diagnostics Engine Log file is located at `$EMGC_JVMDMANAGER1/logs/EMGC_JVMDMANAGER1.out`
- UI related errors are logged in:
 - `$T_WORK/gc_inst/user_projects/domains/GCDomain/servers/EMGC_OMS1/logs/EMGC_OMS1.out`
 - `$T_WORK/gc_inst/user_projects/domains/GCDomain/servers/EMGC_OMS1/logs/EMGC_OMS1.log`
- Communication errors between the JVM Diagnostics Engine and the Console are logged in `$T_WORK/gc_inst/em/EMGC_OMS1/sysman/log/emoms.log`

17.7.2 JVM Diagnostics Engine Status

To check the status of the JVM Diagnostics Engine, follow these steps:

- From the **Setup** menu, select **Middleware Diagnostics**, then click **Setup JVM Diagnostics**.
- Check the JVM Diagnostics Agent log file to verify the connection between Agent and the Manager. If you see an error - JAM Agent ERROR: Cannot connect to Console:Connection refused, this indicates that the JVM Diagnostics Engine is not running.
- Check if the message JAM Console: Agent connection from:[Hostname] is present in the JVM Diagnostics Engine log file. If this message appears, it indicates that the JVM Diagnostics Engine is running and is connected to the Agent.

17.7.3 JVM Diagnostics Agent Status

To check the status of the JVM Diagnostics Agent:

- From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine target. Select the **Live Thread Analysis** option from the Java Virtual Machine menu. Check the JVM Status in the Connected JVMs table.
 - If the status is **Not Active**, this indicates that the Agent is not connected to the Manager. Check the agent logs to verify if it is running and the IP address and port number of the Manager is correct.
 - If the status is **No JVMD Agent Deployed**, the JVM Diagnostics Agent must be deployed on that JVM.

- If the JVM Diagnostics Agent is running, the active threads data must be visible. If the JVM Diagnostics Agent is not running, you will see a message - JVM is inactive, Please try again after some time.

17.7.4 Monitoring Status

To verify if the JVM Diagnostics Engine is monitoring the data:

1. From the **Setup** menu, select **Middleware Diagnostics**, then click **Setup JVM Diagnostics** in the Middleware Diagnostics page. In the JVMD Configuration page, verify that the **Enable Monitoring** check box is checked.
2. Navigate to the Monitoring page under Setup and check if monitoring status is **On** for the Pool to which the JVM being monitored belongs.
3. Navigate to the JVM Pools page under Setup and verify if the **Poll Enabled** check box has been checked for the Pool to which the JVM being monitored belongs. Monitoring should now be enabled.

17.7.5 Running the create_jvm_diagnostic_db_user.sh Script

You can run the `create_jvm_diagnostic_db_user.sh` script if you want to create less privileged users who can only load heaps using the `loadHeap` script.

17.7.6 Usage of the Try Changing Threads Parameter

This parameter should be used only when the JVM is highly active.

17.7.7 Significance of Optimization Levels

The JVM Diagnostics Agent supports three optimization levels:

- Level 0 indicates that the JVM Diagnostics Agent is using a JVMTI based engine. This level is supported for JDK 6 series on almost all supported platforms.
- Level 1 is a hybrid between level 0 and level 2. It is supported only for very few JDKs on selected platforms.
- Level 2 uses Runtime Object Analysis technique for monitoring as it is efficient at run time.

17.7.8 Custom Provisioning Agent Deployment

You can customize the JVMD Agent deployment in the production environment by running custom provisioning scripts.

After the OMS has been installed, the `jvmd.zip` file can be found in the `plugins/oracle.sysman.emas.oms.plugin_12.1.0.0.0` directory in the Middleware installation directory. The zip file contains a set of scripts in the `customprov` directory. Details on using these scripts are described in the `README.TXT` present in the same directory. To use the custom provisioning scripts, follow these steps:

1. From the **Setup** menu, select **Middleware Diagnostics**, then click the **Setup JVM Diagnostics** in the Middleware Diagnostics page. Click the **Register Databases** tab and download the `jamagent.war` file
2. Make a copy of the deployment profile that includes the location of the downloaded `jamagent.war`, domains, and server details.

3. Run the Perl script on the deployment profile which will deploy the JVMD Agent to all the specified servers.

17.7.9 Log Manager Level

The default log manager level is 3. You can temporarily increase this to a higher level if you encounter some issues. Log levels 1 to 5 are supported where:

- 1 - Error
- 2 - Warning
- 3 - Info
- 4 - Debug
- 5 - Trace

17.7.10 Repository Space Requirements

For monitoring data, Oracle recommends 50 MB per JVM per day with the default setting of a 24 hour purge interval. This amount can vary based upon runtime factors (e.g depth of call stacks, etc.) within your environment. Hence, you must check the tablespace growth periodically and if required, you may need to change the space requirements. This will ensure that database growth due to standard monitoring will occur smoothly without sudden spikes. Tablespace sizing can be affected by the following:

- **Heap Dumps:** Analyzing heaps requires a large amount of tablespace. As a standard practice, we recommend that you must have 5 times the size of heap dump file being loaded in your tablespace. Since you know the size of your dump file, make sure that there is adequate space to accommodate the dump file before it is loaded into the database.
- **Thread Traces:** While these are smaller than heaps, they are loaded into the database automatically when a user initiates a trace at the console. The size of these threads can vary dramatically depending on the number of active threads during the trace, the duration of the trace, and the sample interval of the trace. This should usually be under 100MB but if several thread traces have been initiated, it could fill up the database quickly. Before initiating the traces, you must ensure that there is adequate space in the database.

Using Middleware Diagnostics Advisor

The Middleware Diagnostics Advisor analyzes the entire stack and provides diagnostic findings by identifying the root cause of a problem. It correlates and analyzes the input and offers advice on how to resolve the problem. For example, it can help you identify that slow SQL statements or a JDBC connection pool is causing a performance bottleneck.

You can view the diagnostic findings for one or more servers in a WebLogic Domain if the Middleware Diagnostics Advisor has been enabled.

This section covers the following:

- [Diagnosing Performance Issues with Oracle WebLogic Server](#)
- [Diagnosing Performance Issues Using Middleware Diagnostics Advisor](#)
- [Functioning of Middleware Diagnostics Advisor](#)
- [Limiting the Scope of Middleware Diagnostics Advisor](#)
- [Prerequisites](#)
- [Enabling Middleware Diagnostics Advisor](#)
- [Setting the Preferred Frequency for Purging Data](#)
- [Enabling JMS Destination Metrics](#)
- [Using Middleware Diagnostics Advisor to View and Diagnose Performance Issues](#)
- [Troubleshooting Issues Related to Middleware Diagnostics Advisor](#)

18.1 Diagnosing Performance Issues with Oracle WebLogic Server

Oracle WebLogic Server (WLS) is an application server that provides high performance and scalability. WebLogic Server also simplifies deployment and management, and accelerates time to market with a modern, lightweight development platform.

In order to keep up performance, and scalability, it is best to detect violation and provide insight to the cause of the violation, thus enabling faster remedial action. Performance related issues are detected based on the configuration and load of the server. The most common performance issues include slow response times, and application crashes. Using Middleware Diagnostics Advisor (MDA) adds value to the WebLogic Management Pack. To find out more about using MDA for WebLogic Servers, refer [Section 18.2](#).

18.2 Diagnosing Performance Issues Using Middleware Diagnostics Advisor

Middleware Diagnostics Advisor or MDA is a diagnostic module integrated within Enterprise Manager Cloud Control for diagnosing performance issues with targets monitored in Enterprise Manager Cloud Control. Currently, MDA is supported for Oracle WebLogic Server 10g Release 3 (10.3) and higher. MDA monitors JDBC DataSources, EJBs, and JMS Queues.

MDA enables you to easily identify the underlying states in the application server environment that are the causes for degradation in performance. These underlying states can manifest themselves as degradation in performance such as slow response for request, hung server, slow server, high memory utilization and high Disk I/O, and so on.

MDA analyses the overall performance of an aspect in a runtime environment. When the overall performance of the aspect degrades beyond a certain limit, MDA diagnoses the issue to find the underlying cause. However, individual one off issues, which do not affect the overall performance, are not isolated by MDA.

MDA diagnoses performances issues in the following areas:

- JDBC findings:
 - Checks if the SQL execution takes a long time.
 - Checks if the JDBC Pool size is small, and if the wait time for connections is high.
 - Checks if reclaimed connections are found for data source, and if the effective pool size is small.
- JMS findings:
 - Checks if the message processing is slow.
 - Checks if the number of messages reprocessed due to transaction timeout is high.
 - Checks if the number of messages reprocessed due to transaction rollback is high
 - Checks if the message delivery is delayed.
 - Checks if the queue slowed down due to large number of messages.
 - Checks if the queue slowed down due to large size of messages.
- EJB findings:
 - Checks if the remote call made by the EJB takes too long to return.
 - Checks if the EJB takes too long to execute.
- Thread findings:
 - Checks if there are locks that are being waited on by other threads.

18.3 Functioning of Middleware Diagnostics Advisor

Middleware Diagnostics Advisor functions in the following way:

1. Data Collector, located on the target server, collects data at a high frequency. The collected data is then aggregated every 5 minutes.

2. Oracle Management Agent uploads the aggregated data to the Oracle Management Repository periodically.
3. On the OMS, an MDA analysis job runs every 15 minutes, and analyzes the data uploaded to Oracle Management Repository.
4. An analysis is performed every hour. If the analysis is begun at the middle of an hour, it will still be performed for the entire hour. For example, if the analysis is begun at 2:15 PM, the analysis will be performed from 2:00 PM to 3:00 PM.

At the end of the analysis hour, rules are applied to see if there are issues. There are rules which also determine the cause of the problem.

18.4 Limiting the Scope of Middleware Diagnostics Advisor

Using `emctl`, you can disable certain checks performed on targets by MDA. The following table provides the name of the check, the command, and the description.

Check	Command	Description
Thread Analysis	<code>emctl set property -sysman_pwd <sysman password> -name oracle.sysman.emas.mda.disable ThreadAnalysis -value true</code>	Disables Thread Analysis check for targets. Global setting, so cannot be used for individual targets.
EJB Analysis	<code>emctl set property -sysman_pwd <sysman password> -name oracle.sysman.emas.mda.disable EJBAnalysis -value true</code>	Disables EJB Analysis check for targets. Global setting, so cannot be used for individual targets.
JDBC Analysis	<code>emctl set property -sysman_pwd <sysman password> -name oracle.sysman.emas.mda.disable JdbcAnalysisForTargets -value <comma-separated-target-names all></code>	Disables JDBC Analysis check for targets. Enter individual targets in a comma separated list. If disabling for all targets, enter <code>all</code> .
JMS Analysis	<code>emctl set property -sysman_pwd <sysman password> -name oracle.sysman.emas.mda.disable JmsAnalysisForTargets -value <comma-separated-target-names all></code>	Disables JMS Analysis check for targets. Enter individual targets in a comma separated list. If disabling for all targets, enter <code>all</code> .

18.5 Prerequisites

Before you begin using MDA for diagnosing performance issues, meet the following prerequisites.

- Enterprise Manager has discovered the WebLogic Server as a target.
- Before enabling MDA on a target, it is recommended to force the configuration collection for the target.
- The JVMD Manager is configured, and the JVMD Agent is deployed on the target server.
- You have already added preferred credentials for the host, and administrator credentials for WebLogic.
- You have already enabled JVMD.

Note: If the JVMD Agent war file is deployed manually to Oracle WebLogic Server, follow these steps:

1. Copy <MWHOME>/plugins/oracle.sysman.emas.oms.plugin_<pluginversion>/archives/jvmd/javadiagnosticagent.ear to a temporary directory.
 2. Change to the temporary directory.
`cd/scratch/temp`
 3. Extract the files from the .ear file using the following command.
`jar -xvf javadiagnosticagent.ear`
 4. Copy the jamagent.war file that was used to deploy the JVMD agent to this directory.
 5. Create .ear file using the following command.
`jar -cvf javadiagnosticagent.ear *`
 6. Deploy javadiagnosticagent.ear manually, and enable MDA.
-

18.6 Enabling Middleware Diagnostics Advisor

To configure MDA, follow these steps:

1. From the **WebLogic Domain** menu, select **Diagnostics**, then select **Middleware Diagnostics Advisor Configuration**.
2. On the Middleware Diagnostics Advisor Configuration page, select the target that you want to enable MDA for, and click **Enable**. (See [Figure 18–1](#).)

Note: To enable targets to be monitored by MDA, the host and domain credentials have to be provided. If you have already set up the credentials, select the required credentials in the **Host Credentials** and **Domain Credentials** tabs.

If you have not set up the credentials, provide them in the host and domain credentials tabs and proceed with the enabling of targets.

Figure 18–1 *Middleware Diagnostics Advisor Configuration Page*

WebLogic Domain ▾ Page Refreshed Jun 19, 2013 9:11:53 AM PDT ↻

Middleware Diagnostics Advisor Configuration

Use this page to enable or disable analysis of Middleware Diagnostics Advisor for targets under this Oracle WebLogic Domain. A prerequisite to enable or disable a server for analysis is to have the JVMD agent deployed.

To configure a server for analysis, select the corresponding row in the table below and click the Enable or the Disable button.

☒ Credentials for Host and Oracle WebLogic Domain needs to be set prior to enable / disable options.

View ▾ All ▾ Enable Disable

Target Name	Type	Status
▽ /mda_sql_conn_g3_wl1034Domain/wl1034Domain	Oracle WebLogic Domain	
/mda_sql_conn_g3_wl1034Domain/wl1034Domain/AdminServer	Oracle WebLogic Server	Disabled
/mda_sql_conn_g3_wl1034Domain/wl1034Domain/ManagedServer3	Oracle WebLogic Server	Enabled
/mda_sql_conn_g3_wl1034Domain/wl1034Domain/ManagedServer4	Oracle WebLogic Server	Enabled

18.7 Setting the Preferred Frequency for Purging Data

There is an MDA job that runs every 24 hours which purges data from the repository, and by default, it is disabled for all Middleware targets. To set your preferred frequency, follow these steps:

1. From the **Setup** menu, select **Middleware Management**, and then select **Middleware Diagnostics Advisor**.
2. On the Middleware Diagnostics Advisor Setup page, choose whether to purge data or not.
3. If you choose the **Purge Data** option, enter the preferred number of days for which data should be retained.
4. Click **Apply**.

Note: This is a global setting, and will be applied for all targets, and all users.

18.8 Enabling JMS Destination Metrics

If Middleware Diagnostics Advisor is being enabled for the first time, it is recommended that you enable JMS Destination Metrics. However, in order to analyze JMS Queues, it is required to enable the metrics.

There are two methods that you can use to enable JMS Destination Metrics.

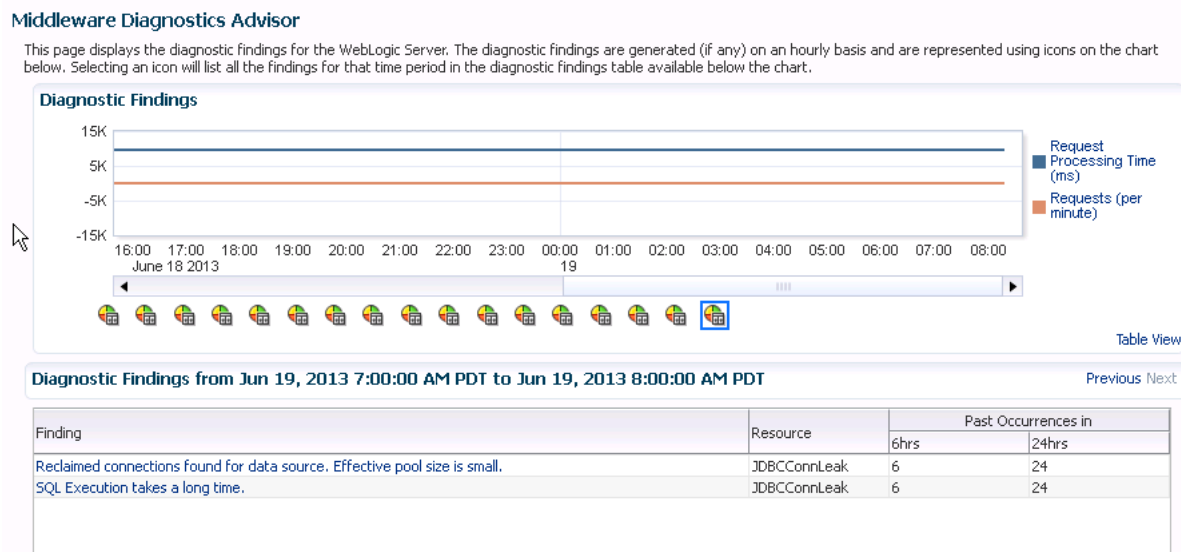
1. First method of enabling JMS Destination Metrics:
 - a. From the **Enterprise** menu, select **Monitoring**, and then select **Monitoring Templates**.
 - b. From the search options, select **Oracle WebLogic Server** from the **Target Type** menu, and select the **Display Oracle Certified Templates** check box.
 - c. Select the **Oracle Certified MDA Template for WebLogic Server** radio button, and click **Apply**.
 - d. On the following page, select the **Template will only override metrics that are common to both template and target** radio button. This option is selected by default.
 - e. From Destination Targets, click **Add**.
 - f. From the Search and Select: Targets dialog box, select the WebLogic server on which to enable JMS Destination Metric, and click **Select**.
 - g. Click **Ok** to apply the template on the selected targets.
2. Second method of enabling JMS Destination Metrics:
 - a. From the **WebLogic Server** menu, select **Monitoring**, and select **Metric and Collection Settings**.
 - b. From the **View** menu, select **All Metrics**, and scroll down the column to find **JMS Destination Metrics**.
 - c. Click **Disabled** from the column.
 - d. On the following page, click **Enable**. Once you click Enable, enter the preferred frequency, and click **Continue**.
 - e. On the following page, click **OK**.

18.9 Using Middleware Diagnostics Advisor to View and Diagnose Performance Issues

To use MDA to view and diagnose performance issues, follow these steps:

1. Navigate to the target home page. From the target specific menu, select **Diagnostics**, then select **Middleware Diagnostics Advisor**.
2. On the Middleware Diagnostics Advisor page, view the graphs for the Request Processing Time, and the Requests per minute.
3. To know more about findings, if any, click on the icon below the graphs. The icons will be visible below the specific hour that the finding appeared in.
4. Once you click the findings icon, the messages appear at the bottom of the screen on the Findings tab. Click each finding to get more information, and suggested workarounds. See [Figure 18–2](#) for an example of the Findings page.

Figure 18–2 Diagnostic Findings



18.10 Troubleshooting Issues Related to Middleware Diagnostics Advisor

To troubleshoot issues related to MDA, follow the steps mentioned in the following table.

Tip	How to
Ensure that MDA is enabled on the target.	<ol style="list-style-type: none"> 1. From the Diagnostics menu, select Middleware Diagnostics Advisor Configuration. 2. Verify whether the status column displays enabled for the target. <p>If the status is disabled, enable MDA for the target.</p>
Ensure that MDA Analysis job is running properly.	<ol style="list-style-type: none"> 1. From the Enterprise menu, select Job Activity. 2. Click Advanced Search. 3. In the Name field, enter MDAANALYSISJOB%. 4. In the Target Type field, enter Targetless. 5. From the Status column, select All. 6. Click Go to check the status of the job. <p>If there are no skipped or failed jobs, the Analysis job is running properly.</p>

Tip	How to
Ensure that the mda_dc_collection metric is getting collected from the EMD browser.	<ol style="list-style-type: none"> 1. Enable the EMD browser on the Enterprise Manager Agent which is monitoring the target. 2. Check the mda_dc_collection metric for WebLogic Server targets. 3. Refresh the page after 5 minutes. <p>Once you refresh the page, the data should be updated.</p>
Ensure that the MDA data collector status check job is running properly, and there is no error in the job output.	<ol style="list-style-type: none"> 1. From the Enterprise menu, select Job Activity. 2. Click Advanced Search. 3. On the Advanced Search page, enter MDADCSTATUSJOB as the Name, Targetless as the Target Type, and select All from the Status Column. 4. Click Go. 5. Click the MDADCSTATUSJOB link. 6. Click Show Link to see the job output.
The top SQLs are not displayed in a SQL execution finding	<p>In a SQL execution finding, if the top SQLs are not displayed, ensure that the Database is registered with JVMD. To do so, follow these steps:</p> <ol style="list-style-type: none"> 1. From the Setup menu, select Middleware Management, and then select Application Performance Management. 2. Click the JVM Diagnostics Engines(..) row, and then click Configure. 3. From the Register Databases tab, click Add, and then select Database Instance. 4. From the Search and Select: Targets dialog box, select the database that you want to register, and click Select.

Part VII

Managing Oracle Coherence

The chapters in this part contain information on discovering and monitoring a Coherence cluster.

The chapters are:

- [Chapter 19, "Getting Started with Management Pack for Oracle Coherence"](#)
- [Chapter 20, "Monitoring a Coherence Cluster"](#)

Getting Started with Management Pack for Oracle Coherence

This chapter describes the procedure to discover and monitor a Coherence cluster using Oracle Enterprise Manager Cloud Control 12c. The following sections are covered in this chapter:

- [About Coherence Management](#)
- [New Features](#)
- [Understanding the Discovery Mechanism](#)
- [Enabling the Management Pack](#)

19.1 About Coherence Management

Oracle Coherence is an in-memory data-grid and distributed caching solution. It is composed of many individual nodes or java processes which work together to provide highly reliable and high speed virtual caching.

Enterprise Manager provides deep visibility into performance of all the artifacts such as caches, nodes, and services. The Cluster Home page displays an overview of the performance hotspots such as nodes with minimum available memory, publisher and receiver success rate, and nodes with maximum send queue size. The Cluster Home page provides immediate visibility into the worst caches in the system based on hits to gets ratio which is an overall health indicator for the cluster.

Nodes and caches can be proactively monitored by the Incident Management feature. You can create a monitoring template by pre-populating the monitoring template with metrics for a Coherence target. You can export and import monitoring templates to share monitoring settings between different Enterprise Manager deployments.

Metric Extensions are the next generation of User-Defined Metrics, which enable you to extend Enterprise Manager to monitor conditions specific to the enterprise's environment by creating new metrics for any target type. By including metric extensions in export or imported monitoring templates, multiple metric extensions can be easily shared at the same time between Enterprise Manager deployments.

You can correlate cluster nodes with the underlying hosts to determine CPU and memory utilization on those hosts in order to make better decisions for scaling your clusters. You can see the association of the caches, nodes, hosts and also Oracle WebLogic targets using Coherence*Web applications.

Highly customizable performance views for monitoring performance charts and trends are available. You can overlay metrics for multiple nodes or caches in the same or different cluster for detail analysis to provide detailed visibility at the desired level.

The drill down views allows you to determine the root cause of performance problems or simply identify performance trends in the Coherence Cluster.

To ensure that node departure is detected and resolved in a timely manner, you can setup a monitoring policy to detect departing nodes and replenish those many nodes.

Enterprise Manager provides a centralized cache data management feature that allows you to perform various cache operations such as add/remove index, view cache data, view query explain plan, and so on.

Enterprise Manager monitors the changing configuration of the nodes over a period of time. You can compare configurations of multiple nodes which helps identify performance bottlenecks caused by configuration changes. The Topology Viewer provides a high level topology of the entire cluster and shows the relation between caches, nodes and hosts. You can customize topology view to show some key performance metrics as well.

All of the Coherence Management features are integrated with JVM Diagnostics and provide real-time visibility into the node JVMs. You can drill down to a Coherence node's JVM from within the context of a cache and a cluster to identify the method or thread that is causing a delay. The JVM Diagnostics feature is part of the WLS Management Pack EE and Management Pack for NonOracle Middleware.

Enterprise Manager provides a complete provisioning solution. You can maintain an Oracle Coherence setup image or gold image in the Software Library and deploy it throughout the infrastructure to create completely new clusters or add nodes an existing cluster. You can use the same deployment procedure to updates nodes as well.

19.2 New Features

This section lists the new features in Oracle Enterprise Manager Cloud Control 12c. The new features are:

- **Cache Data Management:** You can view, export, import, insert, update, purge, add, and remove indexes, view query explain plan and view trace. See [Section 20.5, "Cache Data Management"](#) for details.
- **Topology Viewer:** You can view the top-down hierarchy of the Coherence nodes, caches, clusters, and their related targets in the Topology Viewer. See [Section 20.11, "Viewing Configuration Topology"](#) for details.
- **Configuration Metrics Support:** You can store Coherence configuration metrics and perform operations like compare configurations, view last collected configurations, search for configuration data, view saved configurations, and view configuration history. See the Enterprise Manager Lifecycle Management Administrator's Guide for details on configuration metrics.
- **Performance Summary and Metric Palette Integration:** You can plot charts for critical Coherence cluster component metrics such as node metrics, cache metrics, service metrics, and host metrics. See [Section 20.10, "Viewing Performance Summary"](#) for details.
- **JVM Diagnostics Integration:** You can drill down to a Coherence node's JVM from within the context of a cache and a cluster to identify the method or thread that is causing a delay. The JVM Diagnostics feature is part of the WLS Management Pack EE. See [Section 20.9, "Integration with JVM Diagnostics"](#) for details.
- **Log File Alerts:** You can set up alerts based on a pattern in a log file. Log files are periodically scanned for the occurrence of desired patterns and an alert is raised

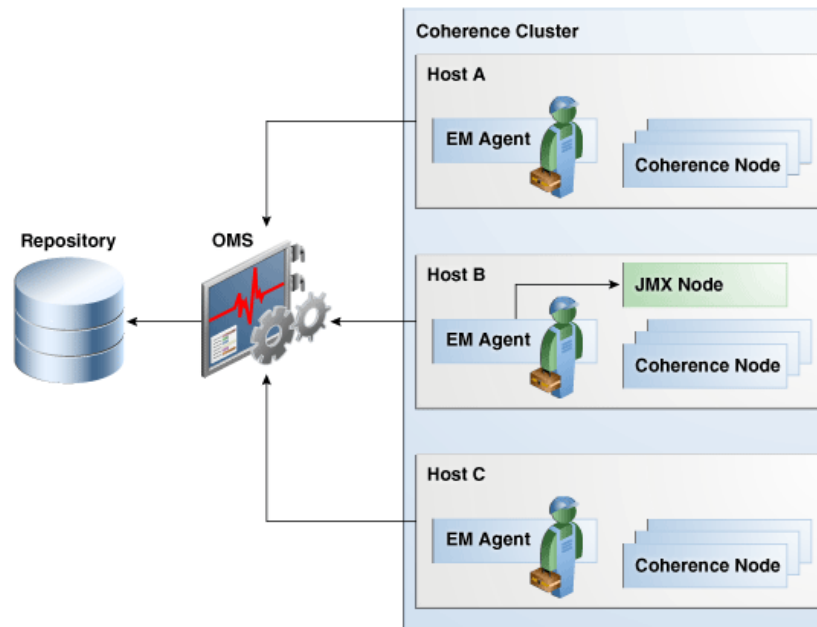
when the pattern occurs during a given scan. See [Section 20.4, "Log File Monitoring"](#) for details.

- **Push Replication:** If a push replication enabled cache is present in your cluster, you can now view the Publisher and Subscriber tables. See [Section 20.7, "Push Replication Pattern"](#) for details.
- **Transaction Cache Support:** You can now view specialized distributed caches or transactional caches in a service. See [Section 20.8, "Transactional Cache Support"](#) for details.
- **Reap Session Support:** You can now view then reap session metrics in the Application Home page. See [Section 20.6, "Reap Session Support"](#) for details.
- **Coherence Provisioning Enhancements:** In this release, the Well Known Address and the Update Nodes features are now available. For more details, see the Enterprise Manager Lifecycle Management Administrator's Guide.
- **Elastic Data Support:** Elastic data is used to seamlessly store data across memory and disk-based devices. To support this feature, new performance and configuration metrics have been added to the Coherence node target in Enterprise Manager 12c.
- **Replenish Storage Enabled Nodes:** You can configure corrective actions to start stopped/crashed storage enabled nodes. This enables cluster operation with minimal disruption.

19.3 Understanding the Discovery Mechanism

Enterprise Manager monitors the entire Coherence cluster and its artifacts. The key targets that can be monitored are Oracle Coherence Cluster, Oracle Coherence Node, and Oracle Coherence Cache. The Oracle Coherence Cluster target provides a high level view of the health of the entire cluster. The Oracle Coherence Node and Oracle Coherence Cache are child targets of the Oracle Coherence Cluster. In addition to monitoring the above target types, additional Coherence components such as Services, Connections, and Applications can also be monitored.

[Figure 19–1](#) shows the Coherence Monitoring Setup.


Figure 19–1 Coherence Monitoring Setup

As shown in [Figure 19–1](#), a typical Coherence deployment has nodes running on one or more hosts. In order to monitor the Coherence cluster in Enterprise Manager, a central Coherence JMX management node must be configured. This JMX management node must expose all Coherence MBeans and attributes. In addition to configuring the JMX management node, the Management Agent must also be installed and configured on the same host as JMX management node. This is required to discover and monitor the Coherence cluster in Enterprise Manager. To provision new Coherence nodes, start, and stop nodes, the Management Agent must be installed on all hosts on which the nodes are running.

19.3.1 Restarting a Node

If a node is restarted or needs to rejoin the cluster due to some network related issues, you must refresh the cluster as described in [Section 19.3.6, "Refreshing a Cluster"](#). Alternatively, you can set up a Refresh job as shown below:

1. From the **Enterprise** menu, select **Job**, then select **Library** to navigate to the Job Library page.
2. Select the Discover New Oracle Coherence Node job type and click **Go**.

Figure 19–2 Job Library Page

3. In the General page, enter a name for the job and click **Add** to add the cluster target.

Figure 19–3 Job General Page

Create 'Discover New Oracle Coherence Nodes' Library Job

Cancel Save to Library

General Parameters Credentials Schedule Access

* Name
 Description
 Target Type Oracle Coherence Cluster

Target
 Add individual targets or one composite target, such as a Group.

Remove Add

Select All Select None

Select	Name	Type	Host	Time Zone
<input type="checkbox"/>	PRODUCTION_BUILD_3.1.26.1245_dslprod	Oracle Coherence Cluster		Greenwich Mean Time

Cancel Save to Library

4. Navigate to the Schedule page. Select option to make it run continuously with interval of few minutes. If the cluster is very dynamic, you can enter 1 minute. Click **Save to Library**.

Figure 19–4 Job Schedule Page

Create 'Discover New Oracle Coherence Nodes' Library Job

Cancel Save to Library

General Parameters Credentials Schedule Access

Type ☐ One Time (Immediately) ☐ One Time (Later) ☒ Repeating

Frequency Type By Minutes

Repeat Every 5 Minutes

Time Zone (UTC-08:00) US Pacific Time (PST)

Start Date Apr 12, 2013

Start Time 7:05 AM PM

Grace Period ☒ Indefinite ☐ End After Hours Minutes

Repeat Until ☒ Indefinite ☐ Specified Date

Date (example: Apr 12, 2013)

Time AM PM

Cancel Save to Library

5. On the Jobs Library page, select the job and click **Submit**.

Figure 19–5 Job Library Page (II)

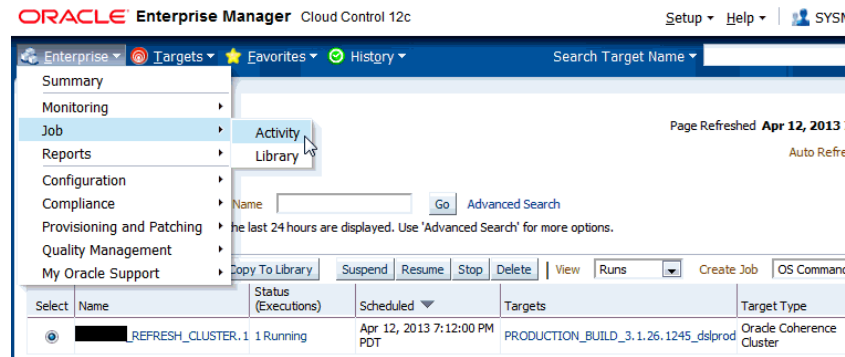
Job Library

Job Type All Name Owner All Go

Submit Create Like Edit Delete Create Library Job OS Command Go

Select	Name	Job Type	Owner	Targets	Target Type
<input type="radio"/>	DISABLE TABLESPACE USED (%) ALERTS FOR UNDO AND TEMP TABLESPACES	SQL Script	SYSMAN		
<input checked="" type="radio"/>	REFRESH_CLUSTER	Discover New Oracle Coherence Nodes	SYSMAN	PRODUCTION_BUILD_3.1.26.1245_dslprod	Oracle Coherence Cluster

6. Navigate to Jobs Activity page and make sure that the job has been scheduled.

Figure 19–6 Job Activity Page

19.3.2 Starting a JMX Management Node

The Management Agent uses the JMX management node (centralized MBean server) to discover and monitor the entire Coherence cluster, including the nodes and caches. As a best practice, it is recommended that the Management Agent is present on the same host as the JMX management node that is used to discover and monitor the Coherence cluster. The Management Agent must be setup on all the machines on which the Coherence nodes are running to monitor and provision the cluster. See the *Using JMX to Manage Coherence* chapter in the *Oracle Coherence Management Guide* for more details on using JMX to manage Oracle Coherence.

The following .jar files must be present in the CLASSPATH along with other application specific .jar files.

```
<OEM_Agent_Home>/plugins/oracle.sysman.emas.agent.plugin_
12.1.1.0.4.0/archives/coherence/coherenceEMIntg.jar

<OEM_Agent_Home>/plugins/oracle.sysman.emas.agent.plugin_
12.1.1.0.4.0/archives/coherence/bulkoperationsmbean.jar
```

Note: The location of the .jar files may change based on the plugin version.

The JMX node must be started using the this Java class (as main class)

```
oracle.sysman.integration.coherence.EMIntegrationServer
```

This class allows you to ensure that the necessary MBeans have been added to the JMX server.

The following mandatory parameters must be added to the Java start command of the JMX node.

- -Dtangosol.coherence.management=all (enables monitoring for all nodes)
- -Dtangosol.coherence.management.remote=true
- -Dcom.sun.management.jmxremote.port=<open TCP port> (required for remote connection)
- -Dtangosol.coherence.cluster=<cluster name>
- -Dtangosol.coherence.member=<Unique member name across the cluster> (required for target name)

- `-Dtangosol.coherence.machine=<fully qualified hostname>`
- `-Doracle.coherence.machine=<fully qualified hostname>` (required to correlate with the host target)
- `-Dtangosol.coherence.distributed.localstorage=false` (disables caching and ensures that the node is a dedicated monitoring node)
- `-server`
- `-Xms4g -Xmx4g` (the JMX node must have adequate resources)
- `-Xms256m -Xmx256m` (you must ensure that there are adequate resources for the management node)
- `-Dtangosol.coherence.management.refresh.expiry=1m` (ensures that the monitoring data is refreshed every minute. If this is not set, the monitoring data is refreshed every second which increases the load on the JVM)

Unless JMX authentication is required, add the following Java options:

- `-Dcom.sun.management.jmxremote.authenticate=false`
- `-Dcom.sun.management.jmxremote.ssl=false`

For more details on JMX authentication, refer to the *Java* documentation.

19.3.2.1 Starting a JMX Management Node Using Shell

If you are using a shell script to start the JMX Management Node, you must use the following startup class for Java.

```
oracle.sysman.integration.coherence.EMIntegrationServer
```

19.3.2.2 Starting a JMX Management Node Using WebLogic Console

This option is supported only for Coherence 3.7 and later versions.

If you are starting the JMX Management Node using the WebLogic Console, you cannot customize the JMX management node start up class. To address this, you must add the `custom-mbeans.xml` to the Coherence server that will act as the JMX management node. Add the following snippet to the `custom-mbeans.xml` file.

```
<mbeans>
  <mbean id="100">
    <mbean-class>oracle.sysman.integration.coherence.CacheDataManager
    </mbean-class>
    <mbean-name>type=Custom,name=CacheDataManager
    </mbean-name>
    <enabled>true</enabled>
  </mbean>
  <mbean id="110">
    <mbean-class>oracle.as.jmx.framework.bulkoperations.BulkOperationsMBeanImpl
    </mbean-class>
    <mbean-name>type=BulkOperations</mbean-name>
    <enabled>true</enabled>
  </mbean>
</mbeans>
```

19.3.3 Starting Other Nodes

To start other nodes in the cluster, you must add the following Java options:

- `-Dtangosol.coherence.management.remote=true`

- `-Dtangosol.coherence.member= <unique member name in the cluster>`
- `-Dtangosol.coherence.cluster= <cluster name>`
- `-Dtangosol.coherence.machine=<fully qualified hostname>`
- `-Doracle.coherence.machine=<fully qualified hostname>`

Note: If you are planning to start or stop Coherence nodes from Enterprise Manager, you must specify the following additional start Java options:

```
-Doracle.coherence.startscript=<node start script>
-Doracle.coherence.home=<coherence home>
```

If you are using the WebLogic Console to start a Coherence node (or Coherence Server as referred to within the WebLogic Console), you can use the WebLogic Console to customize the start arguments of the server.

If you are starting the Coherence node in a WebLogic Managed Server, you can specify the above parameters in the `startWebLogic.sh` or in the script you use to start the WebLogic server.

19.3.3.1 Verifying the JMX Management Node Configuration

After the JMX management node has been configured, you can use tools such as JConsole to verify if all the Coherence MBeans and attributes are available at the management node. Ensure that you are using the Machine Name:Port (and JMX credentials if any) to connect to the JMX management node.

19.3.4 Using JVM Diagnostics with Coherence

JVM Diagnostics provides deep visibility into the runtime of the JVM. JVM Diagnostics allows administrators to identify the root cause of performance problems in the production environment without having to reproduce them in the test or development environment. You can view the JVM Diagnostics data if the JVM Diagnostics Manager and JVM Diagnostics Agent have been deployed on the host machine on which the OMS running. To setup JVM Diagnostics on each Coherence node, you must download the JVM Diagnostics Agent.

To download the JVM Diagnostics Agent, follow the steps listed in the Enterprise Manager Cloud Control Administrator's Guide. When the JVM Diagnostics is downloaded, the `jamagent.war` file is downloaded. You must to copy the `.war` file to all machines on which the Coherence nodes are to be integrated with JVM Diagnostics, and add it to the classpath. See [Section 20.14, "Best Practices"](#) for more information on how to start a Coherence node with JVM Diagnostics.

19.3.5 Discovering Coherence Targets

Prerequisites

To monitor a Coherence cluster in Enterprise Manager, the following prerequisites must be met:

- Install the 12.1.0.3.0 Management Agent on all hosts where Coherence nodes are running.
- Deploy the 12.1.0.4.0 Fusion Middleware Plug-in on all the Management Agents.

- Verify that all Coherence MBeans are available in the Coherence JMX management node as described in [Section 19.3.3.1, "Verifying the JMX Management Node Configuration"](#).

To discover a Coherence target, follow these steps:

1. Login to Enterprise Manager as an administrator with the **Add Target** privilege.
2. From the **Targets** menu, select **Middleware**. You will see a list of Middleware targets.

Note: Alternatively, you can add a Coherence target from the Setup menu. From the **Setup** menu, select **Add Target**, then select **Add Targets Manually**. In the Add Targets Manually page, select the **Add Non-Host Targets Using Guided Process** option. Follow the steps in the wizard to add the Coherence target.

3. Select Oracle Coherence in the **Add** drop-down box and click **Go**. The Oracle Coherence Cluster: Discover Cluster, Node, and Cache Targets page is displayed.

Figure 19–7 Add Coherence Target

Add Oracle Coherence Cluster: Discover Cluster, Node and Cache Targets Continue Cancel

Oracle Coherence Management using Enterprise Manager requires coherenceEMIntg.jar and bulkoperationsmbean.jar to be included in the Coherence Management (MBean Server) Node classpath. Ensure that the Management Node is started with these jars prior to Coherence discovery. Enterprise Manager Oracle Coherence Cluster discovery will discover the Cluster, Nodes and Caches from the specified Management (MBean Server) Node.

MBean Server Connection
Specify the access details for Coherence MBean Server. You may specify either the Service URL or the Host, Port and Service.

☒ Enter Host, Port and Service

* Management Node Host * JMX Remote Port

Service Name

☐ Enter Service URL

MBean Server Credentials
Specify username and password for accessing the Coherence MBean Server if security is enabled.

Username Password

Oracle Coherence Cluster Discovery and Monitoring Agent
Select an Enterprise Manager Agent that is co-located on the host of Management Node or has access to Management Node. This agent will be used to monitor Coherence Cluster target.

* Agent

4. On this page, specify the connection details of the Coherence JMX management node. This is required to discover the Coherence cluster, node and cache targets. You can select either of the following options to provide MBean Server details:
 - Host, Port, and Service: Enter the following details:
 - Select the host on which the Management Node is running.
 - The port used for the JMX RMI connection. If you are using the MBean connector for Coherence MBeans, specify the `tangosol.coherence.management.remote.connectionport` property.

Note: It is recommended that you use the `com.sun.management.jmxremote.port` property.

- The service name used for the connection. The default is `jmxrmi`.
- Service URL that will be used for the connection. You may need to specify the Service URL only in complex cases like when the RMI registry and the MBean Server ports are different. It is recommended that you use the Machine Name and Port option for the MBean server connection.

5. Select the Management Agent that will be used to monitor the Coherence target and click **Continue**.
6. The details of the discovered targets are displayed. Click **Add Targets** to add these targets to Enterprise Manager.

Note: To automatically discover a new node or target in Enterprise Manager, you must refresh the cluster as described in [Section 19.3.6, "Refreshing a Cluster"](#). Alternatively, you can set up a Refresh job as described in [Section 19.3.1, "Restarting a Node"](#).

19.3.6 Refreshing a Cluster

You can manually synchronize the cluster targets with the running Coherence cluster. Click **Refresh** Cluster from the Oracle Coherence Cluster menu. A message indicating that new Coherence nodes and caches that have been discovered will be added as Enterprise Manager targets is displayed. Nodes are updated if there are any changes to their attributes. Click **Continue** to refresh the cluster. This ensures that the latest changes are applied.

Note: Decommissioned nodes and caches will not be removed during the **Refresh** process. You must remove them manually.

19.4 Enabling the Management Pack

You must enable the Management Pack for Oracle Coherence if you want to access additional features beyond Coherence cluster monitoring. To enable the Management Pack, do the following:

1. From the **Setup** menu, select **Management Packs**, then select Management Pack Access.
2. Select **Oracle Coherence** in the Search drop-down list and click **Go**.
3. All the Coherence targets being monitored are displayed. Check the **Pack Access Agreed** check box for the Coherence target and click **Apply** to enable the Management Pack.

Monitoring a Coherence Cluster

After you have discovered the Coherence target and enabled the Management Pack Access, you can start monitoring the health and performance of the cluster. You can monitor the entire cluster or drill down to the various entities of the cluster like nodes, caches, services, connection managers, and connections.

This chapter contains the following sections:

- [Cluster Level Pages](#)
- [Detailed Pages](#)
- [Performance Pages](#)
- [Log File Monitoring](#)
- [Cache Data Management](#)
- [Reap Session Support](#)
- [Push Replication Pattern](#)
- [Transactional Cache Support](#)
- [Integration with JVM Diagnostics](#)
- [Viewing Performance Summary](#)
- [Viewing Configuration Topology](#)
- [Viewing Incidents](#)
- [Troubleshooting Coherence](#)
- [Best Practices](#)

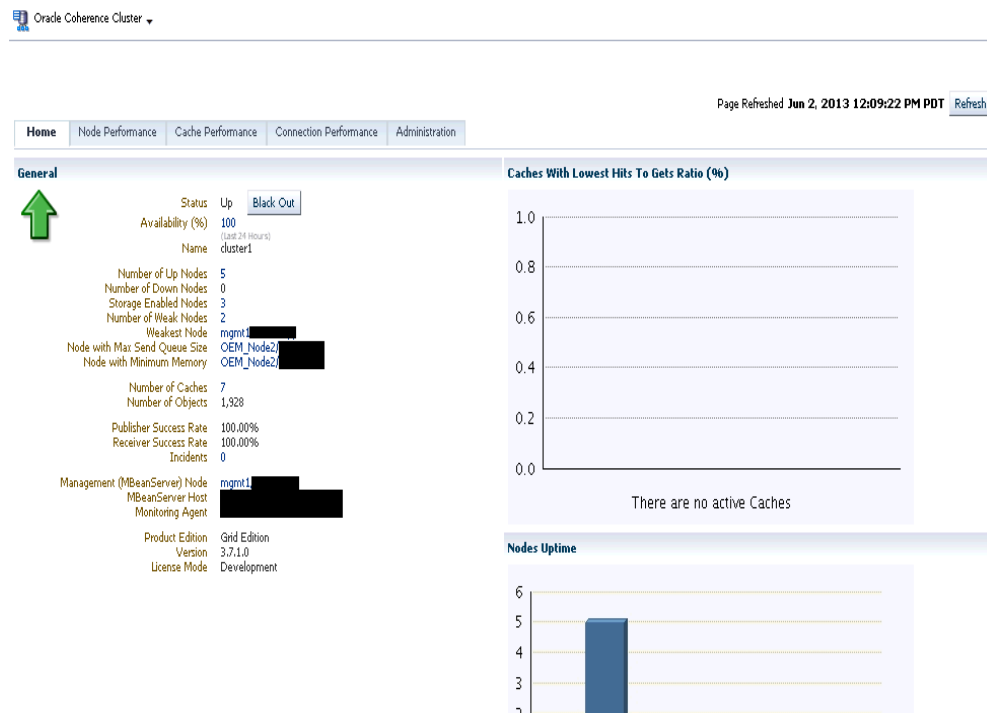
20.1 Cluster Level Pages

At the cluster level, you can view the Home page for the cluster, view the performance of all nodes, caches, and connections in the cluster, and change configuration for the different entities in the cluster.

20.1.1 Cluster Level Home Page

You can get a global view of the cluster from the Home page by following these steps:

1. From the **Targets** menu, select **Middleware**, then click on a Coherence Cluster target. The Coherence Cluster Home page is displayed.

Figure 20–1 Cluster Home Page

2. This page contains the following sections:

- General
- Graphs
- Cluster Management
- Service
- Applications

20.1.1.1 General

This section contains the following details:

- Name and status of the cluster.
- Availability%: This is really the availability of the cluster management node over the last 24 hours.
- Number of Up Nodes: The number of nodes that are Up. Click on the link to drill down to the Selected Nodes Performance page.
- Number of Down Nodes: The number of nodes that are Down. Click on the link to drill down to the Down Nodes page.
- Storage Enabled Nodes: Indicates the number of nodes that are storage enabled.
- Number of Weak Nodes: The number of nodes that are weak and have communication and performance issues. Click on the link to drill down to the Node Performance page.
- Weakest Node: The weakest node in the cluster. Click on the link to drill down to the Node Home page.

- **Node with Max Queue Size:** Indicates the node with the maximum queue size value in the cluster.
- **Node with Minimum Memory:** Indicates the node with the minimum available memory in the cluster.
- **Number of Caches and Objects:** The number of caches in the cluster and the number of objects stored in all caches in the cluster. Click on the Number of Caches link to drill down to the Cache Performance page.
- **Publisher and Receiver Success Rates:** The Publisher and Receiver success rate for this cluster node since the node statistics were last reset.
- **Management Bean Server Node:** This is the management node that contains the Coherence MBeanServer. Click on the link to drill down to the Node Home page.
- **Incidents:** Number of incidents for the target and its members.
- **Monitoring Agent:** The Oracle Management Agent monitoring the cluster.
- **MBean Server Host:** The host on which the management node is running. If the node on the MBean Server Host is not accessible, the monitoring capability of the entire cluster will be affected. To avoid this, we recommend that at least two management nodes are running on the cluster. If a management node departs from the cluster, you must update the host and port target properties to point to the host with the running management node.
- **License Mode:** The license mode that this cluster is using. Possible values are Evaluation, Development or Production.
- **Product Edition:** The product edition that this cluster is running on. Possible values are: Standard Edition (SE), Enterprise Edition (EE), Grid Edition (GE).
- **Version:** The Coherence version.

20.1.1.2 Graphs

Graphs indicating the health of the cluster are displayed here. The following graphs are displayed:

- **Nodes Uptime:** This graph displays groups of nodes according to their uptime. The Node Uptime is calculated as the difference between the Current Time and the Node Timestamp. Nodes that have an uptime of less than a minute are displayed in the seconds bar, nodes with an uptime of less than a hour are displayed in the minutes bar and so on.
- **Caches with Lowest Hits to Gets Ratio:** This graph shows caches (up to a maximum of 5) that have lowest Hits to Gets ratio. Click on the cache name in the legend section to drill down in to the Cache Details page to further investigate the reasons for the low hits to gets ratio.

Note: The data in the General and Graphs section will be refreshed after one collection interval.

20.1.1.3 Cluster Management

In this section, you can start and stop one or more nodes, or stop a cluster.

Prerequisites

To perform Cluster Management operations, you must ensure that:

- The hosts on which the nodes are going to be started or stopped must be monitored targets in Enterprise Manager.
- The Coherence nodes are started with `-Dtangosol.coherence.machine` and `-Doracle.coherence.machine` Java options and the names match the host names monitored by Enterprise Manager.
- The Coherence nodes are started with `-Doracle.coherence.startscript` and `-Doracle.coherence.home` Java options.

The `oracle.coherence.startscript` option specifies the absolute path to the start script needed to bring up a Coherence node. All customizations needed for starting this node should be in this script. The `oracle.coherence.home` option specifies the absolute path to the location in which the coherence folder is present which is `$INSTALL_DIR/coherence`. This folder contains Coherence binaries and libraries.

- Preferred Credentials have been setup for all hosts on which Cluster Management operations are to be performed.

You can do the following:

- **Start New Nodes:** You can start one or more nodes based on an existing node. The new node will have the same configuration as the existing node.
- **Stop Nodes:** You can stop all the nodes on a specific host.
- **Stop Cluster:** You can stop an entire cluster if all the hosts are managed by Enterprise Manager Cloud Control.

Note: You can set up a corrective action to start departed Coherence Nodes automatically. When a node departs, this corrective action is launched, and the new node is automatically started on the same host on which the node was previously running.

- `oracle.coherence.startscript`: The absolute path to the start script needed to bring up a Coherence node. All customizations needed for starting this node should be in this script.
 - `oracle.coherence.home`: The absolute path to the location in which the coherence folder is present which is `$INSTALL_DIR/coherence`. This folder contains Coherence binaries and libraries.
-

20.1.1.4 Services

This section shows all the services in the cluster. You can view the type of service, status of the service (Machine-Safe, Node-Safe, and Endangered), the number of nodes supporting the service, storage enabled nodes, endangered nodes, and active transactions.

20.1.1.5 Applications

This section shows the applications that use this Coherence cluster to cache their HTTPSession Objects. You can view details of the Local Cache, Overflow Cache, and Servlet Context Cache.

20.1.1.6 Cluster Level Operations

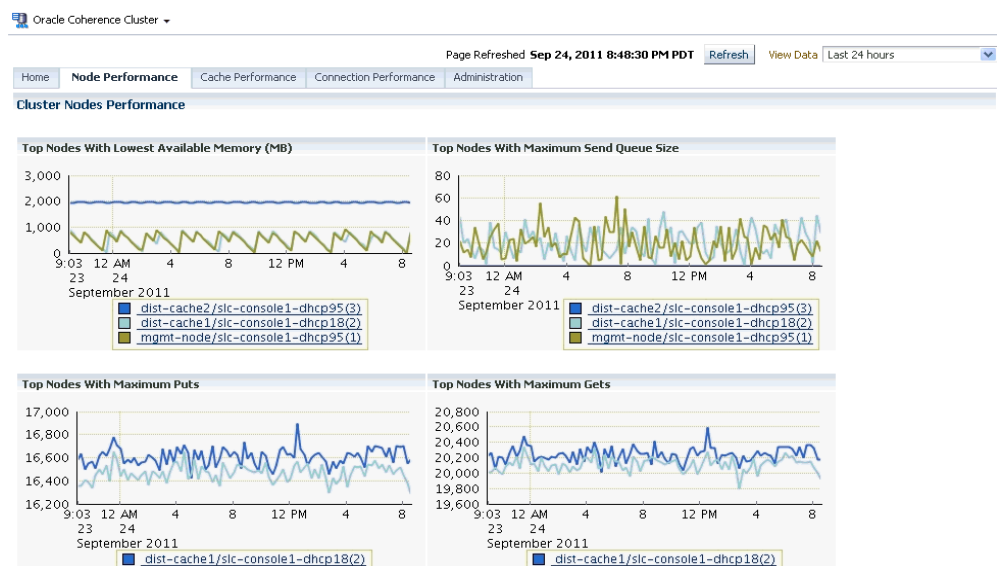
From the Cluster Home page, you can several other operations such as:

- **Viewing The Performance Summary:** From the **Oracle Coherence Cluster** menu, select **Monitoring**, then select **Performance Summary**. You can view the performance of the cluster on this page.
- **Metric and Collection Settings:** From the **Oracle Coherence Cluster** menu, select **Monitoring**, then select **Metric and Collection Settings**. You can set up corrective actions to add nodes and caches as Enterprise Manager targets.
- **Refreshing Cluster:** From the **Oracle Coherence Cluster** menu, select **Refresh Cluster**. You can refresh a cluster to synchronize Coherence targets in Enterprise Manager with a running cluster.
- **Coherence Node Provisioning:** From the **Oracle Coherence Cluster** menu, select **Coherence Node Provisioning**. You can deploy a Coherence node across multiple targets in a farm. See Enterprise Manager Lifecycle Management Administrator's Guide for more details on Coherence Node Provisioning.
- **Last Collected Configuration:** From the **Oracle Coherence Cluster** menu, select **Configuration**, then select **Last Collected**. You can view the latest or saved configuration data for the Coherence cluster.
- **Topology:** From the **Oracle Coherence Cluster** menu, select **Configuration**, then select **Topology**. The Configuration Topology Viewer provides a visual layout of the Coherence deployment.
- **JVM Diagnostics:** From the **Oracle Coherence Cluster** menu, select **JVM Diagnostics** to view the JVM Pool Performance Diagnostics page.

20.1.2 Cluster Level Node Performance Page

This page displays the node related performance over a specified period of time. Click the **Node Performance** tab to view this page.

Figure 20–2 Coherence Node Performance Page



This page displays the performance of all the nodes in the cluster over a specified period of time. You can see charts showing upto 5 top nodes with lowest available memory, maximum send queue size, maximum puts, and maximum gets. By default, you can see the average performance metrics for the last 24 hours in all the

Performance pages. To view the real time charts, select one of the Real Time options in the View Data drop-down list in any of Performance pages. Using the View Data options, you can also view the average performance metrics for the last 7 or 31 days. The Node Performance page tab shows the performance of all nodes in this cluster. If you click on a link that shows multiple nodes like weak nodes, storage nodes, etc., the performance of the selected nodes will be displayed on this page. You can toggle between the two modes to see the performance of the selected nodes or all the nodes.

20.1.3 Cluster Level Cache Performance Page

This page displays the cache related performance over a specified period of time. You can view the performance of the top caches or all the caches.

Figure 20–3 Cluster Level Cache Performance



Select the **All Caches** option from the drop-down list to view the performance of all the caches in the cluster. If you select the All Caches option, you can see the total and average metric values over the selected period of time.

20.1.4 Cluster Level Connection Performance Page

This page shows the performance of all connection managers and connections in the cluster. The following Connection Manager graphs are displayed:

- Top Connection Managers with Most Bytes Sent since the connection manager was last started.
- Top Connection Managers with Most Bytes Received since the connection manager was last started.

A table with the list of Connection Managers is displayed with the following details:

- **Connection Manager:** This is the name of the connection manager. It indicates the Service Name and the Node ID where the Service Name is the name of the service used by this Connection Manager. Click on the link to drill down to the Connection Manager Home page.

- Service: The name of the service. Click on the link to drill down to the Service Home page.
- Node ID: An ID is automatically assigned to any node that is part of the cluster.
- Bytes Sent: The number of bytes sent per minute.
- Bytes Received: The number of bytes received per minute.
- Outgoing Buffer Pool Capacity: The maximum size of the outgoing buffer pool.
- Outgoing Byte Backlog: The number of outgoing bytes in the backlog.

The following Connection related graphs are displayed:

- Top Connections with Most Bytes Sent since the connection was last started.
- Top Connections with Most Bytes Received since the connection was last started.

A table with the list of connections is displayed. Click on the link to drill down to the Details page.

- Remote Client: The host on which this connection exists.
- Up Since: The date and time from which this connection is running.
- Connection Manager: This is the name of the connection manager. Click on the link to drill down to the Connection Manager Home page.
- Service: The name of the service. Click on the link to drill down to the Service Home page.
- Node ID: An ID is automatically assigned to any node that is part of the cluster.
- Bytes Sent: The number of bytes sent per minute.
- Bytes Received: The number of bytes received per minute.
- Connection Time: The connection time in minutes.
- Outgoing Message Backlog: The number of outgoing messages in the backlog.
- Outgoing Byte Backlog: The number of outgoing bytes in the backlog.

20.1.5 Cluster Level Administration Page

This page allows you to change the configuration of nodes, caches, and services.

Figure 20–4 Cluster Level Administration Page

cluster1 ad2170836.us.oracle.com

Oracle Coherence Cluster

Home Node Performance Cache Performance Connection Performance Administration

Node

Change Configuration

Select a node to change its configuration

Node OEM_Node1

Service

Change Configuration

Select a service and its node to change configuration

Service DistributedCache Node store2

Cache

Change Configuration

Select a cache and its node to change configuration

Service DistributedCacheForMessages Cache coherence.messagingpattern.messages Node store2 Tier back Loader None

On this page, you can select an entity (node, cache, or service) for which the configuration needs to be modified and click **Go**. The Change Configuration page is displayed. Enter the new values and click **Update** to save the values and return to the Coherence Cluster Administration page.

20.2 Detailed Pages

From the cluster level pages, you can click on any hyperlink and drill down to the detailed home pages of each entity in the cluster. Hyperlinks from one entity allow you to go to another entity. For example, you can view all the nodes of a cache in the Cache Detailed page, identify the node that is not contributing well for this cache, click on the node hyperlink and drill down to the Node Detailed page for further investigation.

- **Viewing The Performance Summary:** From the **Oracle Coherence Node** or **Oracle Coherence Cache** menu, select **Monitoring**, then select **Performance Summary**. You can view the performance of the node or cache on this page.
- **Metric and Collection Settings:** From the **Oracle Coherence Node** or **Oracle Coherence Cache** menu, select **Monitoring**, then select **Metric and Collection Settings**.
- **Last Collected Configuration:** From the **Oracle Coherence Node** or **Oracle Coherence Cache** menu, select **Configuration**, then select **Last Collected**. You can view the latest or saved configuration data for the Coherence cluster.
- **Topology:** From the **Oracle Coherence Node** or **Oracle Coherence Cache** menu, select **Configuration**, then select **Topology**. The Configuration Topology Viewer provides a visual layout of the relationship of the Coherence cluster with other targets.

20.2.1 Node Home Page

This page shows the details and controls of a specific selected node in the Coherence cluster.

Figure 20–5 Node Home Page



This page contains the following sections:

■ General

- Node ID: When a node becomes part of a cluster, an ID is automatically assigned to the node. This ID can be a new number or the number that was assigned to a Departed Node.
- Machine Name: The name of the machine on which the node is running.
Note: You must set the machine name property when a node is started. To set the machine name, set the `-Dtangosol.coherence.machine=HOST_NAME` flag in the start coherence script. The `HOST_NAME` is the complete name of the machine on which the node is running. This allows you to associate the machine name with the Host Details page. See the Coherence documentation for more details on setting this flag.
- Coherence Cluster: The name of the Coherence Cluster target.
- Role Name: The role could be storage/data, application/process, proxy or management node.
- Site Name: This indicates the location of the coherence node.
- Rack Name: Name of the rack where the machine is located.
- Member Name: The unique name assigned to the Member.
- Up Since: The date and time from this node has been up and running.
- Product Edition: The product edition this Member is running. Possible values are: Standard Edition (SE), Enterprise Edition (EE), Grid Edition (GE).
- Well Known Addresses: If cluster communication has been set to WKA (well known address), the host name and port number are displayed.

- Multicast Port / Address: The IP address and port number of the Multicast Socket used for group communication.
- Quorum Status: The current Quorum state.
- Incidents: Number of incidents for the target and its members.
- Host: Click on the host name to drill down to the Host Home page.
- **Memory Usage:** The maximum memory and the available memory on this node are displayed as line graphs.
- **Node Management:** In this section, you can click **Stop Node** to stop this node and click **Reset Statistics** to reset all the statistics for this node. See **Prerequisites** in [Section 20.1.1.3, "Cluster Management"](#).
- **Components:** The list of components and the type of each component in the cluster are displayed here. Click on the link to drill down to the Component Home page.

20.2.2 Cache Home Page

This page shows the details of the specific selected cache in the Coherence cluster.

Figure 20–6 Cache Home Page



This page contains the following sections:

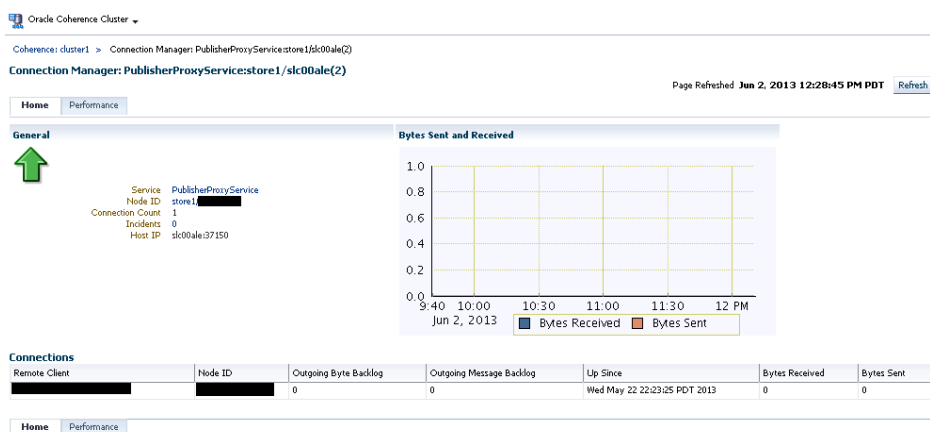
- **General**
 - Status: Indicates the status of the cache.
 - Availability: The percentage of time that the management agent was able to communicate with the cache. Click the percentage link to view the availability details for the past 24 hours.
 - Name: The unique name assigned to the cache.
 - Coherence Cluster: The name of the cluster. Click on the link to drill down to the Cluster Home page.
 - Number of Nodes: The number of nodes on which the cache is running. Click on the link to drill down to the Node Home page.
 - Service Name: The name of the service used by this cache.

- Number of Objects: Shows the number of objects in the cache.
- Memory Consumed: The amount of memory used by the cache in units.
- Queue Size: The size of the write-behind queue size. Applicable only for WRITE-BEHIND persistence type.
- Incidents: Number of incidents for the target and its members
- **Cache Hits and Misses:** This graph displays the total cache hits and misses per minute.
- **Nodes:** This section lists all the nodes supporting this cache.
 - Node ID: The ID number assigned to the node. Click on the link to drill down to the Node Home page.
 - Persistence Type: The persistence type for this cache. Possible values include: NONE, READ-ONLY, WRITE-THROUGH, WRITE-BEHIND.
- **Push Replication Tables:** If a push replication enabled cache is present in your cluster, you can see the Publisher and Subscriber tables.

20.2.3 Connection Manager Home Page

Use this page to view the details and controls of a Connection Manager instance in the Coherence cluster.

Figure 20–7 Connection Manager Home Page



This page contains the following sections:

- **General**
 - Service Name: The unique name assigned to the service.
 - Node ID: An ID is automatically assigned to any node that is part of the cluster. This can be a new number or a number that belongs to a Departed node.
 - Connection Count: The number of connections associated with the connection manager instance.
 - Host IP: The IP address of the host machine.

- Refresh Time: The date and time on which the connection manager instance was last refreshed.
- **Bytes Sent and Received:** This graph displays the number of bytes that were sent and received per minute. Click on the graph to drill down to the Bytes Sent Metric page.
- **Connections**
 - Remote Client: A unique hexadecimal number assigned to each connection.
 - Service Name: The unique name assigned to the service.
 - Node ID: An ID is automatically assigned to any node that is part of the cluster. This can be a new number or a number that belongs to a Departed node.
 - Outgoing Byte Backlog: The number of outgoing bytes in the backlog.
 - Outgoing Message Backlog: The number of outgoing messages in the backlog.
 - Up Since: The date and time from which the connection manager instance is up.
 - Bytes Received: The number of bytes received per minute.
 - Bytes Sent: The number of bytes sent per minute.


20.3 Performance Pages

Entity level performance pages give detailed views of the performance of this particular entity. For caches performance pages, there are two views - Charts and Metrics. You can select the option from the View drop-down list.

20.3.1 Cache Performance Details Page

This page displays the performance of a specific cache over a specific period of time. You can view charts showing the number of cache hits, misses, store reads, and store writes. You can also see the aggregated totals and average metric values over the selected period of time.

Figure 20–8 Cache Performance Details Page



The following graphs are displayed:

- Number of Objects: This graph shows the number of objects in the cache.
- Memory Consumed (Units): This graph shows the amount of memory consumed by the cache per minute.
- Hits: This graph shows the number of cache hits per minute.
- Misses: This graph shows the number of cache misses per minute.
- Puts: This graph shows the total number of put() operations per minute.
- Gets: This graph shows the total number of get() operations per minute.
- Store Reads: This graph shows the total number of load operations per minute.
- Store Writes: This graph shows the total number of store and erase operations per minute.
- Published Batches (per min)
- Replicated Entries
- Publishing Failures (per min)

Metrics: Select the **Metrics** option in the View By drop-down list to view the totals and averages for the cache.

- Totals / Averages: The aggregated total and average values across all nodes and per node total / average values for the cache during the selected period are displayed.
 - Hits: The number of successful fetches of the cached objects per minute.
 - Misses: The number of failed fetches of the cached objects per minute.
 - Puts: The number of addition of objects to a cache per minute.
 - Gets: The number of retrieval of objects from a cache per minute.

- Prunes: The number of prune operations on the cache per minute. A prune operation occurs every time the cache reaches its high watermark.
 - Store Reads: The number of reads from a data store per minute.
 - Store Writes: The number of writes to a data store per minute.
 - Number of Objects: The number of objects in the cache.
 - Memory Consumed (Units): The amount of memory used by the cache in units.
 - Average Time (ms): The average execution time for each operation (hits, misses, puts, gets, store reads, and store writes) are displayed. The values since the last collection divided by the time taken by the operations are displayed. For example, if there are 100 Hits and these 100 Hits took 20 milliseconds since last collection, this value is calculated as $100/20 = 5$.
- Storage Manager
 - Events Dispatched: The total number of events dispatched by the Storage Manager per minute.
 - Eviction Count: The number of evictions from the backing map managed by this Storage Manager caused by entries expiry or insert operations that would make the underlying backing map to reach its configured size limit.
 - Insert Count: The number of inserts into the backing map managed by this Storage Manager. In addition to standard inserts caused by put and invoke operations or synthetic inserts caused by get operations with read-through backing map topology, this counter is incremented when distribution transfers move resources into the underlying backing map and is decremented when distribution transfers move data out.
 - Remove Count: The number of removes from the backing map managed by this Storage Manager caused by operations such as clear, remove or invoke.
 - Listener Filter Count: The number of listener filters.
 - Listener Key Count: The number of listener keys.
 - Listener Registrations: The number of listener registrations.
 - Locks Granted: The number of locks granted.
 - Locks Pending: The number of locks pending.

20.3.2 Connection Manager Performance Page

This page displays the performance of the selected connection manager over a specified period of time. The following graphs are displayed:

- Bytes Sent: This graph shows the number of bytes sent since the connection manager was last started.
- Bytes Received: This graph shows the number of bytes received since the connection manager was last started.

Performance:

The average performance over the selected period is displayed.

- Outgoing Byte Backlog: The number of outgoing bytes in the backlog.
- Outgoing Message Backlog: The number of outgoing messages in the backlog.

- Incoming Buffer Pool Capacity: The maximum size of incoming buffer pool.
- Incoming Buffer Pool Size: The currently used value of the incoming buffer pool.
- Outgoing Buffer Pool Capacity: The maximum size of the outgoing buffer pool.
- Outgoing Buffer Pool Size: The currently used value of the outgoing buffer pool.
- Bytes Received: The number of bytes received per minute.
- Bytes Sent: The number of bytes sent per minute.

20.3.3 Connection Performance Page

This page displays the performance of the specific connection over a specified period of time. The following graphs are displayed:

- Bytes Sent: This graph shows the number of bytes sent within the specified period.
- Bytes Received: This graph shows the number of bytes received within the specified period.

Performance

- Remote Address: The IP address of the remote machine from which the bytes are being received.
- Remote Port: The port number of the remote machine from which the bytes are being received.
- Up Since: The date and time from which the connection has been up and running.
- Bytes Received: The number of bytes received per minute.
- Bytes Sent: The number of bytes sent per minute.
- Total Messages Received: The total number of messages received per minute.
- Total Messages Sent: The total number of messages sent per minute.

20.3.4 Service Performance Page

This page displays the performance of the selected service over a specific period of time. The Request Average Duration and the Request Max Duration charts are displayed. You can also see the average metric values over the selected period of time.

20.3.5 Administration Pages

You can drill down to the Administration page for a specific node, cache, connection, or connection manager. You can change configuration for a node, cache, connection, or connection manager. See the Online Help for configuration parameters that can be changed.

20.4 Log File Monitoring

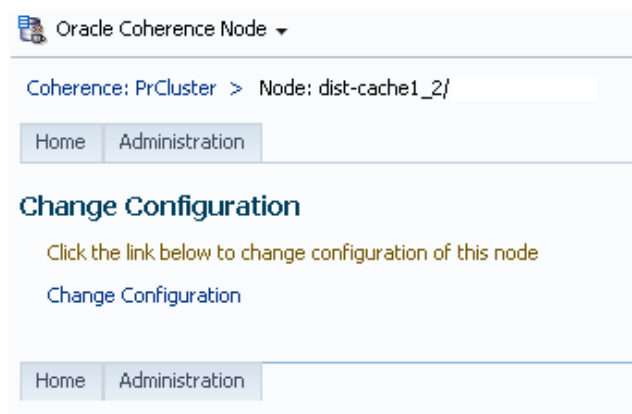
Enterprise Manager can monitor log files for the occurrence of user specified patterns to check for abnormal conditions. Matching patterns are specified as regular expressions. Log files are periodically scanned for occurrence of one or more patterns and an alert is raised when the pattern occurs during a given scan.

You can set up each Coherence node to log all messages into a log file on the host on which the node is running. You must use a specific naming pattern in the log file name

to ensure that it is monitored. To configure the log file monitoring criteria, follow these steps:

1. From the **Targets** menu, select **Middleware**, then click on a Coherence node target.
2. Navigate to the Administration page of the node for which the Log file alerts need to be configured.

Figure 20–9 Node Administration Page



3. Click the **Log Alert Setup** link. The Metric and Policy Settings page of the host on which the Coherence node is running is displayed.
4. Select **Metric with Thresholds** in the **View** drop-down box and search for the **Log File Pattern Matched Line Count** metric.
5. Click the pencil icon in this row to navigate to the Edit Advanced Settings: Log File Pattern Matched Line Count page.

Figure 20–10 Log File Pattern Matched Line Count Page



6. Click **Add** to add a row to the Monitored Objects table.
7. In the **Log File Name** field, specify the name in the format `<coherence_cluster_target_name>_<Node_name>_<...any other optional names>.log`. All node log file names must follow this format.
8. Specify the pattern to be matched for or ignored in the **Match Pattern in Perl** and **Ignore Pattern in Perl** fields. Lines matching the *ignore pattern* will be ignored first, then lines matching specified *match patterns* will result in one record being uploaded to the repository for each pattern.

9. Enter the Critical and Warning thresholds for this metric and click Continue. You will return to the Metrics and Collection Settings page. Click **OK** to update the metric.

20.5 Cache Data Management

The Cache Data Management feature allows you to define indexes and perform queries against currently cached data that meets a specified set of criteria.

Note: This feature is available to users with Administration privileges only if the Cache Data Management MBean has been registered in the Coherence JMX management node.

To perform cache data management operations, navigate to the Cache target Administration page and click **Go** in the Cache Data Management section. In the Cache Data Management page, you can select an operation and a query to perform a data management operation on the cache. You can perform the following operations:

- **Add Indexes:** To create an index, select the **Add Index** option in the Operation field. In the Value Extractor List field, specify a comma separated list of expressions that identify the index, and enter the Host Credentials. The Value Extractor is used to extract an attribute from a given object for indexing.
- **Remove Indexes:** To remove an index, select the **Remove Index** option in the Operation field and specify the Value Extractor List that identifies the index. Specify the Host Credentials and click Execute to remove the index from the cache.
- **Export:** You can export the queried data onto a file. Select a query from the Query section or click **Create** to create a new query. Select the **Export** option in the Operation field and enter the absolute path to the file. This file can be saved on the host machine on which the management node is running.
- **Import:** You can import queried data from a file. This file should be present on the host machine on which the management node is running. Select the **Import** option in the operation field and enter the absolute path to the file.
- **Insert:** Select the Insert option in the Operation field and specify an unique (key value) pair. This key value pair will be inserted into the cache and can be provided from:
 - UI Table on this Page: Select the Type of Keys and Type of Values and the Host Credentials.
 - Text File on Management Host: If the queries are stored in a text file, select this option and specify the location of the file.
 - Database Table: If the queries are stored in a database table, specify the Database URL, Credentials, the SQL Query Statement and Properties.
- **Purge:** Select **Purge** from the Operation drop-down list. Data matching the selected query will be deleted from the cache
- **View:** Select **View** from the Operation drop-down list and specify the number of key-value pairs to be displayed on each page. Data matching the criteria will be displayed.
- **Update:** Select **Update** from the Operation drop-down list. Specify the credentials for the host. Select a query from the Query table or create a new query to update the data in the cache.

- **View Explain Plan:** Select **View Explain Plan** from the Operation drop down list. Select a query that is to be evaluated from the Query table. See View Explain page for details.
- **View Trace:** Select **View Trace** from the Operation drop down list. Select a query that is to be evaluated from the Query table. See View Trace page for details.

20.5.1 View Explain Plan

A query explain record provides the estimated cost of evaluating a filter as part of a query operation. The cost takes into account whether or not an index can be used by a filter. The cost evaluation is used to determine the order in which filters are applied when a query is performed. Filters that use an index have the lowest cost and get applied first.

The View Explain Plan option allows you to estimate the cost of evaluating a filter as part of a query operation. When you select this option, a query record containing details of each step in the query is displayed. After viewing the details, click **Execute** to perform the selected operation or **Return** to return to the previous page.

Note: The View Explain Plan option can be used with Coherence version 3.7.1 or later.

20.5.2 View Trace

The View Trace option allows you to view the actual cost of evaluating a filter as part of a query operation. When you select this option, a query is executed in the cluster and a query record containing details of each step in the query is displayed. After viewing the details, click **Execute** to perform the selected operation or click **Return** to return to the previous page.

Note: The View Trace Plan option can be used with Coherence version 3.7.1 or later.

20.6 Reap Session Support

HTTPSession Objects for one or more Java EE applications deployed on Application Servers can be cached in the Coherence cluster. These HTTP sessions are cleaned by the Session Reaper and the associated memory is freed up. The Session Reaper is responsible for destroying any sessions that are no longer used, which is determined when the session has timed out. It is configured to scan the entire set of sessions over a certain period of time, called the reaping cycle. The Session Reaper scans for sessions that have expired, and when it finds expired sessions it cleans them up.

To view the reap session metrics, click the **Application** link in the Applications table in the Cluster Home page. The following reap session metrics and graphs are displayed in the Application Home page.

- **Reap Duration:** This graph shows the average reap duration in minutes.
- **Reaped Sessions:** This graph shows the average number of reaped sessions in a reap cycle.
- **Reaped Sessions (24 Hour Averages):** This table shows the average reap session data over the last 24 hours. The following details are displayed:

- Module Name: The name of the Coherence cluster on which the HTTP sessions have been cached.
- Node ID: When a node becomes part of a cluster, an ID is automatically assigned to the node. Click on the link to drill down to the Node Home page.
- Average Reap Duration: The average reap duration since the statistics were last reset.
- Average Reaped Sessions: The average number of reap sessions since the statistics were last reset.
- Total Reaped Sessions: The total number of expired sessions that have been reaped since the statistics were last reset.

20.7 Push Replication Pattern

The Push Replication Pattern provides extensible, flexible, high-performance, highly available, and scalable infrastructure to support the replication of `EntryOperations` occurring in one Coherence Cluster to one or more globally distributed Coherence clusters. The Push Replication Pattern advocates that:

- Operations (such as insert, update and delete) occurring on data in one Site should be pushed using one or more Publishers to an associated device.
- A Publisher is responsible for optimistically replicating operations (in the order in which the said Operations originally occurred) on or with the associated device.
- If a device is unavailable for some reason, the operations to be replicated using the associated Publisher will be queued and executed (in the original order) at a later point in time.

Implementation of the Push Replication Pattern additionally advocates that:

- The data on which operations occur are standard Coherence cache entries.
- Operations replicated include inserts, updates, and deletes of cache entries. This includes `NamedCache` put and remove, as well as updates that are artifacts of invoking `AbstractProcessor`.
- The operations that are enqueued for replication are called `EntryOperations`. They contain key and value pairs of updated cache entries in Binary form.
- A Site is a Coherence Cluster.
- A Site may act as both a sender of Operations and receiver of Operations. That is, multi-way multi-site push replication is permitted.
- A Device may be any of the following; a local cluster, a remote cluster, a file system, a database, an i/o stream, a logging system etc.

After the `EntryOperations` are captured, they are sent to the messaging layer, for distribution to the Subscribers. The Publishers then publish the currently queued `EntryOperations`.

The Publisher and Subscriber tables lists all the publishers and subscribers and displays the data used by them. You can Suspend, Drain, or Resume publishing operations on a specific Publisher.

20.8 Transactional Cache Support

Transactional caches are specialized distributed caches that provide transactional guarantees. Transactional caches are required whenever performing a transaction

using the Transaction Framework API. Transactional caches cannot interoperate with non-transactional caches.

The CacheMBean managed resource provides attributes and operations for all caches, including transactional caches. Many of the MBeans attributes are not applicable to transactional cache; invoking such attributes simply returns a -1 value. A cluster node may have zero or more instances of cache managed beans for transactional caches. The object name uses the form:

```
type=Cache, service=service name, name=cache name, nodeId=cluster node's id
```

The following list describes the CacheMBean attributes that are supported for transactional caches.

- **AverageGetMillis:** The average number of milliseconds per get() invocation.
- **AveragePutMillis:** The average number of milliseconds per put() invocation since the cache statistics were last reset.
- **HighUnits:** The limit of the cache size measured in units. The cache will prune itself automatically once it reaches its maximum unit level. This is often referred to as the high water mark of the cache.
- **Size:** The number of entries in the current data set.
- **TotalGets:** The total number of get() operations since the cache statistics were last reset.
- **TotalGetsMillis:** The total number of milliseconds spent on get() operations since the cache statistics were last reset.
- **TotalPuts:** The total number of put() operations since the cache statistics were last reset.
- **TotalPutsMillis:** The total number of milliseconds spent on put() operations since the cache statistics were last reset.

20.9 Integration with JVM Diagnostics

JVM Diagnostics allows administrators to identify the root cause of performance problems in the production environment without having to reproduce them in the test or development environment. JVM Diagnostics is a part of WLS Management Pack EE.

You can drill down to a Coherence node's JVM to identify the method or thread that is causing a delay. This feature allows you to trace live threads, identify resource contention related to locks, and trace the Java session to the database.

You can view the JVM Diagnostics data if the JVM Diagnostics Manager and JVM Diagnostics Agent have been deployed on the host machine on which the OMS is running. You can deploy the JVM Diagnostics Manager and JVM Diagnostics Agent on:

- Standalone Coherence
- CoherenceWeb

To setup JVM Diagnostics on each Coherence node you need to download the JVM Diagnostics Agent. For more details, see the Enterprise Manager Cloud Control Advanced Installation and Configuration Guide.

20.10 Viewing Performance Summary

You can use the Performance Summary page to monitor the performance of the cluster, a node, or a cache. From the **Oracle Coherence Cluster** (Node or Cache) menu, select **Monitoring**, then select **Performance Summary**.

A set of default performance charts that shows the values of specific performance metrics over time are displayed. You can customize these charts to help you isolate potential performance issues. You can also view a series of regions specific to the cluster, node, or cache target. For more details, see the Performance Summary Online Help.

20.11 Viewing Configuration Topology

The Configuration Topology Viewer provides a visual layout of the relationship of the Coherence target with other targets. From the **Oracle Coherence Cluster** (Node or Cache) menu, select **Configuration**, then select **Topology**. A topology graph for the Coherence target is displayed. You can determine the source of its health problem and its impact on other targets. You can also view the members of the cluster and its relationships. For more details, see the Configuration Topology Viewer Online Help.

20.12 Viewing Incidents

The Incident Manager shows incidents for a target and its members. When the Incident Manager is launched from Coherence Cluster target, incidents for Cluster, Node and Cache targets in cluster are displayed. Similarly, when the Incident Manager is launched in the context of Node target, incidents for the Node target and for all Cache targets that are deployed on the node are displayed. When Incident Manager is launched from the Cache target, incidents for that target are displayed.

You can launch the Incident Manager by clicking on the number of Incidents in the General section for Coherence Cluster, Node and Cache targets. Alternatively, from the **Oracle Coherence Cluster** (Node or Cache) menu, select **Monitoring**, then select **Incident Manager** to navigate to the Incident Manager page.

20.13 Troubleshooting Coherence

- **Collecting Metric Data:** If you cannot collect metric data for any of the Coherence targets, check the following to ensure that the steps involved in discovering the target have been followed correctly.
 - Make sure that the management node has been successfully started and the host on which the management node is running is accessible from the Agent host.
 - Specify the appropriate User Name and Password if password authentication is enabled.
 - If you are not using SSL to start the management node, make sure that you have started the JVM using the `com.sun.management.jmxremote.ssl=false` option.
 - If you did not use the bulk operation MBean JAR to start the management node, you must leave the Bulk Operations Mbean field blank during discovery.

- **Node Performance Issues:** If the node is showing down, check if the node was restarted, if yes, refresh the cluster target. If nodes are restarted often then use Refresh job as described in [Section 19.3.1, "Restarting a Node"](#).
- **Dynamic Client Nodes:** If there are dynamic client nodes that are not running all the time, these nodes can be removed from the cluster and proxy service can be used.
- **Target Proliferation of Nodes:** If there is a target proliferation of nodes, this may be due to NULL or duplicate `tangosol.coherence.member` value for the nodes. Verify that each node has a nonNull and unique value for the `tangosol.coherence.member` property.

20.14 Best Practices

This section describes some of the best practices that can be used while setting up and using Oracle Coherence. It covers the following:

- [Mandatory Patches](#)
- [Critical Metrics](#)
- [Refresh Cluster](#)
- [Custom Metrics](#)
- [Monitoring Templates](#)
- [Sample Code Snippets](#)

20.14.1 Mandatory Patches

Ensure that the Coherence Management specific patches (See **Support Note 1526316.1**) have been applied:

20.14.2 Critical Metrics

This section lists some of the key metrics that indicate the overall health of the cluster.

- **Hosts:** The host performance indicates the ability of the nodes running on that host to perform well. The top performance indicators are:
 - CPU Utilization
 - Memory Utilization
- **Caches:** Cache performance can be monitored at the aggregate level (combining all nodes on which the cache is running) or at the individual node level. However, it is useful to set **Alerts** at the node level for cache metrics. This allows you to identify performance degradation of cache even if it happens only one node. The top performance indicators are:
 - Evictions - Aggregate Storage Manager Metrics > Aggregate Evictions Per Minute
 - Number of objects - Aggregate Cache Size Metrics > Aggregate Cache Objects
 - Number of Gets/Puts/Misses - Aggregate Cache Usage Metrics > Aggregate Get/Puts/Misses Per Minute
 - Hit to Gets Ratio% - Aggregate Cache Usage Metrics > Hit to Gets Ratio (%)

- Memory consumed by cache (for Binary operator) - Aggregate Cache Size Metrics > Aggregate Memory Usage in MB
- Queue size (for Write Behind caches) - Aggregate Cache Usage Metrics > Aggregate Queue Size
- Number of non-optimized queries (for applications using CohQL) - Aggregate Query Optimization Metrics > Aggregate Non Optimized Queries Executed Per Minute
- Non Optimized Query Time (ms) (for applications using CohQL) - Aggregate Query Optimization Metrics > Aggregate Non Optimized Query Time (ms) Per Minute
- Nodes: The top performance indicators are:
 - Heap utilization after GC - Garbage Collector After GC > Memory Used After GC (bytes)
 - Packets send/receive success rate - Network Health Details > Send/Receive Success Rate Per Minute (%)
 - Task backlog - Service Requests and Tasks > Total Tasks Backlog
 - Service Status - Service Requests and Tasks > High Availability (HA) Status
 - Endangered Partitions - Service Partitions and Threads > Partitions Endangered
- Cluster: The top performance indicators are:
 - Task backlog - Aggregate Service Requests and Tasks > Aggregate Task Backlog

20.14.3 Refresh Cluster

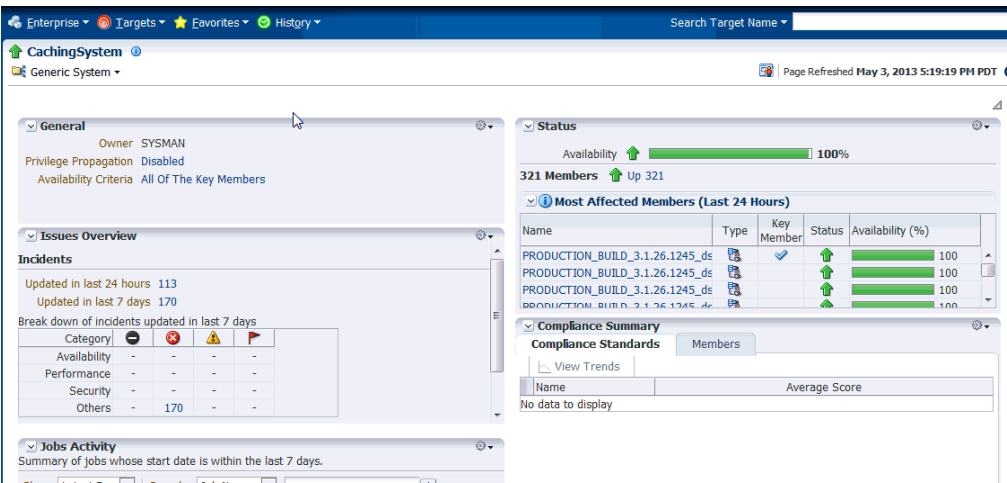
If the nodes are restarted or need to rejoin the cluster for some network related issues then you will have to **Refresh** the cluster target. Alternatively, you can set up a refresh job as described in [Section 19.3.1, "Restarting a Node"](#).

20.14.4 Custom Metrics

If you add custom metrics using Metric Extension at node level, you may need to add aggregate metrics at the cluster level. To create aggregate metric for cluster level monitoring from a node metric, follow these steps:

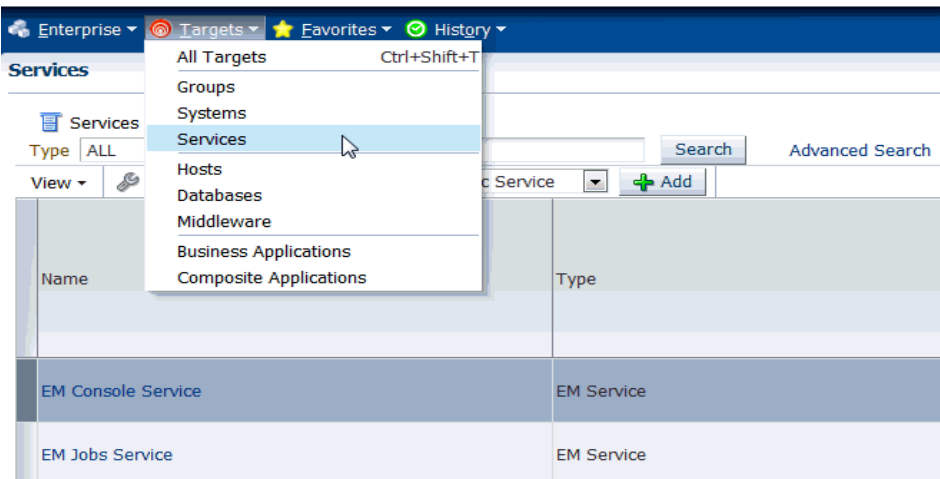
1. Create a Generic System target with all node targets. You can add hosts, caches, and cluster to the target as required.

Figure 20–11 Generic System Page



- 2. Create a Generic Service and associate it with the system created in step 1. From the **Targets** menu, select **Services**.

Figure 20–12 Generic Service Page



- 3. In the Services page, select **Generic Service** from the **Create** menu.

Figure 20–13 Create Generic Service: General

Create Generic Service: General

Define a service to model and monitor a business process or application.

* Name

Time Zone
Select the time zone for this service. Monitored data will be displayed using the selected time zone.

System

Select a system target that hosts this service

System CachingSystem

Time Zone (UTC-08:00) US Pacific Time

4. Define availability based on system rather than service.

Figure 20–14 Create Generic Service: Availability

Create Generic Service: Availability

You can define the service's availability based on either the execution of the service test by the key beacons or the availability of key components of the service or the availability of the system.

Define availability based on:

☒ Availability of System ☐ Availability of selected components of the system

Step 2 of 7

Tip

Select the method of availability testing that best models your business process or application.

Oracle recommends using:

- **Service Test** to define availability for most cases. A service test emulates the way clients access the service.
- **System** when a suitable service test is unavailable.

5. Add a Performance Metric. Make sure you select the right metric from the target(s) you are interested in. In this example, the Node Memory Used (MB) is added. So the service will have Aggregate Node Memory Used (MB).

Figure 20–15 Create Generic Service: Performance Metrics

GeneralAvailabilityService TestBeaconsPerformance MetricsUsage MetricsMore

Add Performance Metric based on System

Define a performance metric by selecting a metric from one or more system components. For multiple components, choose an aggregation function. The name of the pe

System NameCachingSystem

Target TypeOracle Coherence Node

MetricNode Memory Used (MB)

System Components

Use the metric from a specific componentPRODUCTION_BUILD_3.1.26.1245_dslprod/DSL_DTM_

Aggregate the metric across multiple components

Aggregate FunctionSumMinimumMaximumAverageSum

ShowKey Components

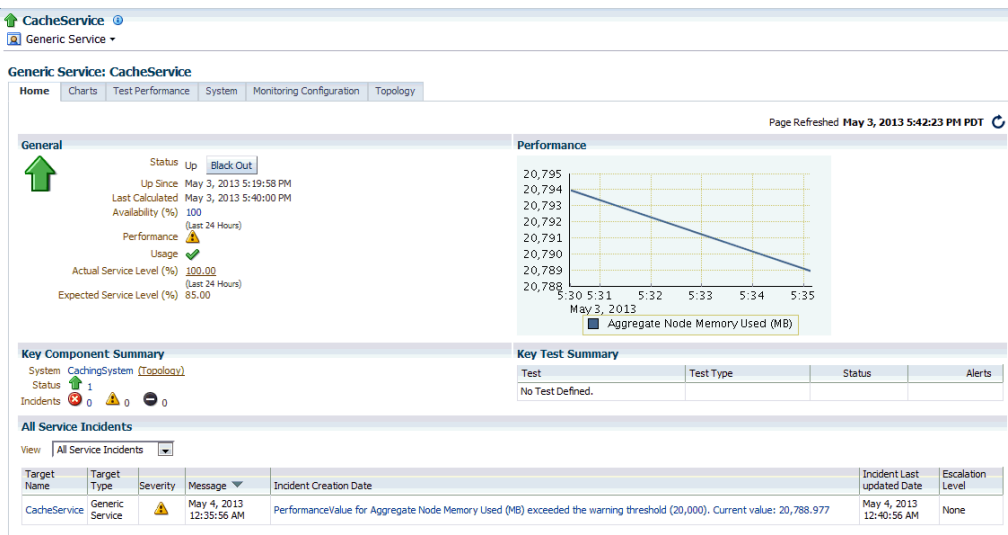
Previous1-10 of 306Next 10

Select AllSelect None

Select	Target
<input checked="" type="checkbox"/>	PRODUCTION_BUILD_3.1.26.1245_dslprod/DSL_DTM_
<input checked="" type="checkbox"/>	PRODUCTION_BUILD_3.1.26.1245_dslprod/DSL_DTM_
<input checked="" type="checkbox"/>	PRODUCTION_BUILD_3.1.26.1245_dslprod/DSL_DTM_
<input checked="" type="checkbox"/>	PRODUCTION_BUILD_3.1.26.1245_dslprod/DSL_DTM_

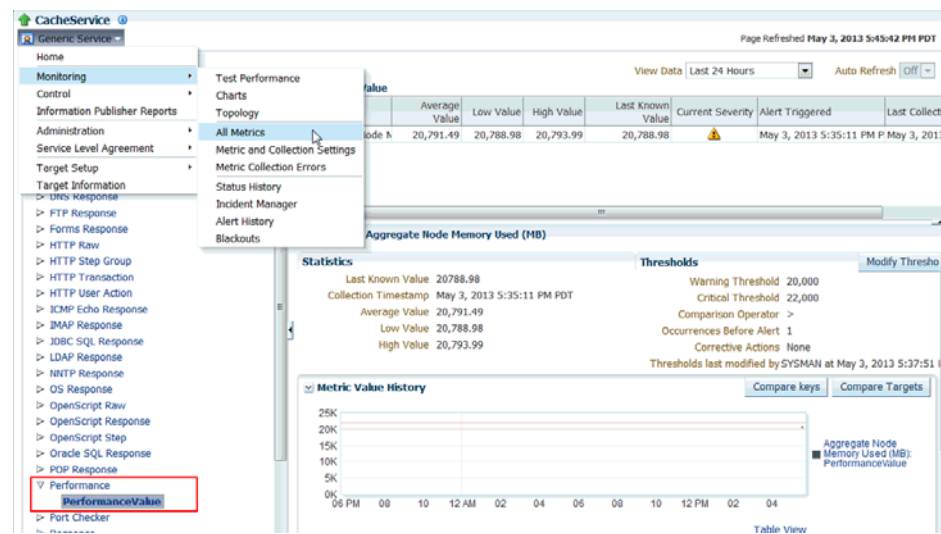
6. The new metric is displayed on the Generic Service Home page.

Figure 20–16 Generic Service Home Page



7. Navigate to the All Metrics page of the service target to view the new performance metric.

Figure 20-17 All Metrics Page



20.14.5 Monitoring Templates

Monitoring templates for each of the Coherence Cluster, Node, and Cache targets are available out-of-the-box. These templates can be used as default monitoring templates for all Coherence targets. Based on specific requirements, you can enable, disable certain metrics or change the collection frequency.

Note: The threshold values provided in the templates are examples and must be changed.

20.14.6 Sample Code Snippets

In this section, you can see some sample code snippets that show you how to use the JVM Diagnostics to start a management node or a storage enabled node.

You must ensure that:

- The Management Agent has been deployed on every machine on which the Coherence node is running.
- Garbage Collection has been enabled on the nodes.
- Heap utilization on the node does not exceed 70%.
- For new nodes and caches or nodes and caches rejoining the cluster, you must set up the Refresh job as described in [Section 19.3.1, "Restarting a Node"](#).
- The caches must be preconfigured on the node to ensure that Enterprise Manager can easily establish associations during discovery.
- Reusable monitoring templates that can be applied on multiple targets must have been created.
- The node has been refreshed after it has been restarted. To refresh the node, from the **Coherence Cluster** menu, select **Configuration**, then select **Last Collected**. In the Latest Configuration page, from the **Actions** menu, select **Refresh** to refresh the node.

Example 20–1 Starting a Management Node

```
JVMD_HOST=<Host_Where_JVMD_Manager_Is_Running>
JVMD_PORT=<Port_of_the_JVMD_Manager>
JVMD_POOL=Contacts
JAM_PARAMS=" jamagent.jamrun jamconshost=$JVMD_HOST jamconspport=$JVMD_PORT
jampool=$JVMD_POOL"
CP=$CP:<path_where_JAMAgent_is_copied>/jamagent.war:
<OEM_Agent_Home>/plugins/oracle.sysman.emas.agent.plugin_
12.1.0.4.0/archives/coherence/coherenceEMIntg.jar:
<OEM_Agent_Home>/plugins/oracle.sysman.emas.agent.plugin_
12.1.0.4.0/archives/coherence/bulkoperationsmbean.jar
CACHE_CONFIG=$CONFIG_DIR/$CACHE_CONFIG_FILE_NAME
POF_CONFIG=$CONFIG_DIR/$POF_CONFIG_FILE_NAME
COH_OPTS="$COH_OPTS -cp $CP"
COH_OPTS="$COH_OPTS -Dtangosol.coherence.cacheconfig=$CACHE_CONFIG"
COH_OPTS="$COH_OPTS -Dtangosol.pof.config=$POF_CONFIG"
# using non-default port to prevent accidentally joining other clusters
COH_OPTS="$COH_OPTS -Dtangosol.coherence.clusterport=<open_port>"
$JAVA_HOME/bin/java $COH_OPTS
-Dtangosol.coherence.management.remote=true
-Dcom.sun.management.jmxremote.authenticate=false
-Dtangosol.coherence.management=all
-Dcom.sun.management.jmxremote.ssl=false
-Dtangosol.coherence.cluster=Contacts
-Dtangosol.coherence.member=MNode1
-Dtangosol.coherence.site=MSite
-Dtangosol.coherence.rack=MRack
-Dtangosol.coherence.machine=<FullyQualifiedHostName_of_localhost>
-Doracle.coherence.machine=<FullyQualifiedHostName_of_localhost>
-Dcom.sun.management.jmxremote.port=<OpenTCP_Port>
-Dtangosol.coherence.distributed.localstorage=false
-Doracle.coherence.jamjvmid=Contacts/MNode1
-Dtangosol.coherence.management.refresh.expiry=1m
-server
-Xms2048m -Xmx2048m
$JAM_PARAMS
jamjvmid=Contacts/MNode1
oracle.sysman.integration.coherence.EMIntegrationServer $*
```

where:

- The string specified in the JVMD_POOL and -Dtangosol.coherence.cluster should match. This is used to correlate the Coherence Cluster target and the JVM Pool.
- The Dtangosol.coherence.member name must be unique. This name is required for JVM Diagnostics integration.
- The JVM Diagnostics integration parameters are:
 - -Doracle.coherence.jamjvmid - <JVM_pool_name>/<member_name>
 - Jamjvmid - <JVM_pool_name>/<member_name>

Note: The location of the .jar files may change depending on the plug-in version.

Example 20–2 Starting a Storage Node

```
JVMD_HOST=<Host_Where_JVMD_Manager_Is_Running>
JVMD_PORT=<Port_of_the_JVMD_Manager>
```



```

JVMD_POOL=Contacts
JAM_PARAMS=" jamagent.jamrun jamconshost=$JVMD_HOST jamconspport=$JVMD_PORT
jampool=$JVMD_POOL"
CP=$CP:<path_where_JAMAgent_is_copied>/jamagent.war
CACHE_CONFIG=$CONFIG_DIR/$CACHE_CONFIG_FILE_NAME
POF_CONFIG=$CONFIG_DIR/$POF_CONFIG_FILE_NAME
COH_OPTS="$COH_OPTS -cp $CP"
COH_OPTS="$COH_OPTS -Dtangosol.coherence.cacheconfig=$CACHE_CONFIG"
COH_OPTS="$COH_OPTS -Dtangosol.pof.config=$POF_CONFIG"

# using non-default port to prevent accidentally joining other clusters
COH_OPTS="$COH_OPTS -Dtangosol.coherence.clusterport=<cluster_port>"
$JAVA_HOME/bin/java $COH_OPTS
-Dtangosol.coherence.cluster=Contacts
-Dtangosol.coherence.member=SNode1
-Dtangosol.coherence.management.remote=true
-Dcom.sun.management.jmxremote.authenticate=false
-Dcom.sun.management.jmxremote.ssl=false
-Dtangosol.coherence.machine=<FullyQualifiedHostName_of_localhost>
-Doracle.coherence.machine=<FullyQualifiedHostName_of_localhost>
-Dtangosol.coherence.log.level=5
-Dtangosol.coherence.log=<path_log_file>
-Doracle.coherence.jamjvmid=Contacts/SNode1
-server -Xms2048m -Xmx2048m
$JAM_PARAMS
jamjvmid=Contacts/SNode1
com.tangosol.net.DefaultCacheServer $*

```


Part VIII

Using Identity Management

The chapters in this part provide a brief introduction to the Management Pack Plus for Identity Management. It guides you through the process of discovering and configuring Oracle Identity Management targets and discusses key features in the Management Pack Plus for Identity Management.

The chapters are:

- [Chapter 21, "Getting Started with Identity Management"](#)
- [Chapter 22, "Prerequisites for Discovering Identity Management Targets"](#)
- [Chapter 23, "Discovering and Configuring Identity Management Targets"](#)

Getting Started with Identity Management

As more and more businesses rely on the Oracle Identity and Access Management Suite to control access to their mission-critical applications (both packaged applications and custom-built web applications) and to provision resources across their organizations, the need to achieve predictable performance and availability for Oracle Identity Management systems has become a top priority for many businesses. An outage or slow performance in access and identity services, for instance, can have negative impacts on the business bottom-line as end-users are unable to log in to mission-critical applications.

To help you maximize the value of Oracle Identity Management systems, and to deliver a superior ownership experience while restraining the systems management costs, Oracle provides Oracle Management Pack Plus for Identity Management (the Identity Management Pack), which leverages the Oracle Enterprise Manager Cloud Control advanced management capabilities, to provide an integrated and top-down solution for your Oracle Identity Management environment.

21.1 Benefits of the Using Identity Management Pack

The benefits of using Identity Management Pack include:

- A centralized systems management solution to efficiently manage multiple Oracle Identity Management deployments including testing, staging, and production environments from a single console
- Gain the ability to monitor a wide range of performance metrics for all critical Identity Management components to find root causes of problems that could potentially slow performance or create outages
- Automated configuration management to accelerate problem resolution
- Record synthetic Web transactions (or service tests) to monitor Identity Management Service availability and analyze end user response times
- Define Service Level Objectives (SLO's) in terms of out-of-box system-level metrics, as well as end user experience metrics to accurately monitor and report on Service Level Agreement (SLA) compliance

21.2 Features of the Identity Management Pack

The features in the Identity Management Pack include:

- Enterprise-Wide View of Oracle Identity Management

- The "Identity and Access" dashboard provides a centralized view of all Oracle Identity Management components - including Identity Management 10g and Identity Management 11g components.
- From the "Identity and Access" dashboard, users can view the performance summary of the associated systems and services based on the underlying dependencies and monitor the overall health of the Identity Management environment.
- Performance Management
 - A wide range of out-of-box performance metrics to find root causes of problems that could potentially slow performance, extend response times, or create outages
 - Customizable performance summaries with a "Metric Palette" that allows users to drag and drop performance charts
- Configuration Management
 - Perform key configuration management tasks like keeping track of configuration changes for diagnostic and regulatory purposes, taking snapshots to store configurations, and comparing component configurations to ensure consistency of configurations within the same environment or across different environments.

21.2.1 New Features for this Release

New features for Identity Management Pack include:

- Performance Management
 - Out-of-box reports for Oracle Internet Directory, Oracle Access Manager, and Oracle Identity Manager
 - Oracle Identity Manager database performance page to analyze the performance of the underlying Oracle Identity Manager database in the context of the OIM-specific tables and user. **Note:** The database target will need to be discovered to take advantage of all the features on the database performance page.
- Configuration Management

Automated compliance monitoring and change detection for Oracle Identity Manager is now available to help customers meet compliance and reporting requirements.

To enable the compliance standard association with the Oracle Identity Manager Cluster target. Perform the following steps:

 1. Click the Oracle Identity Manager Cluster target. From the **Target** menu, select **Compliance**, then select **Standard Associations**.
 2. Click **Edit Association Settings**. Click **Add** and then select **Oracle Identity Manager Cluster Configuration Compliance**.
 3. Click **OK** and then **OK** again to enable the new association setting.
- Monitoring Support

As part of the Oracle Access Management Suite, added monitoring support for the Oracle Mobile and Social, Identity Federation. This includes Up and Down status of Mobile and Social service along with the collection of the select Mobile and Social metrics.

21.3 Monitoring Oracle Identity Management Components in Enterprise Manager

You can use Enterprise Manager to monitor the following Identity Management 11g components ([Table 21–1](#)).

Table 21–1 Licensed Targets for Identity Management 11g Targets

Enterprise Manager Target Type	Purpose
Oracle Adaptive Access Manager Oracle Access Manager Oracle Directory Integration Platform Oracle Identity Federation Oracle Identity Manager Oracle Internet Directory Oracle Virtual Directory	<p>Each component will be presented as a target in Enterprise Manager which provides an interface with access to target overview, customizable performance summary, process control, configuration management, compliance analysis, and Information Publisher reports.</p> <p>For all the Oracle Adaptive Access Managers, Oracle Access Managers, and Oracle Identity Managers that are deployed within the same WebLogic domain, a cluster target will be created for each component:</p> <ul style="list-style-type: none"> ■ Oracle Adaptive Access Manager Cluster ■ Oracle Access Manager Cluster ■ Oracle Identity Manager Cluster <p>Each cluster target is a logically related group of components that are managed as a unit.</p> <p>Every target is part of a WebLogic domain.</p>
Oracle Directory Server Enterprise Edition	<p>The following types of targets will be created for each Oracle Directory Server Enterprise Edition deployment:</p> <ul style="list-style-type: none"> ■ Oracle Directory Server Enterprise Edition Server A target represents the LDAP service and all internal resources ■ Directory Server Group User logical grouping of Oracle Directory Server Enterprise Edition Servers ■ Directory Server Enterprise A set of Oracle Directory Server Enterprise Edition Servers connected through a network that participates in the service, including Directory Server Groups. <p>Each target provides an interface in Enterprise Manager with access to target overview, customizable performance summary, process control, and configuration management.</p>

The following Identity Management 10g components can be monitored by Enterprise Manager ([Table 21–2](#)).

Table 21–2 Licensed Targets for Identity Management 10g Targets

Enterprise Manager Target Type	Purpose
Oracle Delegated Administration Server Oracle Directory Integration Platform Oracle Internet Directory Oracle Single Sign-On	<p>Each component will be presented as a target in Enterprise Manager which provides an interface with access to target overview and performance summary</p>

Table 21–2 (Cont.) Licensed Targets for Identity Management 10g Targets

Enterprise Manager Target Type	Purpose
Oracle Access Manager - Access Server Oracle Access Manager - Identity Server Oracle Identity Federation	<p>Each component will be presented as a target in Enterprise Manager which provides an interface with access to target overview and performance summary.</p> <p>A system target will be created for each component to provide end-to-end system oriented view of the component:</p> <ul style="list-style-type: none"> Access Manager - Access System Access Manager - Identity System Identity Federation System <p>The underlying LDAP servers, database instances and hosts will be monitored within the system.</p>
Oracle Identity Manager	<p>The following types of targets will be created for each Oracle Identity Manager:</p> <ul style="list-style-type: none"> Identity Manager Server A target represents the server tier of Oracle Identity Manager Identity Manager Repository A target represents the data and enterprise integration tier of Oracle Identity Manager <p>A system target will be created for Oracle Identity Manager to provide an end-to-end system oriented view of the component.</p> <ul style="list-style-type: none"> Identity Manager System <p>The underlying LDAP servers, database instances, and hosts will be monitored within the system.</p>

The monitored targets in the Identity Management pack associated with both release 10g and release 11g are summarized in [Table 21–3](#).

Table 21–3 Targets Associated with Both Identity Management 10g and Identity Management 11g Targets

Enterprise Manager Target Type	Purpose
Generic Service	<p>With the Management Pack Plus for Identity Management, users can create targets of type Generic Service associated with any of the monitored Identity Management Systems: Access Manager - Access System, Access Manager - Identity System, Identity Federation System, Identity Manager System, and Identity and Access System. The Generic Service target provides an end-to-end service oriented view of the monitored Oracle Identity Management targets with access to performance and usage metrics, service tests, service level rules, service availability definition, alerts, charts, and topology view.</p>
Host	<p>Representation of hosts running Oracle Identity Management components providing access to metrics, alerts, performance charts, remote file editor, log file alerts, user-defined metrics, host commands and customized reports.</p>

Table 21–3 (Cont.) Targets Associated with Both Identity Management 10g and Identity Management 11g

Enterprise Manager Target Type	Purpose
Oracle Database	Representation of Oracle Database that is used by Oracle Identity Management components providing access to metrics, alerts, performance charts, compliance summary, and configuration management.
Oracle Identity and Access System	System target that can be modeled with any discovered Oracle Identity Management target (including both Identity Management 10g and Identity Management 11g targets) and the underlying hosts and databases as the key components providing an end-to-end system oriented view of the monitored Identity Management environment. The Identity and Access System target provides access to member status, metrics, charts, incidents, and topology view.
Oracle SOA Suite	Representation of Oracle SOA Suite that is used by Oracle Identity Manager 11g providing access to metrics, alerts, performance charts, and configuration management of the SOA infrastructure instance and its service engines.

Prerequisites for Discovering Identity Management Targets

This chapter lists the system requirements and prerequisites needed to discover identity management targets.

22.1 System Requirements

Table 22–1 lists the supported Oracle Identity Management products in the Management Pack Plus for Identity Management in Enterprise Manager Cloud Control 12c Release 2 (12.1.0.2). Please check the Enterprise Manager certification matrix on My Oracle Support (<http://support.oracle.com>) for the most up-to-date list of supported platforms.

Table 22–1 Supported Identity Management Products and Platforms in Enterprise Manager Cloud Control Release 2 (11.1.0.2)

Product	Version	Application Server	Directory Server/Database
Oracle Access Manager	10.1.4.2; 10.1.4.3.0	Not Applicable	Oracle Internet Directory 10.1.4.x; Microsoft Active Directory
Oracle Access Manager	11.1.1.3.0; 11.1.1.5.0	Oracle WebLogic Server 10.3	Oracle Database
Oracle Adaptive Access Manager	11.1.1.3.0; 11.1.1.5.0	Oracle WebLogic Server 10.3	Oracle Database
Oracle Directory Integration Platform	11.1.1.2.0; 11.1.1.5.0	Oracle WebLogic Server 10.3	Oracle Database
Oracle Directory Server Enterprise Edition	6.x; 7.x; 11.1.1.3.0; 11.1.1.5.0	Not Applicable	Not Applicable
Oracle Identity Federation	10.1.4.2; 10.1.4.3.0	Oracle Application Server	Oracle Internet Directory 10.1.4.x
Oracle Identity Federation	11.1.1.2.0; 11.1.1.5.0	Oracle WebLogic Server 10.3	Oracle Internet Directory 11.1.1.2.0; 11.1.1.3.0; 11.1.1.6
Oracle Identity Management Suite - Delegated Administration Services	10.1.4.2; 10.1.4.3.0	Oracle Application Server 10g	Oracle Database
Oracle Identity Management Suite - Directory Integration Platform	10.1.4.2; 10.1.4.3.0	Oracle Application Server 10g	Oracle Database

Table 22–1 (Cont.) Supported Identity Management Products and Platforms in Enterprise Manager Cloud Control Release 2 (11.1.0.2)

Product	Version	Application Server	Directory Server/Database
Oracle Identity Management Suite - Oracle Internet Directory	10.1.4.2; 10.1.4.3.0	Oracle Application Server 10g	Oracle Database
Oracle Identity Management Suite - Single Sign-On Server	10.1.4.2; 10.1.4.3.0	Oracle Application Server 10g	Oracle Database
Oracle Identity Manager	9.1.0.1	Oracle WebLogic Server 10.3; JBoss Application Server	Oracle Database
Oracle Identity Manager	11.1.1.3.0; 11.1.1.5.0	Oracle WebLogic Server 10.3; Oracle SOA Suite 11.1.1.3.0	Oracle Database
Oracle Internet Directory	11.1.1.2.0; 11.1.1.3.0; 11.1.1.4.0; 11.1.1.5.0; 11.1.1.6.0	Oracle WebLogic Server 10.3	Oracle Database
Oracle Virtual Directory	11.1.1.2.0; 11.1.1.3.0; 11.1.1.4.0; 11.1.1.5.0; 11.1.1.6.0	Oracle WebLogic Server 10.3	Not Applicable

22.2 Installing Oracle Enterprise Manager Cloud Control 12c Release 2

Before you begin configuring Cloud Control 12c Release 2 to manage your Identity Management components, you must install and configure Cloud Control 12c Release 2 on at least one host computer on your network. Oracle recommends that you install Cloud Control on dedicated host(s).

For example, if the Identity Management components are installed on host1.us.oracle.com, then install and configure the Oracle Management Service and Oracle Management Repository on host2.us.oracle.com. Install the Cloud Control 12c Management Agent on every host that includes the components you want to manage with Cloud Control.

See Also:

Oracle Enterprise Manager Cloud Control Basic Installation Guide

All documentation files can be accessed on the Oracle OTN website:

http://docs.oracle.com/cd/E24628_01/nav/portal_booklist.htm

22.3 Prerequisites for Discovering Identity Management Targets in Enterprise Manager

Before you start monitoring Oracle Identity Management targets in Enterprise Manager, you must perform the following tasks:

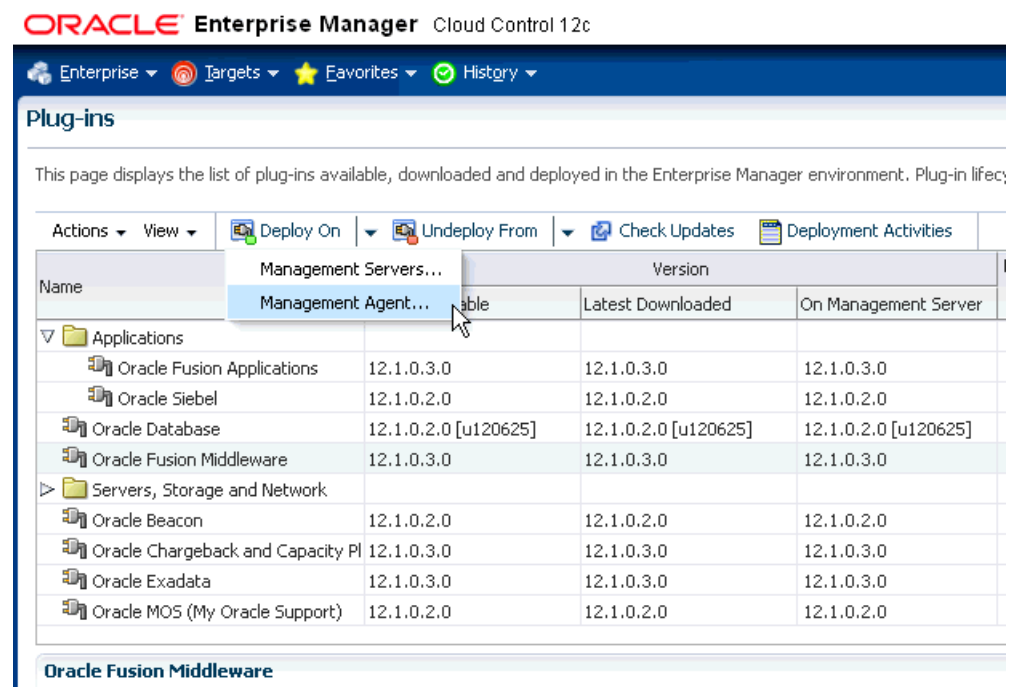
- Install Cloud Control 12c Agent on each of the hosts that run Oracle Identity Management components.

If you would like to monitor additional targets, such as Oracle Application Server, Oracle WebLogic Server, JBoss Application Server, MS Active Directory, MS IIS and databases supporting Oracle Identity Management, and you have the proper

license for monitoring these targets, then install Cloud Control 12c Management Agent on these hosts as well.

- Deploy the "Oracle Fusion Middleware" plug-in on the agents running on the hosts for Oracle Identity Management.
 1. Log in to Enterprise Manager. Navigate to **Setup**, select **Extensibility**, then select **Plugins**.
 2. Select Oracle Fusion Middleware plug-in and ensure that it has been deployed on the agents running on the hosts for Oracle Identity Management. See [Figure 22-1](#).

Figure 22-1 Plug-Ins Deploy On Options



- After Enterprise Manager Cloud Control OMS and Management Agents are installed, complete the following steps before initiating the discovery process:

Oracle Access Manager 10.1.4.2, 10.1.4.3.0

1. Install Oracle Access Manager SNMP Agent on each of the hosts where the Oracle Access Manager Access Server and Identity Server are running. The SNMP Agent collects performance metrics and configuration parameters for the Oracle Access Manager Access Server and Identity Server allowing you to monitor the various Oracle Access Manager components through Enterprise Manager Cloud Control. Refer to the *Oracle Access Manager Installation Guide* for instructions on installing the SNMP Agent (http://download.oracle.com/docs/cd/B28196_01/idmanage.1014/b25353/snmp.htm#CHDFBJJC).
2. Configure the SNMP Agent and specify the Management Agent's UDP and TCP Ports as well as the SNMP Agent Community Name. Make sure that you record the SNMP Agent UDP Port and Community Name; because these details will be needed in the discovery process. Refer to the *Oracle Access Manager Installation Guide* for instructions on configuring the SNMP Agent

(http://download.oracle.com/docs/cd/B28196_01/idmanage.1014/b25353/snmp.htm#CEGEIIFI).

Also, refer to the *Oracle Access Manager Identity and Common Administration Guide* for instructions on setting up the SNMP Agent

(http://download.oracle.com/docs/cd/B28196_01/idmanage.1014/b25343/snmpmntr.htm#CEGHHDBC).

3. Enable SNMP monitoring for both the Oracle Access Manager Access Server and Oracle Access Manager Identity Server by completing the following tasks:

- From the Identity (or Access) System Console, select System Configuration, Identity Server (or Access Server).
- Click a link for a particular server.
- Click **Modify** to display the page where you can turn SNMP monitoring on or off. Click the **SNMP State On** button at the bottom of the page to turn on the collection of SNMP statistics.
- In the SNMP Agent Registration Port field, enter the **port number** to define or change the port on which the SNMP Agent listens.
- Restart the Identity Server (or Access Server).

Refer the *Oracle Access Manager Identity and Common Administration Guide* for instructions on setting up the SNMP Agent

(http://download.oracle.com/docs/cd/B28196_01/idmanage.1014/b25343/snmpmntr.htm#BABFFDDA).

4. Complete all the configuration steps for the Oracle Access Manager Identity Server and Oracle Access Manager Access Server. Ensure that the communication details and the directory server details are defined so that Enterprise Manager can discover the topology of your Oracle Access Manager environment.

Refer to the *Oracle Access Manager Installation Guide* for instructions on configuring the Identity Server

(http://download.oracle.com/docs/cd/B28196_01/idmanage.1014/b25353/id_setup.htm#CHDHIBIB) and the Access Server

(http://download.oracle.com/docs/cd/B28196_01/idmanage.1014/b25353/a_srvr.htm#BGBEFBBD).

5. If you plan to monitor the directory server through Oracle Enterprise Manager Cloud Control, then ensure that the directory server is appropriately discovered in Enterprise Manager before moving on to the discovery of Oracle Access Manager Identity Server and Oracle Access Manager Access Server. Complete the following tasks to discover the supported directory servers:

- Oracle Internet Directory 10.1.4

Discovery of Oracle Identity Management Suite 10g (including Oracle Internet Directory, Directory Integration Platform, Delegated Administration Server, and Single Sign-On Server) can be done using the discovery wizard on the Middleware page. From the Middleware page, select **Oracle Application Server** from the **Add** menu. For more information, refer to the Discovering Oracle Identity Management Suite 10.1.4.2, 10.1.4.3.0 section.

- Self Update

Use the Self Update option in Enterprise Manager to get the plug-in.

- 1) From the **Setup** menu, select **Extensibility**, then select **Self Update**.
- 2) Click **Plug-in** and select the available plug-in for Microsoft Active Directory or other non-Oracle products that need to be monitored. See [Figure 22–2](#).

Figure 22–2 Self Update Screen

Self Update Page Refreshed **Aug 2, 2012 11:03:17 AM**

Oracle periodically provides new functionality and updates for existing features in Enterprise Manager. The Self Update home allows administrators to receive notifications and view, download, and apply updates. While these updates are retrieved automatically, a manual check can be made at any time.

Status Information

Connection Mode: **Offline** Last Download Time: May 22, 2012 3:35:50 AM EDT Last Apply Time: N/A
 Most Recent Refresh Time: **Aug 1, 2012 11:20:31 PM EDT** Last Download Type: Plug-in Last Apply Type: N/A
 Last Successful Refresh Time: N/A

Actions: Open Check Updates Plug-in

Type	Available Updates	Downloaded Updates	Applied Updates	Description
Agent Software	0	0	1	Agent software has to be installed on hosts for managing the host.
Compliance Content	0	0	0	Compliance Content contains Framework, Standard, Rules with support of add a delete on these entities.
Diagnostic Checks	0	0	0	Target side policy checks that identify conditions that may require the attention of target administrators.
EM Deployment Prerequisite Resources	0	0	0	EM Deployment Pre-requisite Checks are the metadata used for checking prerequisites for Install, Upgrade, Patching of EM Platform and Plugins.
Management Connector	0	0	1	Management Connectors are components that integrate different enterprise frameworks into the Enterprise Manager Console
Middleware Profiles and Gold Images	0	0	0	A collection of Software Components used for provisioning of Oracle Application Server homes.
Oracle Database Provisioning Profiles	0	0	0	A collection of Software Components used for provisioning of Oracle Database, Clusterware and Grid Infrastructure homes.
Plug-in	0	12	12	Plug-in extends Enterprise Manager to manage newer target type as well as to extend vertical functionality
Provisioning Bundle	0	0	1	Provisioning bundle is a collection of deployment procedures, software library entities, and other related artifacts that cater to the provisioning and patching of

Oracle Identity Federation 10.1.4.2, 10.1.4.3.0

1. Complete all the configuration steps for the Oracle Identity Federation. Ensure that the Federation Data Store details and User Data Store details are defined so that Enterprise Manager can discover the topology of your Oracle Identity Federation environment.

Refer to the *Oracle Identity Federation Administrator's Guide* for instructions on configuring the Identity Federation

(http://download.oracle.com/docs/cd/B28196_01/idmanage.1014/b25355/configuring.htm#BCGDGAJ).

2. Discover the Oracle Application Server on which Oracle Identity Federation is deployed in Enterprise Manager Cloud Control. Complete the following steps to discover Oracle Application Server in Cloud Control:
 - Log in to Enterprise Manager. Select **Target**, then select **Middleware**.
 - From the **Add** menu, select Oracle Application Server.
 - Enter the information requested for Oracle Application Server. Click **Next** once all the information requested is entered.
3. If you plan to monitor the directory server through Oracle Enterprise Manager Cloud Control, then ensure that the directory server is appropriately discovered in Enterprise Manager before moving on to the discovery of Oracle Identity Federation Server. Complete the following tasks to discover the supported directory servers:
 - Oracle Internet Directory 10.1.4

Discovery of Oracle Identity Management Suite 10g (including Oracle Internet Directory, Directory Integration Platform, Delegated Administration Server, and Single Sign-On Server) can be done using the discovery wizard on the Middleware page.

From the Middleware page, select **Oracle Application Server** from the **Add** menu. For more information, refer to the Discovering Oracle Identity Management Suite 10.1.4.2, 10.1.4.3.0 section.

4. If Oracle Database is used for the User Data Store, ensure that the database instance is discovered in Enterprise Manager Cloud Control before moving on to the discovery Oracle Identity Federation Server. Complete the following steps to discover Oracle Database Instance in Cloud Control:
 - Log in to Enterprise Manager. Select **Targets**, then select **Databases**.
 - Select **Add** from the Search List view.
 - Enter the information requested for the Database Instance. Click **Next** once all the information requested is entered.

Oracle Identity Manager 9.1.0.1

1. Complete all the configuration steps for Oracle Identity Manager. Ensure that the application server and database are appropriately set up and configured for Oracle Identity Manager.

Refer to the *Oracle Identity Manager Installation and Upgrade Guide* for instructions on configuring Oracle Identity Manager (http://download.oracle.com/docs/cd/B31081_01/index.htm).

2. Discover the application server on which Oracle Identity Manager is deployed in Enterprise Manager Cloud Control.

Note: To verify whether the version of your third-party software for Oracle Identity Manager is supported in Oracle Enterprise Manager Cloud Control release 12.1.0.2, refer to the certification matrix located on My Oracle Support (<https://support.oracle.com>).

Complete the following steps to discover the supported application servers:

- JBoss Application Server Version 4.0.2:

Log in to Enterprise Manager. Select **Targets**, then select **Middleware**.

From the **Add** menu, select **JBoss Application Server**.

Enter the information requested for the JBoss Application Server. Click **Next** once all the information requested is entered.
- Oracle WebLogic Application Server Version 7.x and 8.x:

Log in to Enterprise Manager. Select **Targets**, then select **Middleware**.

From the **Add** menu, select **Oracle WebLogic Domain 7.x and 8.x**.

Enter the information requested for the WebLogic Application Server. Click **Next** once all information requested is entered.
- Oracle WebLogic Application Server Version 10.x and later:

Log in to Enterprise Manager. Select **Targets**, then select **Middleware**.

From the **Add** menu, select **Oracle Fusion Middleware/WebLogic Domain**.

Enter the information requested for WebLogic Domain. Click **Continue** once all information requested is entered

– WebSphere Application Server:

Log in to Enterprise Manager. Select **Targets**, then select **Middleware**.

From the **Add** menu, select **IBM WebSphere Application Server**.

Enter the information requested for the WebSphere Application Server. Click **Next** once all information requested is entered.

3. If Oracle Database is used for Oracle Identity Manager, ensure that the database instance is discovered in Enterprise Manager Cloud Control before moving on to the discovery Oracle Identity Manager Server.

Complete the following steps to discover Oracle Database Instance in Cloud Control:

Log in to Enterprise Manager. Select **Targets**, then select **Database**.

Select **Add** from the Search List view.

Enter the information requested for the Database Instance. Click **Next** once all information requested is entered.

Discovering and Configuring Identity Management Targets

This chapter provides the information needed to discover and configure Identity Management targets.

23.1 Discovering Identity Management Targets

This section describes how to discover Identity Management targets.

23.1.1 Discovering Identity Management 11g

Enterprise Manager has a simple Discovery wizard for Oracle Identity Management 11g (including Oracle Internet Directory, Directory Integration Platform, Oracle Virtual Directory, Oracle Identity Federation, Oracle Access Manager, Oracle Adaptive Access and Oracle Identity Manager) targets. The Discovery wizard collects details about Oracle Identity Management 11g targets including information about the host, WebLogic User Name/Password, and other details.

Note: Before discovering the targets associated with Oracle Access Manager 11g, download and install patch 10094106.

To discover Oracle Identity Management 11g (including Oracle Internet Directory, Directory Integration Platform, Oracle Virtual Directory, Oracle Identity Federation, Oracle Access Manager, Oracle Adaptive Access Manager and Oracle Identity Manager), perform the following steps:

1. Log in to Enterprise Manager. Select **Targets**, then select **Middleware**.
2. From the **Add** menu, select **Oracle Fusion Middleware/WebLogic Domain**.
3. Enter the information requested to discover Oracle Identity Management 11g targets.

Field	Description
Administration Server Host	Host on which the WebLogic domain for Identity Management is running. Import the certificates for this WLS domain on the agent if this is a secured domain.
Port	Port used for the WebLogic domain. Enter a number between 1 and 65535.
User Name	WebLogic domain user name.

Field	Description
Password	WebLogic domain password.
Unique Domain Identifier	A unique identifier for the Identity Management domain and is used to create a unique target name. The Unique Domain Identifier can contain only alphanumeric characters and the special character '_' and cannot contain any other special characters.
Agent	Agent that is running on the Identity Management host. Only an agent 12.1 or later can be used for finding targets.
Advanced Fields	Description
JMX Protocol	JMX protocol is used to make a JMX connection to the Administration Server.
Discover Down Servers	This is a signal to discover the servers that are down.
JMX Service URL	JMX Service URL is used to make a JMX connection to the Administration Server. If the URL is not specified, it will be created based on the input parameters. If the URL is specified, the Administration server host and port information must still be provided in the input parameters.
External Parameters	These parameters will be passed to the java process which makes a connection to the Administration Server. All the parameters must begin with -D.
Discovery Debug File Name	The agent side discovery messages for this session will be logged into this file. This file will be generated in the discovery agent's log directory <agent home>/sysman/log. If this file already exists, it will be updated.

- A list of all the Identity Management targets is listed. Click **Add** to complete the discovery. **Note:** If the Configured Agent text-box is blank for one or more of the targets, copy and paste the Management Agent URL before you proceed.
- The status of target discovery is summarized in this screen. Ensure that all targets have been successfully added to Enterprise Manager. Press **OK** to finish the discovery process.
- The discovered targets will now be listed on the Identity and Access dashboard. From the **Targets** menu, select **Middleware**, then select **Middleware Features**.

23.1.2 Discovering Oracle Directory Server Enterprise Edition 6.x, 7.x, 11g

To discover Oracle Directory Server Enterprise Edition 6.x, 7.x, 11g targets, perform the following steps:

- Log in to Enterprise Manager. Select **Targets**, then select **Middleware**.
- From the **Add** menu, select **Oracle Directory Server Enterprise Edition**.
- Enter the information requested.
 - Oracle Directory Server Enterprise Edition Registry Host: Host of the Directory Server Control Center Registry
 - Oracle Directory Server Enterprise Edition Registry Port: Port of the Directory Server Control Center Registry
 - Directory Server User Name - for example CN=Directory Manager
 - Directory Server User Password

- e. Oracle Directory Server Enterprise Edition Install Home: Path under which Directory Server Enterprise Edition is installed.
- f. Unique Deployment Identifier: A unique identifier for ODSEE deployment.

23.1.3 Discovering Oracle Access Manager Access Server 10.1.4.2 and 10.1.4.3.0

Enterprise Manager has a simple Discovery wizard for Oracle Access Manager 10g targets. The Discovery wizard collects details about Oracle Access Manager Targets including information about the host name, host login credentials, SNMP Agent credentials, and other details.

After the Discovery wizard is complete, you can add the discovered targets into an existing System topology or you can create a new System target that stores your topology into the Management Repository.

To discover Oracle Access Manager - Access Server, perform the following steps:

1. Log in to Enterprise Manager. Select **Targets**, then select **Middleware**.
2. From the **Add** menu, select **Oracle Identity Management 10g (OAM, OIF, OIM)**.
3. Select **Access Manager - Access Server** and enter the host name on which your Access Server is running. Click **OK** to continue with the discovery of the Access Server.
4. Enter the information requested for Access Server. (The following table provides descriptions of the fields.) Click **Next** once all information requested is entered.

Field	Description
Host User Name	User name on the operating system with administrator privileges.
Host User Password	Password of host administrator account. <ul style="list-style-type: none"> ■ Save as Preferred Credentials. Select this check box if you would like to save the user name/password for the administrator account. ■ Management Agent is running on Host other than SNMP Host Select this check box if your Cloud Control Management Agent is running on a host other than the SNMP Agent host.
Access Server Home	Enter the home directory of your Access Server (<OAM_HOME>\access) - for example, C:\Program Files\OracleAccessManager\access
Access Server Version	Enter the version of your Oracle Access Manager - Access Server - for example, 10.1.4.0.1
SNMP Agent Host	If your Simple Network Management Protocol (SNMP) Agent is running on a host other than the Cloud Control Management Agent host, then enter the SNMP Agent host name. Otherwise, skip this section.
SNMP Agent Port	Enter the UDP Port of the SNMP Agent - for example, 161
SNMP Agent Community Name	Enter the community name of the SNMP Agent.
LDAP Server Host	Name of the Lightweight Directory Access Protocol (LDAP) host. The host name is available in the LDAPSERVERNAME parameter located in the <AccessServerInstallDir>/config/ldap/ConfigDB.xml file.

Field	Description
LDAP Server Port	Name of the LDAP port. The port name is available in the LDAPSERVERPORT parameter located in the <AccessServerInstallDir>/config/ldap/ConfigDB.xml file.
LDAP User Name	Name of the LDAP user. The user name is available in the LDAPROOTDN parameter located in the <AccessServerInstallDir>/config/ldap/ConfigDB.xml file.
LDAP Password	Password for the LDAP user.
LDAP Base	Name of the LDAP base. The base name is available in the LDAPOBLIXBASE parameter located in the <AccessServerInstallDir>/config/configInfo.xml file.

- Enterprise Manager discovers the topology of your Oracle Access Manager - Access Server deployment including the associated databases and directory servers.

To add this topology into an existing Access Manager - Access System target, select **Use the specified system**, and select an existing target of type Access Manager - Access System.

If you want to create a new Access Manager - Access System target, select the **Create a new system** and enter the name of the new system target. Click **Finish** to complete the discovery.

- The next page shows a message confirming the discovery of Oracle Access Manager - Access Server.

23.1.4 Discovering Oracle Access Manager Identity Server 10.1.4.2 and 10.1.4.3.0

Enterprise Manager has a simple Discovery wizard for Oracle Access Manager 10g targets. The Discovery wizard collects details about Oracle Access Manager Targets including information about the host name, host login credentials, SNMP Agent credentials, and other details.

After the Discovery wizard is complete, you can add the discovered targets into an existing System topology or you can create a new System target that stores your topology into Management Repository.

- Log in to Enterprise Manager. Select **Targets**, then select **Middleware**.
- From the **Add** menu, select **Oracle Identity Management 10g (OAM, OIF, OIM)**.
- Select **Access Manager - Identity Server** and enter the host name on which your Identity Server is running. Click **OK** to continue with the discovery of the Identity Server.
- Enter the information requested for Oracle Access Manager - Identity Server. (The following table describes the fields.) Click **Next** once all information requested is entered.

Field	Description
Host User Name	User name on the operating system with administrator privileges.

Field	Description
Host User Password	Password of host administrator account. <ul style="list-style-type: none"> Save as Preferred Credentials. Select this check box if you would like to save the user name/password for the administrator account. Management Agent is running on Host other than SNMP Host Select this check box if your Cloud Control Management Agent is running on a host other than the SNMP Agent host.
Identity Server Home	Enter the home directory of your Identity Server (<OAM_HOME>\identity) - for example, C:\Program Files\OracleAccessManager\identity
Identity Server Version	Enter the version of your Oracle Access Manager - Identity Server - for example, 10.1.4.0.1
SNMP Agent Host	If your Simple Network Management Protocol (SNMP) Agent is running on a host other than the Cloud Control Management Agent host, then enter the SNMP Agent host name. Otherwise, skip this section.
SNMP Agent Port	Enter the UDP Port of the SNMP Agent - for example, 161
SNMP Agent Community Name	Enter the community name of the SNMP Agent.

- Enterprise Manager discovers the topology of your Oracle Access Manager - Identity Server deployment including the associated databases and directory servers. To add this topology into an existing Access Manager - Identity System target, select **Use the specified system** and select an existing target of type Access Manager - Identity System. If you want to create a new Access Manager - Identity System target, select **Create a new system** and enter the name of new system target. Click **Finish** to complete the discovery.
- The next page shows a message confirming the discovery of Oracle Access Manager - Identity Server.

23.1.5 Discovering Oracle Identity Federation Server 10.1.4.2 and 10.1.4.3.0

Enterprise Manager has a simple Discovery wizard for Oracle Identity Federation targets. The Discovery wizard collects details about Oracle Identity Federation targets including information about the host name, host login credentials, and other details.

After the Discovery wizard is complete, you can add the discovered targets into an existing System topology or you can create a new System target that stores your topology into the Management Repository.

To discover Oracle Identity Federation Server, perform the following steps:

- Log in to Enterprise Manager. Select **Targets**, then select **Middleware**.
- From the Add menu, select **Oracle Identity Management 10g (OAM, OIF, OIM)**.
- Select **Identity Federation Server** and enter the host name on which your Oracle Identity Federation Server is running. Click **OK** to continue with the discovery of the Identity Federation Server.
- Enter the information requested for Oracle Identity Federation Server. Click **Continue** once all required information is entered.

Field	Description
Application Server Target	Select the Application Server target on which Oracle Identity Federation is running.
Host User Name	User name on the operating system with administrator privileges.
Host User Password	Password of host administrator account.

- Enterprise Manager discovers the topology of your Oracle Identity Federation Server deployment including the associated databases and directory servers.

To add this topology into an existing Identity Federation System target, select **Use the specified system** and select an existing target of type Identity Federation System.

If you want to create a new Identity Federation System target, select **Create a new system** and enter the name of new system target. Click **Finish** to complete the discovery.
- The next page shows a message confirming the discovery of Oracle Identity Federation Server.

23.1.6 Discovering Oracle Identity Management Suite 10.1.4.2 and 10.1.4.3.0

Enterprise Manager has a simple Discovery wizard for Oracle Identity Management Suite 10g (including Oracle Internet Directory, Directory Integration Platform, Delegated Administration Server, and Single Sign-On Server) targets. The Discovery wizard collects details about Oracle Identity Management Suite 10g targets including information about the host name, host login credentials, and other details.

To discover Oracle Identity Management Suite 10g (including Oracle Internet Directory, Directory Integration Platform, Delegated Administration Server, and Single Sign-On Server), perform the following steps:

- Log in to Enterprise Manager. Select **Targets**, then select **Middleware**.
- From the **Add** menu, select **Oracle Application Server**.
- Select the host on which Oracle Identity Management Suite 10g targets are running.
- A confirmation page lists Oracle Application Servers found on the host selected. Click **OK** to continue. **Important:** Ensure that the Application Server is up before discovering the Identity Management Suite targets.
- A final confirmation page appears. Click **OK** to finish the discovery process.

23.1.7 Discovering Oracle Identity Manager Server 9.1.0.1

Enterprise Manager has a simple Discovery wizard for Oracle Identity Manager targets. The Discovery wizard collects details about Oracle Identity Manager targets including information about the host name, host login credentials, and other details.

After the Discovery wizard is complete, you can add the discovered targets into an existing System topology or you can create a new System target that stores your topology into Enterprise Manager's Repository.

To discover Oracle Identity Manager Server, perform the following steps:

- Log in to Enterprise Manager. Select **Targets**, then select **Middleware**.
- From the **Add** menu, select **Oracle Identity Management 10g (OAM, OIF, OIM)**.

3. Select **Identity Manager Server** and enter the host name on which your Oracle Identity Manager is running. Click **OK** to continue with the discovery of the Oracle Identity Manager Server.
4. Enter the information requested for Oracle Identity Manager Server. Click **Continue** once all the required information is entered.

Field	Description
Application Server Target	Select the Application Server target on which Oracle Identity Manager is running.
Configured Database Target	Select the configured Database target used by Oracle Identity Manager
Database User Name	Enter the database user name used to access the tablespace reserved for Oracle Identity Manager.
Database Password	Enter the password for the database account reserved for Oracle Identity Manager.
Identity Manager Library Path	Enter the directory path for the Oracle Identity Manager library (<OIM_HOME>\xellerate\lib).
Host User Name	User name on the operating system with administrator privileges
Host Password	Password of host administrator account.

5. Enterprise Manager discovers the topology of your Oracle Identity Manager Server deployment including the associated databases and directory servers.

To add this topology into an existing Identity Manager System target, select **Use the specified system** and select an existing target of type Identity Manager System.

If you would like to create a new Identity Manager System target, select **Create a new system** and enter the name of new system target. Click **Finish** to complete the discovery.

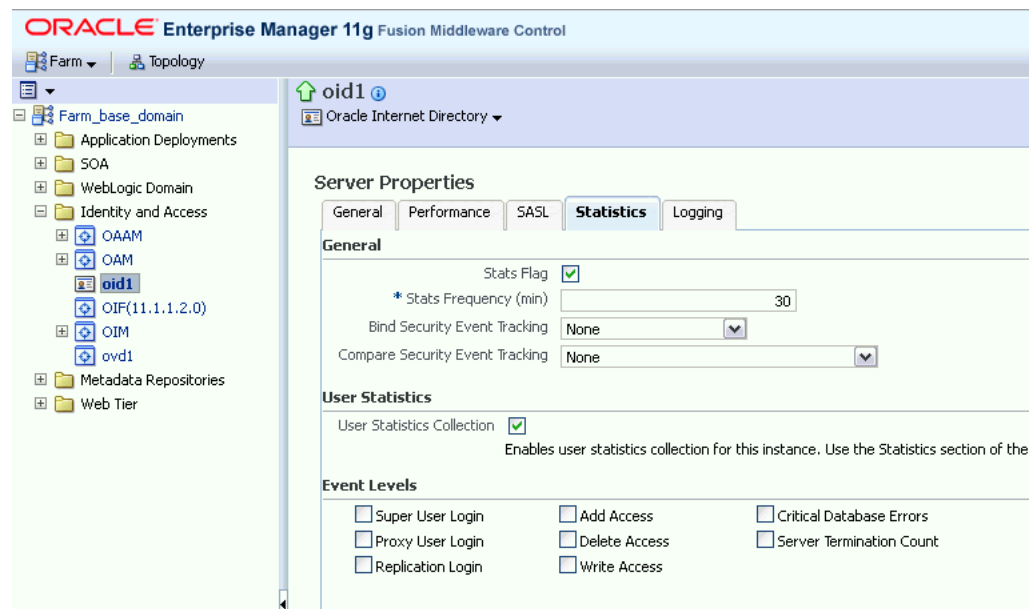
6. The next page shows a message confirming the discovery of Oracle Identity Manager Server.

23.2 Collecting User Statistics for Oracle Internet Directory

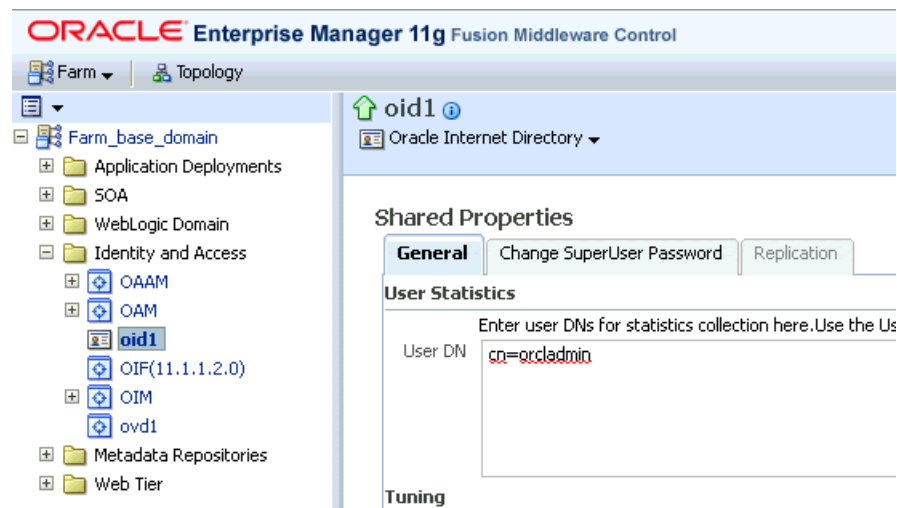
With Enterprise Manager, you can collect user statistics for Oracle Internet Directory allowing you to view charts for failed and completed LDAP operations like Add, Bind, Compare, Delete, Modify, and Search.

To enable the collection of user statistics, perform the following steps:

1. From the **Targets** menu, select **Middleware**. From the **Middleware Features** menu, select **Identity and Access**.
2. Select the discovered Oracle Internet Directory target.
3. From the **Oracle Internet Directory** menu, select **Fusion Middleware Control**.
4. From the **Targets** menu in Fusion Middleware Control, select **Administration**, then select **Server Properties**. Check the box next to **User Statistics Collection** to enable this feature. Click **Apply** to save your changes. See [Figure 23-1](#).

Figure 23–1 Server Properties - Statistics Tab

- From the **Target** menu in Fusion Middleware Control, select **Administration**, then select **Shared Properties**. Enter a valid User DN (for example, cn=orcladmin) to enable user statistics collection for that user. See [Figure 23–2](#).

Figure 23–2 Shared Properties - General Tab

23.3 Creating Identity Management Elements

This section describe how to create Identity Management elements.

23.3.1 Creating Identity and Access System

With Enterprise Manager, you can create an Identity and Access System target that can be modeled with any discovered Oracle Identity Management target (including both Identity Management 10g and Identity Management 11g targets) and the underlying

hosts, databases and LDAP servers as the key components providing an end-to-end system oriented view of the monitored Identity Management environment.

The Identity and Access System target provides access to metrics, alerts, charts, and topology view. In addition to monitoring your Oracle Identity Management environment from a system perspective, you can also monitor your environment from a service-oriented perspective using the Cloud Control Service Level Management framework.

To create a target of type Identity and Access System associated with any of the monitored Identity Management targets, perform the following steps:

1. Log in to Enterprise Manager. Select **Targets**, then select **Systems**.
2. From the **Add** menu, select **Identity and Access System**.
3. Select the Identity Management root target that you would like to include in your system topology. This can be the WebLogic Domain or the ODSEE Registry server.
Click **Next** to continue.
4. Select the targets within the domain that you would like to include in your system topology. You can also add additional targets that are not in the Identity Management domain, for example, databases, non-Oracle middleware, and so on.
Click **Next** to continue.
5. Click **Finish** to complete the creation of Identity and Access System.

23.3.2 Creating Generic Service or Web Application Targets for Identity Management

The Discovery wizard for Oracle Identity and Access Management Suite allows you to create a System target to store the end-to-end topology of monitored Oracle Identity Management components. The Management Pack Plus for Identity Management allows you to create the following System targets:

- Access Manager - Access System
- Access Manager - Identity System
- Identity Federation System
- Identity Manager System
- Identity and Access System

A System target is modeled with all monitored Oracle Identity Management components and the underlying hosts as the key components providing an end-to-end system oriented view of the monitored Oracle Identity Management environment.

A System target provides access to metrics, alerts, charts, and topology view of all the infrastructure components. In addition to monitoring your Oracle Identity Management environment from a system perspective, you can also monitor your environment from a service-oriented perspective using the Cloud Control Service Level Management framework.

With the Management Pack Plus for Identity Management, users can create targets of type Generic Service or Web Application associated with any of the monitored Identity Management Systems: Access Manager - Access System, Access Manager - Identity System, Identity Federation System, and Identity Manager System.

The Web Application or Generic Service target provides an end-to-end service oriented view of the monitored Oracle Identity Management targets with access to performance and usage metrics, service tests, service level rules, service availability definition, alerts, charts, and topology view.

To create a target of type Generic Service associated with any of the monitored Identity Management Systems, perform the following steps:

1. Log in to Enterprise Manager. Select **Targets**, then select **Services**.
2. From the **Add** menu, select **Generic Service**.
3. Enter the general information requested for the new Generic Service.

23.3.3 Creating a Service Dashboard Report

Once you have created Generic Service or Web Application targets associated with your monitored Oracle Identity Management Systems, you can create a Services Monitoring Dashboard that summarizes Service Level Agreement Compliance, Actual Service Level Achieved, Key Performance and Usage Metrics, and Status of Key Components.

Perform the following steps to create a Services Monitoring Dashboard:

1. From the **Enterprise** menu, select **Reports**, then select **Information Publisher Reports**.
2. Click the **Create** button.
3. Enter the general information requested for the new Report. Click the **Elements** tab once all information requested is entered.
 - a. Title
Enter a title for your new dashboard
 - b. Category/Sub-Category
Select a category and sub-category for your dashboard, for example, Category: Monitoring, Sub-Category: Dashboards
 - c. Use the specified target
Leave blank if this report has no report-wide target.
 - d. Options - Visual Style
Select Dashboard for a dashboard-view of your services.
4. Enter the elements information requested for the new Report. Click the **Schedule** tab once all information requested is entered.
 - a. Add
Select **Services Monitoring Dashboard** and click **Continue**.
 - b. Set Parameters
Click **Set Parameters**. Select the available services and click the **Move** button to add them to the Selected Services.
5. Enter the schedule information requested for the new Report. Click the **Access** tab once all information requested is entered.
 - a. Schedule
Enter your scheduling preferences for the report
 - b. E-Mail Report
Enter the email address and preferences for the report recipient.

6. Enter information about your access and security preferences for the new report. Click **OK** to create the new Services Monitoring Dashboard.

Part IX

Discovering and Monitoring Non-Oracle Middleware

The chapters in this part describe how to discover and monitor non-Oracle middleware components.

The chapters are:

- [Chapter 24, "Discovering and Monitoring IBM WebSphere MQ"](#)
- [Chapter 25, "Discovering and Monitoring JBoss Application Server"](#)

Discovering and Monitoring IBM WebSphere MQ

IBM WebSphere MQ is a message oriented middleware and its primary infrastructure is queue based. Message Queue (MQ) clusters are used for high availability, and management and monitoring is supported by a command line tool, user interface, and programmable command format messages.

IBM WebSphere MQ allows administrators to derive instant value, while giving them the flexibility to fine-tune thresholds according to their specific operational requirements.

Note: IBM WebSphere MQ code is bundled with the Oracle Fusion Middleware plug-in.

The following topics are discussed in this document:

- [Introduction](#)
- [Requirements](#)
- [Understanding Discovery](#)
- [Monitoring](#)
- [IBM WebSphere MQ Metrics](#)

24.1 Introduction

IBM WebSphere MQ offers several key benefits, including the following:

- Out-of-box availability and performance monitoring
- Centralized monitoring of all information in a single console
- Enhanced service modeling and comprehensive root cause analysis

Out-of-Box Availability and Performance Monitoring

The plug-in allows you to realize immediate value through out-of-box availability and performance monitoring. Some of the key areas of the more than 60 performance indicators monitored by the plug-in include queue manager status, channel status, queue depth, bytes sent or received, and messages sent or received.

To further aid administrators with critical tasks, such as problem diagnosis, trend analysis, and capacity planning, the System Monitoring for IBM WebSphere MQ includes various out-of-box reports, summarizing key information about availability

and performance. IBM WebSphere MQ System Monitoring Plug-in key features include:

- **Blackout Periods**
Prevent unnecessary alerts from being raised during scheduled maintenance operations, such as hardware upgrade.
- **Monitoring Templates**
Simplify the task of standardizing monitoring settings across the entire IBM WebSphere MQ environment, by allowing administrators to specify the monitoring settings (metrics, thresholds, metric collection schedules and corrective actions) once and applying them to any number of queue manager instances.
- **Corrective Actions**
Ensure that routine responses to alerts are automatically executed, thereby saving administrators time and ensuring problems are dealt with before they noticeably impact users.
- **Notification Rules, Methods, and Schedules**
Define when and how administrators should be notified about critical problems with their applications, ensuring quicker problem resolution.
- **Groups / Systems**
Simplify management of large numbers of components, allowing administrators to “manage many-as-one”.

Centralized Monitoring of all Information in a Single Console

IBM WebSphere MQ provides administrators managing Oracle systems and IBM WebSphere MQ with a consolidated view of the entire enterprise, enabling them to monitor and manage all of their components from a central place.

Having such an integrated tool reduces the total cost of ownership by eliminating the need to manually compile critical information from several different tools, thus streamlining the correlation of availability and performance problems across the entire set of IT components.

In addition, System Monitoring for IBM WebSphere MQ complements existing plug-in support for host systems (requires licensing of System Monitoring Plug-in for Hosts), providing administrators with more depth and greater ability to view IBM WebSphere MQ operations directly in the context of operating system activity.

Enhance Service Modeling and Perform Comprehensive Root Cause Analysis

Oracle Enterprise Manager's Service Level Management functionality provides a comprehensive monitoring solution that helps IT organizations achieve high availability, performance, and optimized service levels for their business services. Administrators can monitor services from the end-users' perspective using service tests or synthetic transactions, model relationships between services and underlying IT components, diagnose root cause of service failure, and report on achieved service levels.

The System Monitoring Plug-in for IBM WebSphere MQ enables IT organizations running applications on top of Oracle and IBM to derive greater value from the Service Level Management features of Enterprise Manager in a number of ways:

- **Enhanced Service Modeling**
Mapping of relationships between services and queue manager instances.

- Complete Service Topology
Including IBM WebSphere MQ as part of the topology view of a service.
- Comprehensive Root Cause Analysis
Identifying or excluding IBM WebSphere MQ as the root cause of service failure.

24.2 Requirements

The following prerequisites must be met before installing the IBM WebSphere MQ plug-in:

- Queue Manager must be running
- TCP listener must be up
- SYSTEM.DEF.SERVERCON channel must be available

Target discovery is performed using Grid Control Agent 10.2.0.3, so this agent must be installed and running. The agent can be installed on the same node as IBM WebSphere MQ (for Local Monitoring) or on another node in the intranet (for Remote Monitoring).

The agent must have a status UP in the Grid Control console and must upload data successfully to the Grid Control OMS (communication from the agent to the OMS). The Preferred Credentials must have been set and successfully tested for the agent node. The Host Credentials to be used for the discovery should be either the Grid Control Agent user or a user part of the same group (primary gid).

24.2.1 About Jar File Requirements for Local Monitoring

For local monitoring, the Grid Control Agent OS user should have read privileges over the following jar files, which are needed for the discovery of the target IBM WebSphere MQ:

- \$MQ_HOME/java/lib/com.ibm.mq.jar
- \$MQ_HOME/java/lib/connector.jar
- \$MQ_HOME/eclipse/plugins/com.ibm.mq.pcf_6.0.0/pcf.jar

Note: MQ_HOME is the directory used to install MQ setup at installation time.

24.2.2 About Jar File Requirements for Remote Monitoring

For Remote Monitoring, the TCP listener port of the Queue Manager must be open to the Agent Host. The appropriate jar files must be copied on the node where the Grid Control Agent is running and the OS user starting this agent must have read privileges on these files. For example, create a directory <AGENT_ORACLE_HOME>/sysman/mq_jar_files and copy the jar files into this directory.

The following sections provide details on the jar files required to support remote monitoring from IBM Websphere MQ V6 and V7:

IBM WebSphere MQ V6

- \$MQ_HOME/java/lib/com.ibm.mq.jar
- \$MQ_HOME/java/lib/connector.jar

- \$MQ_HOME/eclipse/plugins/com.ibm.mq.pcf_6.0.0/pcf.jar

IBM WebSphere MQ V7

- \$MQ_HOME/opt/mqm/java/lib/com.ibm.mq.jar
- \$MQ_HOME/opt/mqm/eclipse/plugins/com.ibm.mq.pcf.event_7.0.0.1.jar
- \$MQ_HOME/opt/mqm/java/lib/com.ibm.mq.jmqi.jar
- \$MQ_HOME/opt/mqm/java/lib/com.ibm.mq.commonservices.jar
- \$MQ_HOME/opt/mqm/java/lib/com.ibm.mq.headers.jar

24.3 Understanding Discovery

IBM WebSphere MQ supports the discovery of entire Queue Manager Clusters from a single Queue manager. Administrators can derive instant value, while giving them the flexibility to fine-tune thresholds according to their specific operational requirements. Some of the key areas of the more than 60 performance indicators being monitored, including queue manager status, channel status, queue depth, bytes sent and/or received, and messages sent and/or received.

To further aid administrators with critical tasks such as problem diagnosis, trend analysis, and capacity planning, plug-in includes various out-of-box reports, summarizing key information about availability and performance.

24.3.1 Discovery Parameters

The following parameters are required for MQ discovery:

- Host name (where the queue manager is running)
- Port number
- JAR path
- MQ administration username
- MQ administration password
- Agent host

The following parameters are required for agent discovery:

- Host name
- Host credentials

24.3.2 Discovery Prerequisites for Local Agent

To enable discovery for the local agent the queue manager must be running and the TCP listener must be up. In addition, the following jar files are required for discovery and should therefore be accessible to agent:

- com.ibm.mq.jar (present under MQ_HOME/java/lib)
- connector.jar (present under MQ_HOME/java/lib)
- pcf.jar (present under MQ_HOME/eclipse/plugins/com.ibm.mq.pcf_<version>)

The Agent Host Credentials to be used for discovery should be either Oracle Agent User or should be part of the same group. The SYSTEM.DEF.SERVERCON channel should also be available.

24.3.3 Discovery Prerequisites for Remote Agent

To enable discovery for the remote agent the queue manager must be running and the TCP listener must be up. The TCP listener port of the Target Queue Manager must also be accessible to the agent. In addition, the following jar files are required for discovery and should therefore be accessible to agent:

- com.ibm.mq.jar (present under MQ_HOME/java/lib)
- connector.jar (present under MQ_HOME/java/lib)
- pcf.jar (present under MQ_HOME/eclipse/plugins/com.ibm.mq.pcf_<version>)

The Agent Host Credentials to be used for discovery should be either Oracle Agent User or should be part of same group. The SYSTEM.DEF.SERVERCON channel should also be available.

24.3.4 Queue Manager Cluster Discovery

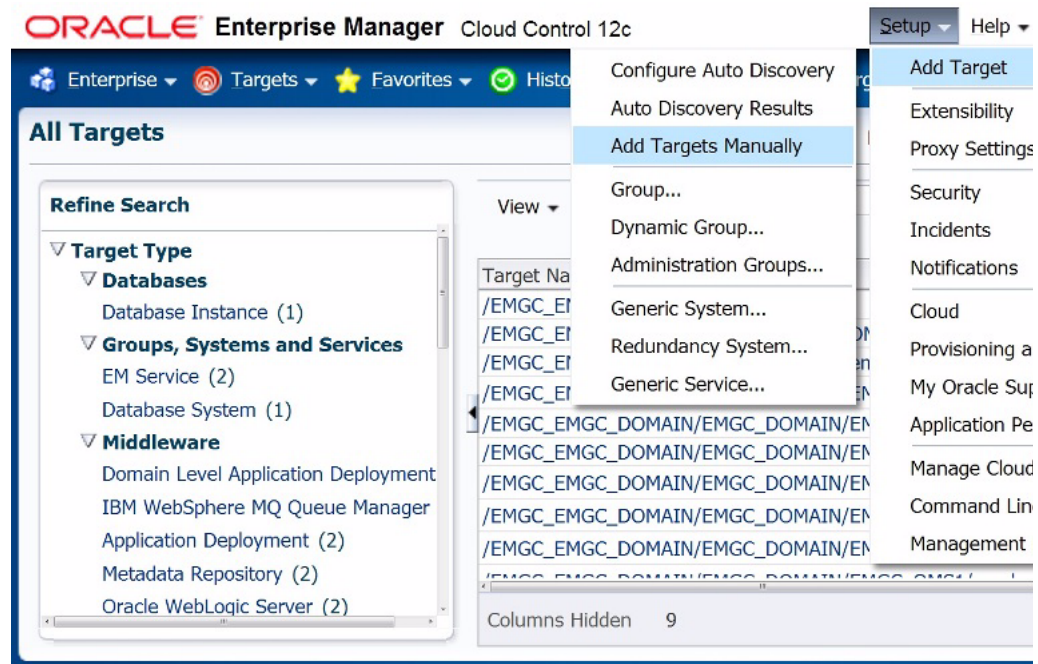
A cluster will be discovered automatically If the queue manager is part of a cluster. To discover other members of cluster, those queue managers should be running. If the queue manager is part of more than one cluster then all clusters will be discovered.

To following points outline the logic involved in cluster discovery:

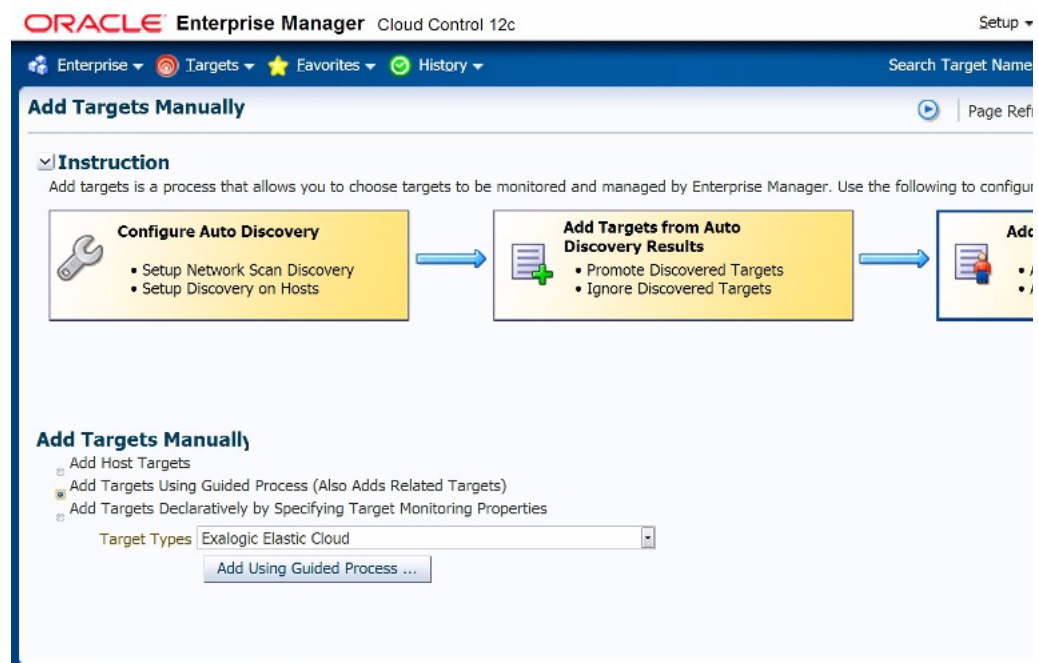
- To discover the entire cluster you first discover any member of cluster
- Get connection details of all other queue managers along with cluster name
If the queue manager is part of multiple clusters then repeat the step for each cluster
- Connect to each queue manager using the connection details
If the queue manager is part of multiple clusters then repeat the step for each cluster
- Get the name of the queue manager (queue manager should be running)
- Once the cluster is added to Grid Control, you cannot explicitly remove any member queue managers from that cluster.

The MQ to be discovered can be either local or remote to the host where the Enterprise Manager Agent resides. To manually add targets, complete the following:

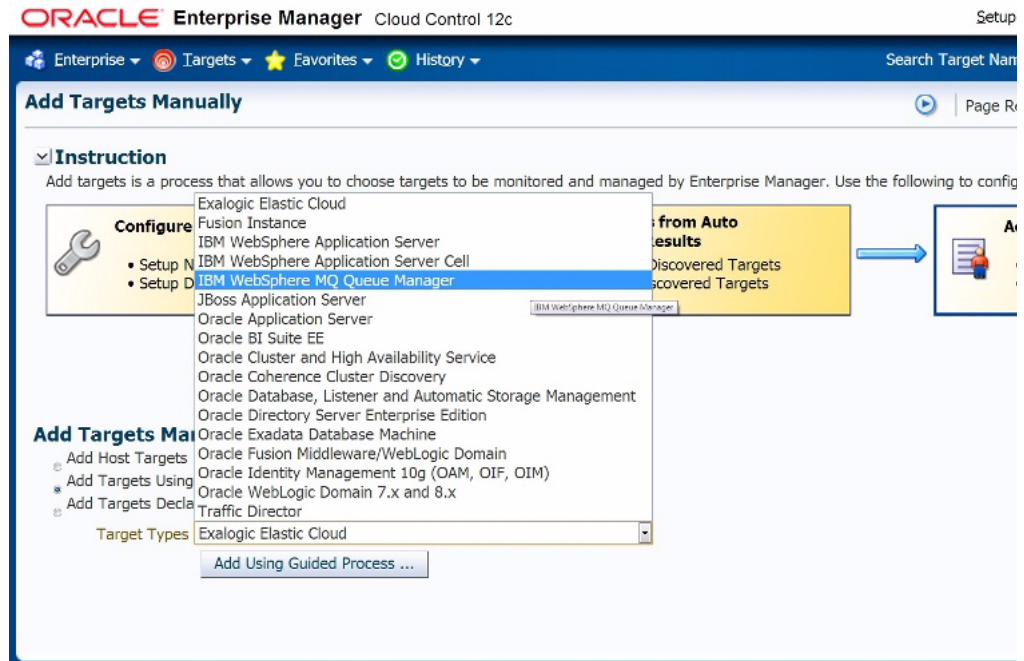
1. From the Enterprise Manager **Setup** menu, select **Add Target**, then select **Add Target Manually**.



2. Select the **Add Targets Using Guided Process (Also Adds Related Targets)** option.



3. Select **IBM WebSphere MQ Queue Manager** from Target Types and click **Add Using Guided Process....**



4. Provide the following required information for discovery and click **Next**.
 - a. Host on which IBM MQ is running.
 - b. Port number of the queue manager you want to discover (associated cluster and all queue manager in this cluster will also be discovered).
 - c. Jar path location (location must be accessible to the OEM agent).
 - d. Provide the agent host details used to monitor IBM WebSphere MQ.
 - e. MQ administration credentials.

ORACLE Enterprise Manager Cloud Control 12c

Authentication Available Targets Confirmation

Add IBM WebSphere MQ Targets: Authentication

In order to add an IBM WebSphere MQ Target to Enterprise Manager, you must first specify details of the host on which the IBM WebSphere MQ Server is running. The server must be running.

WebSphere MQ Host Information

* Host Provide the WebSphere MQ Host name, port and location.

* Port TCP/IP listener port of Queue Manager

Channel Specify the name of the Server Connection channel to be used for monitoring.

* Jar Path Provide the WebSphere MQ Jar path.

Agent Credentials

* Host Provide the Agent URL, user name and password.

WebSphere MQ Admin Credentials

* MQ Admin User Provide the WebSphere MQ Admin user name and password.

* MQ Admin Password

- Select the cluster you want to discover and click **Next**.

ORACLE Enterprise Manager Cloud Control 12c

Authentication Available Targets Confirmation

Add IBM WebSphere MQ Targets: Authentication

In order to add an IBM WebSphere MQ Target to Enterprise Manager, you must first specify details of the host on which the IBM WebSphere MQ Server is running. The server must be running.

WebSphere MQ Host Information

* Host: Provide the WebSphere MQ Host name, port and location

* Port: TCP/IP listener port of Queue Manager

Channel: Specify the name of the Server Connection channel to be used for monitoring.

* Jar Path: Provide the path to the jar file

Agent Credentials

* Host: Provide the Agent URL, user name and password.

WebSphere MQ Admin Credentials

* MQ Admin User: Provide the WebSphere MQ Admin user name and password

* MQ Admin Password:

The following image shows a Queue Manager cluster:

ORACLE Enterprise Manager Cloud Control 12c

Authentication Available Targets Confirmation

Add IBM WebSphere MQ Targets: Available Targets

[Cancel](#) [Back](#) Step 2 of 3

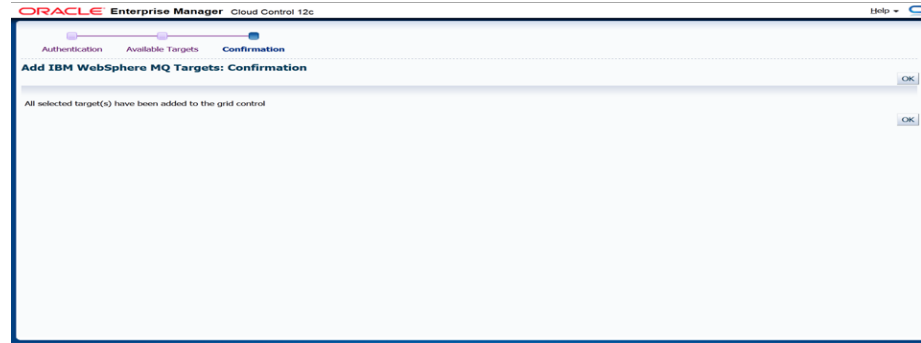
List Of All Available Websphere MQ Target(s)

MQ Targets

Select All | Select None

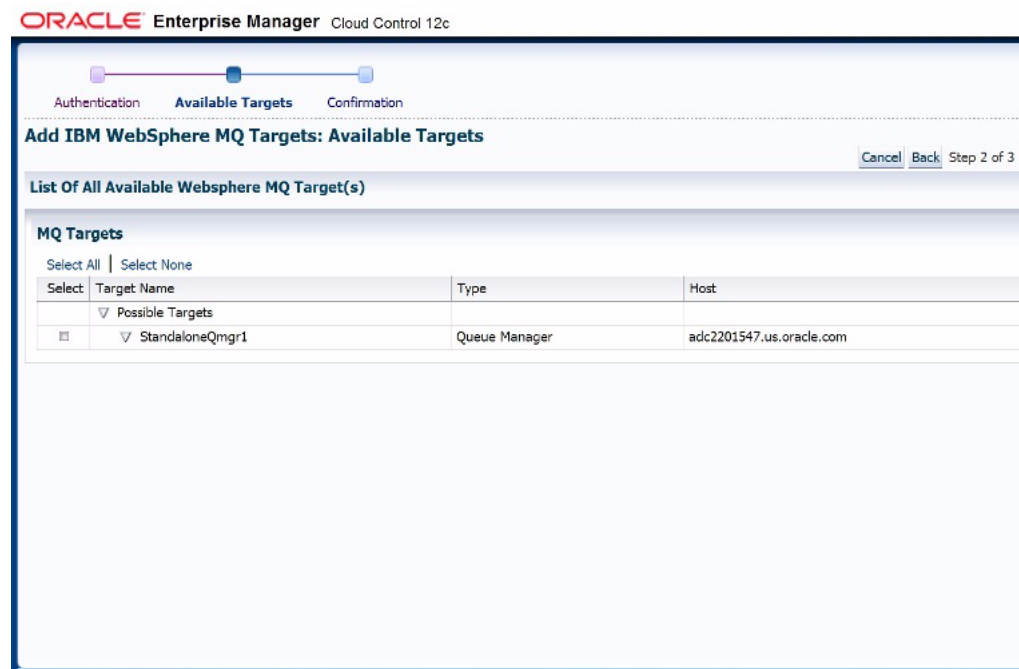
Select	Target Name	Type	Host	Port
<input checked="" type="checkbox"/>	▼ Possible Targets			
<input checked="" type="checkbox"/>	▼ Cluster1			
<input type="checkbox"/>	Qmgr1	Queue Manager	adc2201547	1414
<input type="checkbox"/>	Qmgr2	Queue Manager	adc2201547	1415

- Click **OK** to finish discovery.



24.3.5 Standalone Queue Manager Discovery

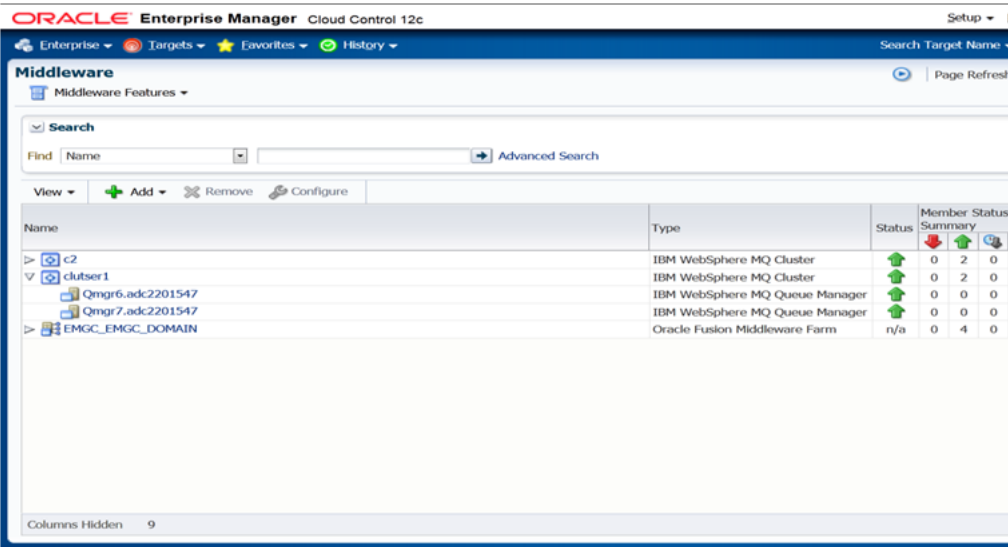
Standalone queue manager discovery is also supported. The steps to manually add a standalone queue manager are the same as those described for cluster queue manager discovery. see [Queue Manager Cluster Discovery](#) for more information. The only difference is that you select a standalone queue manager from the list of possible targets on the Add IBM WebSphere MQ Targets: Available Targets page, as shown below:



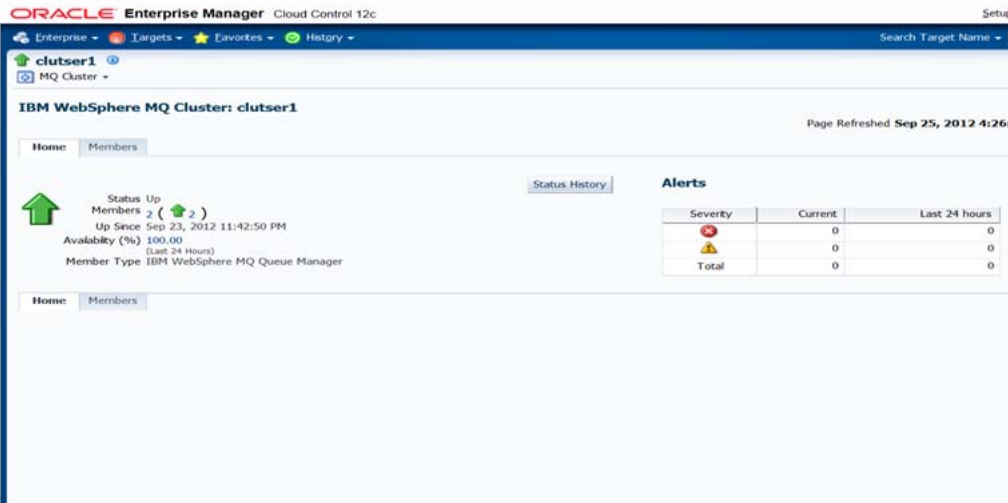
24.4 Monitoring

The following illustrates some of the different methods used to monitor the performance of the IBM WebSphere MQ plug-in:

- Middleware page displaying the discovered IBM MQ clusters and queue-managers.



- Cluster page displaying information about a cluster.



- Cluster Member's page displaying information about members of the cluster.

ORACLE Enterprise Manager Cloud Control 12c

Enterprise Targets Favorites History Search Target Name

clutser1 MQ Cluster

IBM WebSphere MQ Cluster: clutser1

Page Refreshed Sep 24, 2012 11:47:41 P

Home Members

View All Go

Expand All Collapse All

Name	Type	Status	Alerts
clutser1	IBM WebSphere MQ Cluster	Up	0
Qmgr6.adc2201547	IBM WebSphere MQ Queue Manager	Up	0
Qmgr7.adc2201547	IBM WebSphere MQ Queue Manager	Up	0

TIP For an explanation of the icons and symbols used in this page, see the Icon Key.

Home Members

- Queue-Manager page displaying information about queues and channels.

ORACLE Enterprise Manager Cloud Control 12c

Enterprise Targets Favorites History Search Target Name

Qmgr6.adc2201547 WebSphere MQ Queue Manager

Home Performance

General

Status Up Black Out

Availability (%) 100

Host sld01nlu.us.oracle.com

Target Specific Data

Host Platform: Unix

Maximum Handles: 256

Member Of Cluster: yes

Default DeadLetter Queue:

Default Transmission Queue:

Max. Length of Message: 4194304

Cluster Information

Name	Connection	Channel
clutser1	adc2201547(1425)	
clutser1	adc2201547(1426)	

Queue Manager Components

View All Go

Expand All Collapse All

MQ Component

Queues

Local

Queue1

Queue2

Channels

Cluster Receiver

TO:Qmgr6

Cluster Sender

TO:Qmgr7

Incidents

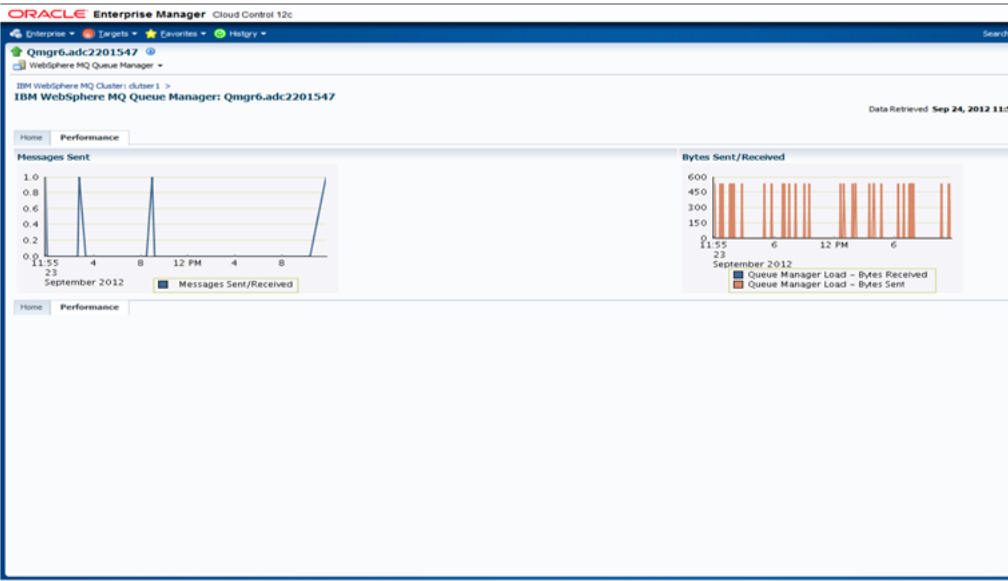
Severity	Message	Created At	Last Updated At	Escalator
No incidents found.				

Host Incidents

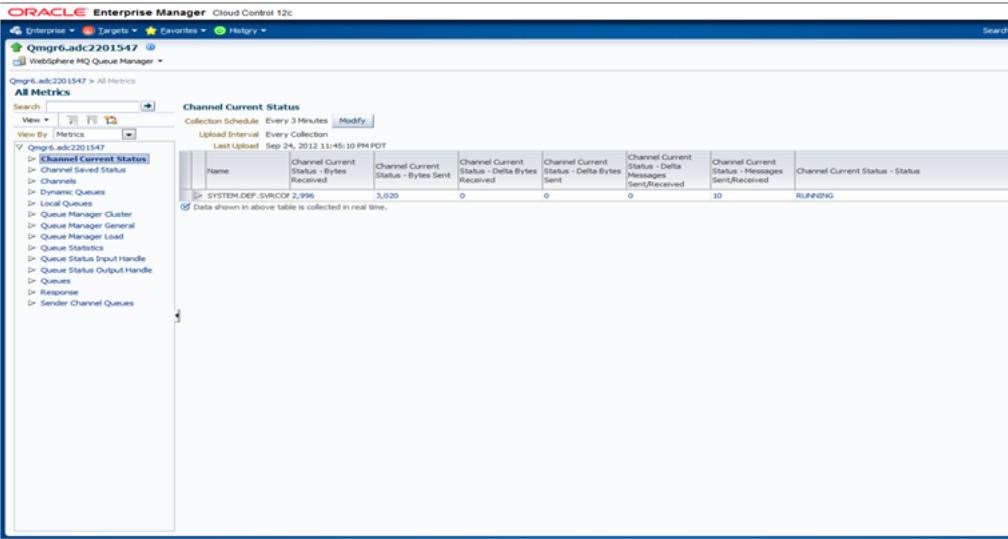
Hosts Collection Errors 1

Severity	Message	Created At	Last Updated At	Escalator
No incidents found.				

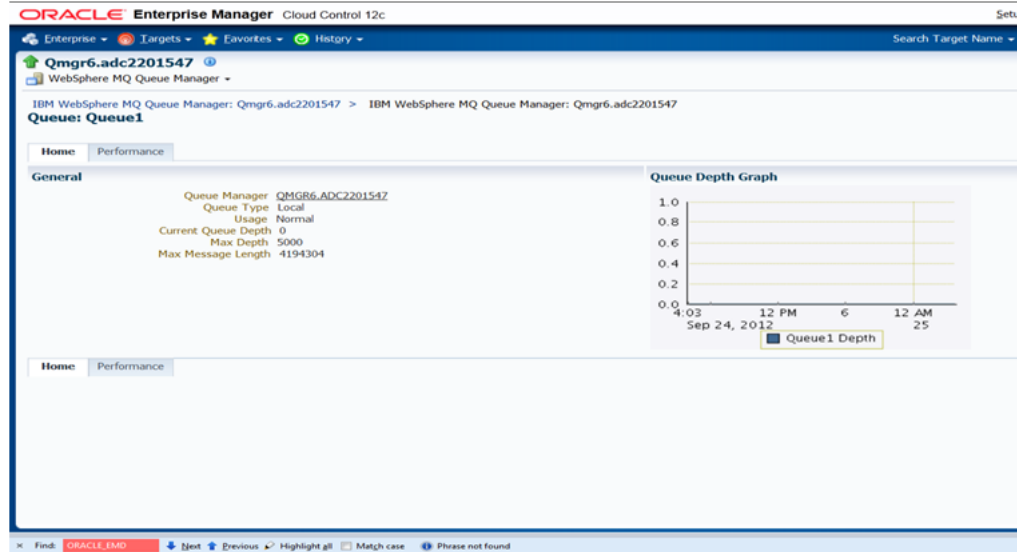
- Queue-manager performance page displaying information about messages.



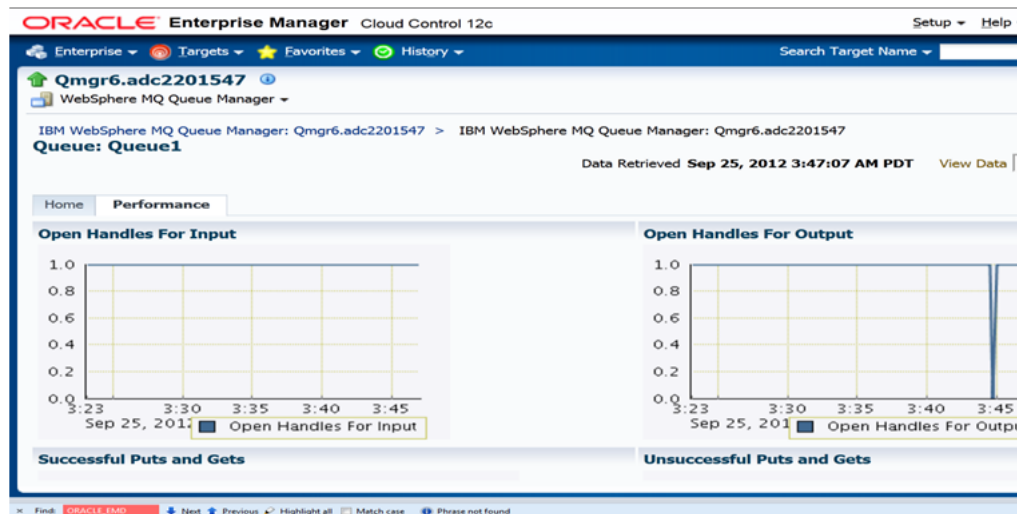
- All metric page for the IBM MQ plug-in displaying all the metrics collected.



- Queue page displaying information about queues discovered.



- Queue performance page displays various information about queue.



24.5 IBM WebSphere MQ Metrics

For information on all IBM WebSphere MQ metric categories, including tables that list and describe associated metrics for each category, see the IBM WebSphere MQ Metrics chapter in *Oracle® Enterprise Manager System Monitoring Plug-in Metric Reference Manual for Non-Oracle Middleware Management*. This can be found at the following location:

http://docs.oracle.com/cd/E11857_01/em.111/e10115/middleware_mq.htm#sthref127

Discovering and Monitoring JBoss Application Server

The JBoss Application Server is an open source application server that is organized in partitions and clusters. JBoss also supports standalone servers that are not part of any cluster. Multicasting is also supported, but needs to be disabled for Enterprise Manager discovery.

The following topics are discussed in this document

- [JBoss Setup](#)
- [JBoss Startup](#)
- [Discovery](#)
- [Monitoring JBoss JVM Metrics](#)
- [About JVM Diagnostics](#)
- [Troubleshooting](#)

25.1 JBoss Setup

To setup JBoss Application Server, download and extract the corresponding zip file and set the JBOSS_HOME and PATH environment variables as follows:

```
setenv JBOSS_HOME <jboss_install_location>
setenv PATH "${PATH}:${JBOSS_HOME}/bin"
```

Copy the MEJB to the appropriate deploy folder. This exposes management APIs to the agent. For example, you can copy the docs/examples/jmx/ejb-management.jar file to the server/default/deploy folder or the server/all/deploy folder.

25.2 JBoss Startup

To start the JBoss Application Server, run the following command from the bin folder:

```
./run.sh -c <deployment profile> -b <binding address>
[-Djboss.partition.name=<Partition name>]
```

The *<deployment profile>* determines the type of server instance that is started, such as standalone or JBoss partition. The *<binding address>* is the hostname/IP running JBoss (for example, server1.us.oracle.com). By default, the servers start as part of DefaultPartition. You can change this by specifying the *<Partition name>*.

To start a standalone JBoss server instance, set a 'default' deployment profile. For example:

```
./run.sh -c default -b <binding address>
```

To start a JBoss partition, the JBoss Clustering service has to be enabled and the deployment profile set to 'all'. For example:

```
./run.sh -c all -b <binding address>
```

To start multiple server instances on the same host, complete the following:

1. Create multiple deployment profiles as per requirements.
2. Use a different port-set to start the individual servers. For example:

```
-Djboss.service.binding.set=ports-01
```

Note that ports-01, ports-02, ports-03, and ports-04 are predefined port-sets. The following is a sample command:

```
./run.sh -c node1 -Djboss.service.binding.set=ports-01 -b example.us.oracle.com
```

25.3 Discovery

The following sections provide details on the steps involved in JBoss local discovery (agent is running on the same host) and JBoss remote discovery (agent is running on a remote host).

- [Understanding JBoss Local Discovery](#)
- [Understanding JBoss Remote Discovery](#)
- [Adding the JBoss Application Server](#)

25.3.1 Understanding JBoss Local Discovery

To enable discovery for the local agent the JNDI service should be configured and Naming Service port should be open. Furthermore, the agent should be able to communicate with the JNDI port. Host user should also have read and execute permissions over the JBOSS_HOME.

25.3.2 Understanding JBoss Remote Discovery

To enable JBoss remote discovery, the JNDI service should be configured and the Naming Service port should be open. In addition, the agent should be able to communicate with the host and JNDI ports, and should also have read and execute permissions on all files in the library path.

The following table provides a list of the JAR files to be copied to the agent machine (location provided as library path during discovery):


Table 25–1 JBoss JAR Files

JBoss Version	JAR files
JBoss 5.x	<ul style="list-style-type: none"> ■ "server/all/lib/dom4j.jar", ■ "client/concurrent.jar", ■ "client/jbossall-client.jar", ■ "client/jnp-client.jar", ■ "client/jboss-client.jar", ■ "client/jboss-common-core.jar", ■ "client/jboss-integration.jar", ■ "client/jboss-javaee.jar", ■ "client/jboss-jsr77-client.jar", ■ "client/jboss-logging-spi.jar", ■ "client/jboss-remoting.jar", ■ "client/jboss-security-spi.jar", ■ "client/jboss-serialization.jar", ■ "client/jbosssx-as-client.jar", ■ "client/jbosssx-client.jar", ■ "server/all/lib/jboss-management.jar",
JBoss 6.x	<ul style="list-style-type: none"> ■ "server/all/lib/dom4j.jar", ■ "client/concurrent.jar", ■ "client/jbossall-client.jar", ■ "client/jnp-client.jar", ■ "client/jboss-client.jar", ■ "client/jboss-common-core.jar", ■ "client/jboss-integration.jar", ■ "client/jboss-javaee.jar", ■ "client/jboss-jsr77-client.jar", ■ "client/jboss-logging-spi.jar", ■ "client/jboss-remoting.jar", ■ "client/jboss-security-spi.jar", ■ "client/jboss-serialization.jar", ■ "client/jbosssx-as-client.jar", ■ "client/jbosssx-client.jar", ■ "client/jboss-logging.jar", ■ "client/jboss-transaction-api_1.1_spec.jar", ■ "client/jboss-ejb-api_3.1_spec.jar", ■ "server/all/lib/jboss-management.jar",

25.3.3 Adding the JBoss Application Server

To manually add the JBoss application server, complete the following:

1. From the Enterprise Manager **Targets** menu, select **Middleware**, then click **Add**.
Select the JBoss Application Server from the list of options.



2. Specify the details of the host on which the JBoss Application Server is running and click **Next**.

The screenshot shows the 'Add JBoss Application Server' wizard in Oracle Enterprise Manager. The 'Host' tab is selected, and the 'Next' button is visible. The wizard is asking for details of the host on which the JBoss Application Server is running. The fields are: 'JBoss Application Server Host' (adc6140210.us.oracle.com), 'JMX Connector Port' (1090), 'JBoss Server Version' (6.x), 'Username' (This is required only if JMX connection requires authentication), 'Password' (This is required only if JMX connection requires authentication), and 'Agent' (slc01hur.us.oracle.com:1838).

3. Select from the list of discovered targets (partition/servers) shown and click **Next**.

Enterprise Manager Cloud Control 12c Help

Targets Review

plays the Standalone JBoss Application Servers and/or JBoss Partitions with their member servers. Select the Partition member servers and/or standalone servers you want Enterprise Manager to monitor. server is selected, associated partition is added automatically. Additional servers of a partition can be added to Enterprise Manager by clicking Refresh Partition on the JBoss Partition Home page.

Cancel Back Step 2 of 3 Next

Expand All Collapse All

Name	Type	Host	JNDI Port	Username	Password
adcs140210.us.oracle.com_DefaultPartition_1099	JBoss Partition	adcs140210.us.oracle.com	1099		

Cancel Back Step 2 of 3 Next

4. Review discovered targets to be saved and click **Finish**.

ORACLE Enterprise Manager Cloud Control 12c

Host **Select Servers** **Review**

Add JBoss Application Server: Review
Click Finish to add the mentioned JBoss Application Server(s) and/or Partition(s) to Enterprise Manager.

Cancel Back Step 3 of 3 Finish

JBoss Application Server Host: adcs140210.us.oracle.com
 Username:
 Agent Host: adcs140210.us.oracle.com
 Selected Partitions: DefaultPartition
 Selected Partition Member Servers: adcs140210.us.oracle.com_1099
 Selected Standalone Servers: None

TIP This wizard allows you to add one/more JBoss Application Server Partitions and/or standalone servers to Enterprise Manager at a time. If you want to add a different partition or a standalone server, complete

5. Click **OK** on the confirmation screen

ORACLE Enterprise Manager Cloud Control 12c Setup Help

Targets Favorites History

Search Target Name

Middleware Features

Page Refreshed Sep 27, 2012 10

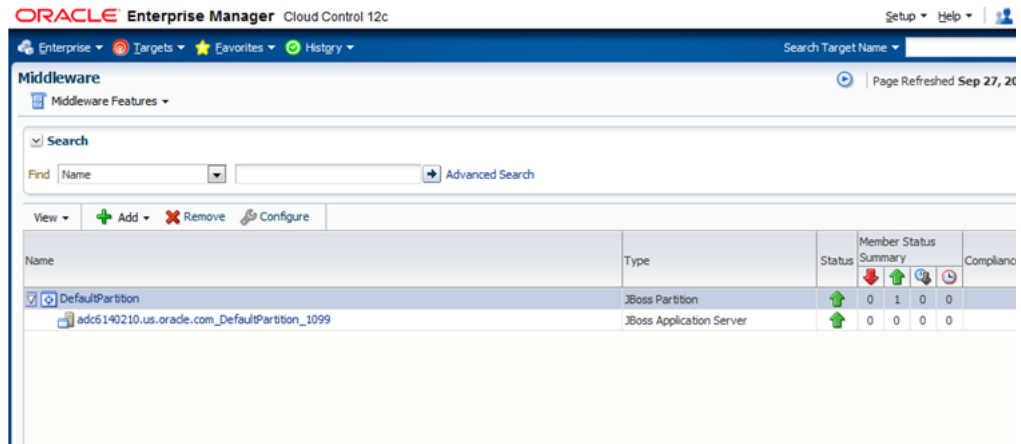
Information

Selected JBoss Partition Members and/or standalone servers on host adcs140210.us.oracle.com have been successfully added, and are now being monitored by Enterprise Manager.

OK

Name	Type	Status	Member Status Summary	Compliance Score
DefaultPartition	JBoss Partition		0 0 0 1	n/a

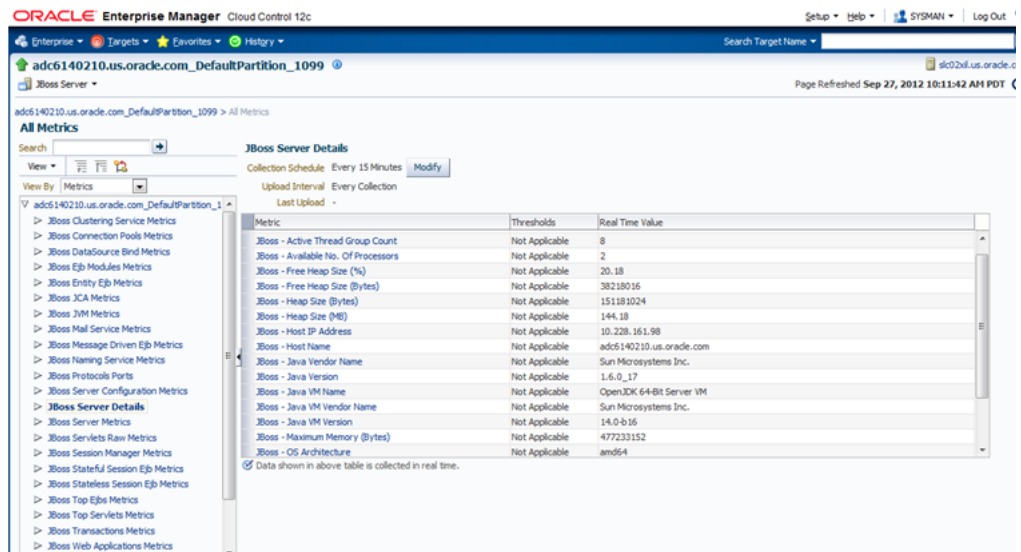
6. From the Enterprise Manager **Targets** menu, select **Middleware**, then verify that the Jboss target is present with status UP.



25.4 Monitoring JBoss JVM Metrics

You can use the All Metrics page on Enterprise Manager to monitor the following JBoss metrics:

- JBoss Server Details



- JBoss JVM Metrics

The screenshot shows the Oracle Enterprise Manager Cloud Control 12c interface. The target is 'adc6140210.us.oracle.com_DefaultPartition_1099'. The left sidebar shows a tree of metrics, with 'JBoss JVM Metrics' selected. The main panel displays a table of JVM metrics.

Metric	Thresholds	Real Time Value
JVM - Free Heap (%)	Set	27.88
JVM - Heap Size (MB)	Not Applicable	152.13
JVM Name	Not Applicable	OpenJDK 64-Bit Server VM
JVM Version	Not Applicable	1.6.0_17
Vendor	Not Applicable	Sun Microsystems Inc.

Collection Schedule: Every 30 Minutes
Upload Interval: Every Collection
Last Upload: -

Data shown in above table is collected in real time.

■ JBoss Web Applications Metrics

The screenshot shows the Oracle Enterprise Manager Cloud Control 12c interface. The target is 'adc6140210.us.oracle.com_DefaultPartition_1099'. The left sidebar shows a tree of metrics, with 'JBoss Web Applications Metrics' selected. The main panel displays a table of web application metrics.

War File Name	Web Apps - Active requests since last agent upload	Web Apps - Average maximum response time (ms)	Web Apps - Average minimum response time (ms)	Web Apps - Average Response Time (ms)	Web Apps - Active
/localhost/invoker	0	0	0	0	0
/localhost/jbossing-http	0	0	0	0	0
/localhost/jbossing	0	0	0	0	0
/localhost/jmx-console	199	2.915	0	235.538	213
/localhost/judd	0	0	0	0	0
/localhost/web-console	199	0.048	0	0.148	0

Collection Schedule: Every 15 Minutes
Upload Interval: Every Collection
Last Upload: -

Data shown in above table is collected in real time.

25.5 About JVM Diagnostics

Oracle Enterprise Manager Cloud Control 12c's JVM Diagnostics enables administrators to diagnose performance problems in Java application in the production environment. By eliminating the need to reproduce problems, it reduces the time required to resolve these problems. This improves application availability and performance.

There is a correlation between JBoss target and the JVMD/JVM target and this allows administrators to navigate to the JVM in context of a JBoss server.

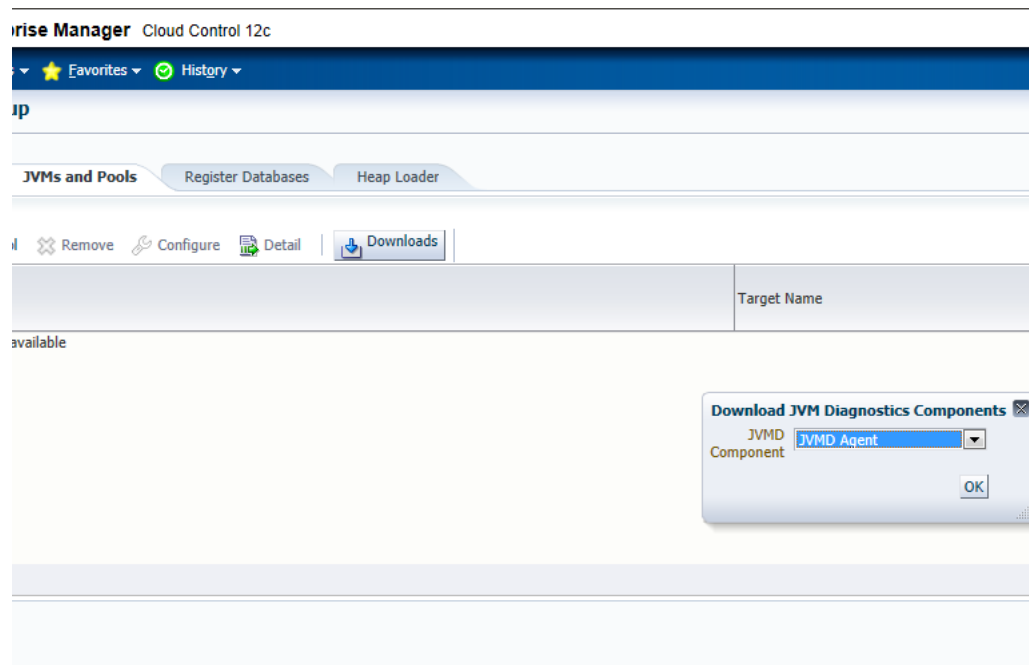
25.5.1 Deploying JVMD on JBoss

To deploy JBoss on JVMD, complete the following:

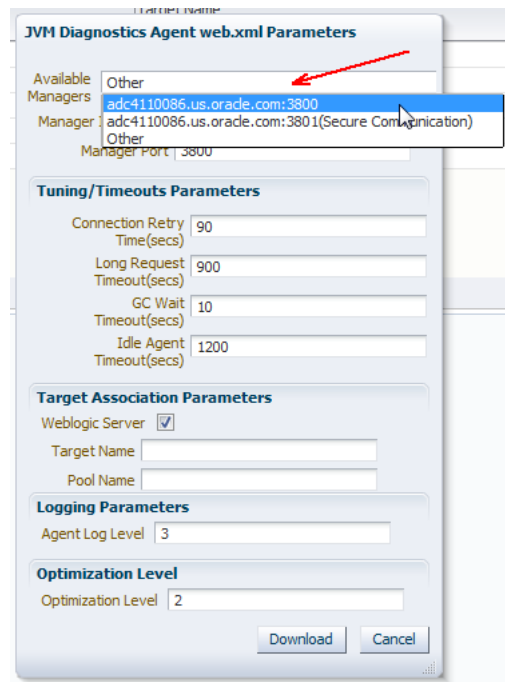
1. Log in to Enterprise Manager

2. From the **Setup** menu, select **Application Performance Management**, then select **JVM Diagnostic Engines**.
3. Click **Configure**.
4. From the JVMs and Pools tab, click **Downloads**.

Figure 25–1 Download JVM Diagnostics Components



5. From the **Download JVM Diagnostics Components** window, select **JVMD Agent**, then click **OK**.
6. From the JVM Diagnostics Agent web.xml Parameters window, select the appropriate JVMD manager from the Available Managers list.

Figure 25–2 JVM Diagnostics Agent web.xml Parameters

7. Uncheck the WebLogic Server checkbox, enter a pool name, then click **Download**.
This will download the jamagent.war file. You need to copy this file to the target machine (the machine the JBoss is running).
8. Copy the script contained at the following location to the same location on the target machine where the jamagent.war. file was copied:


```
https://stbeehive.oracle.com/content/dav/st/NonOracle%20MW/Public%20Documents/makeJAMagents.sh
```
9. Run the script passing the names of the JBoss instances (or whatever names you give to the JVM of each JBoss server). The name has to be unique for each server.
This script will create a separate jamagent WAR for each name you pass to this script.
10. Deploy each of these created WARs on the respective JBoss server. Typically you can use the admin console or copy the WAR file to the deploy folder.

If the hot deployment is not enabled, you may have to restart the JBoss server. Once the WAR is deployed successfully, you will see the respective JVM target on the Middleware page.

25.6 Troubleshooting

This section describes the errors you may encounter while deploying and using JBoss and the workaround steps you can follow to resolve each of them. The following topics are covered:

- [Target Status is Down](#)
- [Target Status is Pending](#)
- [Useful Workarounds](#)

■ [Troubleshooting Resources](#)

Target Status is Down

If the target status is DOWN after discovery, check the Monitoring properties page for JBoss Application Server and verify the following:

- Local discovery: only JBoss home is present
- Remote discovery: only Library path is present

Target Status is Pending

If the target status is PENDING (due to a metric collection error) after discovery, ensure that no other application server is being monitored using the same agent (such as Weblogic).

The following are the various log locations:

- JBoss server logs: \$JBOSS_HOME/server/<config_mode>/log
- OMS logs: emoms.trc (under \$OMS_HOME)
- Agent logs: \$AGENT_STATE_DIR/sysman/log

Useful Workarounds

In the case of JBoss discovery failure, provide the library path along with the install home and try again.

For discovery related issues, manually run the discovery script to check the output, which should look similar to the following:

```
java -Doracle.home=/ade/ankprasa_gc2/oracle/emagent/gcagent/agent
-cp"/ade/ankprasa
gc2/oracle/emagent/gcagent/agent/lib/xmlparserv2.jar:/ade/ankprasa
gc2/oracle/emagent/gcagent/agent/jlib/emConfigInstall.jar:/ade/ankprasa
gc2/oracle/emagent/gcagent/agent/sysman/jlib/log4j-core.jar:/ade/ankprasa
gc2/oracle/emagent/gcagent/agent/modules/oracle.http_client
11.1.1.jar:/scratch/ankprasa/jboss/jboss-4.2.2.GA/client/jbossall-client.jar:/scr
tch/ankprasa/jboss/jboss-4.2.2.GA/server/all/lib/jbosssx.jar:/ade/ankprasa
gc2/oracle/emagent/gcagent/agent/../../plugins/oracle.sysman.emas.discovery.plugi
_12.1.0.4.0/archives/em-as-thirdparty-discovery.jar"
oracle.sysman.emas.thirdparty.discovery.jboss.JBossDiscovery 1099
adc6140210.us.oracle.com "EMPTY"
```

Troubleshooting Resources

Useful troubleshooting information can also be found by checking the following:

- Monitoring Configuration page
- Targets.xml on the agent
- OMS and Agent logs
- Agent metric browser
- JBoss JMX Console
- JBoss server logs
- The JBoss Management Wiki

The wiki can be found at the following location:

<http://aseng-wiki.us.oracle.com/asengwiki/display/ASDevEM/JBoss+Management>

Part X

Managing Oracle Data Integrator

The chapter in this part describes how you can configure and monitor Oracle Data Integrator.

This part contains the following chapter:

- [Chapter 26, "Configuring and Monitoring Oracle Data Integrator"](#)

Configuring and Monitoring Oracle Data Integrator

Oracle Data Integrator (ODI) provides a fully unified solution for building, deploying, and managing complex data warehouses or as part of data-centric architectures in an SOA or business intelligence environment. In addition, it combines all the elements of data integration - data movement, data synchronization, data quality, data management, and data services - to ensure that information is timely, accurate, and consistent across complex systems.

An ODI domain contains the ODI components that can be managed using Enterprise Manager. An ODI domain contains:

- One or several Oracle Data Integrator Console applications. An Oracle Data Integrator Console application is used to browse Master and Work repositories.
- One or several Run-Time Agents attached to the Master Repositories. These agents must be declared in the Master Repositories to appear in the domain. These agents may be Standalone Agents or Java EE Agents.

The Repositories and Agent pages display both application metrics and information about the Master and Work Repositories.

This chapter describes the set of typical tasks you need to perform when managing Oracle Data Integrator:

- [Setting Up ODI](#)
- [Accessing Oracle Data Integrator](#)
- [Managing Oracle Data Integrator](#)
- [Monitoring ODI Components](#)
- [Monitoring Run-Time Agents](#)
- [Configuring Oracle Data Integrator Console](#)

26.1 Setting Up ODI

Before you start managing ODI with Enterprise Manager, you must:

- Deploy the Oracle Management Agent
Oracle Management Agents must be installed on the database hosting the ODI repositories. Optionally, an Oracle Management Agent can also be installed on a machine hosting an ODI Agent.

See Installing the Oracle Management Agent in the *Oracle Enterprise Manager Cloud Control Basic Installation Guide*

- Discover ODI Target

ODI target discovery is part of the Fusion Middleware discovery. Two target types are available: ODI Standalone Agent and ODI Java EE Agent. See [Chapter 2, "Discovering Middleware Targets"](#) for additional information.

Along with ODI specific targets it also discovers domain and admin server.

- Discover the Database

Each Oracle Database Server needs to be discovered because more than one database could be hosting the ODI repositories.

See Performing Additional Configuration Tasks in the *Oracle® Enterprise Manager Cloud Control Advanced Installation and Configuration Guide*

All the operations are available out-of-box in Enterprise Manager.

26.2 Accessing Oracle Data Integrator

To access Oracle Data Integrator (ODI), perform the following steps:

1. From the **Targets** menu on Enterprise Manager, select **Middleware**.
2. In the Middleware Features menu, select **ODI Home**.
3. The ODI Home tab appears.

The regions on this page include:

- Master Repositories Health
- Work Repositories Health
- Agents Health
- Data Servers Health
- Search

26.3 Managing Oracle Data Integrator

The following summarizes the tasks to perform when managing Oracle Data Integrator components.

1. [Starting, Stopping, and Restarting Oracle Data Integrator Agents](#)
2. [Managing Agent Status and Activities](#)
3. Search Sessions and Load Plan Executions.
Navigate to the LPE/Sessions tab under ODI Home Page.
4. [Viewing Log Messages](#)

26.3.1 Starting, Stopping, and Restarting Oracle Data Integrator Agents

To start, stop, and restart Oracle Data Integrator agents, use the Enterprise Manager Fusion Middleware Control Console.

26.3.2 Managing Agent Status and Activities

You can manage the agent status and monitor its activities.

The steps for this process are:

1. Click the target link corresponding to your Java EE Agent either in the target navigation pane or in the ODI Master Repository Home Page. The Java EE Application Page for this agent appears.
2. From the **Agent Page** menu, select **Monitoring** then select **Performance Summary**.
3. The Performance Summary page opens. You can view and customize the displayed metrics and charts. How?

26.3.3 Viewing Log Messages

You can view log messages of Java EE agents in Enterprise Manager Cloud Control.

The steps for this process are:

1. Navigate to the ODI Agent Home page.
2. From the **ODI Agent** menu, select **Logs** then select **View Log Messages**.

You can filter the displayed log messages, for example by date range and message type and search for a search term in the message.

To configure the log configuration settings, select **Logs** then select **Log Configuration** from the **ODI Agent** menu.

26.4 Monitoring ODI Components

Oracle Data Integrator (ODI) provides a fully unified solution for building, deploying, and managing complex data warehouses or as part of data-centric architectures in a SOA or business intelligence environment.

In addition, ODI combines all the elements of data integration - data movement, data synchronization, data quality, data management, and data services - to ensure that information is timely, accurate, and consistent across complex systems.

Click the links for information regarding each of the tabs on this page.

- [Viewing General Health of ODI Components](#)
- [Viewing ODI Agents](#)
- [Viewing Repository Details](#)
- [Viewing Load Plan Executions/Sessions Statistics](#)

26.4.1 Viewing General Health of ODI Components

ODI provides statistics summarizing the general health of the following components:

- Master Repositories
- ODI Agents
- Work Repositories
- Data Servers

Master Repositories Health

This section reports on the:

- Number of master repositories that are either up or down. Click the number for a list of the repositories.
- Master Repositories with Incidents. Click the number to find out which repositories have incidents.

To resolve issues reported in this section:

- If the ODI repositories are down, act based on the statuses by either bringing up the repositories which are down or troubleshoot why they are down and resolve the issues.
- If there are any repositories which are undiscovered, discover the repositories in Cloud Control.
- If there are any repositories with alerts, identify the root cause for those alerts and resolve the issues.

ODI Agents Health

This section reports on the:

- Number of Agents that are either up or down. Click the number for a list of the Agents.
- Number of Agents that are not discovered as targets in Enterprise Manager. Click the number for a list of the Agents that have not been discovered.
- Number of Agents with Incidents. Click the number to find out which repositories have incidents.

To resolve issues reported in this section:

- If the Agents are down, act based on the statuses by either bringing up the Agents which are down or troubleshoot why they are down and resolve the issues.
- If there are any Agents which are undiscovered, discover the Agents in Cloud Control.
- If there are any Agents with alerts, identify the root cause for those alerts and resolve the issues.

Work Repositories Health

This section reports on the:

- Number of work repositories that are either up or down. Click the number for a list of the repositories.
- Number of work repositories that have not been discovered. Click the number of a list of the work repositories that have not been discovered.
- Work Repositories with Incidents. Click the number to find out which repositories have incidents.

To resolve issues reported in this section:

- If the ODI repositories are down, act based on the statuses by either bringing up the repositories which are down or troubleshoot why they are down and resolve the issues.
- If there are any repositories which are undiscovered, discover the repositories in Cloud Control.
- If there are any repositories with alerts, identify the root cause for those alerts and resolve the issues.

Data Servers Health

This section reports on the:

- Number of data servers that are either up or down. Click the number for a list of the servers.
- Number of data servers that have not been discovered. Click the number of a list of the data servers that have not been discovered.
- Data Servers with Incidents. Click the number to find out which data servers have incidents.

To resolve issues reported in these sections:

- If the data servers are down, act based on the statuses by either bringing up the data servers which are down or troubleshoot why they are down and resolve the issues.
- If there are any data servers which are undiscovered, discover the data servers in Cloud Control.
- If there are any data servers with alerts, identify the root cause for those alerts and resolve the issues.

Load Plan Executions/Sessions Quick Links

Click the links for details for each of the following:

Link Name	Description
Sessions With Errors	Number of sessions in error for this agent.
Sessions With Error Records	Number of sessions with error records for this agent.
LPEs With Errors	Number of load plan executions in error for this agent.
LPEs With Error Records	Number of load plan executions with error records for this agent.

26.4.2 Viewing ODI Agents

This page has a table containing the following columns:

Agents

The Agents section displays the Agents declared in the Master Repository.

Element	Description
Name	Displays the name of the Agent. Select an Agent to display the corresponding Agent Home page.
Discovery Status	A blue tick indicates that the ODI Agent is discovered as a custom target in Enterprise Manager. A clock indicates that the ODI Agent is not discovered as a custom target in Enterprise Manager.
View Performance	Click the eye glass icon to view the performance data of the Agent. The metrics include: <ul style="list-style-type: none"> ■ Maximum number of allowed sessions ■ Maximum number of allowed threads ■ Count of active sessions ■ Count of active threads
Active Sessions	Displays the number of active sessions.

Element	Description
Master Repository	A check mark acknowledges that this is a Master Repository.
Version	Displays the version and date of the Agent.
Response Time (ms)	Repository database response time in milliseconds.
User Defined Alerts	Displays the number of Critical and Warning alerts. Click the number to view the alerts in the Incident Manager page.

Search Agents

Use this region to search for agents for all Java EE and Standalone agents.

The latest specified search criteria are always retained. Specify a new criteria and click **Search** to see the updated results. Or, click **Reset** to reset the search form (you must still click **Search** to see the updated results). Note that the search criteria are reset each time you log out or navigate away from all the tabbed pages of the Oracle Data Integrator Console application.

Element	Description
Master Repository	Select the Master Repository.
Execution Agent	Select an Agent from the drop-down list. You can also select All to list all the Agents.
Agent Status	Select the status of the agent: Up, Down, All.
Discovery Status	Select the status of the agent: Discovered, Not Discovered, All.

26.4.3 Viewing Repository Details

The Repositories tab displays the following regions:

- [Repositories](#)
- [Database Details](#)
- [Tablespace Details](#)
- [Search Repositories](#)

Note: Credentials have to be selected for this region to display.

Repositories

Element	Description
Name	Displays the name of the Work Repository.
Status	Displays the status of the Work Repository database. <ul style="list-style-type: none"> ■ Up (green arrow): on ■ Down (red arrow): off ■ Not configured: the Repository is declared in the Master Repository but no connection to this Work Repository is declared in Oracle Data Integrator Console.
Host	Displays the name of the host on which the repository resides.

Element	Description
Port	Displays the port of the host on which the repository resides.
SID	Displays the system identifier of the repository.
Version	Displays the Repository version.
Response Time (ms)	Repository database response time in milliseconds.
External ID	Displays the ODI-specific unique identifier for the repository.
Incidents	Displays the number of incidents associated with this repository: Critical or Warning.
Schema Name	Displays the name of the schema associated with this repository.
LPE/Sessions Tablespace	Total Rows and Segment Size in GB.
Purge	Click the icon to purge the ODI logs. A separate browser window appears where you provide the required information, for example, purge type.

Database Details

By looking at the database details, you have a clear picture of how your database is performing. For example, if the Database Table Space is reaching near full, the Database Administrator can look at extending the table space.

In addition, by taking a look at the database performance chart, Throughput and Wait bottlenecks sections, the Database Administrator can recommend fine tuning the database.

- Wait Bottlenecks

This section provides the following statistics: Average Instance (CPU%), Active Sessions Waiting I/O, and Active Sessions Waiting Others.

- Throughput

This section provides the following statistics: Number of Transactions per second, Physical Writes per transaction, Physical Reads per transaction, and User Commits per transaction.

- Performance

This section provides usage information for CPU, I/O Wait, and Others for the active sessions.

Note: Credential and repository have to be selected for this region to display.

Tablespace Details

This section provides the growth rate for the tablespace by providing Space Used and Space Allocated statistics. Based on the information, you can decide whether to archive or purge the database data, or extend the tablespace.

Search Repositories

Use this region to search for repositories for all master and work repositories.

The latest specified search criteria are always retained. Specify a new criteria and click **Search** to see the updated results. Or, click **Reset** to reset the search form (you must still click **Search** to see the updated results). Note that the search criteria are reset each time you log out or navigate away from all the tabbed pages of the Oracle Data Integrator Console application.

Element	Description
Repository Type	Select the Repository type: Master Repository, Work Repository, All.
Repository Name	Enter the name or a part of the Repository name.
Repository Status	Select the status of the Repository: Up, Down, All.
Discovery Status	Select the status of the repository: Discovered, Not Discovered, All.

26.4.4 Viewing Load Plan Executions/Sessions Statistics

The Load Plan Executions/Sessions displays the execution details of the Load Plan Executions and Sessions executed by the Agent.

Expand a session and review the Steps and Tasks information. For example if an ODI Interface was executed, you can review each task that this interface executed, view the generated code, and drill down to the database execution details.

Note: Oracle Database Diagnostics and Tuning Packs are required to use the SQL Drill Down feature.

Element	Description
Name	Displays the name of the Load Plan Execution or Session.
Execution ID	Load Plan Execution or Session identifier. Every time a Load Plan is executed, a new Load Plan Execution with a unique identifier is created.
Status	Displays the status of the Load Plan Execution run or Session executed. <ul style="list-style-type: none"> Running: The Load Plan Execution/Session is currently running. Done: The Load Plan Execution/Session has terminated successfully. Waiting: The Load Plan Execution/Session is waiting to be executed. Error: The Load Plan Execution/Session has terminated due to an error. Warning: The session has terminated successfully but erroneous rows were detected by an interface during flow control. Queued: The session is waiting for an Agent to be available for its execution.
Started On	Start date and time of the Load Plan Execution/Session run.
Updated On	Displays the last updated date of the Load Plan Execution/Session.
Execution Time	Displays how long it took the Load Plan Execution/Session to run.
Message	Displays the error message of the Load Plan Execution/Session run, if any.
Error Records	Displays the number of error records.
Execution Type	Displays the type LoadPlan or Sessions, for example, Scenario.
Work Repo Name	Displays the name of the Work Repository into which this Load Plan/Session run execution information is stored.
Agent Name	Displays the name of the agent on which the Load Plan Execution/Session ran.
ODI User	Displays the name of the ODI user who started the execution.

Search Sessions/LPEs

Use this region to search for sessions and load plan executions for all master and work repositories.

The latest specified search criteria are always retained. Specify a new criteria and click **Search** to see the updated results. Or, click **Reset** to reset the search form (you must still click **Search** to see the updated results). Note that the search criteria are reset each time you log out or navigate away from all the tabbed pages of the Oracle Data Integrator Cloud Control application.

Element	Description
Master Repository	Select the Master Repository containing the session information.
Work Repository	Select the Work Repository containing the session information.
Execution Agent	Select the Agent used to execute the session.
Context	Select the session's execution context
Execution Type	Select Sessions, Load Plan Executions, or All.
Begin Date	Use the calendar icon to select a date at which to start the search for sessions. Only session started after this date will be returned
End Date	Use the calendar icon to select a date at which to end the search for load plan executions and sessions. Only load plan executions and sessions ended before this date will be returned.
User Name	Name of the ODI user who started the execution.
Status	Select All or narrow the search to display specific statuses: Error, Running, Done, Warning, or Waiting. For example, you can select to view only Running and Warning statuses.
Message	Error message of the Load Plan Execution/Session run.
Keywords	Type keywords to narrow the search. When using multiple keywords, use a comma to separate each keyword, do not include spaces. For example use: lpe1,lpe2.
Execution Name	Type the name of the load plan execution.
Error Records	Select All or narrow the search to display load plan executions and sessions With Error Records or Without Error Records.
Execution ID	Specific Load Plan Execution or Session identifier.

26.5 Monitoring Run-Time Agents

The Agents Home page enables you to monitor the Oracle Data Integrator run-time Agents. Both Standalone agents and Java EE Agents are ODI job executors. The difference between the two agents is that the Standalone Agents are non Java EE based and are managed through Oracle Process Manager and Notification Server (OPMN) or Common Administration Model (CAM) from Enterprise Manager. These run on the Jetty web server. Java EE Agents are Java EE based, that is, they are deployed on Oracle WebLogic Servers or IBM WebSphere. (IBM WebSphere is supported for ODI release 11.1.1.7 only).

Note: OPMN will be used to manage ODI standalone agents until ODI release 11.1.1.7. CAM will be used to manage ODI Standalone agents from ODI release 12c and later.

The Agent Home page is arranged in the following order:

- [General Info](#)
- [Target Incidents](#)
- [Load Balancing Agents](#)
- [Load](#)
- [LPEs/Sessions Execution Incidents](#)

General Info

The General Info region displays general information about this Agent.

Element	Description
Status	Displays the status of the Agent. <ul style="list-style-type: none"> Up (green arrow): on Down (red arrow): off
Response Time (ms)	Displays the repository database response time in milliseconds.
Agent Version	Displays the version of the Agent.
Host and Port	Displays the host (network name or IP address) of the machine where the Agent has been launched on and the port on which the Agent is listening.
Master Repository	Click to access the Database Performance page for the Master Repository.
Incidents	An event or a set of closely correlated events that represent an observed issue requiring resolution through (manual or automated) immediate action or root-cause problem resolution.
More Information	Click to browse for more information in the Oracle Data Integrator Console.

Target Incidents

The Target Incidents region displays notifications raised by the Agents attached to this Repository.

Element	Description
Severity	Seriousness of the incident. <ul style="list-style-type: none"> Fatal - Corresponding service is no longer available. For example, a monitored target is down (target down event). A Fatal severity is the highest level severity and only applies to the Target Availability event type. Critical - Immediate action is required in a particular area. The area is either not functional or indicative of imminent problems. Warning - Attention is required in a particular area, but the area is still functional. Advisory - While the particular area does not require immediate attention, caution is recommended regarding the area's current state. Clear - Conditions that raised the incident have been resolved.
ID	Incident ID.
Summary	Summary description of the incident.
Category	Classification of an incident, for example, Error.

Load Balancing Agents

The Load Balancing Agents region displays (if using ODI Load Balancing) the status and session metrics for the Agents declared as child Agents of the current Agent.

Element	Description
Name	Displays the name of the agent. This is the name you specified when you created the Agent in Oracle Data Integrator. Select an Agent to display the corresponding Agent Home page.

Element	Description
Status	Displays the status of the Agent. <ul style="list-style-type: none"> Up (green arrow): on Down (red arrow): off
Discovered	A blue tick indicates that the ODI Agent is discovered as a custom target in Enterprise Manager. Click the Agent name to access the ODI Console's Agent Detail Page. A clock indicates that the ODI Agent is not discovered as a custom target in Enterprise Manager. Click the Agent name to access the Enterprise Manager Agent Target Page.
Originating LPEs/Sessions	Displays the status of the LPEs and Sessions. <ul style="list-style-type: none"> Error - Number of sessions in error for this agent. Running - Number of sessions currently being executed by this agent. Done - Number of sessions completed by this agent. Warning - Number of sessions in warning state for this agent. Waiting - Number of sessions waiting to be executed. Queued: The session is waiting for an Agent to be available for its execution.
Avg Master Repo Response Time (ms)	Displays the master repository database response time in milliseconds.
Sessions	Maximum and Active number of sessions allowed on this Agent.
Threads	Maximum and Active number of threads allowed on this Agent.

Load

The Load region displays the number of connections supported by the Agent over a period of time.

Elements	Description
Maximum number of allowed sessions	Maximum number of sessions allowed on this Agent.
Maximum number of allowed threads	Maximum number of threads allowed on this Agent.
Count of active sessions	Number of active sessions on this Agent.
Count of active threads	Number of active threads on this Agent.

LPEs/Sessions Execution Incidents

The Load Plan Executions/Sessions Execution Incidents region displays notifications raised by the Agents attached to this Repository.

Element	Description
Severity	Seriousness of the incident. <ul style="list-style-type: none"> Fatal - Corresponding service is no longer available. For example, a monitored target is down (target down event). A Fatal severity is the highest level severity and only applies to the Target Availability event type. Critical - Immediate action is required in a particular area. The area is either not functional or indicative of imminent problems. Warning - Attention is required in a particular area, but the area is still functional. Advisory - While the particular area does not require immediate attention, caution is recommended regarding the area's current state. Clear - Conditions that raised the incident have been resolved.

Element	Description
ID	Incident ID.
Summary	Summary description of the incident.
Category	Classification of an incident, for example, Error.

26.6 Configuring Oracle Data Integrator Console

Oracle Data Integrator Console cannot be configured from Enterprise Manager Cloud Control. To make configuration changes you must use the Fusion Middleware Control Console. For information of how to configure Oracle Data Integrator, see *Oracle Fusion Middleware Developer's Guide for Oracle Data Integrator*.

However, you can configure Oracle Data Integrator Console from Cloud Control to define the linking between Cloud Control and Oracle Data Integrator Console.

By default, the fields on this page are populated with the Oracle Data Integrator Console host, the Oracle Data Integrator Console managed server port, and the default context root. If your Oracle Data Integrator Console must be accessed with a different configuration, you can change the configuration on this page.

The steps for this process are:

1. Navigate to the Agent home page.
2. From the **Agent Page** menu, select **ODI Console Administration** then select **Basic Configuration**.
3. To modify the current configuration, enter new values in the following fields and click **Apply**.
 - **Host:** The name of the server where your application is deployed
 - **Port:** The HTTP listener port number
 - **Context Root:** The Web application's context root
 - **Protocol:** The protocol for the connection

Click **Revert** to revert to the previous settings.

26.6.1 ODI Console Application Configuration Page

Use the ODI Console Application Configuration page to configure the connection to the Oracle Data Integrator Console Application.

This page displays the current configuration for accessing the Oracle Data Integrator Console application. These values are automatically set when the application is discovered by Enterprise Manager and are used to access Oracle Data Integrator Console from Enterprise Manager, for example when clicking **Browse**.

You can modify these values to access Oracle Data Integrator Console in a different way, for example to connect to Oracle Data Integrator Console by using a load balancer.

To modify this configuration, enter new values in the fields and click **Apply**. Click **Revert** to revert to the previous settings.

Element	Description
Host	Displays the name of the server where your application is deployed. If using SSO, enter the Oracle HTTP Server (OHS).
Port	Displays the HTTP listener port number. If using SSO, enter the port of the machine where Oracle HTTP Server 10g or 11g Webgate is installed.
Context Root	Displays the Web application's context root.
Protocol	Displays the protocol of the connection

26.6.2 Configuring an Oracle Data Integrator Domain

Installing and configuring components for an Oracle Data Integrator domain is described in the *Oracle Fusion Middleware Installation Guide for Oracle Data Integrator*.

Part XI

Using Application Dependency and Performance

The chapters and appendixes in this part provide information regarding the usage of Application Dependency and Performance (ADP).

The chapters and appendixes are:

- [Chapter 27, "Introduction to Application Dependency and Performance"](#)
- [Chapter 28, "Exploring Application Dependency and Performance"](#)
- [Chapter 29, "ADP Methodology"](#)
- [Chapter 30, "Frequently Asked Questions About Application Dependency and Performance"](#)
- [Appendix A, "ADP Configuration Directories and Files"](#)
- [Appendix B, "Support Matrix for Application Dependency and Performance"](#)

Introduction to Application Dependency and Performance

The Application Dependency and Performance (ADP) pages within Enterprise Manager Cloud Control analyze Java EE, SOA, and Portal applications to capture the complex relationships among various application building blocks in its Application Schema model - the core of the Oracle intelligent platform.

Using ADP you can:

- Monitor performance of applications deployed in the following type of Servers:
 - Oracle SOA Suite 11g
 - Oracle Service Bus
 - Java EE
 - Oracle WebLogic Portal
- Have visibility into components defined by way of metadata within a framework (for example, components within a composite) with deep dive visibility, where available
- View static relationships defined between components and services, such as OSB business and proxy services, and SOA services and references

This chapter includes the following:

- [Overview](#)
- [Architecture](#)

27.1 Overview

Using the insights stored in Application Schema, ADP is able to deliver an Application Service Management (ASM) environment that self-customizes out-of-the-box, evolves with change, minimizes expert involvement, and delivers a holistic, service-oriented view across heterogeneous environments. ADP enables an enterprise to more efficiently manage distributed applications, attain management agility, and lower total cost of ownership.

See the following sections:

- [Managing Complex Java EE, SOA, OSB, and Portal Applications](#)
- [Delivering a Service-Oriented View Across Environments](#)
- [Avoiding Involvement of Java EE, SOA, OSB, Portal, and Application Experts](#)

- [Eliminating Repetitive Do-It-Yourself \(DIY\) Manual Processes](#)
- [ADP Solution](#)

27.1.1 Managing Complex Java EE, SOA, OSB, and Portal Applications

Today's Java EE, SOA, OSB, and Portal applications enable enterprises to deliver mission-critical business functions to key constituencies - most often their customers, partners, and employees. These composite applications are assembled from many different Java EE components and exposed services distributed across a heterogeneous environment.

To be effective at managing today's complex, distributed Java EE, SOA, OSB, and Portal applications across a heterogeneous environment, enterprises must adopt an intelligent ASM platform with the following characteristics:

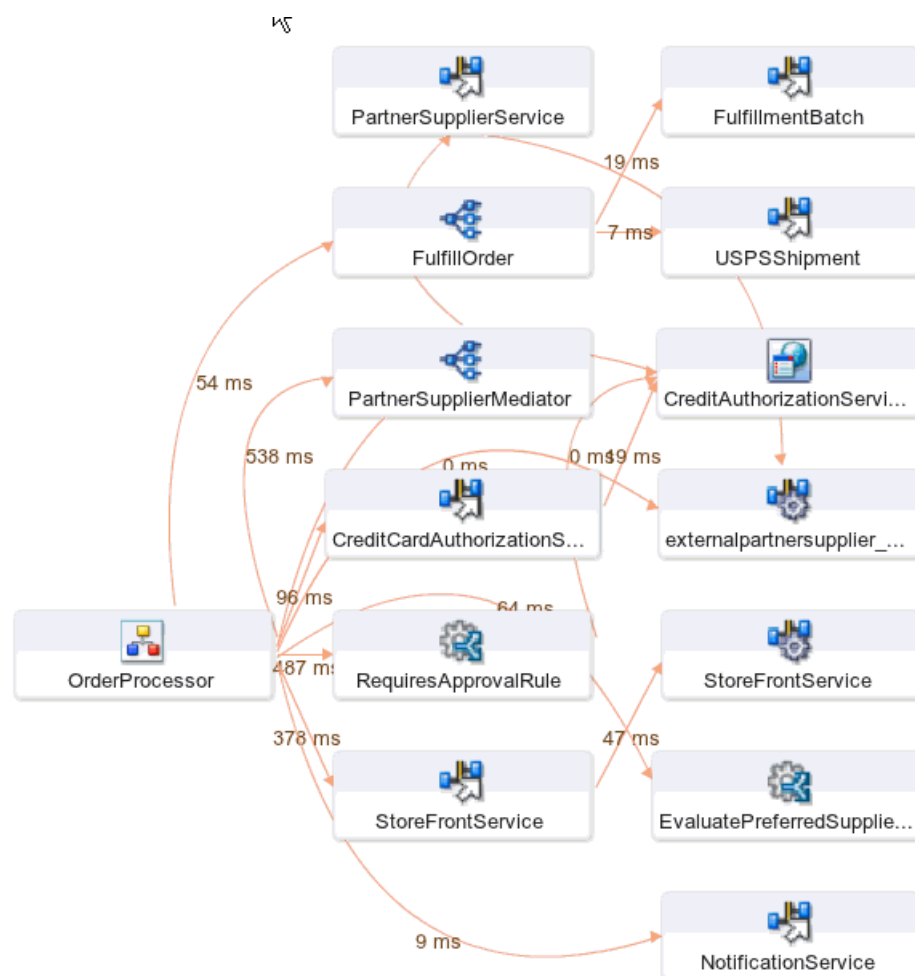
- Provides holistic, service-oriented views across heterogeneous environments
An intelligent ASM platform must provide high-level service-oriented metrics that map to low-level technology-centric metrics. These measurements must be organized in a service-oriented fashion to deliver a unified, holistic view of the numerous interconnected application components deployed across heterogeneous environments.
- Requires minimal Java EE, SOA, OSB, Portal, and application expertise
An intelligent ASM platform must have the ability to capture complex relationships among various interconnected components of today's Java EE, SOA, OSB, and Portal applications. This ability can help minimize reliance on Java EE, SOA, OSB, Portal, and application experts for setting up and maintaining effective APM environments.
- Eliminates repetitive DIY manual processes
An intelligent ASM platform must eliminate repetitive DIY manual processes by delivering the ability to self-customize out-of-the-box and evolve with change. Elimination of these repetitive DIY manual processes is the only way to deal with rising complexity and rapid rate of change with ease.

27.1.2 Delivering a Service-Oriented View Across Environments

Today's mission-critical business functions are powered by Java EE, SOA, OSB, and Portal applications that comprise numerous interconnected components deployed across highly distributed environments. To manage these applications effectively, enterprises must first gain an understanding of the complex relationships among the business functions, associated interconnected components, and the underlying runtime environments. To enable clear and accurate understanding, IT organizations need holistic, service-oriented views that span across heterogeneous environments.

Furthermore, appropriate rendering of these views enables users at different levels of the organization to collaborate with each other and do their respective jobs more efficiently.

Figure 27–1 Application Dependency and Performance Topology View in Enterprise Manager Cloud Control



Application Schema Navigation provides efficient ways for you to access relevant information using techniques like hierarchical traversal, architecture model navigation, string queries, drill down, drill out and more.

27.1.3 Avoiding Involvement of Java EE, SOA, OSB, Portal, and Application Experts

To manage Java EE, SOA, OSB, and Portal performance effectively, IT organizations must adopt an intelligent platform like ADP that requires minimal expertise to set up and maintain. Unlike conventional APM toolkits, ADP does not rely on human expertise to set up and maintain customized APM environments. Instead, ADP uses a unique model-driven approach that leverages the information stored in its Application Schema model to keep the involvement of experts to the minimum. ADP's unique ability to self-customize out-of-the-box and evolve with change makes it the perfect solution for managing not only custom enterprise applications, but also applications developed by external parties.

27.1.4 Eliminating Repetitive Do-It-Yourself (DIY) Manual Processes

Based on a unique model-driven approach, ADP eliminates repetitive DIY manual processes. To achieve this level of self-customization and continuous change adoption, ADP uses its AppsSchema modeling technology to perform the critical task of

analyzing application structure and infrastructure configuration. After capturing these insights in the Application Schema model, ADP leverages this information to establish a fully customized ASM environment. To keep this environment up-to-date, ADP continuously updates the Application Schema model as new applications are deployed and changes are applied. ADP's unique ability to self-customize out-of-the-box and evolve with change enables fast time-to-value, low total-cost-of-ownership (TCO), and maximal return-on-investment (ROI).

27.1.5 ADP Solution

Oracle provides the industry's first intelligent ASM platform for Java EE, SOA, OSB, and Portal. Unlike conventional APM toolkits, ADP analyzes these applications and captures complex relationships among various application building blocks in its Application Schema model - the brain of this intelligent ASM platform.

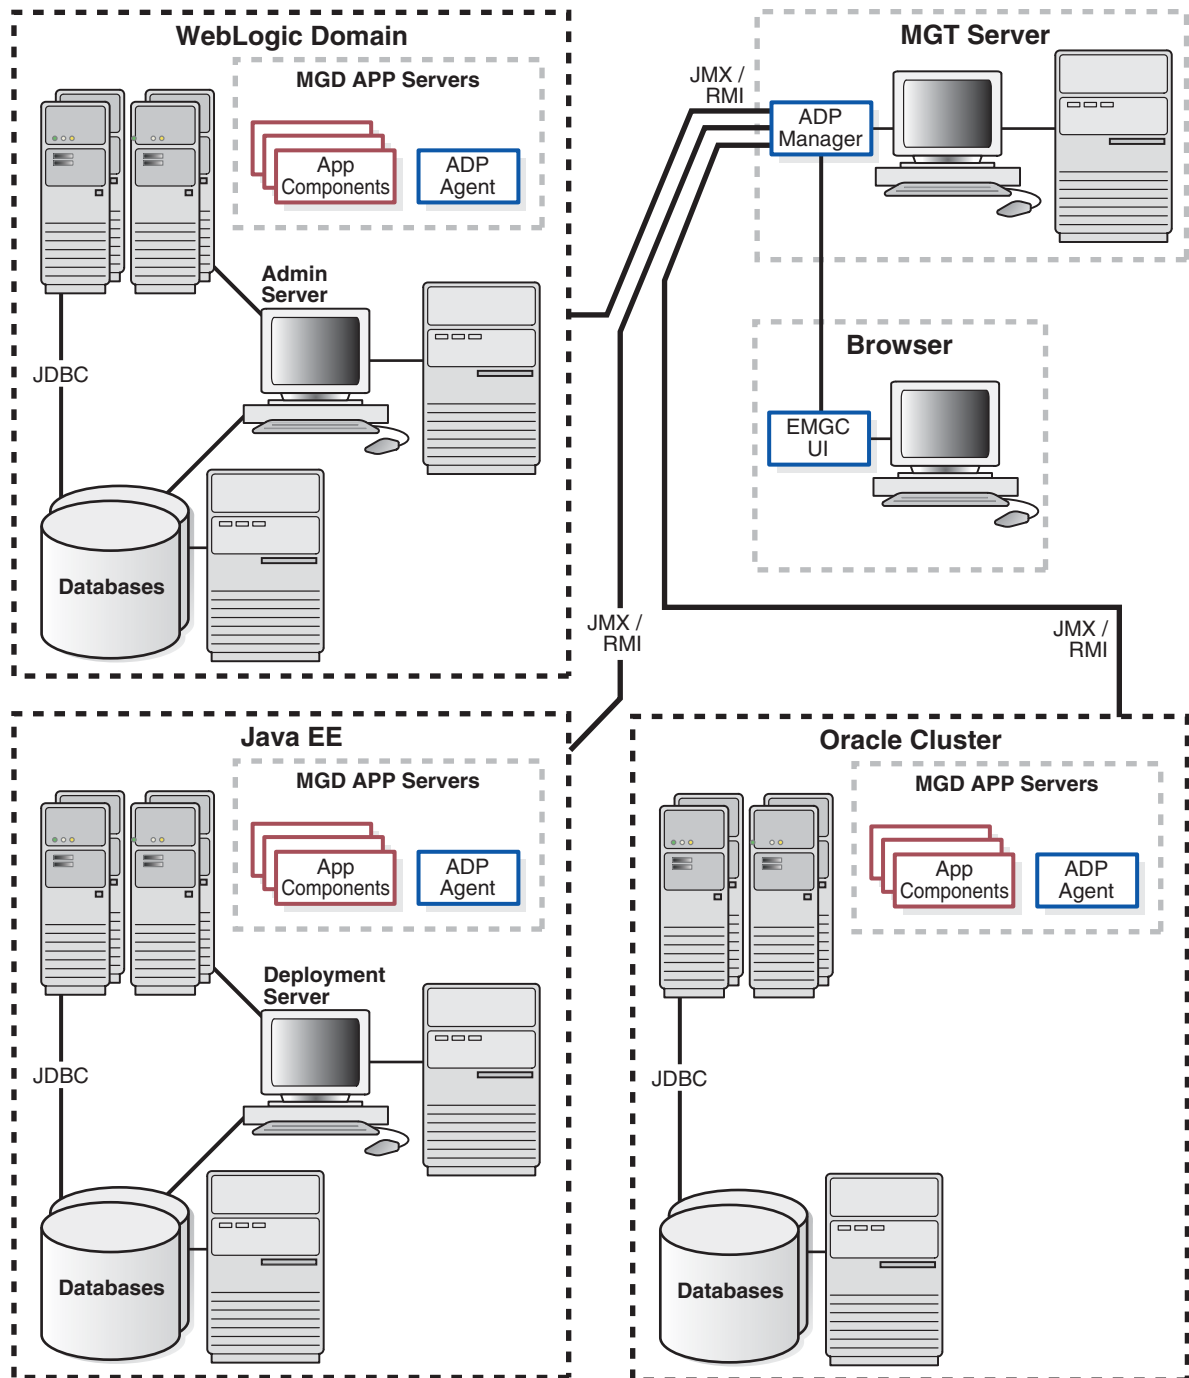
Using the insights stored in the Application Schema model, ADP is able to deliver an ASM solution that self-customizes out-of-the-box, evolves with change, minimizes expert involvement, and delivers a holistic, service-oriented view across heterogeneous environments. Adopting an intelligent platform such as Oracle will enable the enterprise to more efficiently manage distributed applications, attain management agility, and lower total cost of ownership.

27.2 Architecture

ADP employs a multi-tier, fully distributed, configurable architecture to provide the scalability and flexibility to meet the changing needs of enterprise deployments.

ADP operates as a service on the machine and automatically begins running when the machine first boots, and remains on perpetually. ADP is typically installed on its own machine and dedicated to monitor a group of managed application servers.

Figure 27-2 ADP Topology



The following core components are deployed to form the ADP ASM system.

27.2.1 ADP Java Agents

ADP Java Agents are the data collectors of the ADP ASM system. ADP Java Agents are deployed to all managed application servers to perform a series of tasks including collecting performance managements, tracking contextual relationships, and summarizing data in real-time while introducing as little overhead as possible. At the expiration of the predefined aggregation interval, these agents forward the summarized data to ADP for additional analysis. For various Java EE platforms such

as Oracle SOA Suite and Oracle WebLogic, ADP leverages their deployment infrastructures to quickly deploy the ADP Java Agents to all application servers.

27.2.2 ADP Manager

ADP Manager is the core analytical engine of the ADP ASM system. In real-time, ADP Manager performs complex mathematical modeling and statistical calculations with summarized data from all ADP Java Agents. ADP Manager can be configured with a backup to provide higher level of availability.

27.2.2.1 ADP Manager and High Availability

Although the ADP Manager does not have high availability (HA) built into it, administrators can have a backup ADP Manager installed on a separate machine; this backup ADP Manager points to the same database but is disabled. If the production ADP Manager fails, the backup ADP Manager can then be enabled against the same database. The backup ADP Manager then rediscovers the application after the agent is redeployed from the backup ADP Manager, to the managed resources in order to synchronize them. All metrics are preserved, assuming the model does not change in the short time frame it takes to bring the backup ADP Manager online.

The key with this backup procedure is to backup the database that ADP uses as its repository in order to preserve the historical data. On the modeling side, the backup ADP Manager has to rediscover the application which should happen automatically, as long as the resources are configured and the new agent has been deployed.

If historical data preservation is not a necessity, users can simply have another ADP Manager and database and swap agents reducing the backup effort considerably.

27.2.3 ADP User Interface

The ADP User Interface (ADP UI) is the primary user interface for ADP users. Users can use ADP UI to set Service Level Objectives (SLOs), analyze monitoring data, and more. The ADP UI is fully configurable. To access ADP:

1. From the **Targets** menu, select **Middleware**.
2. From the **Middleware Features** menu on the Middleware page, select **Application Dependency and Performance**.

Exploring Application Dependency and Performance

This chapter examines the following:

- [Exploring the User Interface](#)
- [Exploring the Monitoring Tab](#)
- [Exploring the Configuration Tab](#)
- [Exploring the Registration Tab](#)

28.1 Exploring the User Interface

This section explores the ADP User Interface. Topics include:

- [Accessing ADP](#)
- [General ADP UI Elements](#)
- [Drill Down in Operational Dashboard](#)
- [Time Frame](#)
- [Display Interval](#)
- [Graphs and Data Items](#)
- [Custom Metrics](#)
- [Functional View](#)
- [Topology View](#)
- [Architecture View](#)
- [Metric Types](#)

28.1.1 Accessing ADP

To access the Enterprise Manager Application Dependency and Performance (ADP) feature, do the following:

1. From the **Targets** menu, select **Middleware**.
2. From the **Middleware Features** menu on the Middleware page, select **Application Dependency and Performance**.

28.1.2 General ADP UI Elements

ADP UI consists of the following core components:

- Navigation Pane (left)

There are three types of workspaces in the ADP navigation pane: *Monitoring*, *Configuration*, and *Registration*. In the Monitoring workspace, you can navigate the managed environment and monitored applications in a tree.

- Use the Monitoring workspace to traverse the ADP tree model and identify abnormal activities.
- Use the Configuration workspace to create, modify, and review various configuration settings for ADP.
- Use the Registration workspace to enable and disable request monitoring.

- Main Display Window (right)

As you navigate through the ADP tree model and configuration categories, detailed performance information and configuration settings are displayed in the Main Display Window. You can refresh the Main Display Window at anytime by clicking the Refresh icon.

28.1.3 Drill Down in Operational Dashboard

The Operational Dashboard displays the health indicators for various key entities in the managed environment. ADP uses traditional traffic light colors to represent the health of these various key entities.

For each component, ADP uses the following health indicators to provide a comprehensive view. These health indicators are:

- Performance

The performance health indicator depicts the relative responsiveness of the monitored entity to the configured threshold.

- Availability

The availability health indicator informs you to what extent a particular entity is available to service requests. The Availability arrow explains the availability of the particular entity: Red down arrow means the entity is not available whereas the Green up arrow means the entity is available.

- Errors

The errors health indicator informs you if the number of errors and exceptions encountered by this entity are approaching or violating the configured threshold. If there is any errors in the server, the check mark is in red.

- Load

The load health indicator depicts how many operations have been performed and requests have been served by a particular entity.

ADP is aware of clusters. As such, these indicators display overall health of a particular entity across the entire cluster.

28.1.4 Time Frame

In ADP, you can specify the length of the time the window information is to be displayed. To specify the length of this time window, select the appropriate length in the Time Frame list. The following Time Frame values are available:

- 1 hour
- 2 hours
- 4 hours
- 8 hours
- 12 hours
- 24 hours

Note: The ADP default data collection interval is 60 seconds. As you adjust the data collection interval, ADP automatically adjusts the display time frames.

ADP automatically adjusts information displayed to fit the specified time window. You can drill down to see detailed performance information for a specific range of time.

For example, visualize the drill down process with two screen shots of the same graph with different Time Frames of the average response time for Portal campaign. The first graph has a Time Frame of 24 hours. The second graph has a Time Frame of 1 hour. By increasing the granularity of the Time Frame, you are performing a drill down operation.

For example, an IT Operations staff noticed abnormally high response time with Portal campaign subsystem. The person decided to investigate further to evaluate the extent of the problem. By changing the Time Frame from 24 hours to 1 hour, this user is able to see that between 14:17 and 14:18, the Portal campaign response time jumped from an average of 1000 milliseconds to 5000 milliseconds. While the problem did not persist, it may warrant additional investigation.

28.1.5 Display Interval

Display Interval, located above the Main Display window, indicates the start and end time for the data displayed in the Main Display Window. Display Intervals change as you change the following settings:

- Time Frame
- Interval Context
- Turning Off Time Frame Limitation

28.1.5.1 Time Frame

When you select a new Time Frame, the Display Interval automatically changes to fit the selected Time Frame. For example, if you were to change the Time Frame from 1 hour to 2 hours, the Start value of the Display Interval changes.

28.1.5.2 Interval Context

Display Interval can also be changed by setting the Interval Context. The settings for the Interval Context are:

- End Time Is Current System Time

The default Interval Context for ADP is to use the current system time as End value for the Display Interval. In this default setting, you have a sliding Display Interval and can see the latest performance information in the Main Display Window.

- End Time Is Fixed

You can also change the Interval Context setting to use a fixed time as the End value for the Display Interval. By selecting the fixed Interval Context, you can create a fixed time window to display performance data. The fixed time window is particularly useful for performing analytical tasks.

- Date/Time Selector

When you select to fix the End time for the Interval Context, the ADP UI enables a pair of Date/Time Selectors to allow you to set Start or End values for the Display Interval. Click the icon next to the Start and End times to open up the Date/Time Selector.

The Date/Time Selector allows you to set a specific Display Interval to fit your needs. Additionally, the Date/Time Selector enables ADP to compare current performance trends with historical data.

Note: Changing the start and end time do conceptually different things. Users are advised to always change their time frame by modifying the end time first, and then the start time. Changing the end time moves the window in time, whereas changing the start time increases/decreases the size of the window.

28.1.5.3 Turning Off Time Frame Limitation

To support the display of data for more than twenty four hours, ADP allows you to specify your own time frame for data display. To enable this feature, set **Interval Context** to **End time is fixed** and make sure the **Use time frame:** check box is unchecked. Turning off time frame limitation allows ADP to display eight days worth of data.

For example, when you specify the time frame to be eight days by adjusting the start and end times through the Date/Time Selector, ADP then adjusts its view to display eight days worth of data in a single graph. This feature allows you to perform trending analysis over time.

28.1.6 Graphs and Data Items

ADP displays performance information in various formats. Most commonly used display formats in ADP are tables and graphs.

- On graphs, you can gain more information about a data item by pointing the mouse over the interested item.
- Minimum and maximum response time measurements are stored in their database in addition to average response time measurements. The min and max metrics, if present, are displayed visually in the UI.
- For tables, you can perform a table sort by clicking the blue up/down arrow located in the column headings.

28.1.7 Custom Metrics

While ADP intelligently selects relevant performance metrics based on its Application Schema model, you can further customize the monitoring environment by configuring additional custom metrics. In addition, you can use custom metrics in problem diagnostic situations where additional visibility is needed to pinpoint problem root cause.

To configure a new custom metric:

1. Click **Custom Metric Configuration** on the Configuration tab
2. Click the **Create Custom Metric** button.
3. On the Custom Metric File page, either choose an existing custom metric file or provide the name of a new custom metric file. Click **Continue**. ADP walks you through the configuration process.

Custom Metric Configuration page includes the following fields, see [Table 28-1](#).

Table 28-1 Custom Metric Configuration Page

Field	Description
Name	This text field is for defining the display name for the custom metric.
Resource Name	This list is for defining the resource where the custom metric will be collected.
Class Name	This text field is for defining the fully qualified class name (package + class) associated with the custom metric.
Method Name	<p>This optional text field is for defining the method name associated with the custom metric.</p> <p>Usage:</p> <ol style="list-style-type: none"> 1. Type in * - ADP will instrument all methods. 2. Provide comma separated list of methods with no wildcards - ADP will create method entities and only instruments these methods in the agent. 3. Provide comma separated list of methods with wildcard prefixes or suffixes - ADP will instruct the agent to instrument the methods specified along with the wildcards. 4. Provide 1) or 2) preceded by "!" to create an excluded list - ADP will instruct the agent to instrument all methods in the class not defined in the exclude list. <p>Method field examples:</p> <ol style="list-style-type: none"> 1. methodA,methodB,methodC 2. ejb*,*context,methodA 3. !ejb*,*context,methodA

After you define the custom metrics, restart the application server instances associated with these customizations. The new custom metrics will be listed under the Custom Metrics node in the ADP navigation tree.

The newly configured custom metric provides class level performance data, for example invocation count and response time.

28.1.8 Functional View

Functional View is a type of Application Schema Visualization - a visual way for ADP to represent the information stored in its Application Schema model. This view is designed to help you understand how business functions are assembled with various functional building blocks. [Table 28-2](#) provides a list of functional views currently available in ADP.

Table 28–2 Functional View

Entity Type	Function View	Description
Process	Process Workflow View	This functional view depicts the workflow associated with the Oracle BPEL business process. It shows all the process nodes and the relationships among them.
Pageflow	Pageflow Functional View	This functional view depicts the logical flow associated with a JPS or Struts pageflow. It shows all the pages in a pageflow and the relationships among them.
OSB Proxy Service	Proxy Service Functional View	This functional view depicts the pipeline and stage flow associated with an OSB Proxy Service.

Depending on the type of entity selected, ADP displays different functional views. Right-click and select Display Functional View to bring up the relevant Functional View associated with the selected entity.

28.1.9 Topology View

Topology View is another type of Application Schema Visualization - a visual way for ADP to represent the information stored in its Application Schema model. This view is designed to help you understand how application environments are assembled with various applications, application server instances, and shared resources. This information helps you map composite applications and their building blocks to application server instances and share resources.

The highest level topology view graphically depicts domains, external resources, and shared database resources. The applications used in the following examples are CSS and MedRec demos.

For example, you can have a topology with two ADP managed resources, CSS Domain and MedRec Domain, two external resources, and a shared database resource. The lines connecting various entities in the Topology Views depict calls made from one entity to another. You can get more information about a specific call by pointing the mouse over a specific line.

It is possible to hide different types of lines in the Topology View. To the line types, right-click the **Topology View** and highlight the **Edge types** option to reveal a list of different edge (arrow) types associated with the current Topology View.

28.1.10 Architecture View

Architecture View is another type of Application Schema Visualization; a visual way for ADP to represent the information stored in its Application Schema model. This view is designed to help you understand the structure and behaviors of Java EE, SOA, and Portal applications at the module and component level. Some Architecture Views also include built-in delay analysis to help identify potential bottlenecks in a given call path.

The Architecture View in ADP is capable of showing application structure and component relationships at two levels: module and component levels. At each level, ADP can show both active and potential call paths. [Table 28–3](#) describes various types of Architecture Views.

Drill down on a specific application to launch into the Architecture View. This action demonstrates the logical progression of drilling from high level resource-centric topology view, down through application-centric topology view, to module-centric

architecture view. Using this logical drill down, you can understand the structure of your application runtime environments and diagnose problems.

Table 28–3 Various Types of Architecture View

Tab Name	Description
Module Level Execution	This is the default Architecture View at the module level. The Module Level Execution view shows the active calling relationships among various Java EE modules (EAR, WAR, JAR, and so on). Shared resources are also included.
Module Level	The Module Level view shows the potential calling relationships among various Java EE modules. Shared resources are also included. It should also be noted that any object that is not connected within the static view will not be included at this level and if there are no static connections at all between objects, every potential object relationship will be displayed. By default, the Module Level view is not enabled.
Component Level Execution	This is the default Architecture View at the component level. The Component Level Execution view shows the active calling relationships among different Java EE components (EJB, servlet, JSP, and so on). Shared resources are also included.
Component Level	The Component Level view shows the potential calling relationships among various Java EE components. Shared resources are also included. Similar to the module level, any object that is not connected within the static view will not be included at this level and if there are no static connections at all between objects, every potential object relationship will be displayed. By default, the Component Level view is not enabled.

These various types of architecture views are color coded in order to provide additional information. [Table 28–4](#) lists color codes and their meanings.

Table 28–4 Architecture View Color Codes

Background Color	Description
Orange	The orange background color represents entry points into the application or module. The orange color also represents that these entities belong to the same application or module currently selected (in context).
Green	The green background color represents entry points into the application or module. The green color also represents that these entities belong to other applications or modules (out of context). The green color is also used to represent share resources.
White	The white background color represents that these entities belong to the same application and module currently selected (in context).
Blue	The blue background color represents that these entities belong to other applications and modules (out of context). The blue color also represents shared resources.

ADP graphically depicts active calling relationships among various Java EE modules and shared resources.

28.1.10.1 Accessing the Architecture View

There are several ways to access the Architecture View. One way is through the Deployments node associated with a specific application under the Application Node. Application specific Architecture View can be accessed using the Deployments node on the Oracle Tree.

The last way to access the Architecture View is by right-clicking a managed entity and selecting the Architecture View. Right-click and select Architecture View to start the drill down process.

28.1.10.1.1 Arrows in Architecture Views The arrows connecting various entities in the Architecture Views depict calls made from one entity to another. You can get more

information about a specific call by pointing the mouse over a specific arrow. Mousing over arrows shows the details of a specific call in Architecture View

It is possible to hide different types of arrows in the Architecture View. To do this, right-click on the Architecture View and highlight the Edge types option to reveal a list of different edge (arrow) types associated with current Architecture View. Unchecking a specific edge type hides all lines of that type in the Architecture View. Checking a specific edge type makes these lines appear.

To hide all lines not connected with a specific entity, select a monitored entity in the Architecture View, right-click and select **Hide other edges**. Highlight an entity and select **Hide other edges** to hide all arrows not connected to the managed entity.

28.1.10.1.2 Architecture View Summary The Architecture View Summary provides the delay analysis associated with the active call path displayed. The table and pie chart displayed in the right pane guides you to leading delay contributors in the displayed call path. Selecting a specific component in the call path brings up component specific information. You will see the following tabs:

- Summary tab first which includes high-level delay data for both inbound and outbound calls.
- The Instrumentation tab shows detailed method level performance data associated with the selected component. Click the Instrumentation tab to see detailed performance measurements and information at the method level.
- The Errors/Exceptions tab shows the errors metrics associated with the selected portal or BPEL process.
- The SQL Statement tab shows SQL statements and their performance data associated with the selected component.
- The Transactions tab shows the transaction events associated with the selected portal and children below. By default, the Transactions tab is not enabled.

28.1.11 Metric Types

[Table 28–5](#) describes various types of metrics provided by ADP.

Note: All the ADP metric tables have a View drop-down list to change the order of the columns in the tables.

Table 28–5 Metric Types

Examples	Metric Type	Metric Description
Active Sessions	Snapshot Count	A count of the monitored entity at a point in time. ADP plots these snapshot counts in trend graphs.
Completions		
Pending Requests		
Running Instances		
Max Capacity		
Messages High		

Table 28–5 (Cont.) Metric Types

Examples	Metric Type	Metric Description
Requests Serviced Total Sessions [Processes] Aborted [Processes] Terminated [Method] Invocation Count Bytes Received	Aggregated Count	A count of the monitored entity incrementally aggregated from the beginning of display time window. ADP shows these aggregated counts in summary tables.
Response Time Elapse Time Connection Delay	Average Timing	<p>Calculated every sampling period (default 60 seconds), the average timing is calculated by dividing the total amount of time needed to complete the monitored business unit of work by the number of completed business units of work.</p> <p>ADP uses this data in the following two ways:</p> <ol style="list-style-type: none"> 1. Plot the average timings in trend graphs. 2. Calculate average timing of this business unit of work for the display time window and display in a summary table.
Min/Max	Minimum and Maximum Response Time Measurement	Minimum and maximum response time measurements found per collection sampling intervals. These are stored in their embedded database in addition to average response time measurements. The default is 60 seconds.

28.2 Exploring the Monitoring Tab

When ADP is pointed to an Oracle WebLogic domain or an Oracle SOA Suite cluster, it automatically discovers information about this particular domain including all deployed applications, configuration, resources, and others. ADP displays this information in the Monitoring tab under Oracle Enterprise Manager.

Each node represents a construct in the platforms monitored by ADP. Each construct is described in this section.

Note: Promote to dashboard can be configured in ADP to incorporate ADP metrics tables in the ADP dashboard page. The dashboard configuration can be selected for the entity type which is discovered by ADP and to display the entity metrics table on the dashboard.

This section includes the following topics:

- [Monitoring ADP Entities](#)
- [WebLogic](#)
- [Oracle WebLogic Portals](#)
- [Oracle BPEL Processes](#)
- [Oracle ESB](#)
- [Oracle WebCenter](#)
- [Processes](#)
- [Web Services](#)
- [Pageflows](#)
- [Services](#)

- [WSRP Producers](#)
- [Integration](#)
- [Applications](#)
- [Oracle WebLogic Resources](#)
- [Oracle Resources](#)
- [Custom Metrics](#)
- [Status](#)
- [Service Component Architecture \(SCA\)](#)

28.2.1 Monitoring ADP Entities

Performance can be monitored for the ADP entities:

- SOA Suite 11g
- OSB Applications
- Java EE Applications
- ADF Applications

28.2.1.1 Monitoring SOA Suite 11g Performance

To monitor the performance of service-oriented architecture applications (SOA), perform the following steps:

1. Navigate to **Application Dependency and Performance**.

From the **Targets** menu, select **Middleware**. On the Middleware page, click the SOA Infrastructure target. On the **Home** tab, select the **Summary** region and click the **Application Dependency and Performance** link.

2. Click the **Monitoring** tab.
3. ADP discovers all the deployed Composites on the configured Oracle WebLogic Domain.
4. A Composite node appears under the configured ADP Manager (for example, select Oracle Enterprise Manager, select SOAServer, then select Composites).

Under the Composites node, the following nodes appear:

- SCA Partition
- Composite
 - Services
 - Components
 - References
 - Wires

The performance metrics are displayed on the right-hand side panel when the respective node is selected.

28.2.1.2 Monitoring OSB Performance

To monitor the performance of Oracle Service Bus (OSB) applications, perform the following steps:

1. Navigate to **Application Dependency and Performance**.

From the **Targets** menu, select **Middleware**. On the Middleware page, click the OSB target. On the **Home** tab, select the **Summary** region and click the **Application Dependency and Performance** link.

2. Click the **Monitoring** tab, then select OSB in the tree.

3. ADP discovers all the deployed OSB proxy and business services on the configured Oracle WebLogic Domain.

4. An OSB node appears under the configured ADP Manager (for example, select Oracle Enterprise Manager, select servicbusServer, then select OSB).

Under the OSB node, the following nodes appear:

- Business Services
- Proxy Services
 - Pipeline
 - References

The performance metrics are displayed on the right-hand side panel when the respective node is selected.

28.2.1.3 Monitoring Java EE Application Performance

To monitor the performance of Java EE applications, perform the following steps:

1. Navigate to **Application Dependency and Performance**.

From the **Targets** menu, select **Middleware**. On the Middleware page, click the Java EE Application target. On the **Home** tab, select the **Summary** region and click the **Application Dependency and Performance** link.

2. Click the **Monitoring** tab, then select the application in the tree.

3. ADP discovers Java EE artifacts like Servlet/JSPs, EJBs, and Web services.

4. You can find Java EE Application related data under following nodes:

- Applications
- Web Services
- Services

The performance metrics for related components are displayed on the right-hand side panel when the respective component is selected.

28.2.1.4 Monitoring ADF Application Performance

To monitor the performance of Application Development Framework (ADF) applications, perform the following steps:

1. Navigate to **Application Dependency and Performance**.

From the **Targets** menu, select **Middleware**. On the Middleware page, click the ADF target. On the **Home** tab, select the **Summary** region and click the **Application Dependency and Performance** link.

2. Click the **Monitoring** tab, then select the application in the tree.

3. ADP discovers all the deployed ADF artifacts on the configured Oracle WebLogic Domain release 11gR1.

4. An ADF node appears under the configured ADP Manager (for example, Oracle Enterprise Manager, select Server, then select ADF). The ADF node contains the following:
 - ADF taskflows
 - JSF Pages
 - Managed Beans
 - Business Components

The performance metrics for related components are displayed on the right-hand side panel when the respective component is selected.

28.2.2 WebLogic

You can perform the following on WebLogic using Application Dependency and Performance.

- Monitor a WebLogic portal
- Add a WebLogic domain to be monitored by an existing ADP manager
- Remove a WebLogic Domain from monitoring

28.2.2.1 Monitoring WebLogic Portal Performance

Perform the following steps:

1. Navigate to **Application Dependency and Performance**.
From the **Targets** menu, select **Middleware**. In the Related Links section, click **Application Dependency and Performance** link.
2. Click the **Monitoring** tab, then select the application in the tree.
3. ADP discovers portal on the configured Oracle WebLogic Domain.
4. A Portals node appears under the configured ADP Manager (for example, Oracle Enterprise Manager, select PortalServer, then select Portals). The Portals node contains the following:
 - Portal Desktop, which is made up of Books, which are in turn made up of Pages, which in turn are made up of Portlets

The performance metrics for related components are displayed on the right-hand side panel when the respective component is selected.

28.2.2.2 Adding WebLogic Domain To Be Monitored By Existing ADP Manager

To add a WebLogic domain to be monitored by an existing ADP manager, perform the following steps:

1. Navigate to **Application Dependency and Performance**.
From the **Targets** menu, select **Middleware**. On the Middleware page, click the WebLogic target. On the **Home** tab, select the **Summary** region and click the **Application Dependency and Performance** link.
2. Click the **Configuration** tab and expand the Configuration (Manager Name). Click **Resource Configuration** and provide the details.

28.2.2.3 Removing a WebLogic Domain From Monitoring

To remove a WebLogic domain from being monitored, perform the following steps:

1. Navigate to **Application Dependency and Performance**.

From the **Targets** menu, select **Middleware**. In the **Related Links** section, click **Application Dependency and Performance**.

2. Click the **Configuration** tab and expand the Configuration (Manager Name). Click **Resource Configuration**, choose the resource, and click **Delete Resource**.

28.2.3 Oracle WebLogic Portals

The Portals node under Oracle Enterprise Manager contains information about all deployed WebLogic Portal applications in the managed domains. The Portals node is organized hierarchically using the same framework developers use to build these Portal applications. The minimum and maximum response time measurements are stored in the database in addition to the average response time measurements. These metrics, if present, display visually in the Main Display Window on the right.

For WebLogic Portal, this hierarchy contains the following (Table 28–6):

Table 28–6 WebLogic Portal Hierarchy

Component	Description
Portals	The Portal is the logical containment unit for a Portal application. A typical Portal can contain a few desktops, several books, tens of pages, and hundreds of portlets.
Desktops	The desktop is the top-level container for the portal components included in that specific view of the portal. Portal administrators can create new desktops beyond what portal developers create in WebLogic Workshop.
Books	The top-level book contains all sub-books, pages, and portlets. The top-level book defines the initial menu navigation style used for the desktop. For each sub-book you add to a desktop you can select a different navigation style.
Pages	Pages and sub-books are the navigable containers used for organizing portlets.
Portlets	Portlets are the containers that surface Web content and applications in your desktops.

When you click the Portals node under Oracle Enterprise Manager, ADP displays summary information on active portal applications. This summary includes the following (Table 28–7):

Table 28–7 Tree Summary

Metric	Description
Portal web application activity	A summary of user sessions for a specific portal application
Portal completions	Total number of requests fulfilled by a specific portal application
Portal response time (ms)	Average response time for a specific portal application
Portal entitlement response time	Average response time of WebLogic Portal entitlement subsystem for a specific portal application
Portal campaign response time	Average response time of WebLogic Portal campaign subsystem for a specific portal application

For Portal web application activity and Portal performance, ADP displays information in both table and graph formats. For the other metrics, ADP shows the information in graph format. When you click the plus (+) icon next to the Portals node, ADP expands the tree to show all managed portal applications currently deployed on the WebLogic domain.

You can also see information specific to a particular portal application. By selecting a specific portal application, all information displayed in the Main Display Window changes to only show data relevant to this new context. For example, when a user selects a particular portal application under the Portals node, the Main Display Window only shows information specific to that portal application.

At the Portal level, you can navigate to different levels of the portal application by using different tabs. Use the tabs available to quickly access lower level components. [Table 28–8](#) provides a list of the tabs available for portal level nodes and their descriptions.

Table 28–8 Portal Level Tab Description

Tab	Description
Summary	Performance summary specific to the selected portal.
Desktops	Performance summary for all the desktops associated with the selected portal.
Headers	Performance summary for all the headers associated with the selected portal.
Books	Performance summary for all the books associated with the selected portal.
Pages	Performance summary for all the pages associated with the selected portal.
Portlets	Performance summary for all the portlets associated with the selected portal.
Footers	Performance summary for all the footers associated with the selected portal.
WSRP Topology	View WSRP consumer-producer relationships and WSRP deployment topology.
Analysis	Two performance analytics - Multi-Point Regression Analysis performed at the portal level and Entity Performance Ranking performed at the portlet level.
Events	SLO violation events associated with the selected portal.
Errors/Exceptions	Errors metrics associated with the selected portal.
Transactions	Transaction events associated with the selected portal and children below. By default, the Transactions tab is not enabled. To enable the Transactions tab, enable the <code>UIProvider.Modes=product,trace</code> property in the <code>Acsera.properties</code> file before starting the manager.

28.2.3.1 Desktops

Expand a particular portal application further to reach the Desktops node. By selecting the Desktops node, ADP provides a list of currently active desktops associated with that portal application.

This Desktop Summary includes the following metrics:

Table 28–9 Desktop Summary Metrics

Metrics	Description
Desktop arrivals	Total number of requests for a specific desktop
Desktop completions	Total number of requests fulfilled by a specific desktop
Desktop response time (ms)	Average response time for a specific desktop.

Note: Portal desktops are end-user facing entities. Metrics such as Desktop hits and response time represents request arrival rate and application performance respectively. Violations in thresholds set on these metrics would indicate unacceptable end-user experience.

ADP displays these metrics in both table and graph formats.

For example, when you have two active desktops, you can drill down further to a specific desktop by expanding the Desktops node. Again, clicking on the plus (+) icon expands the tree view for you.

When you select a node in the expanded tree to get more information specific for that desktop, ADP changes information in the Main Display Window to reflect the new context.

ADP not only shows the performance metrics associated with a specific node, but it also displays other relevant settings for that node. For example, there can be pre-configured Service Level Objectives (s). These SLOs are displayed in the graphs as red lines.

Expand the desktop node to see Header, Footer, and Books. You can see detailed information for these components by clicking on the appropriate nodes.

28.2.3.1.1 Display Portal Desktop - Desktop Structure Viewer One of the unique capabilities of ADP is its automatic discovery and modeling of deployed applications. The Desktop Structure Viewer provides visibility into how a portal desktop is organized. To activate the Desktop Structure Viewer, right-click on a specific desktop. Select the Display Portal Desktop menu option to access the Desktop Structure Viewer.

After the Desktop Structure Viewer appears, you can navigate through the portal desktop structure by clicking on the appropriate book, page, or portlet. The ability to see portal desktop structure using the same perspective as portal end-users is a unique value especially for the IT support staff.

With the Desktop Structure Viewer, the IT support staff can speak the same language with end-users while at the same time looking at performance oriented information for a specific component. The IT support staff can also use the Desktop Structure Viewer to isolate a particular performance problem. By drilling down from the top-level desktop to individual portlets, the IT support staff can get more insight into which components are having performance problems.

The Desktop Structure Viewer consists of two main panes. The pane on the left is the Desktop Structure pane. This pane allows you to graphically navigate the portal desktop. The pane on the right is called the Main Display Window. The Main Display window displays performance information in the context of the selected component in the Desktop Structure pane. As you navigate through the portal desktop and click different components, the Main Display Window provides information relevant for that selected context.

The Main Display Window shows relevant performance metrics for different portal desktop components - desktop, books, pages, and portlets.

Since ADP understands the WebLogic Portal framework and knows that a pageflow can be associated with a portlet, it is designed to allow easy access to the Pageflow Viewer from the Desktop Structure Viewer.

To activate the Pageflow Viewer, double-click the interesting portlet. In turn you can double-click the portlet in the Desktop Structure pane to open the appropriate pageflow in the Pageflow Viewer.

28.2.3.2 Portlet Drill Down

You can drill down on a portlet in the portal desktop view to activate the Display Architecture View.

1. Select a portlet under a node.

2. Double-click on a name to see the Portal Desktop Status page.
3. In the Portal Desktop Status window, right-click on a service box to select Display Architecture View.

28.2.3.3 Pageflow Viewer

The Pageflow Viewer has two panes. The pane on the right is the Main Display Window. The Main Display Window shows information corresponding to the item selected in the left pane. The left pane shows either the Flow View or the Component View. You can choose to see either the Flow View or the Component View by selecting the appropriate tab.

The Main Display Window changes to show information relevant to the selected item in either the Flow View or the Component View.

28.2.3.4 Books

Expand a particular portal desktop further to see the Books node. By selecting the Books node, ADP provides a list of currently active books associated with the specific desktop.

This Books Summary includes the following metrics ([Table 28–10](#)):

Table 28–10 Book Summary Metrics

Metrics	Description
Book completions	Total number of requests fulfilled by a specific book
Book response time (ms)	Average response time for a specific book

ADP displays these metrics in both table and graph formats. For example, you can have two active books for the portal desktop. These active books are listed in the table and plotted in the graphs.

You can drill down further to a specific book by expanding its node. Click the plus (+) icon to expand the tree view. Expand the Books node to see a list of specific books configured.

When you select a particular active book, the Main Display Window shows the relevant information in that context.

28.2.3.5 Pages

Expand a particular book to see the Pages node. By selecting the Pages node, ADP provides a list of currently active pages associated with the specific book.

This Pages Summary includes the following metrics ([Table 28–11](#)):

Table 28–11 Pages Summary Metrics

Metrics	Description
Page completions	Total number of requests fulfilled by a specific page
Page response time (ms)	Average response time for a specific page

ADP displays these metrics in both table and graph formats. For example, you can have one active page for a book. The active page is listed in the table and plotted in the graphs.

You can drill down further to a specific page by expanding the Pages node. Click the plus (+) icon to expand the tree view. This reveals the next level of components - Portlets.

28.2.3.6 Portlets

Expand a particular page to see the Portlets node. Select a Portlets node to view a list of currently active portlets associated with the specific page.

This Portlets Summary includes the following metrics ([Table 28–12](#)):

Table 28–12 Portlet Metrics

Metrics	Description
Portlet completions	Total number of requests fulfilled by a specific portlet.
Portlet response time (ms)	Average response time for a specific portlet.

ADP displays these metrics in both table and graph formats. For example, you can have four active portlets for a particular page. These active portlets are listed in the table and plotted in the graphs.

Drill down further to a specific page by expanding the Portlets node. Click the (plus) + icon to expand the tree view. This provides additional information about the page.

28.2.4 Oracle BPEL Processes

The BPEL Processes node in the navigation tree contains information about all deployed Oracle BPEL processes within the managed domain. ADP organizes information for various process nodes into domains.

In the right-hand pane, you can view the minimum and maximum response time measurements stored in the database in addition to the average response time, arrivals, errors, and completions measurements. These metrics, if present, display visually in the window on the right pane.

When you select the root of the BPEL Processes tree, ADP displays the BPEL Processes Summary in the Main Display Window.

The BPEL Process Summary includes the following ([Table 28–13](#)):

Table 28–13 BPEL Process Summary Metrics

Metrics	Description
Domain	Name of the OC4J domain container
Process	Name of the BPEL process
Arrivals	Total number of currently running instances for a specific BPEL process
Response Time (ms)	Average response time in milliseconds for a specific BPEL process
Completions	Total number of fulfilled requests for a specific BPEL process. A Completed status represents a BPEL process instance that has finished normally.
Errors	Total number of aborted instances of a specific BPEL process
Min Response Time (ms)	Minimum average response time in milliseconds for a specific BPEL process
Max Response Time (ms)	Maximum average response time in milliseconds for a specific BPEL process

ADP presents these metrics in a table format in the Main Display Window when you select the BPEL Processes node. Graphical representations of two metrics, Arrivals and Completions, are displayed below the table.

When you click the plus (+) icon next to the domains sub-node under the main BPEL Processes node, ADP expands the tree to show all managed BPEL domains currently deployed on that particular Oracle SOA Suite instance.

You can see information specific to a particular process. By selecting a specific process, all information displayed in the Main Display Window changes to only show data relevant to this new context.

To see the BPEL process work flow associated with a BPEL process, select the node, right-click and select the Display Functional View option. ADP displays the appropriate functional work flow diagram and associated performance data in a new pop-up window.

See [Table 28–14](#) for BPEL Functional View summary.

Table 28–14 BPEL Functional View Summary

Column/Metric	Description
Activity	Name of a specific activity in the BPEL process
Type	Control Type for a specific node
Arrivals	Number of requests that have arrived for a specific node
Response Time (ms)	Average response time for a specific node
Completions	Number of completed requests for a specific node
Errors	Number of aborted instances for a specific node
Response Time Min (ms)	Minimum response time for a specific node
Response Time Max (ms)	Maximum response time for a specific node

By looking at this summary table, you can determine which BPEL process node is running slowly and whether there are errors.

In addition to the summary, the following views are available for a node:

- Delay Analysis view
- Metadata view
- Partner Links view
- Partner Link Type Role view
- Partner Link Bindings view
- Modeled Entities view
- Topology view

You can get to these views by selecting the appropriate tab.

28.2.4.1 Delay Analysis View

Delay Analysis gives you a bird's eye view of a specific BPEL process. You can see what nodes in the BPEL process are taking up a majority of the average elapsed time. The red bar indicates the slowest BPEL process group or BPEL process node. The blue represents the time spent for the particular nodes.

28.2.4.2 Metadata View

The Metadata view displays the tables containing specific metadata associated with the selected active BPEL process being displayed in the left-hand pane. Information provided in this view includes caller and called class metadata information as well as general summarized metadata in relation to the BPEL process and the associated web services. [Table 28–15](#) explains the metadata.

Table 28–15 Metadata View Summary

Column/Metric	Description
SummaryTable -Process	Name of the BPEL process node
SummaryTable -Web Service	Name of the web service being called from the BPEL process
SummaryTable -Version	Version of the web service being called from the BPEL process
SummaryTable -Location	Location of the web service being called from the BPEL process
Caller Table - Caller Class	Class name for the caller class that is calling the BPEL process
Caller Table - Caller Method	Class method for the caller class that is calling the BPEL process
Caller Table -Target Host	Target host that the caller class targeted to instantiate the BPEL process
Caller Table -Target Port	Target port that the caller class targeted to instantiate the BPEL process
Caller Table -Target URL	Target URL that the call class targeted to instantiate the BPEL process
Caller Table - Invocation Count	Number of invocations of the BPEL process instantiated by the caller class
Caller Table - Response Time	Average response time of the BPEL process instantiated by the caller class
Called Clients Table - Called Class	Class name of the class that was called by the BPEL process
Called Clients Table - Target URL	Target URL of the class that was called by the BPEL process
Called Clients Table - Invocation Count	Number of invocations made from the BPEL Process to the called class.
Called Clients Table - Response Time	Response time of the called class

28.2.4.3 Partner Links View

The partner links view provides detailed information on the various roles related to how and why the partner link service is being utilized. The information provided includes both the caller and callee roles, as well as the partner link type. See [Table 28–16](#).

Table 28–16 Partner Links View Summary

Column/Metric	Description
Partner Link	Name of the partner link
My Role	Role in regards to the BPEL process calling the partner link service
Partner Role	Role of the partner link service
Partner Link Type	Partner link category (type) of the service being called

28.2.4.4 Partner Link Type Role View

See [Table 28–17](#) describes the columns in the Partner Link Type Role view.

Table 28–17 Partner Link Type Role View Summary

Column/Metric	Description
Name	Name of the partner link
Link Type Name	Category (type) of the partner link
Port Type	Partner link service URL

28.2.4.5 Partner Link Bindings View

The Partner Link Bindings view provides insight into the actual roles and types of the partner link instances which represent web services that have been bound by the BPEL process. See [Table 28–18](#).

Table 28–18 Partner Link Bindings View Summary

Column/Metric	Description
Partner Link Role	Defines the web service role that the BPEL process will communicate with
Partner Link Type	Defines the web service type that the BPEL process will communicate with
WebService PortType	Name of the web service
WebService Port Namespace ID	URL of the webservice instance

28.2.4.6 Modeled Entities View

The modeled entities view consist of a list and count of the general entities as catalogued during the discovery phase of the resource configuration. The tables contain both a total entity count as well as a breakdown of the entity count by entity type. See [Table 28–19](#).

Table 28–19 Modeled Entities Summary

Column/Metric	Description
Total Entities Modeled Table - Total	Total entities (static label)
Total Entities Modeled Table - Count	Total number of entities catalogued during the discovery phase of the BPEL process
Modeled Entities Table - Entity Type	Entity type being catalogued as part of the discovery phase of the BPEL process
Modeled Entities Table - Count	Total number of entities catalogued during the discovery phase of the BPEL process for a particular entity type

28.2.4.7 Topology View

The Topology View utilizes the modeled entities that were captured during the discovery process to provide a bird's eye view of all of the various high-level relationships between BPEL processes, web services, and business services. You can toggle between static and dynamic relationship views using the tabs at the top of the Topology pane.

28.2.4.8 Node Hierarchy

Expanding a particular BPEL process further, the first item you see is the Node Hierarchy node. By selecting the Node Hierarchy node, ADP provides a list of nodes associated with the specific process.

When you click the plus (+) icon next to a specific Node Hierarchy node, ADP expands the tree to show BPEL process nodes in the Node Hierarchy. Click an individual BPEL

process node to see the load and performance of the selected node in the Main Display Window.

The BPEL process node information also includes the name of the method invoked. This information is displayed as part of the summary table at the top of the main view window.

28.2.5 Oracle ESB

The Oracle ESB node under Oracle Enterprise Manager contains information about all of the deployed Oracle ESB servers running in the managed domain. ADP organizes the information for various Oracle ESB nodes into various categories.

When you select the root of the ESB tree, ADP displays the ESB Summary in the Main Display Window.

The ESB Summary includes the following (Table 28–20):

Table 28–20 ESB Summary Metrics

Metric	Description
ESB System	Name of ESB System
ESB Service	Name of the ESB Service identifier
Arrivals	Total number of ESB service instance arrivals
Completions	Total number of ESB service instance completions
Response Time	Total number of completed instances for a specific BPEL process. A Completed status represents a BPEL process instance that has finished normally.

ADP presents these metrics in a table format in the Main Display Window when you select the ESB node. When you click the plus (+) icon next to the ESB Systems sub-node under the main ESB node, ADP expands the tree to show all managed ESB Systems currently deployed on that particular Oracle SOA Suite instance.

You can see information specific to a particular ESB System. By selecting a specific ESB System, all information displayed in the Main Display Window changes to only show data and the topology relevant to this new context.

By looking at the summary table, you can find out which ESB node is running slowly and whether there are errors.

Besides the summary, the following views are available for the Node Hierarchy node:

- Service Details view
- Service Parent Details view
- Service Definition view
- Service Operations view
- Operation Routing Rules view
- Topology view

You can get to these views by selecting the appropriate tab.

28.2.5.1 Service Details View

The Service Details view provides specific information related to the details of the bound service process instances. Instance IDs and other descriptive details are included as part of this view. See [Table 28–21](#).

Table 28–21 Service Details View Summary

Column/Metric	Description
Service Name	Name of the ESB service
GUID	GUID of the ESB service
Qname	Canonical qualified name for the bound ESB service
Description	Description of the ESB service

28.2.5.2 Service Parent Details View

The Parent Service Details view provides specific information related to the details of the parent of the bound service process instances. Instance IDs, roles, and other descriptive details are included as part of this view. See [Table 28–22](#).

Table 28–22 Service Parent Details View Summary

Column/Metric	Description
Service Name	Name of the parent ESB service
ParentGUID	GUID of the parent ESB service
ParentQname	Canonical qualified name for the parent of the bound ESB service
ParentType	Parent type of the parent ESB service
MyRole	Role of the caller of the parent ESB service instance
ParentRole	Role of the callee of the parent ESB service instance

28.2.5.3 Service Definition View

The Service Definition view contains information regarding the bound ESB service including the Business Service (ESB) WSDL and Port Type as well as the associated URLs. See [Table 28–23](#).

Table 28–23 Service Definition View Summary

Column/Metric	Description
Service Name	Name of the ESB service
BusinessServiceWSDL	URL of the Business Service WSDL
BusinessServicePortType	Port type of the Business Service
ConcreteServiceWSDL	URL of the Concrete Service WSFL
ConcreteServiceURI	URI for the concrete service

28.2.5.4 Service Operations View

The Service Operations views provides details regarding the various method operations being executed. All information is provided in regards to the metadata associated with a specific business service instance. See [Table 28–24](#).

Table 28–24 Service Operations View Summary

Column/Metric	Description
Service Name	Name of the ESB service
Name	Service operation name being executed
GUID	GUID of the ESB service
Qname	Canonical qualified name for the bound ESB service
Element	Associated element within the ESB Service
SchemaLocation	Schema location for the associated ESB service
Type	Type of ESB service operation

28.2.5.5 Operation Routing Rules View

The Operation Routing Rules view provides various details regarding the operation routing rules for Business Service operations. This includes the specific instance business service names being utilized for operations. See [Table 28–25](#).

Table 28–25 Operation Routing Rules View Summary

Column/Metric	Description
Service Name	Name of the ESB service
Name	Instance name ID of the ESB service instance
GUID	GUID of the ESB service instance

28.2.6 Oracle WebCenter

Oracle WebCenter provides a set of features and services (for example, portlets, customization, and content integration) that simplify the process of reaching a solution with JSF applications. This solution brings information from multiple sources into a single interface, simplifying transactions and providing everything users need to support a given task within the application itself.

Table 28–26 WebCenter Tree Summary

Component	Description
ADF Business Components	ADF business component
ADF Data Controls	ADF data controls
ADF Taskflows	ADF task flows provide a modular approach for defining control flow in an application. See Section 28.2.6.1, "ADF Task Flows" .
JSF Pages	JSF page definition files define the binding objects that populate the data in UI components at runtime. See Section 28.2.6.2, "JSF Pages" .
Portlets	Portlets are the containers that surface Web content and applications on desktops. See Section 28.2.6.3, "Portlets" .

28.2.6.1 ADF Task Flows

Instead of representing an application as a single large JSF page flow, you can break it up into a collection of reusable task flows. Each task flow contains a portion of the application's navigational graph. The nodes in the task flows are activities. An activity node represents a simple logical operation such as displaying a view, executing application logic, or calling another task flow. The transactions between the activities

are called control flow cases. A task flow consists of activities and control flow cases that define the transitions between activities.

28.2.6.1.1 User-Defined Taskflows The following taskflows are available in WebCenter.

Table 28–27 Taskflow Activities

Activity Name	Description
Managed Beans	A backing bean that is managed by the JSF framework and used during the JSF page lifecycle.
Taskflow Method Calls	Invokes a method, typically a method on a managed bean.
Taskflow Views	Displays a JSF page or page fragment. Multiple view activities can represent the same page or same page fragment.
Taskflow URL Views	Redirects the root view port (for example, a browserpage) to any URL-addressable resource, even from within the context of an ADF region.
Taskflow Calls	Calls an ADF bounded task flow from an ADFunbounded task flow or another bounded task flow
Routers	Evaluates an EL expression and returns an outcome based on the value of the expression. For example, a router in a credit check task flow might evaluate the return value from a previous method call and generate success, failure, or retry outcomes based on various cases. These outcomes can then be used to route control to other activities in the task flow.

28.2.6.1.2 Web 2.0 Service Oracle WebCenter provides a wide range of Web 2.0 capabilities, including discussion forums, wikis, blogs, content services, RSS, presence, instant messaging, linking, tagging, and search. Both developers and business users can easily add these services to their pages to maximize productivity.

Table 28–28 Taskflow Activities

Activity Name	Description
Managed Beans	A backing bean that is managed by the JSF framework and used during the JSF page lifecycle.
Taskflow Method Calls	Invokes a method, typically a method on a managed bean.
Taskflow Views	Displays a JSF page or page fragment. Multiple view activities can represent the same page or same page fragment.
Taskflow URL Views	Redirects the root view port (for example, a browserpage) to any URL-addressable resource, even from within the context of an ADF region.
Taskflow Calls	Calls an ADF bounded task flow from an ADFunbounded task flow or another bounded task flow
Routers	Evaluates an EL expression and returns an outcome based on the value of the expression. For example, a router in a credit check task flow might evaluate the return value from a previous method call and generate success, failure, or retry outcomes based on various cases. These outcomes can then be used to route control to other activities in the task flow.

28.2.6.2 JSF Pages

A typical JSF application couples a backing bean with each page in the application. The backing bean defines properties and methods that are associated with the UI components used on the page. The UI component's value is bound to the bean's property.

A Managed Bean is a backing bean that is managed by the JSF framework and used during the JSF page lifecycle.

28.2.6.3 Portlets

Portlets can display excerpts of other Web sites, generate summaries of key information, perform searches, and access assembled collections of information from a variety of data sources. You can use the portlets that Oracle or third parties provide, or create your own programmatically. Oracle WebCenter supports WSRP 1.0, WSRP 2.0, JSR 168, and Oracle PDK-Java. You can include any portlets adhering to those standards in your WebCenter applications.

28.2.6.4 Monitoring WebCenter Performance

Perform the following steps:

1. Navigate to **Application Dependency and Performance**.

From the **Targets** menu, select **Middleware**. On the Middleware page, click the WebCenter target. On the **Home** tab, select the **Summary** region and click the **Application Dependency and Performance** link.

2. Click the **Monitoring** tab, then select the application in the tree.
3. ADP discovers the deployed WebCenter ADF artifacts on the configured Oracle WebLogic Domain.
4. A WebCenter node appears under the configured ADP Manager (for example, Oracle Enterprise Manager, select Server, then select WebCenter). The WebCenter node contains the ADF Application Name and the following subnodes:
 - ADF Taskflows
 - JSF Pages
 - Managed Beans

The performance metrics for related components are displayed on the right-hand side panel when the respective component is selected.

28.2.7 Processes

The Processes node under Oracle Enterprise Manager contains information about all deployed WebLogic business processes in the managed domain. ADP organizes information for various process nodes into the following major categories:

- Node Hierarchy
- Persistent Containers
- Instrumentation

The minimum and maximum response time measurements are stored in the database in addition to the average response time measurements. These metrics, if present, display visually in the window on the right pane.

When you select the root of the Processes tree, ADP displays the Processes Summary in the Main Display Window. See [Table 28–29](#).

Table 28–29 Process Summary Metrics

Metrics	Description
Process	Name of process
Running	Total number of currently running instances for a specific process
Suspended	Total number of suspended instances for a specific process. A Suspended request from a user is a common cause for a process instance to go into a Suspended state.

Table 28–29 (Cont.) Process Summary Metrics

Metrics	Description
Frozen	Total number of frozen instances for a specific process
Completed	Total number of completed instances for a specific process. A Completed status represents a process instance that has finished normally.
Aborted	Total number of aborted instances for a specific process
Terminated	Total number of terminated instances for a specific process. An external Terminate request would terminate a process instance.
Average Execution Time (ms)	Average execution completion time for a specific process

Tip: Statistics on the number of process instances with Terminated, Aborted, and Frozen states can indicate abnormal operation of the WebLogic Integration application or container. It is possible to unfreeze Frozen process instances from WLI Console.

ADP presents these metrics in a table format in the Main Display Window when you select the Processes node. Graphical representations of two metrics, Running Instances and Average Execution Time, are displayed below the table.

When you click the plus (+) icon next to the Processes node, ADP expands the tree to show all managed processes currently deployed on the WebLogic domain.

You can see information specific to a particular process by selecting a specific process. All information displayed in the Main Display Window changes to only show data relevant to this new context.

To see the process work flow associated with a particular process, select the process node, right-click and select the Display Functional View option. ADP displays the appropriate functional work flow diagram and associated performance data in a new pop-up window.

28.2.7.1 Node Hierarchy

When expanding a particular process further, the first item you see is the Node Hierarchy node. By selecting the Node Hierarchy node, ADP provides a list of nodes associated with the specific process. See [Table 28–30](#).

Table 28–30 Node Hierarchy Summary

Column/Metric	Description
Node	Name of a specific node
ID	Process Node ID for a specific node
Type	Control Type for a specific node
Method	Node Method Name for a specific node
Arrivals	Number of Requests Arrived for a specific node
Active	Number of Active Instances for a specific node
Elapsed Time (ms)	Average Time Elapsed to Complete an Instance for a specific node
Completions	Number of Completed Instances for a specific node
Aborts	Number of Aborted Instances for a specific node
Exceptions	Number of Exception Encountered for a specific nod.

By looking at this summary table, you can determine which process node is running slowly and whether there are aborts or exceptions.

The following additional views are available for the Node Hierarchy node:

- Delay Analysis view
- Events view

You can get to these views by selecting the appropriate tab.

28.2.7.1.1 Delay Analysis View Delay Analysis gives you a bird's eye view of a specific process. You can see what nodes in the process are taking up a majority of the average elapsed time. The red bar indicates the slowest process group or process node. The blue represents the time spent for the particular nodes.

28.2.7.1.2 Events View The Events view shows a list of SLO violations events relevant to this process in a table format. The Events view table includes the following information ([Table 28–31](#)):

Table 28–31 Events View Summary

Column/Metric	Description
Start Time	Start time for the process instance that violated a SLO
Entity Name	Name of the process node that violated a SLO
SLO Name	Name of the violated SLO
Service URI	URI of the process that violated a SLO
Application	Name of the application that violated a SLO
Event Type	Violation type (violation or cautionary)
Entity Type	Violation Metric type
SLO Threshold	Type of threshold (high or low)
SLO Trigger Value	Value that triggered a SLO violation

When you click the plus (+) icon next to a specific Node Hierarchy node, ADP expands the tree to show process nodes in the Node Hierarchy. Click an individual process node to see the load and performance of the selected node in the Main Display Window.

The process node information also includes the name of the method invoked. This information is displayed as part of the summary table at the top of the main view window.

28.2.7.2 Persistent Containers

When you expand a particular process further, the Persistent Containers node is included. By selecting the Persistent Containers node, ADP provides a list of persistence performance statistics relevant to the selected process.

As you select the root of the Persistent Containers tree, a summary of all Persistent Containers relevant to the selected process is presented. For example, a summary can contain the following high level items:

- Container persistence invocations
- Container persistence response time (milliseconds)
- Entity EJB activity

- Entity EJB cache
- Entity EJB transactions
- Entity EJB locking

These items are displayed in both table and graph formats.

The Persistent Containers Summary includes different tables:

- Entity EJB Activity
- Entity EJB Cache
- Entity EJB Transactions
- Entity EJB Locking

28.2.7.2.1 Entity EJB Activity Table Entity EJB Activity table ([Table 28–32](#)) includes the following information:

Table 28–32 Entity EJB Activity Table

Metrics	Description
EJB	Name of the Entity EJB
In Use	Number of instances for a specific Entity EJB currently being used from the free pool. [Snapshot Count]
Idle	Number of instances for a specific Entity EJB currently in the idle state in the free pool. These bean instances are available for use. [Snapshot Count]
Waits	Number of Threads currently waiting for a specific Entity EJB bean instance from the free pool [Snapshot Count]
Timeouts	Total number of Threads that have timed out waiting for an available bean instance from the free pool [Aggregated Count]

Tip: Pay attention to Waits and Timeouts metrics. Activities in the Waits metric and increasing count in the Timeouts metric are signs that requests waiting to be serviced by the EJB container. Ideally, 0 should be indicated for these metrics.

28.2.7.2.2 Entity EJB Cache Table Entity EJB Cache table ([Table 28–33](#)) includes the following information:

Table 28–33 Entity EJB Cache Table

Metrics	Description
EJB	Name of the Entity EJB
Hits	Total number of times an attempt to access the Entity EJB instance from the cache succeeded [Aggregated Count]
Accesses	Total number of attempts to access the Entity EJB instance from the cache [Aggregated Count]
Size	Number of beans instances from this EJB Home currently in the EJB cache [Snapshot Count]
Activations	Total number of beans from this EJB Home that have been activated [Aggregated Count]
Passivations	Total number of beans from this EJB Home that have been passivated [Aggregated Count]

Tip: Passivation (serializing EJB state information to disk) and activation (reconstituting EJB state information from disk) are resource intensive operations. Ideally, it is preferable to see low level of activity in these metrics.

28.2.7.2.3 Entity EJB Transactions Table Entity EJB Transactions table ([Table 28–34](#)) includes the following information:

Table 28–34 Entity EJB Transactions Table

Metrics	Description
EJB	Name of the Entity EJB
Commits	Total number of transactions that have been committed for this EJB. [Aggregated Count]
Rollbacks	Total number of transactions that have been rolled back for this EJB. [Aggregated Count]
Timeouts	Total number of transactions that have timed out for this EJB. [Aggregated Count]

Tip: High number of EJB Transaction Rollbacks may indicate problems with the data used - for some reason the target database is unable to commit the change. High number of EJB Transaction Timeouts may indicate problems accessing the database including network outage, database lock contention, database outage, and more.

28.2.7.2.4 Entity EJB Locking Table Entity EJB Locking table ([Table 28–35](#)) includes the following information:

Table 28–35 Entity EJB Locking Table

Metrics	Description
EJB	Name of the Entity EJB
Entries	Number of Entity EJB instances currently locked [Snapshot Count]
Lock Accesses	Total number of attempts to obtain a lock on an Entity EJB instance [Aggregated Count]
Current Waiters	Number of Threads that currently waiting for a lock on an Entity EJB instance [Snapshot Count]
Total Waiters	Total number Threads that have waited for a lock on an Entity EJB instance [Aggregated Count]
Timeouts	Total number Threads that have timed out waiting for a lock on an Entity EJB instance [Aggregated Count]

Tip: Pay attention to Current Waiters and Timeouts. These metrics can indicate possible performance problems caused by EJB Locking. Ideally, 0s should be displayed for these metrics.

By looking at the activities related to Persistence Containers, you can determine if EJB persistence calls are causing performance problems.

28.2.7.3 Instrumentation

When expanding a particular process further, the last item you see is the Instrumentation node. Click the plus (+) icon next to Instrumentation to expand the tree to reveal the following categories of instrumentation:

- Class

- Methods
- Errors/Exceptions
- Transactions

The Class node in the Instrumentation tree provides the following information ([Table 28–36](#)):

Table 28–36 Class Node

Column/Metric	Description
Probe Point	Class name in which instrumentation probe point is inserted
Response Time (ms)	Average response time for a specific class
Invocation Count	Number of times a specific class is called

The Method node in the Instrumentation tree provides the following information ([Table 28–37](#)):

Table 28–37 Method Node

Column/Metric	Description
Probe Point	Method name in which instrumentation probe point is inserted
Response Time (ms)	Average response time for a specific method
Invocation Count	Number of times a specific method is called

The Errors/Exceptions and Transactions are described in [Section 28.1.10, "Architecture View"](#).

28.2.8 Web Services

The Web Services node in the navigation tree contains information about all deployed Web Services in the managed domain. By selecting the Web Services node under Oracle Enterprise Manager, ADP shows the Web Services Summary in the Main Display Window.

This summary view lists all discovered web services and their associated URL entry points. Below this list, ADP lists out all active web services and their performance data (invocation count and response time).

When you click the plus (+) icon next to the Web Services node, ADP expands the tree to show all monitored web services currently deployed on the WebLogic domain.

When you select a specific web service, ADP displays performance data associated with the selected web service. Click the plus (+) icon next to a specific web service to expand the tree to show all public operations associated with that web service.

The Operations table provides the following information ([Table 28–38](#)):

Table 28–38 Operations Table

Column/Metric	Description
Operation	Name of the web service operation
Invocation Count	Number of times the operation is called
Response Time (ms)	Average response time for the operation in milliseconds

Table 28–38 (Cont.) Operations Table

Column/Metric	Description
Delay (ms)	Overall delay contributed by the operation in milliseconds

28.2.9 Pageflows

The Pageflows node in the navigation tree contains information about all deployed pageflows in the managed domain. By selecting the Pageflows node under Oracle Enterprise Manager, ADP shows the Pageflows Summary in the Main Display Window.

28.2.10 Services

The Services node in the navigation tree contains information about all external entry points into the managed domain. ADP currently monitors the following types of services:

- HTTP
- EJBs
- JDBC

Selecting each service type reveals service summary in the Main Display Window.

The minimum and maximum response time measurements are stored in the database in addition to the average response time measurements. These metrics, if present, display visually in the window in the right pane.

ADP displays entry point activity summary associated with the selected EJB service.

Tip: Setting thresholds at some of these entry points enables ADP to monitor the performance of key business services. When a violation event occurs, you can begin investigating from the Service node.

28.2.10.1 HTTP

Expanding the HTTP node under the Services node reveals a list of discovered HTTP based entry points into the managed domain. HTTP service end points include JSPs, struts actions, and servlet mappings. These discovered HTTP entry points are listed by their root context. When you select a specific HTTP entry point, ADP displays the associated summary in the Main Display Window.

When a specific file is selected, ADP displays more detailed performance data.

Method level performance data is displayed when you select a specific HTTP service entry point.

Table 28–39 HTTP Performance Summary

Column/Metric	Description
Servlet	Name of the servlet associated with the selected service
Method	Name of the method invoked by external call
Arrivals	Total number of requests received by this method
Invocation Count	Total number of method invocations
Response Time (ms)	Average method response time in milliseconds

28.2.10.2 EJBs

To view the performance summary for EJBs invoked from outside the JVM, click the EJBs node.

Table 28–40 EJB Performance Summary

Column/Metric	Description
EJB	Name of the EJB
Invocation Count	Number of times the EJB is called
Response Time (ms)	Average response time for the EJB in milliseconds
Delay (ms)	Overall delay contributed by the EJB in milliseconds

Tip: As a general rule, external calls that terminate in EJBs are RMI calls. Web services calls that ultimately terminate in EJBs use SOAP and enter the application server via HTTP.

28.2.10.3 JDBCs

To bring up the performance summary for JDBC operations invoked from outside of the JVM, click the JDBC node.

Table 28–41 JDBC Performance Summary

Column/Metric	Description
SQL Statement	Generalized SQL Statement executed by the JDBC operation
Class	Name of the class used in the JDBC operation
Method	Name of the method used in the JDBC operation
Invocation Count	Number of times the JDBC operation is called
Response Time (ms)	Average response time for the JDBC operation in millisecond
Delay (ms)	Overall delay contributed by the JDBC operation in milliseconds

28.2.11 WSRP Producers

The Web Services Remote Portlet (WSRP) Producers node in the navigation tree contains information about the WebLogic WSRP consumer - producer relationships in the managed domain. By selecting an entity in the WSRP node, ADP displays the performance measurements for the associated WSRP consumer or producer.

WebLogic Portal can act as either a WSRP remote producer or as a consumer. When acting as a consumer, WebLogic Portal's remote--or proxy--portlets are WSRP-compliant. These portlets present content that is collected from WSRP-compliant producers, allowing you to use external sources for portlet content, rather than having to create its content or its structure yourself.

The following types of portlets can be exposed with WSRP inside a WebLogic portal:

- Page flow portlets
- JavaServer Pages (JSP) portlets
- Struts portlets
- Java portlets (JSR168; supported only for complex producers)
- JavaServer Faces (JSF) portlets

The minimum and maximum response time measurements are captured in addition to the average response time measurements. These metrics, if present, display visually in the window in the right pane.

28.2.11.1 WSRP Summary

To view the WSRP Producers Summary:

1. Select the WSRP Producers node to show the WSRP Producers Summary tab.

The WSRP Producers summary includes the following table ([Table 28–42](#)):

Table 28–42 WSRP Producers Summary

Column	Description
WSRP Producer	Name of the producer portlet
WSDL URL	URL of the WSD.

2. To view the portlet details, click the Consumer Portlets node under the WSRP Producers.

The following tables are in this view:

- WSRP Producer Information
- WSRP Consumer Portlet Performance

Select the portlet name and right click. Select **Open** to drill down to more detailed view.
- WSRP Producer Portlets

Select the portlet name and right click. Select Open to drill down to more detailed view.

Table 28–43 WSRP Producers Information

Column	Description
TestPortlets	Defined by the user for the Producer, for example description, handle, and more
URL	Lists the details of each item under the TestPortlet column

Table 28–44 WSRP Consumer Portlet Performance

Column	Description
Portal	The Portal is the logical containment unit for a Portal application. A typical Portal can contain a few desktops, several of books, tens of pages, and hundreds of portlets.
Desktop	The desktop is the top-level container for the portal components included in that specific view of the portal. Portal administrators can create new desktops beyond what portal developers create in WebLogic Workshop.
Book	The top-level book contains all sub-books, pages, and portlets. The top-level book defines the initial menu navigation style used for the desktop. For each sub-book, you add to a desktop you can select a different navigation style.
Page	Pages and sub-books are the navigable containers used for organizing portlets.
Portlet	Portlets are the containers that surface Web content and applications in your desktops.
Response Time (ms)	Average response time in milliseconds.

Table 28–44 (Cont.) WSRP Consumer Portlet Performance

Column	Description
Completions	Number of Completed Instances for a specific node.
Response Time Min (ms)	Minimum response time in milliseconds.
Response Time Max (ms)	Maximum response time in milliseconds.

Table 28–45 WSRP Producer Portlets

Column	Description
Producer Portlet	Name of the producer portlet
Producer	Name of the producer

3. Click a portlet name in the tree view to see the performances associated with the consumer and producer portlets.

28.2.11.2 WSRP Topology

Use this option to visually explore WSRP consumer - producer relationships and the WSRP deployment topology.

To view the WSRP Topology:

1. Select the WSRP Producers node to show the WSRP Topology tab.
2. Click the WSRP Topology tab to view the details.

28.2.11.3 Display Portal Desktop

The portal desktop is described in [Section 28.2.3.1.1, "Display Portal Desktop - Desktop Structure Viewer"](#).

Access the Architecture View:

1. To view the portal desktop for a specific portlet, right-click the portlet name under the Consumer Portlet node.
2. Select Display Portal Desktop.
3. You can drill down to view the Architecture View from this view. See the instructions in [Section 28.2.3.2, "Portlet Drill Down"](#).

28.2.12 Integration

The Integration node under Oracle Enterprise Manager contains information about the WebLogic Integration resources in the managed domain. By selecting the Integration node under Oracle Enterprise Manager, ADP displays the Integration Summary.

The Integration Summary includes the following ([Table 28–46](#)):

Table 28–46 Integration Summary

Metric	Description
Process	Name of process
Running	Total number of currently running instances for a specific process
Suspended	Total number of suspended instances for a specific process

Table 28–46 (Cont.) Integration Summary

Metric	Description
Frozen	Total number of frozen instances for a specific process
Completed	Total number of completed instances for a specific process
Aborted	Total number of aborted instances for a specific process
Terminated	Total number of terminated instances for a specific process
Average Execution Time	Average execution completion time for a specific process

Tip: Statistics on the number of process instances with Terminated, Abort, and Frozen states can indicate abnormal operation of WebLogic Integration application or container. It is possible to unfreeze Frozen process instances from WLI Console.

ADP presents these metrics in a table format in the Main Display Window when you select the Integration node. Graphical representations of these metrics, Running Instances, Completed Instances, and Average Execution Time, are displayed below the table.

Expand the Integration tree by clicking on the plus (+) icon next to Integration node.

The expanded Integration tree allows you to look at various components of WebLogic Integration and help identify performance bottlenecks. This section explains the nodes under the Integration Tree.

28.2.12.1 Health

In the expanded Integration tree, the first node you see is the Health node. Under the Health node, ADP lists various subsystems in WebLogic Integration. By expanding the Health node, you can see the following:

- Execute Queues
- Async Dispatchers
- Sync Dispatchers
- JMS Destinations
- Stateless Containers
- Persistent Containers

You can get to the health information specific to each of these subsystems by clicking the appropriate node. Also, you can get to a particular instance of a subsystem.

28.2.12.1.1 Execute Queues In the Execute Queues node, ADP provides operational statistics of each execute queues configured for WebLogic Integration. Select the Execute Queues node in the Monitor Workspace to display the Execute Queues Summary in the Main Display Window.

The Execute Queues Summary provides the following information ([Table 28–47](#)):

Table 28–47 Execute Queues Summary

Metric	Description
Execute Queue	Execute Queue ID
Aggregated Execute Queue	Aggregated execute queue statistics per resource

Table 28–47 (Cont.) Execute Queues Summary

Metric	Description
Idle Threads	Current number of idle threads in a specific Execute Queue
Pending Threads	Current number of pending threads in a specific Execute Queue
Requests	Total number of requests serviced for a specific Execute Queue
Total Threads	Total number of threads configured in a specific Execute Queue

Tip: Pay attention to Idle Threads and Pending Threads counts. Rapidly decreasing Idle Threads count combined with rapidly increasing Pending Threads count can indicate a backup in the Execute Queue.

Use the following guidelines to adjust the Execute Queue Thread Count ([Table 28–48](#)):

Table 28–48 Guidelines to Adjust the Execute Queue Thread Count

Execute Queue Is Backed Up?	Application Is CPU Bound?	Adjustment Guideline
Yes	No	Increase execute queue thread count.
Yes	Yes	Decrease thread count and explore JVM or application issues that may be causing high CPU utilization.

ADP presents these metrics in a table format in the Main Display Window when you select the Health node. Graphical representations of these metrics, Idle Treads, Pending Threads, and Requests, are displayed below the table.

Expand the Health tree by clicking on the plus (+) icon next to Health node. You can get the same summary as previously described for a specific execute queue.

28.2.12.1.2 Async Dispatchers In the Async Dispatcher node, ADP provides operational statistics of each of the Async Dispatchers configured in WebLogic Integration. Select the Async Dispatchers node in the Monitor Workspace to show the Async Dispatchers Summary in the Main Display Window.

The Async Dispatcher Summary includes the following information ([Table 28–49](#)):

Table 28–49 Async Dispatcher Summary

Metric	Description
EJB	Name of the Message Driven EJB
In Use	Number of instances for a specific Message Driven EJB currently in use
Idle	Number of instances for a specific Message Driven EJB currently in the idle state
Waits	Number of instances for a specific Message Driven EJB currently in the wait state
Timeouts	Number of instances for a specific Message Driven EJB currently in the timeout state
Commits (Transaction)	Total number of commits performed for a specific Message Driven EJB
Rollbacks (Transaction)	Total number of transaction rollbacks performed for a specific Message Driven EJB
Timeouts (Transaction)	Total number of transaction timeouts performed for a specific Message Driven EJB

Tip: Rapidly increasing counts in MDB Waits and Timeouts metrics may indicate a tuning opportunity for the MBD container. Furthermore, increasing numbers in the Transaction Rollbacks and Timeouts metrics may indicate issues interacting with the database. Ideally, these metrics should not increase rapidly.

ADP presents these metrics in a table format in the Main Display Window when you select the Async Dispatchers node. Graphical representation of one metrics, Message Driven EJB in use, is displayed below the table.

Expand the Async Dispatchers tree by clicking the plus (+) icon next to Async Dispatchers node. You can get the same summary as previously described for a specific async dispatcher.

28.2.12.1.3 Sync Dispatchers In the Sync Dispatchers node, ADP provides operational statistics of each of the Sync Dispatchers used by WebLogic Integration. Select the Sync Dispatchers node in the Monitor Workspace to show the Sync Dispatchers Summary in the Main Display Window.

The Sync Dispatcher Summary includes the following information ([Table 28–50](#)):

Table 28–50 Sync Dispatcher Summary

Metric	Description
EJB	Name of the Stateless EJB
In Use	Number of instances for a specific Stateless EJB currently in use
Idle	Number of instances for a specific Stateless EJB currently in the idle state
Waits	Number of instances for a specific Stateless EJB currently in the waits state
Timeouts	Number of instances for a specific Stateless EJB currently in the timeouts state

Tip: Rapidly increasing counts in Stateless EJB Waits and Timeouts metrics may indicate performance issues and a tuning opportunity for the EJB container. Ideally, these metrics should not increase at a rapid pace.

ADP presents these metrics in a table format in the Main Display Window when you select the Sync Dispatchers node. Graphical representation of one metrics, Stateless EJB in use, is displayed below the table.

Expand the Sync Dispatchers tree by clicking on the plus (+) icon next to Sync Dispatchers node. You can get the same summary as previously described for a specific sync dispatcher.

28.2.12.1.4 JMS Destinations In the JMS Destination node, ADP provides operational statistics of each of the JMS Destinations used by WebLogic Integration. Select the JMS Destinations node in the Monitor Workspace to show the JMS Destinations Summary in the Main Display Window.

JMS Destination Summary includes the following tables: JMS destination message statistics and JMS destination byte statistics. The JMS destination message statistics table includes the following information ([Table 28–51](#)).

Table 28–51 JMS Destination Message Statistics

Column/Metric	Description
JMS Destination	Name of the JMS destination
Message Current	Number of JMS messages currently at a specific JMS destination
Message High	Maximum number of JMS messages at a specific JMS destination
Message Pending	Number of JMS messages pending to be delivered to a specific JMS destination
Message Received	Total number of JMS messages at a specific JSM destination

Tip: Pay attention to Message Pending metric. Too many pending messages in a specific JMS destination could result in a performance slowdown. Rapidly increasing count for the Message Pending metric may indicate a performance problem and a JMS destination tuning opportunity.

The JMS destination byte statistics table includes the following information ([Table 28–52](#)).

Table 28–52 JMS Destination Byte Statistics

Column/Metric	Description
JMS Destination	Name of the JMS destination
Byte Current	Byte count of JMS messages currently at a specific JMS destination
Byte High	Maximum byte count of JMS messages at a specific JMS destination
Byte Pending	Byte count of JMS messages pending to be delivered to a specific JMS destination
Byte Received	Total Byte count of JMS messages at a specific JMS destination

ADP presents these metrics in table format in the Main Display Window when you select the JMS Destinations node. Graphical representations of the following metrics, Message pending and Byte pending, are displayed below the table.

Expand the JMS Destinations tree by clicking on the plus (+) icon next to JMS Destinations node. You can get the same summary as described above for a specific JMS destination.

28.2.12.1.5 Stateless Containers In the Stateless Containers node, ADP provides operational statistics of each of the Stateless Containers used by WebLogic Integration. Select the Stateless Containers node in the Monitor Workspace to show the Stateless Containers Summary in the Main Display Window.

The Stateless Containers Summary includes the following information ([Table 28–53](#)):

Table 28–53 Stateless Containers Summary

Metric	Description
EJB	Name of the Stateless EJB
Stateless EJB Transactions	Runtime statistics. You can monitor stateless session EJBs using the metrics in this table.
In Use	Number of instances for a specific Stateless EJB currently being used from the free pool [Snapshot Count]

Table 28–53 (Cont.) Stateless Containers Summary

Metric	Description
Idle	Number of instances for a specific Stateless EJB currently in the idle state in the free pool. These bean instances are available for use. [Snapshot Count]
Waits	Number of Threads currently waiting for a specific Stateless EJB instance from the free pool [Snapshot Count]
Timeouts	Total number of Threads that have timed out waiting for an available bean instance from the free pool [Aggregated Count]

ADP presents these metrics in a table format in the Main Display Window when you select the Stateless Containers node. Graphical representation of one metrics, Stateless EJB in use, is displayed below the table.

Expand the Stateless Containers tree by clicking on the plus (+) icon next to Stateless Containers node. You can get the same summary as previously described for a specific stateless container.

28.2.12.1.6 Persistent Containers In the Persistent Containers node, ADP provides operational statistics of each of the Persistent Containers used by WebLogic Integration. Select the Persistent Containers node in the Monitor Workspace to show the Persistent Containers Summary in the Main Display Window.

The Persistent Containers Summary includes the following tables:

- Entity EJB Activity
- Entity EJB Cache
- Entity EJB Transactions
- Entity EJB Locking

Entity EJB Activity table includes the following information ([Table 28–54](#)):

Table 28–54 Entity EJB Activity

Metrics	Description
EJB	Name of the Entity EJB
In Use	Number of instances for a specific Entity EJB currently being used from the free pool [Snapshot Count]
Idle	Number of instances for a specific Entity EJB currently in the idle state in the free pool. These bean instances are available for use. [Snapshot Count]
Waits	Number of Threads currently waiting for a specific Entity EJB bean instance from the free pool [Snapshot Count]
Timeouts	Total number of Threads that have timed out waiting for an available bean instance from the free pool [Aggregated Count]

Tip: Pay attention to Waits and Timeouts metrics. Activities in the Waits metric and increasing count in the Timeouts metric are signs that requests waiting to be serviced by the EJB container. Ideally, 0 should be indicated for these metrics.

Entity EJB Cache table includes the following information ([Table 28–55](#)):

Table 28–55 Entity EJB Cache

Metrics	Description
EJB	Name of the Entity EJB
Hits	Total number of times an attempt to access the Entity EJB instance from the cache succeeded [Aggregated Count]
Accesses	Total number of attempts to access the Entity EJB instance from the cache [Aggregated Count]
Size	Number of beans instances from this EJB Home currently in the EJB cache [Snapshot Count]
Activations	Total number of beans from this EJB Home that have been activated [Aggregated Count]
Passivations	Total number of beans from this EJB Home that have been passivated [Aggregated Count]

Tip: Passivation (serializing EJB state information to disk) and activation (reconstitute EJB state information from disk) are resource intensive operations. Ideally, Oracle recommends low level of activity in these metrics.

Entity EJB Transactions table includes the following information ([Table 28–56](#)):

Table 28–56 Entity EJB Transactions

Metric	Description
EJB	Name of the Entity EJB
Commits	Total number of transactions that have been committed for this EJB [Aggregated Count]
Rollbacks	Total number of transactions that have been rolled back for this EJB [Aggregated Count]
Timeouts	Total number of transactions that have timed out for this EJB [Aggregated Count]

Tip: High number of EJB Transaction Rollbacks may indicate problems with the data used - for some reason the target database is unable to commit the change. High number of EJB Transaction Timeouts may indicate problems accessing the database including network outage, database lock contention, database outage, and more.

Entity EJB Locking table includes the following information ([Table 28–57](#)):

Table 28–57 Entity EJB Locking

Metric	Description
EJB	Name of the Entity EJB
Entries	Number of Entity EJB instances currently locked [Snapshot Count]
Lock Accesses	Total number of attempts to obtain a lock on an Entity EJB instance [Aggregated Count]
Current Waiters	Number of Threads that currently waiting for a lock on an Entity EJB instance [Snapshot Count]
Total Waiters	Total number Threads that have waited for a lock on an Entity EJB instance [Aggregated Count]
Timeouts	Total number Threads that have timed out waiting for a lock on an Entity EJB instance [Aggregated Count]

Tip: Pay attention to Current Waiters and Timeouts. These metrics can indicate possible performance problems caused by EJB Locking. Ideally, 0s should be displayed for these metrics.

ADP presents these metrics in a table format in the Main Display Window when you select the Persistent Containers node. Graphical representations of three metrics, Entity EJB in use, Entity EJB cache access, and Entity EJB lock access, are displayed below the table.

Expand the Persistent Containers tree by clicking on the plus (+) icon next to Persistent Containers node. You can get the same summary as previously described for a specific persistent container.

28.2.12.2 Performance

In the expanded Integration tree, the second node you see is the Performance node. ADP provides the Performance Summary for WebLogic Integration in the Main Display Window when the Performance node is selected.

The Performance Summary includes the following tables: Process Node and Events. The Process Node table provides performance information for various process nodes running in WebLogic Integration. It includes the following information ([Table 28–58](#)):

Table 28–58 Performance - Process Node Summary

Column/Metric	Description
Node	Name of a specific node
ID	Process Node ID for a specific node
Type	Control Type for a specific node
Method	Node Method Name for a specific node
Arrival	Number of Requests Arrived for a specific node
Active	Number of Active Instances for a specific node
Elapsed Time	Average Time Elapsed to Complete an Instance for a specific node
Completions	Number of Completed Instances for a specific node
Aborts	Number of Aborted Instances for a specific node.
Exceptions	Number of Exception Encountered for a specific node

Tip: You can use Arrivals and Elapsed Time data collected by ADP to characterize the performance of your installation. Since ADP measures performance at cluster level, you are capturing the actual performance of your configuration. You can also perform simple capacity planning analysis by plotting Arrivals versus Elapsed Time (arrival rate versus response time). Ask your Oracle consultant for more information.

The Events table provides a list of SLO violations triggered relevant to WebLogic Integration. It includes the following information ([Table 28–59](#)):

Table 28–59 Performance - Events Node Summary

Column/Metric	Description
Start Time	Start time for the process instance that violated a SLO

Table 28–59 (Cont.) Performance - Events Node Summary

Column/Metric	Description
Entity Name	Name of the process node that violated a SLO
SLO Name	Name of the violated SLO
Service URI	URI of the process that violated a SLO
Application	Name of the application that violated a SLO
Event Type	Violation type (violation or cautionary)
Entity Type	Violation Metric type
SLO Threshold	Type of threshold (high or low)
SLO Trigger Value	Value that triggered a SLO violation

28.2.12.3 Channels

In the expanded Integration tree, the third node you see is the Channels node. ADP shows the Channels Summary for various channels configured for WebLogic Integration.

The Channels Summary includes the following information ([Table 28–60](#)):

Table 28–60 Channels Summary

Column/Metric	Description
Channel	Name of channel
Type	Channel type
Message Count	Total number of messages processed for a specific channel
Dead Message Count	Total number of dead messages for a specific channel

Tip: Increasing count in the Dead Message Count metric may indicate a configuration issue. When the Message Broker is unable to determine the URI to send a message to, the message is sent to the appropriate deadletter channel. Ensure the URI configured for the channel is reachable.

Expand the Channels tree by clicking the plus (+) icon next to Channels node. You can get the same health summary as previously described for a specific channel.

28.2.12.4 Subscribers

In the expanded Integration tree, the fourth node you see is the Subscribers node. ADP shows the Subscribers Summary for various subscribers configured for WebLogic Integration.

Expand the Subscribers tree by clicking the plus (+) icon next to Subscribers node. You can get specific information about an individual subscriber.

28.2.13 Applications

The Applications node in the navigation tree contains information about all deployed applications in the managed domain. By selecting the Applications node, ADP displays the Applications Summary.

The Applications Summary includes the following information ([Table 28–61](#), "Applications Summary"):

Table 28–61 Applications Summary

Column/Metric	Description
Application	Name of application
Status	Operations status for a specific application
Response Time (ms)	Average response time in milliseconds for a specific application. This is the average of response times of all JSPs and servlets contained in the deployment archive.
Invocation Count	Total number of invocations for a specific application. This is the total invocation count of all JSPs and servlets contained in the deployment archive.

Tip: Application is a packaging unit in Java EE. Each EAR, WAR, and JAR files deployed to the application server is considered an individual application. These metrics track performance and arrival rate of these entities.

ADP presents these metrics in a table format in the Main Display Window when you select the Applications node. Graphical representations of the following metrics, Response Time, Invocation Count, and Active Sessions, are displayed below the table.

Expand the Applications tree by clicking the plus (+) icon next to Applications node. You can get more information about a specific application.

ADP displays performance summary for the selected application in the Main Display Window. You can obtain additional performance data by clicking different tabs in the Main Display Window.

The Applications Summary includes the following tabs ([Table 28–62](#)):

Table 28–62 Applications Summary Tabs

Tab Name	Description
Summary	Includes performance data at the application level including time-based trend graphs of Application Response Time, Application Invocation Count, and Application Active Sessions. The invocation count and response time for the top 10 slowest servlets, the usual application entry points, are also included.
Response Times	Includes time-based trend graphs of component response times. Graphs include Servlet Response Time, EJB Response Time, and JDBC Response Time.
Invocations	Includes time-based trend graphs of component invocation counts. Graphs include Servlet Invocation Count, EJB Invocation Count, and JDBC Invocation Count.
Errors/Exceptions	Errors metrics associated with the selected portal.
Transactions	Transaction events associated with the selected portal and children below. By default, the Transactions tab is not enabled.
Modeled Entities	Includes a catalog of entities modeled by ADP. Only the modeled entities associated with the selected application are included.
Instrumentation	Includes performance data by different types of instrumentation probe points. There are different tabs available: Class, Method, and SQL. Each tab includes basic information such as Probe Point Name, Invocation Count, and Response Time. This detailed performance data can help you identify low-level bottlenecks.
Topology	Includes the topology view associated with the selected application.

Under each named application node, ADP displays performance and other relevant information specific to that application. For example, by clicking the children nodes, the relevant data is displayed in the Main Display Window. Application response time and invocations measurements can be reached by clicking the panes in the Main Display Window.

In this section, we will further expand on the following nodes:

- Services
- Dependencies
- Deployments
- Workshop Projects
- Web Applications
- Stateless Beans
- Stateful Beans
- Entity Beans
- Message Driven Beans

Note: The number of children nodes available under each application node depends solely on the complexity of the selected application. Simple Java EE web applications will not have nodes like Workshop Projects, Stateless Beans, Stateful Beans, Entity Beans, and Message Driven Beans.

28.2.13.1 Services

The Services node includes all the external entry points associated with the selected application. When this node is selected, ADP displays a summary view in the Main Display Window. ADP displays the performance data associated with various entry points associated with the selected application.

Tip: The children nodes under the Services node include entry point specific performance data.

28.2.13.2 Dependencies

The Dependencies node shows a list of internal and external components and share resources that a specific application depends on for its normal operation. When the Dependencies node is selected, ADP displays all external references made by the application in the Main Display Window. The following is a list of columns and their descriptions ([Table 28–63](#)):

Table 28–63 *Dependencies Column Descriptions*

Column/Metric	Description
Name	Display name of the component or resource used by the application. If this is undefined in the Deployment Descriptor, the reference name for the component is used.
Reference	Reference name of the component or resource used by the application.

Table 28–63 (Cont.) Dependencies Column Descriptions

Column/Metric	Description
Reference Type	Component or resource type.
Referer Component	Name of the component that is part of the application which obtained the reference to external component or resource.
Referer Module	Name of the module that is part of the application which obtained the reference to external component or resource.

ADP displays all the references associated with components in the selected application.

The Dependencies node can be further expanded by clicking the plus (+) icon. The children nodes of the Dependencies node are organized by type. Here are the list of dependency types and their descriptions ([Table 28–64](#)):

Table 28–64 Dependency Types

Dependency Type	Description
Data Sources	All shared data sources used by the application
Entity Beans	All entity beans used by the application
Session Beans	All session beans used by the application
JMS Queues	All JMS queues used by the application for publishing JMS messages
JMS Topics	All JMS topics subscribed by the application
Web Services	All web services used by the application

When a specific node is selected, ADP displays relevant performance summary. These nodes can also be expanded by clicking the plus (+) icons. The expanded tree includes specific components and share resources used by the application.

The Performance summary view associated with the Data Sources node under Dependencies provides information on both connection pools and SQL statements.

For more information on the metric description, refer to [Section 28.1.11, "Metric Types"](#).

28.2.13.3 Deployments

The Deployments node shows the architecture of the deployed application. When this node is selected, ADP shows all the modules deployed as part of this application. The default view in the Main Display Window shows the active module-level call path.

[Table 28–65](#) lists the tabs available as part of this summary view and their descriptions.

Table 28–65 Deployment Tabs

Tab Name	Description
Module Level Execution	Shows the active calling relationships among various Java EE modules (EAR, WAR, JAR, and more). Shared resources are also included. This is the default Architecture View at the module level.

Table 28–65 (Cont.) Deployment Tabs

Tab Name	Description
Module Level	Shows the potential calling relationships among various Java EE modules. Shared resources are also included. By default, the Module Level tab is not enabled.
Instrumentation	Includes detailed performance data at the method level. The table includes caller components, caller method, callee (target) component, callee module, invocation count, and response time.
SQL Statement	Includes all SQL statements executed as part of this application. It also includes performance information such as invocation count and response time.

Active module-level call path is displayed as the default view for the Deployments node of a selected application.

Double-click a specific module to trigger ADP to display the architecture of the selected module.

Expand the Deployments node by clicking the plus (+) icon to reveal all the deployed modules in this application. Further expanding the nodes at the *module* level reveals components associated with the selected module. Further expanding the nodes at the *component* level reveals methods associated with the selected component.

When you select one of these children nodes (module, component, and method levels), ADP displays associated tabs for active call path diagram, static call path diagram, instrumentation and SQL statements.

Tip: Use the active call path diagram as a guide to identify entities with performance data. If an entity does not have performance data, ADP displays *No data available for the selected time frame* in the Main Display Window.

28.2.13.4 Workshop Projects

The Workshop Projects node includes performance information about modules and components created using the Oracle WebLogic Workshop. These modules and components include WebLogic Integration processes, WebLogic Integration web services, and WebLogic Portal pageflows.

Workshop Project node and its children nodes provide performance data associated with WLI processes, web services, and WLP pageflows.

When you select a specific children node, ADP displays detailed performance information.

28.2.13.5 Web Applications

The Web Applications node includes performance information related to the Web Applications modules and components associated with the selected application. Click the Web Applications node to reveal a performance summary in the Main Display Window. Click the plus (+) icon to expand the Web Applications node to reveal various web modules deployed as part of this application.

Click the plus (+) icon to expand on a specific web module and reveal different groupings for web components, for example, Pageflows, Struts Modules and Servlets. Clicking one of these nodes triggers ADP to display rolled up performance summary for the entire grouping. You can further expand these nodes by clicking the plus (+) icon to reveal more detailed information. Fully expanded Web Applications node contains all web modules organized by type.

Detailed performance information at the individual pageflow, struts action, and servlet levels will be displayed when you click the lowest level nodes.

28.2.13.6 Stateless Beans

The Stateless Beans node includes activity information related to the stateless EJB components associated with the selected application. Click the Stateless Beans node to reveal an activity summary in the Main Display Window. Click the plus (+) icon to expand the Stateless Beans node to reveal various stateless EJBs deployed as part of this application.

You can further select individual nodes to obtain detailed activity information. Selecting a specific Stateless Bean node triggers ADP to display detailed activity metrics.

The detailed view contains the following activity metrics ([Table 28–66](#)):

Table 28–66 *Stateless Beans Detail View*

Column/Metric	Description
EJB	Name of the stateless EJB.
In Use	Number of instances for a specific stateless EJB currently being used from the free pool. [Snapshot Count]
Idle	Number of instances for a specific stateless EJB currently in the idle state in the free pool. These bean instances are available for use. [Snapshot Count]
Waits	Number of threads currently waiting for a specific stateless EJB bean instance from the free pool. [Snapshot Count]
Timeouts	Total number of threads that have timed out waiting for an available bean instance from the free pool. [Aggregated Count]

Note: The metrics reported in the Stateless Beans node are reported by the MBean (Management Bean) of the EJB container. These activity metrics can be used for checking the overall health of the EJB container. When the EJB container is restarted, these metrics are reset.

28.2.13.7 Stateful Beans

The Stateful Beans node includes activity information related to the stateful EJB components associated with the selected application. Click the Stateful Beans node to reveal an activity summary in the Main Display Window. Click the plus (+) icon to expand the Stateful Beans node to reveal various stateful EJBs deployed as part of this application.

You can further select individual nodes to obtain detailed activity information.

The Stateful EJB Summary includes the following tables:

- Stateful EJB Cache
- Stateful EJB Transactions
- Stateful EJB Locking

28.2.13.7.1 Stateful EJB Cache Stateful EJB Cache table includes the following information ([Table 28–67](#)):

Table 28–67 Stateful EJB Cache

Metrics	Description
EJB	Name of the Stateful EJB
Hits	Total number of times an attempt to access the Stateful EJB instance from the cache succeeded [Aggregated Count]
Accesses	Total number of attempts to access the Stateful EJB instance from the cache [Aggregated Count]
Size	Number of beans instances from this Stateful Home currently in the EJB cache [Snapshot Count]
Activations	Total number of beans from this Stateful Home that have been activated [Aggregated Count]
Passivations	Total number of beans from this Stateful Home that have been passivated [Aggregated Count]

Tip: Passivation (serializing EJB state information to disk) and activation (reconstitute EJB state information from disk) are resource intensive operations. Ideally, Oracle recommends low level of activity in these metrics.

28.2.13.7.2 Stateful EJB Transactions Stateful EJB Transactions table includes the following information ([Table 28–68](#)):

Table 28–68 Stateful EJB Transactions

Metrics	Description
EJB	Name of the Stateful EJB
Commits	Total number of transactions that have been committed for this Stateful [Aggregated Count]
Rollbacks	Total number of transactions that have been rolled back for this Stateful [Aggregated Count]
Timeouts	Total number of transactions that have timed out for this EJB [Aggregated Count]

Tip: High number of EJB Transaction Rollbacks may indicate problems with the data used; for some reason the target database is unable to commit the change. High number of EJB Transaction Time-outs may indicate problems accessing the database including network outage, database lock contention, and database outage.

28.2.13.7.3 Stateful EJB Locking Stateful EJB Locking table includes the following information ([Table 28–69](#)):

Table 28–69 Stateful EJB Locking

Metric	Description
EJB	Name of the Stateful EJB
Entries	Number of Stateful EJB instances currently locked [Snapshot Count]
Lock Accesses	Total number of attempts to obtain a lock on an Stateful EJB instance [Aggregated Count]
Current Waiters	Number of Threads that currently waiting for a lock on an Stateful EJB instance [Snapshot Count]
Total Waiters	Total number Threads that have waited for a lock on an Stateful EJB instance [Aggregated Count]
Timeouts	Total number Threads that have timed out waiting for a lock on an Stateful EJB instance [Aggregated Count]

Tip: Pay attention to Current Waiters and Time-outs. These metrics can indicate possible performance problems caused by EJB Locking. Ideally, 0s should be displayed for these metrics.

ADP presents these metrics in a table format in the Main Display Window when you select the Stateful Beans node. Graphical representations of two metrics, Stateful EJB cache access, and Stateful EJB lock access, are displayed below the table.

By looking at the activities related to Stateful EJBs, you can determine if there any abnormal activities associated with Stateful EJBs.

Note: The metrics reported in the Stateful Beans node are reported by the MBean (Management Bean) of the EJB container. These activity metrics can be used for checking the overall health of the EJB container. When the EJB container is restarted, these metrics are reset.

28.2.13.8 Entity Beans

The Entity Beans node includes activity information related to the Entity EJB components associated with the selected application. Click the Entity Beans node to reveal an activity summary in the Main Display Window. Click the plus (+) icon to expand the Entity Beans node to reveal various Entity EJBs deployed as part of this application.

You can further select individual nodes to obtain detailed activity information. Selecting a specific Entity Bean node triggers ADP to display detailed activity metrics.

The Entity EJB Summary includes the following tables:

- Entity EJB Activity
- Entity EJB Cache
- Entity EJB Transactions
- Entity EJB Locking

28.2.13.8.1 Entity EJB Activity Entity EJB Activity table includes the following information ([Table 28–70](#)):

Table 28–70 Entity EJB Activity

Metrics	Description
EJB	Name of the Entity EJB.
In Use	Number of instances for a specific Entity EJB currently being used from the free pool. [Snapshot Count]
Idle	Number of instances for a specific Entity EJB currently in the idle state in the free pool. These bean instances are available for use. [Snapshot Count]
Waits	Number of Threads currently waiting for a specific Entity EJB instance from the free pool. [Snapshot Count]
Timeouts	Total number of Threads that have timed out waiting for an available bean instance from the free pool. [Aggregated Count]

Tip: Pay attention to Waits and Timeouts metrics. Activities in the Waits metric and increasing count in the Timeouts metric are signs that requests are waiting to be serviced by the EJB container. Ideally, 0 should be indicated for these metrics.

28.2.13.8.2 Entity EJB Cache Entity EJB Cache table includes the following information (Table 28–71):

Table 28–71 Entity EJB Cache

Metrics	Description
EJB	Name of the Entity EJB
Hits	Total number of times an attempt to access the Entity EJB instance from the cache succeeded [Aggregated Count]
Accesses	Total number of attempts to access the Entity EJB instance from the cache [Aggregated Count]
Size	Number of beans instances from this EJB Home currently in the EJB cache [Snapshot Count]
Activations	Total number of beans from this EJB Home that have been activated [Aggregated Count]
Passivations	Total number of beans from this EJB Home that have been passivated [Aggregated Count]

Tip: Passivation (serializing EJB state information to disk) and activation (reconstituting EJB state information from disk) are resource intensive operations. Ideally, Oracle recommends a low level of activity in these metrics.

28.2.13.8.3 Entity EJB Transactions Entity EJB Transactions table includes the following information (Table 28–72):

Table 28–72 Entity EJB Transactions

Metric	Description
EJB	Name of the Entity EJB
Commits	Total number of transactions that have been committed for this EJB [Aggregated Count]
Rollbacks	Total number of transactions that have been rolled back for this EJB [Aggregated Count]
Timeouts	Total number of transactions that have timed out for this EJB [Aggregated Count]

Tip: High numbers of EJB Transaction Rollbacks may indicate problems with the data used; for some reason the target database is unable to commit the change. High numbers of EJB Transaction Timeouts may indicate problems accessing the database including network outage, database lock contention, database outage, and more.

28.2.13.8.4 Entity EJB Locking Entity EJB Locking table includes the following information (Table 28–73):

Table 28–73 Entity EJB Locking

Metric	Description
EJB	Name of the Entity EJB
Entries	Number of Entity EJB instances currently locked [Snapshot Count]
Lock Accesses	Total number of attempts to obtain a lock on an Entity EJB instance [Aggregated Count]

Table 28–73 (Cont.) Entity EJB Locking

Metric	Description
Current Waiters	Number of Threads that currently waiting for a lock on an Entity EJB instance [Snapshot Count]
Total Waiters	Total number Threads that have waited for a lock on an Entity EJB instance [Aggregated Count]
Timeouts	Total number Threads that have timed out waiting for a lock on an Entity EJB instance [Aggregated Count]

Tip: Pay attention to Current Waiters and Timeouts. These metrics can indicate possible performance problems caused by EJB Locking. Ideally, 0s should be displayed for these metrics.

When you select the Entity Beans node, ADP presents these metrics in a table format in the Main Display Window. Graphical representations of the following metrics, Entity EJB in use, Entity EJB cache access, and Entity EJB lock access, are displayed below the table.

Expand the Entity Beans tree by clicking the plus (+) icon next to Entity Beans node. You can get the same summary as previously described for a specific Entity EJB.

By looking at the activities related to Entity EJBs, you can determine if there any abnormal activities associated with Entity EJBs.

Note: The metrics reported in the Entity Beans node are reported by the MBean (Management Bean) of the EJB container. These activity metrics can be used for checking the overall health of the EJB container. When the EJB container is restarted, these metrics are reset.

28.2.13.9 Message Driven Beans

The Message Driven Beans node includes activity information related to the message driven EJB components associated with the selected application. Click the Message Driven Beans node reveals an activity summary in the Main Display Window. Click the plus (+) icon to expand the Message Driven Beans node to reveal various message driven EJBs deployed as part of this application.

You can further select individual nodes to obtain detailed activity information.

The Message Driven EJB Summary includes the following tables:

- Message Driven EJB Activity
- Message Driven EJB Transactions

28.2.13.9.1 Message Driven EJB Activity Message Driven EJB Activity table includes the following information ([Table 28–74](#)):

Table 28–74 Message Driven EJB Activity

Metric	Description
EJB	Name of the Message Driven EJB.
In Use	Number of instances for a specific Message Driven EJB currently being used from the free pool. [Snapshot Count]
Idle	Number of instances for a specific Message Driven EJB currently in the idle state in the free pool. These bean instances are available for use. [Snapshot Count]
Waits	Number of Threads currently waiting for a specific Message Driven EJB instance from the free pool. [Snapshot Count]
Timeouts	Total number of Threads that have timed out waiting for an available bean instance from the free pool. [Aggregated Count]

Tip: Pay attention to Waits and Timeouts metrics. Activities in the Waits metric and increasing count in the Timeouts metric are signs that requests are waiting to be serviced by the EJB container. Ideally, 0 should be indicated for these metrics.

28.2.13.9.2 Message Driven EJB Transactions Message Driven EJB Transactions table includes the following information ([Table 28–75](#)):

Table 28–75 Message Driven EJB Transactions

Metric	Description
EJB	Name of the Message Driven EJB
Commits	Total number of transactions that have been committed for this EJB [Aggregated Count]
Rollbacks	Total number of transactions that have been rolled back for this EJB [Aggregated Count]
Timeouts	Total number of transactions that have timed out for this EJB [Aggregated Count]

Tip: High numbers of EJB Transaction Rollbacks may indicate problems with the data used; for some reason the target database is unable to commit the change. High numbers of EJB Transaction Timeouts may indicate problems accessing the database including network outage, database lock contention, database outage, and more.

ADP presents these metrics in a table format in the Main Display Window when you select the Message Driven Beans node. Graphical representation of the Message Driven EJB in use metric is displayed below the table.

By looking at the activities related to Message Driven EJBs, you can determine if there are any abnormal activities associated with Message Driven EJBs.

Note: The metrics reported in the Message Driven Beans node are reported by the MBean (Management Bean) of the EJB container. These activity metrics can be used for checking the overall health of the EJB container. When the EJB container is restarted, these metrics are reset.

28.2.14 Oracle WebLogic Resources

The Resources node under Oracle Enterprise Manager contains information for the managed domain organized by logical clusters, machines, servers, and more. You can look for low-level technology metrics organized by technology subsystems for a specific WebLogic Server.

The Resources tree includes the following nodes ([Table 28–76](#)):

Table 28–76 WebLogic Resources Tree

Example Node	Description
CSS Domain	Name of the WebLogic Domain configured
b-15/192.168.128.15	ID of the physical machine
cgServer	Name of the WebLogic Server configured
Applications	Performance measurements of all deployed applications running on this server
JDBC	Information of all configured JDBC resources for this server
JMS Servers	Information of all JMS destinations configuration for this server
Execute Queues	Information of all Execute Queues configured for this server
JVM	JVM information including Heap Size for this server
JRockit	JRockit information including Heap Size for this server
Modeling Status	Entities modeled by ADP for this server
ADP Modules	Status of the ADP Java Agent Module for this server

Expand these nodes by clicking the plus (+) icon next to the node name to get more information.

If the ADP OS Agent is deployed on the machine, clicking on the physical machine ID would show OS metrics collected by the OS Agent. These OS metrics include CPU Usage, Disk Usage, and Physical Memory Usage.

28.2.15 Oracle Resources

The Resources node under Oracle Enterprise Manager contains information for the managed domain organized by logical clusters, machines, servers, and more. You can look for low-level technology metrics organized by technology subsystems for a specific Oracle AS Server.

The Resources tree includes the following nodes ([Table 28–77](#)):

Table 28–77 Oracle Resources Tree

Example Node	Description
Managed System Resource Name	Top-level Resource name, for example, oc4j_soa
Oracle AS Server	Machine name which can be navigated to both within or outside a cluster, for example, oc4j_soa@192.168.1.119 which includes both the server name and the host server IP address
Applications	Performance measurements of all deployed applications running on this server
JDBC	Information of all configured JDBC resources for this server
JMS Servers	Information of all JMS destinations configuration for this server

Table 28–77 (Cont.) Oracle Resources Tree

Example Node	Description
Thread Pools	Performance information about all threads used by the container to process requests
JVM	JVM information including Heap Size for this server
BPEL Processes	Performance measurements about BPEL Processes deployed in the container
ESB	Performance measurements about ESB services deployed in the container
Modeling Status	Modeled entities for the container
ADP Modules	Status of the ADP Java Agent Module for this server
Applications	Performance information about the applications deployed in the container

Clicking the physical machine ID would show OS metrics. These OS metrics include CPU Usage, Disk Usage, and Physical Memory Usage.

28.2.16 Custom Metrics

The Custom Metrics node under Oracle Enterprise Manager contains all the custom metrics you defined. Currently ADP supports custom metrics for Java classes. When Custom Metrics node is selected, ADP displays various summaries. You can select individual entities to get more detailed performance information.

Expanding the Custom Metrics node reveals a list of Java classes with custom metrics configured.

The following is a list of columns in the Custom Class Performance table and their descriptions ([Table 28–78](#)):

Table 28–78 Custom Class Performance

Column/Metric	Description
Caller Class	Fully qualified name of the class that is making the inbound call
Caller Method	Method name in the class that is making the inbound call
Class	Fully qualified name of the class that is the destination of the inbound call
Invocation Count	Total number of times the inbound call is made
Response Time (ms)	Average response time of the inbound call in milliseconds

28.2.17 Status

Status in the navigation tree contains information for the ADP environment for the monitored WebLogic domain, WebSphere cell, or Oracle AS cluster. Select Status to see the ADP Java Agent status for the WebLogic domain.

The ADP Java Agent status includes the following ([Table 28–79](#)):

Table 28–79 ADP Java Agent Status

Column/Metric	Description
Server	Name of the WebLogic server, WebSphere cell, or Oracle AS cluster
Container Status	Operational status of the WebLogic, WebSphere, or Oracle AS server (running or not)
Agent In Sync	Version synchronization between ADP and ADP Agent status (true or false)
EJB Installed	ADP EJB installation status (true or false)

Table 28–79 (Cont.) ADP Java Agent Status

Column/Metric	Description
Agent Installed	ADP Java Agent installation status
Agent Activated	ADP Java Agent activation status
Agent Status	ADP Java Agent operational status
Server Type	Identifies server as administration, individual, or clustered server
Admin URI	Location of the domain admin server
Manager RMI Registry Host	Host name of the ADP RMI registry
Manager RMI Registry Port	Port number of the ADP RMI registry
EJB Major Version	ADP EJB major version
EJB Minor Version	ADP EJB minor version
EJB Build ID	ADP EJB build number - for version synchronization check
Agent Major Version	ADP Java Agent major version
Agent Minor Version	ADP Java Agent minor version
Agent Build ID	ADP Java Agent build number - for version synchronization check

Click the Modeling Status node under Status to see a table of all modeled entities in the managed domain. This table shows all the managed clusters, servers, and applications in the ADP environment. Mismatches between the Modeling Status table and your environment are indications of configuration problems.

You can use this information to debug and resolve ADP configuration issues.

28.2.18 Service Component Architecture (SCA)

Service Component Architecture (SCA) provides a set of features and services that simplify the process of detecting the presence of Service-Oriented Architecture (SOA) components.

Table 28–80 SCA Composites

Composite	Description
Services	Metrics related to Services defined on the SOA composite.
Wires	Metadata related to Wires defined in the SOA composite
References	Metrics related to References defined in the SOA composite
Components	Metrics related to Components within the SOA composite

28.2.18.1 Components

The following components make up the Service Component Architecture:

Table 28–81 Components in SCA

Component	Description
Decision Services	Metrics related to components in the Decision Services engine
Mediators	Metrics related to components in the Mediator engine

Table 28–81 (Cont.) Components in SCA

Component	Description
Human Workflows	Metrics related to components in the Human Workflow engine
BPEL	Metrics related to components in the BPEL engine

28.3 Exploring the Configuration Tab

Using the Configuration tab you can set up the resources you want to monitor using ADP.

The configurations explained in this section are:

- [Database Configuration](#)
- [Resource Configuration](#)
- [Service Level Objective Configuration](#)
- [Custom Metric Configuration](#)

A running ADP manager must be registered in Enterprise Manager. After the registration, Enterprise Manager continues to keep the manager as a valid manager even if it is down. When this occurs, the Enterprise Manager UI displays the ADP manager as Unreachable.

28.3.1 Database Configuration

The Database Configuration page lists the databases accessible to ADP which you want to monitor. You can configure a database to be used by ADP, edit an existing database configuration, delete a database configuration, and enable a configuration.

28.3.2 Resource Configuration

The Resource Configuration node in the Configuration tree enables you to create resources (for example, target application server domains) that can be monitored by ADP.

28.3.3 Service Level Objective Configuration

In ADP, thresholds configured for various measurements are called Service Level Objectives (SLOs). A service level objective is a measurable attribute, for example, availability. Service Level Agreements (SLA) are made up of SLOs.

Configuring SLOs is a key activity for establishing and maintaining an effective performance monitoring system. To configure a SLO, click the **Configuration** tab and select the **Service Level Objective Configuration** option.

ADP categorizes SLOs into the following types:

- **Performance**
Depicts the relative responsiveness of the monitored entity to the configured threshold.
- **Availability**
Informs you to what extent a particular entity is available to service requests.
- **Errors**

Informs you if the number of errors and exceptions encountered by this entity are approaching or violating the configured threshold.

- **Load**

Depicts how many operations have been performed and requests have been served by a particular entity.

ADP is aware of clusters. As such, these indicators display overall health of a particular entity across the entire cluster.

To configure a SLO, perform the following steps:

1. From the **Targets** menu, select **Middleware**. On the Middleware page, select **Application Dependency and Performance** from the Middleware Features menu. Ensure the Configuration tab is highlighted.
2. Select **SLO Blackout Configuration**.
3. You can view any existing SLO blackout.
4. Use this window to create, delete, or view the details of existing blackouts.

28.3.3.1 Creating a New SLO

When you select Service Level Objectives Configuration, ADP displays the Service Level Objective Configuration window. This window allows you to apply existing SLOs or create new ones. When you click **Create New SLO**, ADP guides you through the process of setting up a new SLO.

The steps for SLO creation are as follow:

1. Either select a SLO file or create a new SLO file. ADP can store SLO configurations in different files to improve configuration portability.
2. Define the SLO Entity Type. ADP automatically selects the appropriate entity type for you based on the selected monitoring element. For example, if you want to set a SLO on a Portal Desktop element, ADP automatically sets the Entity Type for you.
3. Other information is filled in by default. Normally, there is no need to modify the SLO Entity values.
4. When you are done setting the SLO Entity Type values, click **Create New SLO** to go to the second step of the SLO creation process, Defining the SLO Parameters.
Note: The (*) character means Select All. It is recommended that you do not use the (*) character.

Note: SLOs are hierarchical which allows you to set service levels at any level within the modeled hierarchy of an application.

28.3.3.2 Defining SLO Parameters

Follow these steps to define the SLO parameters:

1. Navigate to **Application Dependency and Performance**.
From the **Targets** menu, select **Middleware**. On the Middleware page, click the target of choice. On the **Home** tab, select the **Summary** region and click the **Application Dependency and Performance** link.
2. Expand the **Configuration** tab. Select **Service Level Objective Configuration**.
3. Either create a new SLO or edit an existing SLO.
4. Select the performance metric.

5. Define the monitoring window size, which determines how long the condition must persist before generating an alert.
6. Set threshold values for the SLO.
7. Select what actions to take when a trigger is fired. A list of preconfigured actions is available in the view pane.
8. Add new actions by going to the Action Configuration node in the Configuration Workspace.
9. Click **Save** to set the SLO for this monitored element.
10. You can delete unwanted SLOs for any element from this window.

Types of SLOs

ADP categorizes SLOs as Performance, Availability, Error, and Load.

SLO Events Viewer

Right-click on any tree node and select **View Service Level Objective Events** to open a new window. You can see all the SLO violation events triggered for the selected entity. ADP automatically applies a filter to show only relevant events.

Once new SLOs are added, ADP updates the relevant graphs to visually display these new thresholds. [Table 28–82](#) explains the different line types.

Table 28–82 SLO Line Types

Line Description	Description
Solid Red Line	A violation threshold that triggers on high.
Solid Yellow Line	A cautionary threshold that triggers on high.
Dashed Red Line	A violation threshold that triggers on low.
Dashed Yellow Line	A cautionary threshold that triggers on low.

28.3.3.3 SLO Blackout Configuration

You can prevent having unwanted alerts being fired during planned or unplanned down time. The SLO Blackout Configuration node in the Configuration tree enables you to create time periods when information will not be monitored for a specific SLO. You can define blackouts by a SLO file, an individual SLO, or by entity.

28.3.3.4 Creating and Maintaining SLO Blackouts

You can prevent having unwanted alerts being fired during planned or unplanned down time. SLO Blackout Configuration enables you to create time periods when information will not be monitored for a specific SLO. You can define blackouts by a SLO file, an individual SLO, or by entity.

To create and maintain SLO blackouts, perform the following steps:

1. Navigate to Application Dependency and Performance.
From the **Targets** menu, select **Middleware**. On the Middleware page, click the target of choice. Select **Application Dependency and Performance** from the Middleware Features menu.
2. Expand the **Configuration** tab. Select **SLO Blackout Configuration**.
3. You can view any existing SLO blackout.

4. Use this window to create, delete, or view the details of existing blackouts.

Creating SLO Blackout

1. Click **Create SLO Blackout** to view the detail window.
2. On the SLO Blackout File page, type the name of the blackout in the New SLO Blackout File field. Click **Continue**.
3. On the SLO Blackout Configuration page, fill in the fields. Refer to [Table 28–83](#) for details.

Table 28–83 SLO Blackout Configuration

Column/Metric	Description
Blackout Name	Type in the name.
Description	Type in the description of the SLO you are creating.
Blackout By SLO File	Use to blackout at the file level. The SLO files display in a list where you can select them or cancel out of the window. This option restricts the blackout to the SLO file name.
Blackout By SLOs	Use to blackout at the SLO level. The SLOs display in a list where you can select them or cancel out of the window. This option restricts the blackout to the SLO name.
Blackout By Entity	Use to blackout at the entity type level. Click the Blackout by Entity: button to view the list of entity types. Select the entity. This option restricts the blackout to the entity type selected.
year, month, date, hour, minute, duration	Use the guidelines to the right of these columns to enter the appropriate information.
recurring	Select how often you would like to run this blackout event from the list.

Viewing SLO Blackout Summary List

1. Click **SLO Blackout Summary List**.
2. View the details on the existing SLO Blackout events.
3. Click **Show SLO Blackout List** to return to the previous window.

Deleting SLO Blackout

1. Select an existing event on the list.
2. Click **Delete SLO Blackout**.
3. Confirm that you want to delete the entry and click **Yes**.

28.3.3.5 Propagating Threshold Violation Events

ADP is designed to propagate threshold violation events up the hierarchy. Therefore, when a SLO is set on a lower level metric, the higher level health indicator light becomes activated. Additionally, the health indicator light for the application server that hosts this component also becomes active. Oracle calls this *containment approach* to SLO event propagation. When a lower level SLO is violated, the violation event propagates all the way up the hierarchy and changes the status of all containers for this event.

28.3.4 Event Integration

Use the Enterprise Manager Incident console to check for events fired as a result of SLO violations in ADP.

To access the ADP alerts:

1. From the Enterprise menu, select **Monitoring**, then select **Incident Manager**.
2. In the Views region, select **Events without incidents**.

Look for events with target type "Application Deployment" and "Application Dependency and Performance Alert". These are the ADP alerts.

28.3.5 Custom Metric Configuration

There are cases where additional instrumentation is needed based on your specialized requirements. Custom metrics allow you to instrument a class or method of your choice and receive performance metrics collected by the ADP agent.

To create a metric configuration, do the following:

1. From the **Targets** menu, select **Middleware**. In the Related Links sections, select **Application Dependency and Performance**.
2. Click the **Configuration** tab, choose the configuration in which you are interested. Click **Custom Metric Configuration**.
3. In the right pane, click the **Create Custom Metric** button.
4. On the Custom Metric File page, choose whether to use an existing .xml file or a new file. If you choose a new file, the ADP Manager will create the new .xml file. Click **Continue**.
5. On the Custom Metric Configuration page, provide the following information:
 - Resource name is a monitored WebLogic domain or Oracle Application Server or WebSphere cell.
You created a name when you configured ADP to monitor. The same name is used here during custom metric configuration.
 - Class name is the name of the implementation class in the code. You are required to enter a fully qualified class name.
 - Method name is the name of the implementation method in the code.

After you define the custom metrics, restart the application server instances associated with these customizations. The new custom metrics will be listed under the Custom Metrics node in the ADP navigation tree.

The newly configured custom metric provides class level performance data, for example, invocation count and response time.

28.4 Exploring the Registration Tab

The managers perform complex mathematical modeling and statistical calculations with summarized data from all Java Agents.

Using the Registration tab, you can add, edit, and remove Managers configured to Enterprise Manager. By accessing ADP through Remote Method Invocation (RMI), you can manipulate all the managers configured to Enterprise Manager through a secured protocol.

28.4.1 Using RMI Configuration for Managers

In ADP, the Configuration tab lists all the managers currently configured to Enterprise Manager. By using the Configuration for Managers feature, you can access Application Dependency and Performance through Remote Method Invocation (RMI). You can then manipulate all the managers configured to Enterprise Manager through a secured protocol. The following sections provide additional information.

- [Adding a New Manager \(RMI Configuration\)](#)
- [Editing a Previously Configured Manager \(RMI Configuration\)](#)
- [Removing or Disabling a Previously Configured Manager](#)

The Configuration tab displays only if the Enterprise Manager user is an Administrator as defined by examining the user's role.

28.4.2 Adding a New Manager (RMI Configuration)

The first time the Registration tab displays there are no managers in the Managers tree. To add a new manager, perform the following steps:

1. Navigate to the **Application Dependency and Performance** feature.

From the **Targets** menu, select **Middleware**. In the **Related Links** section, click **Application Dependency and Performance**.

2. In the **Registration** tab, click the **Managers** node in the tree.
3. Type the new manager information in the Main Display window.
4. Decide whether this manager should be monitored.

Request monitoring provides end-to-end visibility into requests, localizes end-user performance problems to specific application deployments, and provides a platform for context-based drill down diagnostics.

When you select **Enable Request Monitoring**, ADP creates and sets up targets for collecting request performance data. If you do not select Enable Request Monitoring, the ADP manager is only registered in Enterprise Manager.

Note: The grayed out information represents configuration data for connecting to the ADP manager by way of a secure protocol, for example Key Store, Trust Store, and passwords. This information is extracted from the ADP manager by way of the RMI call.

5. If you enable request monitoring on an existing manager, click **Upload** to populate the manager configuration properties to the ADP target in Request Monitoring.
6. Click **Test Connect** to test the connection to the new manager. Should the test connection fail, this may be because the manager is not running or the manager is not yet installed.
7. Click **Add**.

Once the manager is added, the name of the manager will display in the Configuration tab under the Managers node in the tree.

28.4.3 Editing a Previously Configured Manager (RMI Configuration)

To add a previously configured manager, perform the following steps:

1. Click + (plus sign) next to the Managers node in the tree, then select the subnode for the manager you want to edit.
2. After you make changes to the manager information, click **Update**. This results in the manager entries in the Enterprise Manager repository to be updated with the new values.

If a manager is configured before using this Enterprise Manager configuration page, Enterprise Manager continues to keep the manager as a valid manager even though the manager may be down or permanently removed.

The list of managers is not refreshed.

28.4.4 Removing or Disabling a Previously Configured Manager

To remove a configured manager, perform the following steps:

1. Navigate to **Application Dependency and Performance**.
From the **Targets** menu, select **Middleware**. In the **Related Links** section, click **Application Dependency and Performance**.
2. Click the **Registration** tab.
3. Click + (plus sign) next to the Managers node in the tree, then select the subnode for the manager you want to remove.
4. Click **Remove** in the main pane.

Deleting a manager from Enterprise Manager does not uninstall and remove the manager from the remote host where the manager is located and may be running. Remove only deletes the manager entry from the Enterprise Manager repository.

To shut down the manager after the Remove operation, execute the `acshut.sh/.bat` command from the command line.

To disable a configured manager:

1. Click + (plus sign) next to the Managers node in the tree, then select the subnode for the manager you want to disable.
2. Deselect **Enable Request Monitoring**.
3. Click **Update**.

When you deselect the Enable Request Monitoring option, the manager settings are preserved. The UI displays these managers as disabled. There will not be any further information under the disabled manager in the tree.

28.5 Using emctl to Manage the ADP Diagnostics Engine

You can use `emctl` commands to start and stop the ADP Diagnostics Engine. This is useful when a running ADP Diagnostics Engine may cause an upgrade process to fail. See [Table 28–84](#) for the a list of `emctl` commands.

Table 28–84 Extended ADP emctl Commands

Command	Description
<code>emctl extended oms adp list</code>	Queries and lists all the ADP Managed servers from the Repository.

Table 28–84 (Cont.) Extended ADP emctl Commands

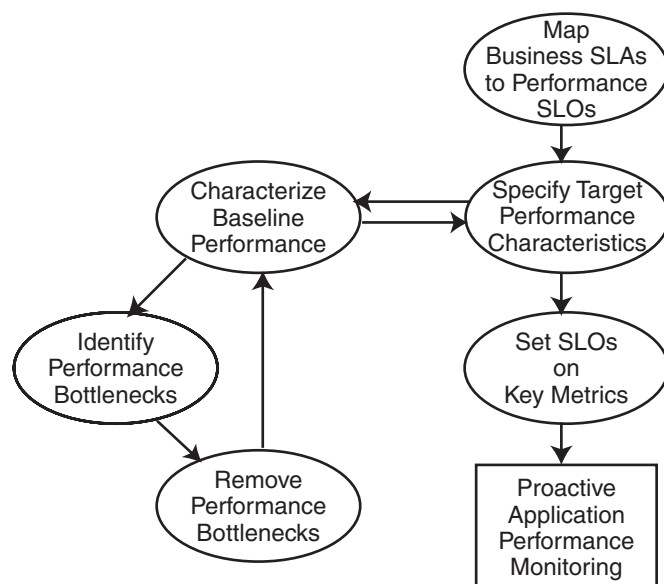
Command	Description
emctl extended oms adp start -server=<server_name1>,<server_name2>... For example: emctl extended oms adp start -server=EMADPENGINE,MYADPDMGR	Starts the ADP Managed servers mentioned in command line arguments. The servers could be running on the same local host on which the OMS is running or can be running on a remote host.
emctl extended oms adp start -all	Starts all ADP Managed servers on the same local host on which the OMS is running.
emctl extended oms adp start -global	Starts all ADP Managed servers, even if they are running on remote hosts (remote to this OMS host).
emctl extended oms adp stop -server=<server_name1>,<server_name2>... For example: emctl extended oms adp stop -server=EMADPENGINE,MYADPDMGR	Stops the ADP Diagnostics Managed servers mentioned in command line arguments. The servers could be running on the same local host on which the OMS is running or can be running on a remote host.
emctl extended oms adp stop -all	Stops all ADP Managed servers that are running on the same local host on which the OMS is running.
emctl extended oms adp stop -global	Stops all ADP Managed servers, even if they are running on remote hosts (remote to this OMS host).
emctl extended oms adp status -server=<server_name1>,<server_name2>... For example: emctl extended oms adp status -server=EMADPENGINE,MYADPDMGR	Shows the status of the ADP Diagnostics Managed servers mentioned in command line arguments. The servers could be running on the same local host on which the OMS is running or can be running on a remote host.
emctl extended oms adp status -all	Status of all the ADP Engines in this domain.
emctl extended oms adp -help	Shows the online help for the ADP Diagnostics commands.

ADP Methodology

The Enterprise Manager Application Dependency and Performance (ADP) capabilities automatically select performance metrics and track contextual relationships for various applications. The methodology focuses on other important activities to allow you to setup and maintain an effective application performance monitoring environment.

These activities include the following:

Figure 29–1 Steps of ADP Methodology



This methodology describes a series of steps for users to establish and maintain a proactive application performance monitoring environment leveraging the Enterprise Manager Application Dependency and Performance capabilities. [Figure 29–1](#) illustrates these steps in a sequential order.

Methodology steps:

1. Map business Service Level Agreements (SLAs) to performance Service Level Objectives (SLOs).

The process of using agreed business SLAs to determine the value of performance SLOs.

2. Specify target performance characteristics.

Specify the ideal application performance characteristics using performance SLOs identified in step 1.

3. Characterize baseline performance.
4. Identify performance bottlenecks.
5. Remove performance bottlenecks.

Steps 3, 4, and 5 should be grouped together to form a process of incremental performance improvement. Iterations of this process may be required to improve the application performance to meet the performance target as specified in step 2.

6. Set SLOs on key metrics.

Once application performance reaches the targeted goal, you need to set performance SLOs on key metrics to establish a proactive monitoring environment. This environment provides you with warnings when key performance metrics start to report abnormalities. These warnings enable you to proactively solve potential problems before they begin to impact business.

This chapter explains the following activities in more detail:

- [ADP Methodology Activities](#)
- [Mapping Business SLAs to Performance SLOs](#)
- [Characterizing Baseline Performance](#)
- [Identifying Performance Bottlenecks](#)
- [Setting SLOs on Key Metrics](#)

29.1 ADP Methodology Activities

The ADP methodology activities include the following:

- [Mapping Business SLAs to Performance SLOs](#)
- [Specifying Target Performance Characteristics](#)
- [Improving Performance](#)

29.1.1 Mapping Business SLAs to Performance SLOs

To successfully setup a proactive application performance monitoring environment, the first step is to map a set of business objectives to a set of performance thresholds for you to monitor. These business objectives are often referred to as business service level agreements (SLAs). These business SLAs provide the basic application performance requirements at a high level. As such, mapping these high level SLAs to low level performance thresholds is often a very difficult activity to do well.

Using tools that only measure performance at technology levels (EJB, JSP, servlet, portlet, SQL calls, and so on) to perform this type of activity continues to be very difficult as the correlations between low-level metrics and high-level objectives are often fuzzy at best. Consequently, the mapping activity is considered by many as an art rather than a science.


By measuring performance at both technology and functional levels, Enterprise Manager makes this mapping activity significantly less complicated. Since functional metrics measure performance for high level constructs such as business processes or portal desktops, mapping business SLAs directly to performance SLOs (Service Level Objectives) is straightforward.

29.1.2 Specifying Target Performance Characteristics

Defining the target performance characteristics for the monitored applications is the next step after mapping business SLAs to performance SLOs. Since these SLOs represent absolute minimal performance requirements for these applications, using these *violation* thresholds as target performance characteristics makes little sense. Instead, you need to define what performance range is acceptable for normal operation and when to send out cautionary alerts for abnormal activities.

For some applications, it may be sufficient to just specify a set of *cautionary* performance thresholds. Application performance monitoring tools, such as ADP, will send out cautionary alerts if these thresholds have been breached. Since these thresholds are cautionary, it may be acceptable to have a few violations before an alert is sent out. By defining the minimal violation duration, you can minimize the number of duplicate alerts generated. [Figure 29–2](#) illustrates this concept.

Figure 29–2 Control Number of Alerts



In [Figure 29–2](#), the minimal violation duration is defined to be 15 seconds. So if the cautionary state does not persist for more than 15 seconds, no cautionary alert would be fired.

For other applications, it is necessary to define both a high and low performance thresholds. Having both thresholds would effectively define a normal range of operation for these applications. With ADP, both high and low triggers can be set for any SLO.

With a set of clearly defined target performance characteristics, you are able to determine how much performance tuning is needed to achieve the ideal performance range. You will also have a set of cautionary performance thresholds to enable a proactive application performance monitoring environment to be established.

29.1.3 Improving Performance

The activities explained in this section should be grouped together as a single performance improvement process. This process would start with characterizing the baseline performance of our application, move on to identifying performance bottlenecks, and finish with removing performance bottlenecks. We would continue to perform these activities in iterations until the performance of our application meets the target characteristics.

29.1.3.1 Characterizing Baseline Performance

Once the specification of the target performance characteristics is completed, the next activity is to capture the performance baseline for our application. The performance baseline will be compared with the set of target performance characteristics to determine if further performance improvement is needed. If so, you will improve application performance iteratively through the next two steps until the performance meets the target characteristics.

29.1.3.2 Identifying Performance Bottlenecks

To identify performance bottlenecks, you must first isolate performance abnormalities in your performance baseline. Once you isolate a performance abnormality, you need to determine if this issue is a localized occurrence or a systematic problem. By using monitoring and diagnostic tools available to you, you can perform the analysis needed to identify the cause of the performance bottleneck.

29.1.3.3 Removing Performance Bottlenecks

Once these performance bottlenecks are identified, you need to determine how to remove them. The strategies for bottleneck removal vary by cause. A few examples follow:

- If the cause of the bottleneck is an application defect, the strategy would involve the application development team.
- If the bottleneck is caused by a configuration problem, you would request assistance from system administrators.
- If the bottleneck is in the application server or framework, you would seek help from those vendors.

The following is a list of possible bottleneck removal activities:

- Change application code to fix defects
- Modify environment setting to fix configuration problem
- Install patches to fix software defects
- Replace defective hardware
- Upgrade network infrastructure
- Add computing resources
- Remove resource hogging programs
- Tune back-end connectivity and response time

As you can see from the list, the Remove Performance Bottlenecks activity varies widely by cause. Correctly and quickly finding the appropriate groups to help resolve performance bottlenecks is the key for success for this activity. Once the performance bottleneck removal is completed, you must redo the Characterize Baseline Performance activity to confirm the fix implemented indeed improved performance.

29.1.3.4 Setting SLOs on Key Metrics

In addition to setting application specific performance thresholds, it is also important to set performance thresholds on key system metrics and on some selective component metrics. Setting these thresholds will help you establish an early warning system and alert you to smaller issues before they manifest into big production problems.

Setting SLOs on key system metrics involves some basic understanding of how the system behaves under load. If the system becomes unstable or performs poorly when it runs out of free JDBC connections or idle ExecuteQueue threads, these system metrics should be monitored.

To determine which system metrics to monitor, it is critical to figure out the correlations between overall system performance and specific system metrics. You would use this information to decide which of the system metrics to monitor. Once appropriate system metrics are identified, you will then determine the performance

range for normal operation and figure out the cautionary as well as violation thresholds.

While it is fairly straightforward to determine which system metrics to monitor and what system metric performance thresholds to set, setting SLOs on key component metrics is significantly more difficult. In theory, you can assume performance degradation at the component level would negatively impact application level performance. However, this assumption may not accurately reflect reality.

To predict application performance by monitoring component level performance metrics, there must be a very strong correlation between the performance of a specific component and that of the application. Sometimes, a drop in performance in one component is compensated by a jump in performance in another. These performance changes in opposite directions at the component level would essentially result in little change at the application level. Therefore, you must be careful not to draw conclusions by monitoring the performance of a few components unless there are strong correlations.

The last task to perform is to associate various actions and responses for various threshold violations. Once these associations are completed, you can begin to use your proactive application performance monitoring environment.

29.2 Mapping Business SLAs to Performance SLOs

One of the primary reasons companies purchase solutions to establish proactive application performance monitoring is the demand to meet business SLAs. Business SLAs for enterprise applications are a set of service level expectations defined by internal or external customers. In most cases, these business SLAs are defined at such a high level, they are not useful for setting thresholds in application performance monitoring tools.

As a result, the process of mapping business SLAs to performance SLOs is extremely important for companies to meet the service requirements set forth by their customers. Since Enterprise Manager monitors performance at both functional and technology levels, it is easy to perform this mapping exercise.

In this section, our example explains how to use the Enterprise Manager Application Dependency and Performance features to determine the proper performance threshold values for a set of business SLAs.

In the example, we were given the following high-level business SLAs:

Table 29–1 Example - Guidelines for Business SLAs

Business SLA	SLA Requirement
Fast customer self-service portal.	On average, pages in customer self-service portal should load within 2 seconds. This SLA must be fulfilled 99% of the time.
Customer service representative portal must be as fast as mainframe system.	All pages in customer service rep. portal must load within 6 seconds. This SLA must be fulfilled 99.9% of the time.
Fast to schedule a service call.	On average, scheduling a service call should take less than 30 seconds. This SLA must be fulfilled 99.99% of the time.

Let's map the first business SLA to a performance SLO. This SLA requirement states that the average response time for customer self-service portal (desktop) should be less than 2 seconds. In Enterprise Manager, we would set a high-level performance SLO at the desktop. Using the hierarchy in the ADP UI, you would select the *customer* desktop and right-click to set the SLO.

Because ADP monitors performance at both functional and technology levels, you can directly translate business SLAs to SLOs on functional metrics. In our example, it is the response time for portal desktop *customer*. For our example, we would proceed to set a violation SLO and a warning SLO.

We can calculate how often violations occur to figure out whether or not our current system is able to meet the SLA requirement 99% of the time. With Enterprise Manager, we can see whether there are any obvious violations. If there are any violations or close calls, we should confirm by examining actual data. If we have data for at least 24 hours, we would use the export capability within Enterprise Manager Application Dependency and Performance to prepare raw data for this calculation.

For the other two business SLAs, we would set performance SLOs on the appropriate metrics.

29.3 Characterizing Baseline Performance

Also known as performance base-lining, characterize baseline performance involves a set of activities to capture the baseline performance of a system under specific level of load. For example, you can measure the baseline performance of a portal application deployed to a WebLogic cluster.

For example, you can display the performance data during the first four hours of a load test. The number of active sessions grows at a steady pace for the first ninety minutes. Eventually, the number of active sessions stays at approximately seven hundred as the number of new and expiring sessions reach an equilibrium.

Visualize the portal performance following a typical pattern of slow performance initially and gradually reaching a steady state. The initial slow performance is expected as the application server load components into memory and populates data into its caching mechanism. The performance improves gradually and reaches a steady state after approximately thirty minutes. The performance pattern during this initial thirty-minute period can be characterized as *startup performance during increasing load*. After thirty minutes, performance of the portal application stabilizes.

Enterprise Manager's ability to quickly establish an application performance monitoring environment allows you to carry out *characterize baseline performance* painlessly. Because Enterprise Manager is able to monitor at cluster level as well as at individual server level, it can characterize performance for the entire cluster or individual servers.

By verifying that a less loaded server has lower resource utilization and faster performance, you can draw the following observations about the performance characteristics of a portal application running on this environment:

- Since Server A's resource utilization is near maximum, we can use the load on that server as the maximum limit for individual servers. We can calculate individual server maximum load limit by using the load metric provided by the Enterprise Manager Application Dependency and Performance features.
- We should examine the load balancing algorithm and the configuration of the load balancer.

Comparing load and resource usage of two servers in a cluster confirms resource usage is inversely correlated to the load.

Note: This is a very basic performance characterization of an individual server. Performance of a multi-server cluster cannot be calculated by multiplying performance characteristics of individual servers because of the overhead involved with a clustered configuration. True cluster level performance must be measured with application performance monitoring tools like Enterprise Manager.

29.4 Identifying Performance Bottlenecks

Enterprise Manager can also be used to quickly identify performance bottlenecks in QA, staging, and production settings. You can use the hierarchical model within the Enterprise Manager Application Dependency and Performance feature to identify an application performance bottleneck. Furthermore, you can use Enterprise Manager to track down an application performance problem caused by resource starvation.

29.4.1 Determining System Level Performance

Since the performance problem seems to affect all components in the same way, we should suspect there is some type of performance degradation at the system level. To view system level performance data, we would look under the *Resources* hierarchy. Performance metrics under the Resources hierarchy provides the raw data for us to perform correlation analysis. This type of analysis is needed to determine whether resource starvation is the cause of the performance slowdown.

Graphs can reveal some interesting patterns.

- OS Agent Abnormalities

We noticed a sudden drop in CPU utilization and a sudden increase in disk utilization during the time period in question. This pattern indicates a large amount of virtual memory paging activities on this machine. Memory paging to disk is extremely expensive and slows down request processing as indicated by lowered CPU utilization. To understand why page is occurring, we will take a look at performance metrics on the JVM.

- JVM Heap Size

We noticed that for the initial twelve hours of this example, both total JVM heap size and free JVM heap size grew at a steady pace. The growth of the total JVM heap size stopped at 512 MB - an expected behavior since we configured WebLogic to have a maximum heap size of 512 MB. While we expect the free JVM heap size to stop growing after total heap size reached 512 MB, the free JVM heap size actually starts to drop. Combining this information with some previously obtained information such as no sudden increase in load, we can conclude that there is a high likelihood of a memory leak.

This abnormal consumption of memory caused the total JVM heap to reach its pre-defined maximum. It is also very likely this memory leak caused the increase in virtual memory paging activities and corresponding reduction in CPU utilization. This reduction in CPU utilization impacts the response time for all components running in this machine including JDBC Connections.

- Memory Leak

We were able to identify a memory leak in the WebLogic JVM gradually caused resource starvation and eventually impacts application performance. In order to further diagnose this problem, a deep-level memory profiling tool is required to understand memory usage of the JVM.

29.5 Setting SLOs on Key Metrics

This step in the Oracle Methodology allows you to proactively monitor key system metrics to avoid catastrophic failures such as server hangs (non-responsive), server crashes, cluster hangs, and more. The ability to recognize signs leading up to these catastrophic failures is a must to maintain quality of service for your WebLogic infrastructure.

You can proactively set thresholds and actions for key WebLogic system metrics.

[Table 29–2](#) lists the key system metrics for the WebLogic Platform:

Table 29–2 List of Key System Metrics for WebLogic

Key System Metric	Reason to Monitor
ExecuteQueue Idle Thread Count	Running out of ExecuteQueue threads is often a precursor to application server hangs (non-responsive). In some severe cases, when the application server runs out of ExecuteQueue threads, all of its operations would stop working.
ExecuteQueue Pending Request Count	A steady increase in the number of ExecuteQueue pending requests is also a precursor to server hangs. This metric is inversely correlated with the ExecuteQueue Idle Thread Count metric.
Total JVM Heap Size	There are two reasons to monitor this metric: <ol style="list-style-type: none"> 1. If total JVM heap size grows to predefined maximum, a cautionary event should be fired notifying the administrator. 2. If total JVM heap size suddenly drops to 0, this may be an indication of a JVM crash or a non-operational application server.
Free JVM Heap Size	A steady decrease in the free JVM heap size is an indicator of either a memory leak or misconfigured application server. A JVM running out of heap will experience instability and performance degradation as garbage collector and JVM competes for resources to perform cleanup and object creation respectively.
Open Sessions Count	If open session count drops to 0 and remains at 0 for a period of time, some investigation is warranted. Often this pattern indicates a network or load balancing problem.
Application Invocation Count	If application invocation count drops to 0 and remains at 0 for a period of time, some investigation is warranted. While this pattern often indicates a network or load balancing problem, it could also be a symptom of a hanged server.

Understanding these key WebLogic system metrics, setting the SLO thresholds and assigning appropriate responses are critical to establishing a proactive monitoring. In this example, we will configure SLOs and actions with ADP.

The first task is to set a cautionary and a violation SLO for ExecuteQueue Idle Thread Count metric so the appropriate person can be alerted when available ExecuteQueue is running low. To configure SLOs, right-click on the Execute Queues metric and select **Configure service level objects**. In this example, we will create the following SLOs for ExecuteQueue Idle Threads:

Table 29–3 SLOs for ExecuteQueue Idle Threads

SLO Name	Metric	Threshold Type	Threshold Value	Trigger On
Low ExecuteQueue Idle Threads	Metric.J2EE.Dispatcher.IdleThreads	Cautionary	3	Low
ExecuteQueue Idle Threads Exhaustion	Metric.J2EE.Dispatcher.IdleThreads	Violation	0	Low

When SLO trigger is set to Low, ADP will fire an alert when current measurement reaches the threshold value AND the previous measurement has a higher value than the threshold.

For example, we would create the following actions for the SLOs previously configured:

Table 29–4 Action for SLO

SLO Name	Action Name	Action Type
Low ExecuteQueue Idle Threads	Enter Low ExecuteQueue Idle Threads event into server log	Log
ExecuteQueue Idle Threads Exhaustion	Email ExecuteQueue Idle Threads Exhaustion alert	Email
ExecuteQueue Idle Threads Exhaustion	Send ExecuteQueue Idle Threads Exhaustion SNMP trap to HP Overview	SNMP
ExecuteQueue Idle Threads Exhaustion	Enter ExecuteQueue Idle Threads Exhaustion event into server log	Log

After configuring these SLOs and actions, we now have a proactive monitoring environment to detect ExecuteQueue resource starvation related problems before a catastrophic event occurs. We would use this approach to establish proactive monitoring for other key WebLogic system metrics.

The following is the Oracle recommendation:

Table 29–5 ExecuteQueue Pending Requests

SLO Name	Metric	Threshold Type	Threshold Value	Trigger On
ExecuteQueue Pending Request Warning	Metric.J2EE.Dispatcher.PendingRequests	Cautionary	5 ~ 10 ¹	High
ExecuteQueue Pending Request Violation	Metric.J2EE.Dispatcher.PendingRequests	Violation	10 ~ 20	High

¹ **Threshold values for these SLOs vary by environment.** Figuring out what threshold values to use is an iterative process. Users should gather information about the performance characteristic of their WebLogic environment as the first step. Based on this information, users can set SLOs accordingly. As users continue to improve the performance of their WebLogic environment, they should re-evaluate these threshold values and change them as needed.

When SLO trigger is set to High, ADP will trigger an alert when current measurement hits the threshold value.

Table 29–6 Total JVM Heap Size

SLO Name	Metric	Threshold Type	Threshold Value	Trigger On
JVM Heap Reached Max	Metric.J2EE.JVM.HeapSizeCurrent	Cautionary	512 MB ¹	High
JVM Heap Reached 0	Metric.J2EE.JVM.HeapSizeCurrent	Violation	0 MB	Low

¹ **Threshold value for this SLO varies by environment.** Users would set this value to the maximum heap size specified in the WebLogic configuration file.

Table 29–7 Free JVM Heap Size

SLO Name	Metric	Threshold Type	Threshold Value	Trigger On
Low JVM Free Heap Warning	Metric.J2EE.JVM.HeapFreeCurrent	Cautionary	72 MB	Low
Low JVM Free Heap Violation	Metric.J2EE.JVM.HeapFreeCurrent	Violation	24 MB	Low

Table 29–8 Open Session Count

SLO Name	Metric	Threshold Type	Threshold Value	Trigger On
No user session in system for 5 minutes	Metric.J2EE.WebApplication.OpenSessionCurrentCount	Cautionary	0 ¹	Low

¹ In the example, this SLO would have a measurement window of 5 minutes. By setting the measurement window to 5 minutes, ADP will fire an alert only if this condition persists for at least 5 minutes.

Table 29–9 Application Invocation Count

SLO Name	Metric	Threshold Type	Threshold Value	Trigger On
No application invocation in system for 5 minutes	Metric.J2EE.Servlet.InvocationTotalCount	Cautionary	0	Low

Setting these SLOs and corresponding actions establishes a proactive monitoring environment for your WebLogic deployment. This proactive monitoring approach allows you to identify problems leading up to catastrophic problems before they impact your system's performance and availability.

29.6 Conclusion

The ADP Methodology is a critical aspect of your application performance management strategy. By following this methodology carefully, you will be able to use ADP to improve your ability to proactively monitor the performance and availability of your deployed applications and WebLogic infrastructure. ADP's automation reduces time, effort, and errors associated with manual processes. This allows ADP users to focus on other crucial activities such as the ones listed in the Oracle Methodology.

Frequently Asked Questions About Application Dependency and Performance

This chapter answers frequently asked questions regarding Application Dependency and Performance.

This chapter includes the following sections:

- [Can I Erase the darchive Directory?](#)
- [How Do I Undeploy the Agent?](#)

30.1 Can I Erase the darchive Directory?

In general, you should not erase the darchive directory while the ADP manager is running. If you erase the darchive directory after shutting down the ADP manager, the darchive directory will be re-populated when the ADP manager next restarts.

If the size of the darchive directory is of concern, determine which subdirectories are occupying the most space. If there are large directories that do not contain Java class files (for example `.`, `..`, or `wlserver10.3`), modify the `Acsera.properties` file to exclude them. Search for `Model.StaticAnalysis.ExcludeClassPaths` and add the path to the list, separated by comma.

The darchive directory is located under `$ACSERA_HOME`:

```
$GCDomain/EMGC_ADPMANAGER1/ADPManager.ear/ADPManager.war
```

30.2 How Do I Undeploy the Agent?

There is a drop-down menu in the last step in the deployment process that, by default, is set to *deploy*. Change it to *disable* to remove the agent startup arguments from the application servers, and to *remove* to erase the agent files from the application servers.

Note: On some platforms, for example Windows, you need to restart the application servers before the Remove command can be run. Otherwise the Remove command may run into File Still In Use errors.

To undeploy the agent, follow these steps:

1. From the Targets menu, select **Middleware**, then select **Middleware Features** (drop-down list). Select **Application Dependency and Performance**.
2. Click the **Configuration** tab.

ADP Configuration Directories and Files

This appendix lists and defines the files and directories available in ADP. Topics include:

- [Configuration Directories](#)
- [Acsera.properties File](#)
- [UrlMap.properties](#)

A.1 Configuration Directories

After ADP is installed, all the components of the application package are located in the EMGC_ADPMANAGER1 directory. This directory is in the GC domain home, for example:

```
/net/abcdef1234/scratch/jdoe/view_storage/jdoe_aug21/work/user_
projects/domains/EMGC_DOMAIN/EMGC_ADPMANAGER1
```

A.1.1 Directory Structure

The path where the ADP Manager is installed is similar to:

```
/scratch/Middleware0712/gc_inst/user_projects/domains/GCDomain/EMGC_ADPMANAGER1
```

where domain.home=/scratch/Middleware0712/gc_inst/user_projects/domains/GCDomain
and ORACLE_HOME=/scratch/Middleware0712/oms

The directory structure is as follows:

```
ADPManager.ear/
ADPManager.ear/APP-INF/
ADPManager.ear/APP-INF/lib/
ADPManager.ear/META-INF/
ADPManager.ear/ADPManager.war/
ADPManager.ear/ADPManager.war/bin/
ADPManager.ear/ADPManager.war/config/
ADPManager.ear/ADPManager.war/mcconfig/
ADPManager.ear/ADPManager.war/deploy/
ADPManager.ear/ADPManager.war/lib/
ADPManager.ear/ADPManager.war/lib/bea/
ADPManager.ear/ADPManager.war/lib/oracle/
ADPManager.ear/ADPManager.war/META-INF/
ADPManager.ear/ADPManager.war/WEB-INF/
```

Table A–1 ADP Manager Directories

Directory	Description
bin	Contains all the executable files to start and stop ADP, run deployer for Agent and ADP EJB, run export utility.
config	Contains all the ADP runtime configuration parameters that control execution logic, ADP schemas enablement, ADP GUI functionality, Service Level Objectives definition, export logic and many more.
deploy	Contains agent libraries and configuration files, as well as <i>ADP EJB</i> and <i>ADP Admin Web Application</i> . These components are deployed on the remote host (Web or Application servers) using <i>deployer</i> utility found in bin directory of ADP package.
lib	Has all the libraries required for ADP's proper functionality
mcconfig	Contains internal base instrumentation configuration. Do not modify these files.

A.1.2 Config Directory

Config directory has many files that potentially can be configured and make ADP to run in a particular way. Any changes applied to files in this directory require restarting ADP server.

Most of the files never get touched directly by user. The following are the main three files which can be configured manually to achieve desired effect:

File	Description
Acsera.properties	This file is the main ADP configuration file customization of which helps to tune up ADP.
configuration.xml	In this file you define location of Administration Server and credentials to access it. Usually you do not touch this file. The entire configuration is done through ADP GUI.
export.xml	This file contains information that drives proper data export logic. It is used for manual and automatic export of performance metric and events data from the ADP Data Repository.
UrlMap.properties	This file is used to map server addresses to load balancer addresses. By default, this file is not available; it must be created by the user.

It is worth mentioning that Service Level Objective definitions and Actions associated with the SLOs are described in *slo.xml* and *event.xml* respectively. The content of these files is completely controlled by definitions applied from ADP GUI (configuration tab).

A.1.3 Deploy Directory

The /deploy directory contains the ADP Java Agent distributable, including configuration files as well as corresponding libraries. These files are copied to the target systems hosting the Managed Servers when running the deployer utility. Rarely one needs to modify configuration files in this directory. Remember though if you modify the files they will be distributed to ALL targets within single server/cluster.

A.2 Acsera.properties File

The *acsera.properties* file contains global configuration parameters that define the operation of the ADP Manager.

A.2.1 Log Files Management

This section of Acsera.properties file defines log rotation policies. Log.MaxFiles indicates max number of log files available at any given moment, whereas the Log.MaxFileSizeMB indicates maximum size of the log file.

Example A-1 Log Files Management Section

```
Log.CopyOut = false
Log.MaxFiles = 10
Log.MaxFileSizeMB = 30
Log.MergeLogs = true

Debug.CopyOut = false
Debug.LogLevel = all
Debug.MaxFiles = 10
Debug.MaxFileSizeMB = 30
```

Log files are stored in the log directory.

A.2.2 Multi-Domain Monitoring Configuration

One can limit number of domains to be monitored by setting resource limit parameter: ConfigurationManager.ResourceLimit=4

Example A-2 Multi-Domain Monitoring Configuration

```
ConfigurationManager.ResourceLimit=4
```

A.2.3 ADP RMI Port Assignment

ADP uses RMI ports for communication with the agents and collects incoming performance metrics from a particular RMI port. By default, the RMI port is set on the same machine that hosts ADP. RMI.Registry.Host needs to be un-commented and have a value other than localhost if the host is multi-homed (such as, many network interfaces or has any ipv6 addresses) and you need to make sure that ADP listens to the incoming traffic on the particular interface.

You may need to change RMI.Registry.Port value in case the default 51099 port number has been allocated to an other application. Also if ADP is running in multi-instance mode, the port number will be different from instance to instance.

Example A-3 ADP RMI Port Assignment

```
#RMI.Registry.Host = localhost
RMI.Registry.Port = 51099
```

A.2.4 ADP Aggregation and Data Life Time Configuration

ADP has sophisticated multi-tiered logic for aggregation (or compression) of performance data. This helps to optimize performance of interaction with the internal data repository both when querying data for presentation or inserting new performance metrics.

Users who want to store longer term data should look for this section in Acsera.properties:

```
#####
# Production setting
# NOTE: use Model.GlobalSamplingRateSecs to configure Metric.Grain.0
```

```
#####  
Metric.Grain.0 0s  
Metric.TableInterval.0 = 4h  
Metric.DataLife.0 = 2d  
  
Metric.Grain.1 = 3m  
Metric.TableInterval.1 =1d  
Metric.DataLife.1 = 8d  
  
#Metric.Grain.2 = 30m  
#Metric.TableInterval.2 = 7d  
#Metric.DataLife.2 = 420d
```

and uncomment the last 3 lines for the Metric.*.2 properties

A.2.5 Aggregating Incoming Metrics On the Fly

ADP by default aggregates data coming from multiple cluster members by application thus minimizing rate of insertion in to the data repository. This greatly improves performance of ADP in heavily loaded environments.

As a side effect of this approach though, the user is unable to see metrics from instrumentation (processes and portals) on per server level. If you need to enable this then set the JavaMIP.AggregateInserts to *false*.

A.2.6 Listing Applications to Be Monitored or Excluded From Monitoring

To avoid overhead of unnecessary monitoring of certain applications, you can explicitly state which applications to monitor, or which applications to exclude from monitoring.

Users should append the name of their application to the property ComponentProvider.Application.Exclude.

Example A-4 Specifying Which Applications to Monitor

```
# Control which applications to analyze  
#  
ComponentProvider.Application.Exclude=WLI System EJBs,WLI-AI  
Design-time,B2BDefaultWebAppApplication,WLI  
Worklist,JWSQueueTransport,Deployer,BEA_WLS_DBMS_ADK,  
Acsera,ClearApp,HttpDeployer,ServiceBus_Console,em
```

A.2.7 Firewall Mitigation (for Internal RMI Ports)

If there is a firewall between the ADP Manager and the monitored application servers, ports need to be opened between them especially in the case where multiple resources are configured. For example, if two resources are configured and the first one uses 55006 as the port, then the next resource must use 55007 as the port. Each additional resource increments the port by 1.

In addition to the application server's JMX access ports, the following two properties in Acsera.properties indicate the ports used specifically by ADP:

- RMI.Registry.Port (51099 by default)
- RMI.JavaProvider.ServerPort (55003 by default)

A.2.8 SLO Dampening

There are times when you deliberately want to cut down on the number of repeated notifications should SLO violation persist for a given period of time. To suppress notifications of the same violation in a short period of time, ADP provides the SLO Dampening feature. Once enabled, should a SLO violation occur and be repeated several times in a short period, ADP will not fire the SLO violation notification for the time period defined in SLO.RearmDelay. To disable this feature, set the value of this parameter to 0.

SLO.SuppressDelayedAsserts indicates that if the violation still persists upon time period expiry ADP, should fire the SLO notification. By default it is *false*, for example, fire the notification.

Example A-5 SLO Dampening

```
# The following property is specified in units of
# minutes (m), hours (h) or days (d)
SLO.RearmDelay = 15m
SLO.SuppressDelayedAsserts = false
```

A.3 UrlMap.properties

The UrlMap.properties file should be created in the ADP Manager's config directory and used to provide address mappings between load balancers and application servers. The format of this file is:

```
# Format:
#   $app_server_ip = $load_balancer_id
# E.g:
#   http://localhost:7001 = http://localhost:7005
#
# Note: ":" character need to be escaped with "\"
#
http://192.168.128.53:7002 = http://192.168.3.187:80
http://192.168.128.53:7003 = http://192.168.3.187:80
http://192.168.128.54:7005 = http://192.168.128.54:7011
http://192.168.128.54:7006 = http://192.168.128.54:7011
```

Support Matrix for Application Dependency and Performance

For information about the platforms supported by the Application Dependency and Performance (ADP) feature in Enterprise Manager Cloud Control, search the My Oracle Support (<https://support.oracle.com>) knowledge base for “Platform Support List ADP” to locate the current certification article.

Index

A

- accessing architecture view, 28-7
- acsera.properties file, A-2
- activity
 - entity EJB, 28-39
- ADF task flows
 - Oracle WebCenter, 28-23
- ADP
 - architecture, 27-4
 - how to use, 29-1
 - Java agents, 27-5
 - manager, 27-6
 - node, 28-54
 - starting, 28-1
 - user interface, 27-6
- aggregated count
 - metric type, 28-9
- aggregating
 - incoming metrics, A-4
 - performance data, A-3
- Application Replay
 - analyzing replay results, 4-20
 - creating captures, 4-8
 - introduction, 4-1
 - monitoring capture process, 4-12
 - prerequisites, 4-4
 - replaying captures, 4-13
 - restarting database, 4-6
 - synchronized vs non-synchronized replay, 4-3
 - testing against real workloads, 4-2
 - troubleshooting, 4-23
- Application Replay, Importing Divergences, 4-22
- Application Replay, OpenScript, 4-23
- Application Schema model
 - functional view, 28-5
- Application Server
 - extensible monitoring, 3-12
 - managing configurations, 3-18
- applications, 28-42 to 28-52
 - Web, 28-46
- architecture
 - ADP, 27-4
- architecture view, 28-6
 - accessing, 28-7
 - summary, 28-8

- types of, 28-7
- async dispatchers, 28-36
- automatic domain member discovery, 2-13
- automatic target discovery, 2-1
 - changes in domain, 2-13
- average timing
 - metric type, 28-9

B

- baseline performance, 29-3, 29-6
- beans
 - entity, 28-49
 - message driven, 28-51
 - stateful, 28-47
 - stateless, 28-47
- blackouts
 - configuring for Service Level Objectives, 28-58
 - creating for Service Level Objectives, 28-59
 - deleting Service Level Objectives, 28-59
 - Service Level Objectives summary list, 28-59
- books
 - Oracle WebLogic portal, 28-16
- bottlenecks
 - identifying performance, 29-4, 29-7
 - removing performance, 29-4
- BPEL Process Manager
 - Configuring, 11-10
 - Discovery, 11-2, 11-3, 11-4
 - Software Library, 11-4
 - Supported Versions, 11-1
 - Troubleshooting, 11-11
- business SLAs guidelines, 29-5
- byte statistics, JMS destinations, 28-38

C

- cache
 - entity EJB, 28-40
 - stateful EJB, 28-47
- channels node, 28-42
- class node
 - in processes, 28-30
- Coherence
 - Discovery, 19-3
 - Refresh Cluster, 19-10

- Enable Management Pack, 19-10
- JVM Diagnostics Integration, 20-20
- Monitor
 - Administration, 20-15, 20-17
 - Cluster, 20-1, 20-5, 20-6, 20-7
 - Configuration Topology, 20-21
 - Detail, 20-8, 20-10, 20-11
 - Performance, 20-12
 - Performance Summary, 20-21
 - Push Replication Pattern, 20-19
 - Reap Session, 20-18
 - Transactional Cache, 20-19
- New Features, 19-2
- Troubleshooting, 20-21
- component level
 - architecture view, 28-7
- component level execution
 - architecture view, 28-7
- config directory, A-2
- configuration
 - data life time, A-3
 - directories and files, A-1
 - multi-domain monitoring, A-3
- configurations
 - managing, 3-18
- configuring
 - list of applications excluded from monitoring, A-4
 - list of applications to be monitored, A-4
- containers
 - persistent, 28-39
 - persistent in processes, 28-27
 - stateless, 28-38
- creating
 - Service Level Objectives, 28-57
 - Service Level Objectives blackouts, 28-59
- custom metrics, 28-54
 - monitoring environment, 28-5

D

- dashboards
 - health indicators, 28-2
- data
 - performance graphs and data items, 28-4
- data displayed
 - display interval, 28-3
 - time frame, 28-3
- data items
 - performance data, 28-4
- data life time configuration, A-3
- delay analysis
 - view
 - in processes, 28-27
 - Oracle BPEL processes, 28-18
- dependency
 - node, 28-44
 - types, 28-45
- deploy directory, A-2
- deployments nodes, 28-45

- desktops
 - displaying portal, 28-34
 - Oracle WebLogic portal, 28-14
- destinations
 - JMS, 28-37
- directory
 - config, A-2
 - configuration, A-1
 - deploy, A-2
 - structure, A-1
- dispatchers
 - async, 28-36
 - sync, 28-37
- display interval
 - context, 28-3
 - data displayed, 28-3
 - time frame, 28-3
- do-it-yourself manual processes, avoiding, 27-3

E

- EJB
 - activity, 28-39
 - activity table in processes, 28-28
 - cache table in processes, 28-28
 - entity cache, 28-40, 28-50
 - entity locking, 28-40, 28-50
 - entity transactions, 28-40, 28-50
 - entity transactions table in processes, 28-29
 - locking table in processes, 28-29
 - message driven activity, 28-51
 - message driven transactions, 28-52
 - services, 28-32
 - stateful cache, 28-47
 - stateful locking, 28-48
 - stateful transactions, 28-48
- entity
 - beans, 28-49
 - EJB activity, 28-39
 - EJB activity table in processes, 28-28
 - EJB cache, 28-40, 28-50
 - EJB cache table in processes, 28-28
 - EJB locking, 28-40, 28-50
 - EJB locking table in processes, 28-29
 - EJB transactions, 28-40, 28-50
 - EJB transactions table in processes, 28-29
- environments
 - service-oriented views, 27-2
- events
 - node performance, 28-41
 - view in processes, 28-27
- execute queues
 - health, 28-35
- execute thread
 - thread count, 28-36
- extensible monitoring, 3-12

F

- firewall mitigation, A-4

functional view of Application Schema model, 28-5
Fusion Middleware plug-in, 14-3

G

graphs
 performance data, 28-4
guidelines
 business SLAs, 29-5

H

health
 integration, 28-35
health indicators
 dashboard, 28-2
how to use ADP, 29-1
HTTP services, 28-31

I

improving performance, 29-3
instrumentation in processes, 28-29
integration, 28-34 to 28-42
interval context
 display interval, 28-3

J

Java agents
 ADP, 27-5
Java EE application
 managing, 27-2
Java EE application responsiveness, monitoring, 3-9
JDBC services, 28-32
JMS destinations, 28-37
JSF pages
 Oracle WebCenter, 28-24
JVM Diagnostics
 Features
 Cross-Tier Correlation, 15-2
 In-depth Visibility, 15-2
 JVM Pooling, 15-3
 Low Overhead, 15-2
 Memory Leak Detection, 15-2
 New Features, 15-3
 Real Time and Historical Diagnostics, 15-3
 Real Time Transaction Tracing, 15-2
 Supported Platforms and JVMs, 15-4
 User Roles, 15-4
 Heap Snapshots, 16-30
 Compare, 16-35
 Heap Object Information, 16-33
 Heap Usage by Roots, 16-32
 Top 40 Objects, 16-33
 Upload Trace Diagnostics Images, 16-39
 Installing, 16-1
 Manage JVMs
 Offline Diagnostics, 16-40, 16-41, 16-42
 Managing JVM Pools, 16-2
 JVM Pool Home Page, 16-10

 Live Thread Analysis, 16-12
 Performance Diagnostics, 16-11
Managing JVMs, 16-16
 JVM Home Page, 16-16
 Live Heap Analysis, 16-25
 Live Time Thread Analysis, 16-21
 Performance Diagnostics, 16-18
 Performance Summary, 16-20
Overview, 15-1
Setting Up, 16-2
Thread Snapshots
 Analyze Trace Diagnostic Images, 16-38
Threshold Violations, 16-42
Troubleshooting
 Cross Tier Functionality Errors, 17-1
 Deployment Script Execution Errors, 17-6
 FAQ, 17-10
 LoadHeap Errors, 17-9
 Trace Errors, 17-5
 User Interface Errors, 17-10

L

locking
 entity EJB, 28-40
 stateful EJB, 28-48
log files management
 acsera.properties file, A-3

M

managing
 configurations, 3-18
manual domain member discovery, 2-13
mapping
 SLAs to SLOs, 29-2, 29-5
maximum response time measurement
 metric type, 28-9
message driven
 beans, 28-51
 EJB activity, 28-51
 EJB transactions, 28-52
message statistics
 JMS destinations, 28-38
metadata view
 Oracle BPEL processes, 28-19
method node
 in processes, 28-30
methodology, ADP, 29-1 to 29-10
metrics
 aggregating incoming, A-4
 custom, 28-54
 customizing for monitoring, 28-5
 setting Service Level Objectives on, 29-4
 setting SLOs on, 29-8
 types, 28-8
 WebLogic, 29-8
Middleware
 management, 1-1, 3-1
Middleware Diagnostics Advisor

- diagnosing performance issues, 18-2, 18-5
- functions, 18-2
- limiting scope of, 18-3
- overview, 18-1
- prerequisites, 18-3
- purging data, 18-4
- troubleshooting issues, 18-6
- minimum response time measurement
 - metric type, 28-9
- modeled entities view
 - Oracle BPEL processes, 28-20
- module level
 - architecture view, 28-7
- module level execution
 - architecture view, 28-7
- Monitor
 - Job System, 3-22
 - Lifecycle Management, 3-18
 - Compliance Management, 3-19
 - Manage Configurations, 3-18
 - Patch Management, 3-20
 - Provisioning, 3-20
 - Manage Service Levels, 3-21
 - Support Workbench, 3-23
 - Topology Viewer, 3-23
- Monitoring
 - Administer, 3-17
 - Middleware Targets, 3-1
 - Non-Oracle Middleware Components, 3-5
 - Oracle Application Server Components, 3-4
 - Out-of-Box Monitoring
 - Blackouts, 3-11
 - Extending, 3-12
 - Historical Performance, 3-10
 - Metric Thresholds, 3-11
 - Monitoring Templates, 3-11
 - Out-of-box Metrics, 3-9
 - Request Monitoring, 3-12
 - Performance Problems, 3-15
 - Diagnostics Snapshots, 3-16
 - Home Page, 3-15
- monitoring
 - configuring custom metrics, 28-5
 - extensible, Application Server, 3-12
 - multi-domain configuration, A-3

N

- node
 - hierarchy in processes, 28-26
 - method in processes, 28-30
- nodes
 - ADP, 28-54
 - class in processes, 28-30
 - dependencies, 28-44
 - deployments, 28-45
 - hierarchy in Oracle BPEL processes, 28-20
 - pageflows, 28-31
 - services, 28-44

O

- operation routing rules view
 - Oracle ESB, 28-23
- Oracle
 - Application Server
 - Web Cache, 3-9
 - Oracle BPEL Processes, 28-17 to 28-21
 - Oracle Business Analytics, 14-1
 - Oracle Business Intelligence, 14-1
 - Oracle Business Intelligence Instance, 14-1
 - component failovers, 14-22
 - dashboard reports, 14-15
 - Discovering, 14-4
 - Monitoring, 14-5
 - monitoring credentials, 14-22
 - scheduler reports, 14-16
 - Oracle Business Intelligence Instance
 - components, 14-2
 - BI Cluster Controller, 14-2
 - BI Java Host, 14-2
 - BI Presentation Server, 14-2
 - BI Scheduler, 14-2
 - BI Server, 14-2
 - Oracle Business Intelligence targets, 14-5
 - alerts, 14-11
 - availability, 14-7
 - blackouts, 14-20
 - compliance, 14-14
 - configuration, 14-13
 - health, 14-10
 - incidents, 14-11
 - job activity, 14-14
 - logs, 14-12
 - metrics, 14-9
 - monitoring configuration, 14-21
 - performance, 14-8
 - resource usage, 14-8
 - Oracle ESB, 28-21 to 28-23
 - Oracle Essbase, 14-2
 - applications, 14-18
 - Discovering, 14-4
 - Monitoring, 14-5
 - Oracle Fusion Middleware Components, 3-2
 - Oracle Business Intelligence, 3-4, 14-1, 14-2
 - Oracle Coherence, 3-4
 - Oracle Forms Services, 3-3
 - Oracle Identity Management, 3-3
 - Oracle Portal, 3-3
 - Oracle SOA Suite, 3-2
 - Oracle Web Tier, 3-2
 - Oracle WebCenter, 3-2
 - Oracle WebLogic Server Domains, Clusters, and Managed Servers, 3-2
 - Oracle GlassFish
 - before getting started, 9-1
 - cluster
 - home page, 9-12
 - how to access, 9-13
 - creating configuration comparison template, 9-14
 - domain, 9-2

- adding, 9-5
 - displaying results, 9-8
 - finding and assigning targets, 9-5
 - how to access, 9-4
 - refreshing, 9-9
- overview, 9-1
- roles and privileges, 9-1
- server
 - home page, 9-10
 - how to access, 9-11
- start and stop procedures, 9-2
- viewing configuration data, 9-13
- Oracle HTTP Server session volumes, 3-9
- Oracle resources, 28-53
- Oracle Service Bus
 - Discovery, 12-2, 12-4
 - Enabling Management Packs, 12-6
 - Supported Versions, 12-1
 - Troubleshooting, 12-8
- Oracle Traffic Director
 - adding to Exalogic target, 7-2
 - configuration, 7-2
 - instance, 7-6
 - overview, 7-1
 - refresh flow, 7-6
- Oracle WebCenter, 28-23 to 28-25
- Oracle WebLogic
 - portals, 28-13 to 28-17
 - resources, 28-53

P

- pageflow viewer
 - Oracle WebLogic portal, 28-16
- pageflows node, 28-31
- pages
 - JSF in Oracle WebCenter, 28-24
 - Oracle WebLogic portal, 28-16
- parameters
 - defining Service Level Objectives, 28-57
- partner link bindings view
 - Oracle BPEL processes, 28-20
- partner link type role view
 - Oracle BPEL processes, 28-19
- partner links view
 - Oracle BPEL processes, 28-19
- performance
 - baseline, 29-3, 29-6
 - characteristics of target, 29-3
 - graphs and data items, 28-4
 - identifying bottlenecks, 29-4, 29-7
 - improvement process, 29-3
 - node, 28-41
 - removing bottlenecks, 29-4
 - system level, 29-7
- persistent containers, 28-39
 - in processes, 28-27
- platform
 - intelligent, 27-3
- port assignment

- RMI, A-3
- portal desktop
 - displaying, 28-34
- portal hierarchy
 - Oracle WebLogic portal, 28-13
- portals
 - Oracle WebLogic, 28-13
- portlet drill down
 - Oracle WebLogic portal, 28-15
- portlet performance
 - Web Services Remote Portlet producers, 28-33
- portlets
 - Oracle WebCenter, 28-25
 - Oracle WebLogic portal, 28-17
 - Web Services Remote Portlet producer, 28-34
- process nodes
 - performance, 28-41
- processes, 28-25 to 28-30
 - deployed WebLogic, 28-25
 - node hierarchy, 28-26
- properties file
 - acsera.properties, A-2
 - UrlMap.properties, A-5

R

- refresh WebLogic Domain targets, 2-13
- resources
 - Oracle, 28-53
 - Oracle WebLogic, 28-53
- RMI port assignment, A-3
- RMI ports
 - firewall mitigation, A-4

S

- SCA *See* Service Component Architecture
- Service Component Architecture, 28-55
- service definition view
 - Oracle ESB, 28-22
- service details view
 - Oracle ESB, 28-22
- Service Level Objectives
 - blackouts summary list, 28-59
 - configuring blackouts, 28-58
 - creating, 28-57
 - creating blackouts, 28-59
 - defining parameters for, 28-57
 - deleting blackouts, 28-59
 - setting on key metrics, 29-4
- service mode
 - ADP topology, 27-5
- service operations view
 - Oracle ESB, 28-22
- service parent details view
 - Oracle ESB, 28-22
- service-oriented views, 27-2
- services, 28-31 to 28-32
- services node
 - applications, 28-44

- setting
 - Service Level Objectives on key metrics, 29-4
 - SLOs on key metrics, 29-8
- SLO dampening, A-5
- snapshot count
 - metric type, 28-8
- SOA application
 - monitoring, 27-2
- SOA Management Pack Enterprise Edition, 10-1
 - BPEL Process Manager, 10-1
 - Central Management Console, 10-1
 - Service Bus, 10-1
 - SOA Composite, 10-1
 - SOA Infrastructure, 10-1
- SOA Suite
 - Dehydration Store, 13-10
 - Discovery, 13-3
 - Generating IP Reports, 13-14
 - New Features, 13-1
 - Service Topology, 13-11
 - SOA Artifacts and Composites, 13-14
 - SOA Instance Tracing, 13-7
 - Supported Versions, 13-2
 - Troubleshooting, 13-31
 - UDDI Publishing, 13-12
- starting ADP, 28-1
- stateful
 - beans, 28-47
 - EJB cache, 28-47
 - EJB locking, 28-48
 - EJB transactions, 28-48
- stateless
 - beans, 28-47
 - containers, 28-38
- steps in using ADP, 29-1
- subscribers node, 28-42
- sync dispatchers, 28-37
- system level performance, 29-7

T

- tab descriptions
 - Oracle WebLogic portal level, 28-14
- target
 - performance characteristics, 29-3
- target discovery
 - automatic, 2-1
 - using EMCLI, 2-12
 - WebLogic 10.x, 2-9
 - WebLogic 9.x, 2-9
- task flows
 - ADF in Oracle WebCenter, 28-23
 - user-defined in Oracle WebCenter, 28-24
- time frame
 - data displayed, 28-3
 - display interval, 28-3
 - specifying, 28-4
- topology
 - service mode, 27-5
 - Web Services Remote Portlet, 28-34

- topology view, 28-6
 - Oracle BPEL processes, 28-20
- transactions
 - entity EJB, 28-40
 - stateful EJB, 28-48
- tree summary
 - Oracle WebLogic portal, 28-13
- types
 - dependency, 28-45
 - metric, 28-8

U

- UrlMap.properties file, A-5
- user
 - interface in ADP, 27-6
- user-defined task flows
 - Oracle WebCenter, 28-24

V

- view
 - architecture, 28-6
 - delay analysis in Oracle BPEL processes, 28-18
 - delay analysis in processes, 28-27
 - events in processes, 28-27
 - functional of Application Schema model, 28-5
 - metadata in Oracle BPEL processes, 28-19
 - operation routing rules in Oracle ESB, 28-23
 - partner link bindings in Oracle BPEL
 - processes, 28-20
 - partner link type role in Oracle BPEL
 - processes, 28-19
 - partner links in Oracle BPEL processes, 28-19
 - service definition view in Oracle ESB, 28-22
 - service details in Oracle ESB, 28-22
 - service operations in Oracle ESB, 28-22
 - service parent details in Oracle ESB, 28-22
 - topology, 28-6
 - topology of Oracle BPEL processes, 28-20
- view modeled entities in Oracle BPEL
 - processes, 28-20
- viewers
 - pageflow for Oracle WebLogic portal, 28-16

W

- Web 2.0 service
 - Oracle WebCenter, 28-24
- Web applications, 28-46
- Web Services, 28-30
- WebLogic
 - key metrics, 29-8
- WebLogic domain, 14-1
 - Oracle Business Intelligence Instance, 14-1
 - Oracle Essbase, 14-2
- WebLogic Domain Refresh job, 2-13
- what's new in this manual, xxiii
- workshop projects, 28-46
- WSRP producers, 28-32 to 28-34