

Oracle® Enterprise Manager Cloud Control Middleware Management Guide



13c Release 3
E92977-02
November 2018

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

ORACLE®

Copyright © 2015, 2018, Oracle and/or its affiliates. All rights reserved.

Contributing Authors: Pavithra Mendon

Contributors: Enterprise Manager Cloud Control 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, then the following notice is applicable:

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

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

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

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

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

Contents

Preface

Audience	xxiv
Documentation Accessibility	xxiv
Related Documents	xxiv
Conventions	xxiv

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 Managing Middleware Targets

2.1	Middleware Targets in Enterprise Manager	2-1
2.1.1	Oracle Fusion Middleware Components	2-1
2.1.2	Oracle Application Server Components	2-4
2.1.3	Non-Oracle Middleware Components	2-4
2.2	Monitoring Middleware Targets	2-5
2.2.1	Middleware Summary Page	2-5
2.2.1.1	Heat Map	2-6
2.2.1.2	Searching Middleware Managed Targets	2-7
2.2.2	Target Home Page	2-8
2.2.3	Predefined Performance Metrics	2-10
2.2.4	Analyzing Historical Performance	2-11
2.2.5	Setting Metric Thresholds for Alert Notifications	2-11
2.2.6	Monitoring Templates	2-12
2.2.7	Managing and Creating Blackouts and Notification Blackouts	2-12
2.2.8	Extend Monitoring for Applications Deployed to WebLogic Server	2-13
2.2.9	Using Multi-Tenancy	2-14

2.3	Diagnosing Performance Problems	2-15
2.3.1	Using Home Pages to Diagnose Performance Issues	2-15
2.3.2	Diagnostic Snapshots	2-15
2.3.3	Log File Viewer	2-16
2.6	Administering Middleware Targets	2-16
2.6.1	Shutting Down, Starting Up, or Restarting a Middleware Target	2-17
2.7	Auditing WebLogic-specific Operations	2-21
2.8	About Lifecycle Management	2-22
2.8.1	Managing Configurations	2-23
2.8.2	Compliance Management	2-24
2.8.3	Patch Management	2-24
2.8.4	Provisioning	2-25
2.8.4.1	Deploying / Undeploying Java EE Applications	2-25
2.9	Managing Service Levels	2-25
2.9.1	Service Dashboard	2-26
2.10	Job System	2-27
2.11	Routing Topology Viewer	2-27
2.4	Analyzing Middleware Problems Using Problem Analysis	2-27
2.5	Managing Problems with Support Workbench	2-30
2.5.1	Accessing and Logging In To Support Workbench	2-31
2.5.1.1	Accessing Support Workbench	2-31
2.5.1.2	Logging In	2-31
2.5.2	Using Fusion Middleware Support Workbench	2-32
2.5.2.1	Viewing Diagnostics	2-32
2.5.2.2	Viewing an Aggregated Diagnostic Summary	2-33
2.5.2.3	Searching for Problems	2-34
2.5.2.4	Annotating a Problem	2-34
2.5.2.5	Adding More Files	2-35
2.5.2.6	Creating a Package	2-35
2.5.2.7	Providing Additional Files	2-36
2.5.2.8	Uploading a Package to Oracle Support	2-36
2.5.2.9	Creating a Service Request	2-37
2.5.2.10	Managing Problem Resolution	2-37

3 Testing Application Load and Performance

3.1	Introduction to Application Replay	3-1
3.2	Testing Against Real-World Application Workloads	3-1
3.3	Capturing Application Workload Using RUEI	3-2
3.4	Prerequisites and Considerations When Using Application Replay	3-3
3.4.1	Using RUEI to Capture Application Workloads	3-4

3.4.2	Configuring Required User Privileges in Enterprise Manager	3-5
3.4.3	Setting up the Test System Database for Application Replay	3-5
3.4.4	Setting up the Capture Directory for Application Replay	3-5
3.5	Understanding the Application Capture and Replay Process	3-6
3.6	Creating Application Workload Captures	3-6
3.7	Monitoring the Application Capture Process	3-12
3.8	Replaying Application Workload Captures	3-13
3.8.1	Preparing to Replay Workload Captures	3-14
3.8.2	Understanding Application Replays and Replay Tasks	3-14
3.8.3	Resolving References to External Systems for Application Replays	3-14
3.8.4	Remapping URLs for Application Replays	3-15
3.8.5	Substituting Sensitive Data for Application Replays	3-15
3.8.6	Replaying Workload Captures	3-15
3.8.7	Analyzing Application Replay Results	3-21
3.9	Importing Replay Session Divergences into OpenScript	3-22
3.10	Troubleshooting Application Replay	3-23

4 Composite Applications

4.1	Viewing the Composite Application Dashboard	4-1
4.2	Creating a Composite Application	4-3
4.3	Editing a Composite Application	4-4
4.4	Editing a Composite Application Home Page	4-5
4.5	Using Composite Applications	4-5

Part II Monitoring Exalytics Target and Traffic Director

5 Monitoring an Exalytics Target

6 Oracle Traffic Director

6.1	Before Discovering Traffic Director 11g	6-2
6.2	Adding a Traffic Director to an Exalogic Target	6-2
6.3	About Traffic Director Configuration	6-2
6.3.1	Using the Traffic Director Configuration Page	6-3
6.3.2	Adding Traffic Director Target Configuration	6-3
6.3.2.1	Finding Configurations and Instances	6-4
6.3.2.2	Discovered Targets	6-5
6.3.2.3	Viewing Results	6-5
6.4	About Traffic Director Instance	6-6

6.5	About Traffic Director Refresh Flow	6-6
6.5.1	Adding New Targets to Newly Added Configurations	6-7
6.5.2	Adding New Targets for Newly Added Instances of Configurations	6-7
6.5.3	Deleting Targets of Configurations That Have Been Removed	6-7
6.5.4	Deleting Targets of Instances That Have Been Removed	6-8

Part III Monitoring Oracle WebLogic Domains and Oracle GlassFish Domains

7 Monitoring WebLogic Domains

7.1	Updating the Agent Truststore	7-1
7.1.1	Importing a Demo WebLogic Server Root CA Certificate	7-2
7.1.2	Importing a Custom Root CA Certificate	7-2
7.1.3	Prerequisites for Domain Discovery When in TLS Mode	7-2
7.2	Changing the Default AgentTrust.jks Password Using Keytool	7-2
7.3	Collecting JVM Performance Metrics for WebLogic Servers	7-3
7.3.1	Setting the PlatformMBeanServerUsed Attribute	7-3
7.3.2	Activating Platform MBeans on WebLogic Server 9.x to 10.3.2 versions	7-3

8 Overview of Oracle GlassFish Server Management

8.1	Before Getting Started	8-1
8.1.1	GlassFish Roles and Privileges	8-1
8.1.2	Adding Domain Certificate to Activate Start and Stop Operations	8-1
8.2	Understanding the Oracle GlassFish Domain	8-2
8.2.1	How to Add an Oracle GlassFish Domain To Be Monitored	8-4
8.2.2	Adding an Oracle GlassFish Domain: Finding and Assigning Targets	8-5
8.2.3	Adding an Oracle GlassFish Domain: Displaying Results	8-9
8.2.4	How to Access an Oracle GlassFish Domain	8-9
8.2.5	Refreshing an Oracle GlassFish Domain	8-10
8.3	Understanding the Oracle GlassFish Server Home Page	8-11
8.3.1	How to Access an Oracle GlassFish Server	8-12
8.4	Understanding the Oracle GlassFish Cluster Home Page	8-13
8.4.1	How to Access an Oracle GlassFish Cluster	8-14
8.5	Viewing Collected Configuration Data for Oracle GlassFish Members	8-15
8.6	Creating an Oracle GlassFish Server Configuration Comparison Template	8-15

Part IV Managing Oracle SOA

9	Overview of Oracle SOA Management	
9.1	About Oracle SOA Management Pack Enterprise Edition	9-1
10	Discovering and Monitoring Service Bus	
10.1	New Features in This Release	10-1
10.2	Supported Versions	10-1
10.3	Understanding the Discovery Mechanism	10-1
10.4	Understanding the Discovery Process	10-2
10.5	Discovering Service Bus	10-3
10.5.1	Discovering Service Bus Deployed to WLS Not Monitored by Enterprise Manager	10-3
10.5.2	Discovering Service Bus Deployed to WLS Monitored by Enterprise Manager	10-4
10.6	Enabling Management Packs	10-5
10.7	Monitoring Service Bus in Cloud Control	10-5
10.7.1	Enabling Monitoring for Service Bus Services	10-5
10.8	Generating Service Bus Reports Using BI Publisher	10-6
10.9	Troubleshooting Service Bus	10-7
10.9.1	System and Service	10-7
10.9.2	SOAP Test	10-8
11	Discovering and Monitoring the SOA Suite	
11.1	List of Supported Versions	11-1
11.2	Monitoring Templates	11-2
11.3	Discovering the SOA Suite	11-2
11.3.1	Discovering the SOA Suite Using a Local Agent	11-2
11.3.2	Discovering the SOA Suite Using a Remote Agent	11-5
11.3.3	Discovering the SOACS Instance Using the Hybrid Cloud Agent	11-5
11.4	Configuring the SOA Suite with Target Verification	11-6
11.4.1	Running Functionality-Level Diagnostic Checks	11-7
11.4.2	Running System-Level Diagnostic Checks	11-7
11.4.3	Repairing Target Monitoring Setup Issues	11-7
11.5	Metric and Collection Settings	11-8
11.6	Integration Workload Statistics (IWS)	11-10
11.6.1	Statistics in an IWS Report	11-11
11.6.2	Enabling, and Configuring, or Disabling IWS	11-11
11.6.3	Generating an IWS Report	11-12
11.7	Setting Up and Using SOA Instance Tracing	11-13
11.7.1	Configuring Instance Tracing (SOA 11g Targets Only)	11-13

11.7.2	Setting Search Criteria for Tracing an Instance	11-13
11.7.2.1	Instance Tracing for SOA 11g Targets	11-13
11.7.2.2	Instance Tracing for SOA 12c Targets	11-14
11.7.3	Tracing an Instance Within a SOA Infrastructure	11-17
11.7.4	Tracing Instance Across SOA Infrastructures	11-17
11.8	Viewing Composite Heat Map	11-18
11.9	Monitoring Dehydration Store	11-18
11.9.1	Enabling Monitoring of the SOA Dehydration Store	11-19
11.9.2	Viewing the SOA Dehydration Store Data	11-19
11.10	Publishing a Service to UDDI	11-20
11.11	Generating SOA Reports	11-20
11.11.1	Generating SOA Reports Using BI Publisher	11-20
11.11.2	Generating SOA Reports Using Information Publisher	11-22
11.11.3	Generating SOA Diagnostic Reports	11-22
11.11.4	Viewing SOA Diagnostics Jobs	11-23
11.12	Exporting a Composite .jar File	11-24
11.13	Provisioning SOA Artifacts and Composites	11-24
11.14	Diagnosing Issues and Incidents	11-25
11.15	Searching Faults in the SOA Infrastructure	11-25
11.15.1	Overview of Faults and Fault Types in SOA Infrastructure	11-25
11.15.2	Overview of the Recovery Actions for Resolving Faults	11-26
11.15.3	Prerequisites for Searching, Viewing, and Recovering Faults	11-27
11.15.4	Searching and Viewing Faults	11-27
11.15.4.1	Setting Search Criteria	11-28
11.15.4.2	Finding Total Faults in the SOA Infrastructure	11-29
11.15.4.3	Limiting Faults Searched and Retrieved from the SOA Infrastructure	11-30
11.15.4.4	Searching Only Recoverable Faults	11-30
11.15.4.5	Searching Faults in a Particular Service Engine	11-30
11.15.4.6	Searching Faults by Error Message	11-31
11.15.4.7	Filtering Displayed Search Results	11-31
11.15.5	Recovering a Few Faults Quickly (Simple Recovery)	11-31
11.16	Recovering Faults in Bulk	11-32
11.16.1	Performing Bulk Recovery from the Bulk Recovery Jobs Page	11-33
11.16.1.1	Setting Fault Details for Recovering Faults in Bulk	11-34
11.16.1.2	Setting Recovery and Batch Details for Recovering Faults in Bulk	11-34
11.16.1.3	Scheduling Bulk Recovery Jobs to Run Once or Repeatedly	11-35
11.16.2	Performing Bulk Recovery from Faults and Rejected Messages Tab	11-36
11.16.3	Performing Bulk Recovery from the Error Hospital Tab	11-37
11.16.4	Tracking Bulk Recovery Jobs	11-38

11.16.4.1	Tracking Bulk Recovery Jobs, and Viewing Their Results and Errors	11-38
11.16.4.2	Creating Bulk Recovery Jobs Using EMCLI and Web Services	11-40
11.16.5	WorkFlow Examples for Bulk Recovery	11-42
11.16.5.1	Running Bulk Recovery Job Every Night	11-43
11.16.5.2	One Time Job with Specific Time Interval to Recover Faults	11-43
11.17	Generating Error Hospital Reports	11-44
11.17.1	Generating an Error Hospital Report	11-47
11.17.2	Customizing the Error Hospital Report	11-48
11.18	Recovering BPMN Messages	11-48
11.19	Troubleshooting	11-49
11.19.1	Discovery	11-49
11.19.2	Monitoring	11-50
11.19.3	Instance Tracing Errors	11-50
11.19.4	Recent Faults	11-50
11.19.5	Fault Management	11-51
11.19.5.1	Bulk Recovery	11-51
11.19.5.2	Fault Search and Recovery	11-52
11.19.5.3	Fault Management and Instance Tracing Errors	11-53
11.19.6	Information Publisher Reports	11-53
11.19.7	BI Publisher Reports	11-54
11.19.8	Systems and Services	11-55
11.19.9	BPEL Recovery	11-56
11.19.10	SOA License Issue	11-56
11.19.11	Dehydration Store Issue	11-56

Part V Managing Oracle Business Intelligence

12 Discovering and Monitoring Oracle Business Intelligence Instance and Oracle Essbase

12.1	Overview of Oracle Business Intelligence Targets You Can Monitor	12-1
12.1.1	Oracle Business Intelligence Instance	12-2
12.1.2	Oracle Essbase	12-2
12.2	Understanding the Monitoring Process	12-3
12.3	Discovering Oracle Business Intelligence Instance and Oracle Essbase Targets	12-4
12.3.1	Discovering Targets of an Undiscovered WebLogic Domain	12-4
12.3.2	Discovering New or Modified Targets of a Discovered WebLogic Domain	12-5
12.4	Monitoring Oracle Business Intelligence Instance and Essbase Targets	12-5

12.4.1	Performing General Monitoring Tasks	12-6
12.4.1.1	Viewing Target General and Availability Summary	12-6
12.4.1.2	Viewing Target Status and Availability History	12-7
12.4.1.3	Viewing Target Performance or Resource Usage	12-8
12.4.1.4	Viewing Target Metrics	12-10
12.4.1.5	Viewing or Editing Target Metric and Collection Settings	12-10
12.4.1.6	Viewing Target Metric Collection Errors	12-10
12.4.1.7	Viewing Target Health	12-11
12.4.1.8	Viewing Target Alert History	12-11
12.4.1.9	Viewing Target Incidents	12-12
12.4.1.10	Viewing Target Logs	12-12
12.4.1.11	Viewing Target Configuration and Configuration File	12-14
12.4.1.12	Viewing Target Job Activity	12-15
12.4.1.13	Viewing Target Compliance	12-15
12.4.2	Performing Target-Specific Monitoring Tasks	12-15
12.4.2.1	Viewing Oracle Business Intelligence Dashboard Reports	12-16
12.4.2.2	Viewing Oracle Business Intelligence Scheduler Reports	12-18
12.4.2.3	Viewing Oracle Business Intelligence Instance Key Metrics	12-18
12.4.2.4	Viewing Oracle Essbase Applications Summary	12-19
12.4.2.5	Viewing Oracle Essbase Application Data Storage Details	12-20
12.5	Administering Oracle Business Intelligence Instance and Essbase Targets	12-21
12.5.1	Performing General Administration Tasks	12-21
12.5.1.1	Starting, Stopping, or Restarting the Target	12-21
12.5.1.2	Administering Target Access Privileges	12-22
12.5.1.3	Administering Target Blackouts	12-22
12.5.1.4	Viewing Target Monitoring Configuration	12-22
12.5.2	Performing Target-Specific Administration Tasks	12-23
12.5.2.1	Viewing Oracle Business Intelligence Component Failovers	12-23
12.5.2.2	Editing Oracle Business Intelligence Monitoring Credentials	12-24
12.6	Scaling Out Oracle Business Intelligence Domains	12-24
12.7	Creating Oracle Business Intelligence Instance Provisioning Profiles	12-27
12.8	Cloning Oracle Business Intelligence Instances	12-28

Part VI Monitoring Application Performance

13 Monitoring Performance

13.1	Monitoring Views and Dimensions	13-1
13.2	Using ECIDs to Track Requests	13-4
13.2.1	ECIDs for Components Other Than Oracle Fusion Middleware Components	13-5

13.3	Setting up End-to-end Monitoring	13-6
13.3.1	Set up Enterprise Manager	13-8
13.3.2	Set up Java Virtual Machine Diagnostics	13-8
13.3.3	Set up Real User Experience Insight	13-8
13.3.4	Create the Business Application	13-9
13.4	User Roles and Privileges	13-9

14 Understanding the User Experience

14.1	What Does RUEI Discover?	14-1
14.2	Viewing and Analyzing RUEI Data	14-3
14.2.1	Dashboards	14-3
14.2.2	Reports	14-4
14.2.3	Session Diagnostics	14-5
14.2.4	User Flows	14-5
14.2.5	KPIs and Service Level Agreements	14-7
14.3	What Questions Can RUEI Answer?	14-8
14.4	What Aspects of RUEI Can You Access from the EM Console?	14-9
14.5	How Does RUEI Work with JVM Diagnostics?	14-10

15 Getting Detailed Execution Information

15.1	Using JVM Diagnostics	15-1
15.2	Using Request Instance Diagnostics	15-3

16 Monitoring Business Applications

16.1	Introduction to Business Applications	16-1
16.1.1	Systems, Services, and Business Applications	16-2
16.1.2	MyBank: An Example Business Application	16-2
16.2	Prerequisites and Considerations	16-3
16.2.1	Requirements for Using RUEI	16-4
16.2.1.1	Registering RUEI Installations with Self-Signed Certificates	16-4
16.3	Registering RUEI Systems	16-5
16.3.1	Setting Up a Connection Between RUEI and the Oracle Enterprise Manager Repository	16-9
16.4	Creating Business Applications	16-10
16.5	Monitoring Business Applications	16-14
16.6	Monitoring End User Experience	16-16
16.6.1	Monitoring End User Experience Data	16-17
16.6.1.1	Key Performance Indicators	16-17
16.6.1.2	Usage Data	16-17

16.6.1.3	Violations Data	16-18
16.6.2	Working With Session Diagnostics	16-19
16.6.2.1	Creating an Enterprise Manager User for Session Diagnostics	16-20
16.6.2.2	Getting Started with Session Diagnostics	16-20
16.6.2.3	Customizing Session Diagnostics Reporting	16-23
16.6.2.4	Exporting Full Session Information	16-24
16.6.2.5	Exporting Session Pages to Microsoft Excel	16-25
16.6.3	Monitoring End User Experience Metrics	16-26
16.6.4	Monitoring User Flows	16-28
16.6.5	Monitoring Logs	16-28
16.7	Monitoring an End User Service	16-29
16.7.1	Troubleshooting an End User Service	16-30
16.8	Monitoring KPI and SLA Alert Reporting	16-30
16.9	Upgrading End User Service	16-32
16.10	Upgrading Business Applications	16-32

17 Monitoring End-to-end Performance

17.1	Troubleshooting: A Case Study	17-1
17.2	Finding Solutions	17-6

18 Troubleshooting Middleware Applications Using Enterprise Manager

18.1	Introduction to Troubleshooting Middleware Applications	18-1
18.2	Preparing the Environment to Troubleshoot Applications	18-3
18.2.1	Document the Topology of the Systems in the Environment	18-3
18.2.2	Install Management Agents on All Systems in the Environment	18-4
18.2.3	Install RUEI to Help Troubleshoot Web Applications	18-4
18.3	Configure the Environment to Help Troubleshoot Applications	18-5
18.3.1	Discover All Targets in the Environment	18-5
18.3.2	Deploy JVM Agents in the Environment	18-5
18.3.3	Define Composite Applications to Help Troubleshoot Multiple Tier Applications	18-6
18.3.4	Define Synthetic Monitoring Beacons in Enterprise Manager	18-6
18.3.5	Define Thresholds in Enterprise Manager	18-6
18.3.6	Set up Compliance Management in Enterprise Manager	18-7
18.3.7	Create a RUEI Application to Help Troubleshoot Web Applications	18-7
18.3.8	Define RUEI Service Level Agreements	18-8
18.3.9	Create a Business Application in Enterprise Manager	18-8
18.4	Analyzing Issues Using Enterprise Manager and RUEI	18-8
18.4.1	Analyzing Incidents using Log Files	18-10

18.4.2	Analyzing Incidents using Business Applications	18-11
18.4.3	Analyzing Incidents	18-11
18.4.3.1	Check EM Dashboards to Analyze Incidents	18-11
18.4.3.2	Use RUEI to Check Pages Affected by an Incident	18-11
18.4.3.3	Use JVMD to Isolate Issue	18-12
18.4.3.4	Use Thresholds and Compliance to Analyze Incidents	18-12
18.5	Resolving Issues Using Enterprise Manager	18-12
18.5.1	Resolve an Issue Using Configuration Tools	18-12
18.5.2	Work with Application Developers or DBAs to Resolve Application Issues	18-13
18.5.3	Resolve a Capacity Issue Using Provisioning Tools	18-13

Part VII Using JVM Diagnostics and MDA Advisor

19 Introduction to JVM Diagnostics

19.1	Overview	19-1
19.1.1	Java Activity Monitoring and Diagnostics with Low Overhead	19-1
19.1.2	In-depth Visibility of JVM Activity	19-2
19.1.3	Real Time Transaction Tracing	19-2
19.1.4	Cross-Tier Correlation with Oracle Databases	19-2
19.1.5	Memory Leak Detection and Analysis	19-2
19.1.6	JVM Pooling	19-3
19.1.7	Real-time and Historical Diagnostics	19-3
19.2	New Features in this Release	19-3
19.3	Supported Platforms and JVMs	19-3
19.4	User Roles	19-4

20 Using JVM Diagnostics

20.1	Setting Up JVM Diagnostics	20-1
20.1.1	Configuring the JVM Diagnostics Engine	20-2
20.1.2	Configuring JVMs and JVM Pools	20-4
20.1.3	Registering Databases	20-5
20.1.4	Configuring the Heap Analysis Hosts	20-6
20.1.5	Viewing Registered JVMs and Managers	20-7
20.2	Accessing the JVM Diagnostics Pages	20-7
20.3	Managing JVM Pools	20-8
20.3.1	Viewing the Java Virtual Machine Pool Home Page	20-8
20.3.1.1	Promoting JVM Diagnostics Events to Incidents	20-9
20.3.2	Viewing the JVM Pool Live Thread Analysis Page	20-9

20.3.3	Configuring a JVM Pool	20-12
20.3.3.1	Updating Pool Thresholds	20-12
20.3.4	Removing a JVM Pool	20-13
20.3.5	Adding a JVM Pool to a Group	20-13
20.4	Managing JVMs	20-13
20.4.1	Viewing the JVM Home Page	20-14
20.4.2	Viewing the JVM Diagnostics Performance Summary	20-14
20.4.3	Viewing the JVM Live Thread Analysis Page	20-15
20.4.3.1	Performing Cross Tier Analysis	20-18
20.4.3.2	Establishing Cross-Tier Correlation in Oracle RAC Databases	20-20
20.4.4	Viewing Memory Diagnostics	20-23
20.4.5	Working with Class Histograms	20-26
20.4.5.1	Saving a Class Histogram	20-26
20.4.5.2	Viewing Saved Histograms	20-26
20.4.5.3	Scheduling a Histogram Job	20-26
20.4.5.4	Comparing Class Histograms	20-27
20.4.5.5	Deleting Class Histograms	20-27
20.4.6	Taking a Heap Snapshot	20-27
20.4.7	Taking a Heap Snapshot and Loading Into the Repository	20-28
20.4.8	Analyzing Heap Snapshots	20-29
20.4.8.1	Viewing Heap Usage by Roots	20-30
20.4.8.2	Viewing Heap Usage by Objects	20-33
20.4.8.3	Memory Leak Report	20-34
20.4.8.4	Anti-Pattern Report	20-34
20.4.9	Managing JFR Snapshots	20-34
20.4.10	Configuring a JVM	20-35
20.4.11	Removing a JVM	20-35
20.4.12	Adding a JVM to a Group	20-35
20.5	Managing Thread Snapshots	20-35
20.5.1	Tracing Active Threads	20-36
20.6	Analyzing Trace Diagnostic Images	20-37
20.7	Viewing Heap Snapshots and Class Histograms	20-37
20.8	JVM Offline Diagnostics	20-38
20.8.1	Creating a Diagnostic Snapshot	20-38
20.8.2	Using the Diagnostic Snapshots Page	20-39
20.8.3	Analyzing a Diagnostic Snapshot	20-39
20.8.4	Viewing a Diagnostic Snapshot	20-40
20.9	Viewing JVM Diagnostics Threshold Violations	20-40
20.10	Using Java Workload Explorer	20-40
20.10.1	Accessing Java Workload Explorer	20-40
20.10.2	Performance Analysis and Search Criteria	20-41

20.10.3	Graph Highlights	20-43
20.10.4	Diagnostics	20-43
20.11	Managing and Troubleshooting JVMD (Globally)	20-50
20.12	Managing and Troubleshooting JVMD (Specific Agent)	20-51
20.13	Enable or Disable Monitoring of JVM Targets using EMCLI	20-52

21 Troubleshooting JVM Diagnostics

21.1	Cross Tier Functionality Errors	21-1
21.2	Trace Errors	21-4
21.3	Deployment Execution Errors	21-4
21.4	LoadHeap Errors	21-7
21.5	Heap Dump Errors on AIX 64 and AIX 32 bit for IBM JDK 1.6	21-8
21.6	Errors on JVM Diagnostics UI Pages	21-8
21.7	Frequently Asked Questions	21-9
21.7.1	Location of the JVM Diagnostics Logs	21-9
21.7.2	JVM Diagnostics Engine Status	21-10
21.7.3	JVM Diagnostics Agent Status	21-10
21.7.4	Monitoring Status	21-10
21.7.5	JVMD SLB Configuration	21-10
21.7.6	Running the create_jvm_diagnostic_db_user.sh Script	21-12
21.7.7	Usage of the Try Changing Threads Parameter	21-12
21.7.8	Significance of Optimization Levels	21-12
21.7.9	Custom Provisioning Agent Deployment	21-13
21.7.10	Log Manager Level	21-13
21.7.11	Repository Space Requirements	21-13

22 Using Middleware Diagnostics Advisor

22.1	Diagnosing Performance Issues with Oracle WebLogic Server	22-1
22.2	Diagnosing Performance Issues Using Middleware Diagnostics Advisor	22-2
22.3	Functioning of Middleware Diagnostics Advisor	22-2
22.4	Prerequisites for Configuring Middleware Diagnostics Advisor	22-3
22.5	Configuring Middleware Diagnostics Advisor	22-3
22.6	Enabling Middleware Diagnostics Advisor for a Target	22-4
22.7	Setting Up Middleware Diagnostics Advisor (MDA)	22-5
22.8	Limiting the Scope of Middleware Diagnostics Advisor	22-6
22.9	Using Middleware Diagnostics Advisor to View and Diagnose Performance Issues	22-7
22.10	Running an Unscheduled Middleware Diagnostics Advisor Analysis on a Target	22-8

Part VIII Managing Oracle Coherence

23 Getting Started with Management Pack for Oracle Coherence

23.1	About Coherence Management	23-1
23.2	New Features for Oracle Coherence	23-2
23.3	Configuring a Coherence Cluster	23-3
23.3.1	Creating and Starting a JMX Management Node	23-4
23.3.1.1	Specifying Additional System Properties	23-4
23.3.1.2	Including the Additional Class Path	23-5
23.3.1.3	Using the Custom Start Class	23-5
23.3.1.4	Example Start Script for the Coherence Management Node	23-5
23.3.2	Configuring All Other Nodes	23-6
23.3.2.1	Additional System Properties for All Other Coherence Nodes	23-6
23.3.2.2	Example Start Script for All Other Coherence Nodes	23-6
23.3.3	Testing the Configuration	23-7
23.3.3.1	Verifying Remote Access for the MBean Objects Using JConsole	23-7
23.4	Discovering Coherence Targets	23-8
23.4.1	Discovering a Standalone Coherence Cluster	23-8
23.4.1.1	Refreshing a Cluster	23-12
23.4.1.2	Managing Mis-configured Nodes	23-13
23.4.2	Discovering a Managed Coherence Cluster	23-14
23.5	Enabling the Management Pack	23-16

24 Monitoring a Coherence Cluster

24.1	Understanding the Page Layout	24-1
24.1.1	Navigation Tree	24-2
24.1.2	Personalization	24-3
24.2	Viewing the Home Pages	24-4
24.2.1	Coherence Cluster Home Page	24-4
24.2.1.1	General Tab	24-6
24.2.1.2	Heatmap	24-8
24.2.1.3	Cluster Menu Navigation	24-9
24.2.2	Node Home Page	24-10
24.2.2.1	Node Menu Navigation	24-13
24.2.3	Cache Home Page	24-13
24.2.3.1	Near Cache	24-16

24.2.3.2	Cache Menu Navigation	24-17
24.2.4	Partition Cache Home Page	24-18
24.2.4.1	Cache Menu Navigation	24-19
24.2.5	Application Home Page	24-19
24.2.6	Service Home Page	24-20
24.2.7	Connection Manager Home Page	24-23
24.3	Viewing the Summary Pages	24-24
24.3.1	Nodes Page	24-24
24.3.2	Caches Page	24-26
24.3.3	Services Page	24-27
24.3.4	Applications Page	24-29
24.3.5	Proxies Page	24-29
24.4	Log Viewer	24-30
24.4.1	Configuring the Log Location Settings	24-31
24.4.2	Viewing the Log Messages	24-31
24.5	Viewing the Performance Pages	24-32
24.5.1	Performance Summary Page	24-32
24.5.1.1	Customizing the Performance Page Charts	24-32
24.5.2	Connection Manager Performance Page	24-32
24.6	Removing Down Members	24-33
24.7	Topology Viewer	24-33
24.8	Viewing Incidents	24-34

25 Administering a Coherence Cluster

25.1	Cluster Administration Page	25-1
25.1.1	Changing the Node Configuration	25-1
25.1.2	Changing the Cache Configuration	25-3
25.1.3	Changing the Service Configuration	25-3
25.2	Node Administration Page	25-3
25.3	Cache Administration Page	25-4
25.4	Service Administration Page	25-4
25.5	Cache Data Management	25-4
25.5.1	Explain Plan	25-7
25.5.2	Trace	25-7

26 Troubleshooting and Best Practices

26.1	Troubleshooting Coherence	26-1
26.2	Best Practices	26-1

26.2.1	Monitoring Templates	26-1
--------	----------------------	------

27 Coherence Integration with JVM Diagnostics

27.1	Overview	27-1
27.2	Configuring Coherence Nodes for JVM Diagnostics Integration	27-1
27.2.1	Example Start Script for Coherence Management Node	27-2
27.2.2	Example Start Script for All Other Nodes	27-2
27.3	Accessing JVM Diagnostics from Coherence Targets	27-3
27.3.1	Accessing JVM Diagnostics from Oracle Coherence Node Menu	27-3
27.3.2	Accessing JVM Diagnostics from Oracle Coherence Cache Menu	27-3
27.3.3	Accessing JVM Diagnostics from Oracle Coherence Cluster Menu	27-3
27.4	Including the JVM Diagnostics Regions in the Coherence Target Home Pages	27-3

Part IX Using Identity Management

28 Getting Started with Oracle Identity Management

28.1	Benefits of Using the Identity Management Pack	28-1
28.2	Features of the Identity Management Pack	28-1
28.2.1	New Features for this Release	28-2
28.3	Monitoring Oracle Identity Management Components in Enterprise Manager	28-3

29 Prerequisites for Discovering Oracle Identity Management Targets

29.1	System Requirements	29-1
29.2	Installing Oracle Enterprise Manager Cloud Control 13c	29-3
29.3	Prerequisites for Discovering Identity Management Targets in Enterprise Manager	29-3

30 Discovering and Configuring Oracle Identity Management Targets

30.1	Discovering Identity Management Targets	30-1
30.1.1	Discovering Identity Management 11g and 12c	30-1
30.1.2	Discovering Oracle Directory Server Enterprise Edition 11g and 12c	30-2
30.2	Collecting User Statistics for Oracle Internet Directory	30-3
30.3	Creating Identity Management Elements	30-4
30.3.1	Creating Identity and Access System Target	30-4
30.3.2	Creating Generic Service or Web Application Targets for Identity Management	30-5

Part X Discovering and Monitoring Non-Oracle Middleware

31 Discovering and Monitoring IBM WebSphere MQ

31.1	Introduction	31-1
31.1.1	Out-of-Box Availability and Performance Monitoring	31-1
31.1.2	Centralized Monitoring of all Information in a Single Console	31-2
31.1.3	Enhance Service Modeling and Perform Comprehensive Root Cause Analysis	31-2
31.2	Prerequisites	31-3
31.2.1	Basic Prerequisites	31-3
31.2.2	JAR File Requirements (for Local Monitoring and Remote Monitoring)	31-3
31.3	Understanding Discovery	31-4
31.3.1	Discovery Prerequisites for Local Agent	31-4
31.3.2	Discovery Prerequisites for Remote Agent	31-4
31.3.3	Queue Manager Cluster Discovery	31-5
31.3.4	Standalone Queue Manager Discovery	31-6
31.4	Monitoring	31-7

32 Discovering and Monitoring IBM WebSphere Application Servers, Clusters, and Cells

32.1	About Managing IBM WebSphere Application Servers, Clusters, and Cells	32-1
32.2	Supported Versions for Discovery and Monitoring	32-3
32.3	Prerequisites for Discovering IBM WebSphere Application Servers, Clusters, and Cells	32-3
32.4	Discovering IBM WebSphere Application Servers, Clusters, and Cells	32-7
32.5	Monitoring IBM WebSphere Application Servers	32-9
32.5.1	Monitoring IBM WebSphere Application Servers	32-9
32.5.1.1	General Section	32-10
32.5.1.2	Monitoring and Diagnostics Section	32-10
32.5.1.3	Response and Load Section	32-11
32.5.1.4	Applications Tab	32-11
32.5.1.5	Servlets and JSPs Tab	32-11
32.5.1.6	EJBs Tab	32-11
32.5.2	Administering IBM WebSphere Application Servers	32-11
32.5.3	Monitoring the Performance of IBM WebSphere Application Servers	32-12
32.5.4	Monitoring the Applications Deployed to IBM WebSphere Application Servers	32-13

32.5.5	Viewing the Top EJBs of IBM WebSphere Application Servers	32-13
32.5.6	Viewing the Top Servlets and JSPs of IBM WebSphere Application Servers	32-13
32.5.7	Viewing IBM WebSphere Application Server Metrics	32-14
32.6	Monitoring IBM WebSphere Application Server Clusters	32-14
32.6.1	Monitoring IBM WebSphere Application Server Clusters	32-14
32.6.1.1	Summary Section	32-14
32.6.1.2	Monitoring and Diagnostics Section	32-15
32.6.1.3	Servers Section	32-15
32.6.1.4	Resource Usage Section	32-15
32.6.2	Administering IBM WebSphere Application Server Clusters	32-16
32.6.3	Viewing IBM WebSphere Application Server Cluster Members	32-16
32.6.4	Viewing IBM WebSphere Application Server Cluster Metrics	32-17
32.7	Monitoring IBM WebSphere Application Server Cells	32-17
32.7.1	Monitoring IBM WebSphere Application Server Cells	32-17
32.7.1.1	General Section	32-18
32.7.1.2	Incidents Summary Section	32-18
32.7.1.3	Clusters Section	32-19
32.7.1.4	Servers Section	32-19
32.7.2	Administering IBM WebSphere Application Server Cells	32-19
32.7.3	Viewing IBM WebSphere Application Server Cell Members	32-20
32.8	Troubleshooting IBM WebSphere Application Server Discovery and Monitoring Issues	32-21
32.8.1	Troubleshooting Discovery Issues	32-21
32.8.2	Troubleshooting Monitoring Issues	32-25

33 Discovering and Monitoring JBoss Application Server

33.1	About Managing JBoss Application Servers, JBoss Domains, and JBoss Partitions	33-1
33.2	Finding Out the Supported Versions for Discovery and Monitoring	33-3
33.3	Prerequisites for Discovering JBoss Application Servers, Domains, and Partitions	33-3
33.4	Discovering JBoss Application Servers 7.x and JBoss Domains	33-4
33.5	Discovering JBoss Application Servers 6.x and JBoss Partitions	33-6
33.6	Monitoring JBoss Application Servers	33-8
33.6.1	Monitoring JBoss Application Servers 7.x	33-9
33.6.1.1	General	33-9
33.6.1.2	JVM Threads	33-10
33.6.1.3	Transaction Metrics	33-10
33.6.1.4	Response and Load Section	33-10
33.6.1.5	Deployments Section	33-11

33.6.2	Monitoring JBoss Application Servers 6.x	33-11
33.6.2.1	General	33-11
33.6.2.2	Servlets/JSPs	33-12
33.6.2.3	JVM Threads	33-12
33.6.2.4	Datasource	33-12
33.6.2.5	Response and Load Section	33-13
33.6.2.6	Most Requested Servlets/JSPs Section	33-13
33.6.3	Administering JBoss Application Servers 7.x and 6.x	33-13
33.6.4	Monitoring Applications Deployed to JBoss Application Servers 7.x and 6.x	33-14
33.6.5	Monitoring the Performance of JBoss Application Servers 7.x and 6.x	33-14
33.6.6	Monitoring Servlets and JSPs Running on JBoss Application Servers 6.x	33-15
33.6.7	Viewing JBoss Application Server Metrics	33-15
33.6.8	Analyzing Problems Using Metric Correlation	33-16
33.7	Monitoring JBoss Domains	33-17
33.7.1	Monitoring JBoss Domains	33-17
33.7.1.1	Summary Section	33-18
33.7.1.2	Servers Section	33-18
33.7.1.3	Incidents Section	33-18
33.7.2	Monitoring JBoss Server Groups	33-18
33.7.2.1	Summary Section	33-19
33.7.2.2	Servers Section	33-19
33.7.2.3	Incidents Section	33-20
33.7.3	Administering JBoss Domains	33-20
33.7.4	Viewing JBoss Domain Members	33-20
33.7.5	Refreshing JBoss Domains	33-21
33.8	Monitoring JBoss Partitions	33-21
33.8.1	Monitoring JBoss Partitions	33-21
33.8.1.1	Summary Section	33-22
33.8.1.2	Servers Section	33-23
33.8.1.3	Incidents Section	33-23
33.8.2	Administering JBoss Partitions	33-23
33.8.3	Viewing JBoss Partition Members	33-24
33.8.4	Refreshing JBoss Partitions	33-25
33.9	Deploying JVMD on JBoss Application Server 7.x and 6.x to Diagnose Issues	33-25
33.10	Troubleshooting JBoss Application Server Discovery and Monitoring Issues	33-29
33.10.1	Troubleshooting Monitoring Issues	33-29
33.10.2	Troubleshooting Discovery Issues	33-29
33.10.3	Additional Useful Resources	33-30

34 Discovering and Monitoring Apache HTTP Server

34.1	Introduction to HTTP Servers	34-1
34.2	Supported Versions of Apache HTTP Server for Discovery and Monitoring	34-1
34.3	Prerequisites for Discovering and Monitoring Apache HTTP Server	34-2
34.4	Discovering Apache HTTP Servers	34-2
34.5	Monitoring Apache HTTP Servers	34-3
34.6	Configuration Management for Apache HTTP Servers	34-4
34.7	Troubleshooting Apache HTTP Server Issues	34-4

Part XI Managing Oracle Data Integrator

35 Configuring and Monitoring Oracle Data Integrator

35.1	Prerequisites for Monitoring Oracle Data Integrator	35-1
35.2	Monitoring Oracle Data Integrator	35-2
35.2.1	Monitoring Oracle Data Integrator	35-2
35.2.1.1	Master Repositories Health	35-3
35.2.1.2	ODI Agents Health	35-3
35.2.1.3	Work Repositories Health	35-4
35.2.1.4	Data Servers Health	35-4
35.2.1.5	Sessions/Load Plan Executions	35-5
35.2.2	Monitoring ODI Agents	35-5
35.2.2.1	Search Agents	35-5
35.2.2.2	ODI Agents	35-5
35.2.3	Monitoring Repositories	35-6
35.2.3.1	Search Repositories	35-7
35.2.3.2	Repositories	35-7
35.2.3.3	Cluster Databases	35-8
35.2.3.4	Database Details	35-9
35.2.3.5	Tablespace/File Group Details	35-9
35.2.4	Monitoring Load Plan Executions and Sessions	35-9
35.2.4.1	Search Sessions/LPEs	35-10
35.2.4.2	Load Plan Executions/Sessions	35-11
35.2.4.3	Load Plan Executions/Session Detail	35-12
35.3	Administering Oracle Data Integrator	35-13
35.3.1	Starting Up, Shutting Down, and Restarting Oracle Data Integrator Agents	35-13
35.3.2	Managing Agent Status and Activities	35-14
35.3.3	Searching Sessions and Load Plan Executions	35-14
35.3.4	Viewing Log Messages	35-14

35.4	Creating Alerts and Notifications	35-14
35.5	Monitoring Run-Time Agents	35-15
35.5.1	Agent Home Page	35-16
35.5.1.1	General Info	35-16
35.5.1.2	Load	35-16
35.5.1.3	Target Incidents	35-17
35.5.1.4	LPEs/Sessions Execution Incidents	35-17
35.5.1.5	Load Balancing Agents	35-17
35.6	Configuring Oracle Data Integrator Console	35-18
35.7	Configuring an Oracle Data Integrator Domain	35-19

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 that have purchased support 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 13c Release 3 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.

Part I

Managing Oracle Fusion Middleware

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

The chapters are:

- [Introduction to Middleware Management](#)
- [Managing Middleware Targets](#)
- [Testing Application Load and Performance](#)
- [Composite Applications](#)

1

Introduction to Middleware Management

This section introduces the use of Oracle Enterprise Manager Cloud Control to monitor and manage middleware software, including Oracle Fusion Middleware and Oracle WebLogic Server.

This section 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 environments 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 Domains, Partitions, clusters, and single server instances
- Oracle GlassFish Domains, Clusters, and Servers
- Partitioned, Clustered, and standalone Java EE applications
- Oracle HTTP Server (Collocated and Standalone)
- Oracle Traffic Director
- Service-Oriented Architecture (SOA) components
- Oracle Identity Management
- Metadata Services repositories
- Oracle WebCenter
- Oracle Portal

- Oracle Business Intelligence
- Oracle Forms Services
- Oracle Reports
- Directory Server Enterprise Edition
- Oracle Coherence
- Oracle Exalogic Elastic Cloud
- Java EE

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 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 Farms and WebLogic Domains.
- Manage third party products such as IBM WebSphere Application Server, JBoss Application Server, Apache HTTP 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.
- 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.
- Perform critical configuration and administration operations such as the following:
 - Start, stop, or restart Fusion Middleware components and processes
 - Configure domain, clusters, managed servers, resources, and multitenancy
 - Schedule and track execution of WLST scripts
- View Business Applications to access RUEI as well as information about the application's supporting infrastructure.

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 Farm and a single WebLogic Domain. Unlike Cloud Control, this is current information only. There is no storage of historical data when using Fusion Middleware Control.
- Monitor the availability and performance of Fusion Middleware software in real time mode.

- Perform routine administration tasks such as deploying applications, configuring parameters, and so on.

Note: In Fusion Middleware Control, you cannot analyze historical metric data, and the real-time analysis is limited to a single domain

For more details, see the *Oracle Fusion Middleware Administrator's Guide 11g Release 2* and *Oracle Fusion Middleware Administering Oracle Fusion Middleware 12c*.

2

Managing Middleware Targets

This section describes how to use Enterprise Manager to monitor Middleware software.

This section covers the following:

- [Middleware Targets in Enterprise Manager](#)
- [Monitoring Middleware Targets](#)
- [Diagnosing Performance Problems](#)
- [Analyzing Middleware Problems Using Problem Analysis](#)
- [Administering Middleware Targets](#)
- [Managing Problems with Support Workbench](#)
- [About Lifecycle Management](#)
- [Managing Service Levels](#)
- [Job System](#)
- [Routing Topology Viewer](#)

For more information, see *Discovering and Adding Middleware Targets in the Enterprise Manager Cloud Control Administrator's Guide*.

2.1 Middleware Targets in Enterprise Manager

After you have added a Middleware target (for example, Oracle Fusion Middleware, Oracle WebLogic Domain, 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 monitor the following middleware software using Oracle Enterprise Manager Cloud Control:

- Oracle Fusion Middleware software
- Non-Oracle Middleware software

2.1.1 Oracle Fusion Middleware Components

You can monitor the following Oracle Fusion Middleware components using Enterprise Manager:

- **Oracle WebLogic Domains, Partition, Clusters, Managed Servers, and Node Managers:** A WebLogic 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 and WebLogic Node Managers.

A domain partition (partition) is an administrative and runtime slice of a WebLogic domain. You can create one or more partitions in the domain. Each partition will contain its own apps and resources.

Clusters are groups of WebLogic Servers instances that work together to provide scalability and high-availability for applications.

A Node Manager is a WebLogic Server utility used to start, shut down, and restart Administration Server and Managed Server instances from a remote location. In addition, the Node Manager target enables you to determine whether a Node Manager is up or down. Although Node Manager is optional, it is recommended if your WebLogic Server environment hosts applications with high availability requirements. Ensure that the Node Manager has been discovered as part of the discovery of the Oracle WebLogic Domain.

With Oracle Enterprise Manager, you can monitor and manage the farm, domains, clusters, servers, node managers, 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. Enterprise Manager supports WebCenter Portal and WebCenter Content.
- **Oracle Web Tier:** This consists of:
 - **Oracle Traffic Director:** Oracle Traffic Director is a fast, reliable, and scalable layer-7 software load balancer. You can set up Oracle Traffic Director to serve as the reliable entry point for all HTTP, HTTPS and TCP traffic to application servers and web servers in the back end. Oracle Traffic Director distributes the requests that it receives from clients to servers in the back end based on the specified load-balancing algorithm, routes the requests based on specified rules, caches frequently accessed data, prioritizes traffic, and controls the quality of service. The architecture of Oracle Traffic Director enables it to handle large volumes of application traffic with low latency. The product is optimized for use in Oracle Exalogic Elastic Cloud and Oracle SuperCluster. It can communicate with servers in the back end over Exalogic's InfiniBand fabric.

- **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.x 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:
 - * Discover and monitor Oracle HTTP 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.
 - * Start, stop, and restart Oracle HTTP Servers.

Note: Cloud Control console supports both managed, as well as standalone HTTP Servers.
- **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. For the list of the IDM components monitored by Enterprise Manager, see [System Requirements](#).
- **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 standalone and managed Coherence clusters and their various entities. See [New Features for Oracle Coherence](#) and [Discovering a Managed Coherence Cluster](#) for details.
- 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 WebCenter Content** 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 WebCenter Content 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.

2.1.2 Oracle Application Server Components

Discovering and monitoring Oracle Application Server targets outside of Oracle E-Business Suite is no longer supported as of Enterprise Manager release 13.x. Enterprise Manager release 13.1 supports Oracle Application Server targets only in the context of Oracle E-Business Suite. When discovering Oracle E-Business Suite releases 12.1.x and 12.0.x, Oracle Application Server targets such as OC4J 10.1.3 and Oracle HTTP Server 10.1.3 are automatically discovered. You cannot discover and monitor Oracle Application Server targets in any other context.

2.1.3 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

For additional third-party middleware software that can be monitored, check the Enterprise Manager certification matrix on My Oracle Support (<http://support.oracle.com>).

2.2 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.

2.2.1 Middleware Summary Page

Enterprise Manager provides centralized monitoring across domains, configuration management, provisioning, real time and historical performance analysis. Beginning with the Fusion Middleware Plug-in release 12.1.0.4 and continuing with release 13c, administration features are exposed within the Cloud Control console. These features enable you to perform configuration changes directly from the Cloud Control console rather than drilling down to administration consoles such as the WebLogic Server Administration Console or the Oracle Enterprise Manager Fusion Middleware Control console. Some examples of the administration features exposed from Cloud Control include: management of JDBC data sources (for example create, edit, delete, test, control data sources), configure multitenancy, domain, clusters, servers, access to the System MBean Browser to view, edit, and invoke MBeans. However, not all administration and configuration operations can be made from Cloud Control; in some cases, you still need to drill down to the administration consoles.

The Middleware summary page, accessed from the Targets menu, provides two different views of the middleware components configured as managed targets.

These two views are referred to as the Table view and the Heat Map view. While the more traditional Table view provides a detailed summary of status across middleware-related targets, the Heat Map view provides a graphical and more efficient way to analyze the same data. On the Heat Map view, targets are represented as boxes and the size and color of each box depicts potential problem areas. This view enables administrators to quickly analyze a large amount of data, customize the filtering, and pinpoint problems more efficiently.

You can use the Table tab to add or remove middleware targets, as well as set certain monitoring configuration properties for targets.

By default, the Name, Type, Status, and Member Status Summary are listed for middleware targets. You can also add any of the global target properties such as Department and Line of Business as columns in this table. From the **View** menu, select **Columns**, then select **Manage Columns**.

Columns of particular interest are:

- **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.

- **Member Status Summary:** The availability of the middleware components associated with the target. The counts only reflect the components that meet all search criteria.
- **Version:** The target version.
- **Compliance Score:** An overall evaluation of the target's compliance with compliance standard rules defined in your enterprise, presented as a percentage of compliance. A compliance score of 100% indicates full compliance with a policy. For additional information about compliance management, see *Managing Compliance in the Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

2.2.1.1 Heat Map

You can use the Heat Map tab to view the Middleware Targets Heat Map, a graphical representation of a set of targets depicted as boxes in the heat map which are the root targets that are shown in the table tab. They can be grouped and optionally summarized by attributes such as 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 Status or the WebLogic Servers Only: CPU Usage. You can hover or click on graph elements to see more detail.

If you choose WebLogic Servers Only: CPU Usage, the graph displays boxes that are root targets containing WebLogic servers. If a root target does not contain any WebLogic servers, it is not displayed in the view. The box size is based on the number of WebLogic servers it contains. The box color is based on the average CPU value of all servers it contains. The Properties area in the lower right corner shows the number of WebLogic servers it contains as well as the average CPU value. You can also use tooltips to display this information.

The color of the boxes is meaningful. If you choose Status, red means that several members of the target are down. If you choose WebLogic Servers Only: CPU Usage, then the color represents CPU Usage for the WebLogic Servers. Red would indicate high CPU usage values while green would indicate low.

The slider enables you to set which CPU usage values are red and which are green.

Status and CPU Usage

You can use the Show drop-down menu to change to either of two displays: Status or WebLogic Servers Only: CPU Usage.

The default view is by Status and organized by target version. While this is the default view, you can modify the default and organize the data in a variety of ways using the Options region. For instance, you can organize the data by location of the target or lifecycle status of the target. You can also provide multiple levels of organization; for example, you may want to first organize by location and then by version to gain an understanding of the health of different versions of middleware targets in different geographic areas.

The WebLogic Servers Only: CPU Usage option supports only WebLogic Servers. Each box represents a WebLogic Server or the parent of a WebLogic Server (a cluster, for example). A WebLogic Server will be excluded from the graph if it is down or if its CPU metric data has not yet been collected.

Organizing Data Using Options Region

Each box in the Heat Map view represents a target or set of targets; for example, a farm or domain target. The size of the box represents the number of member targets; therefore, the larger the box, the more members the target contains.

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 heat map that was grouped using the Organize By menus. To focus on one section of the heat map, drill in by double-clicking on the section header. This displays only the boxes that are in that box and hides all others. To drill out from the view, use the locator links available above the heat map.

Using the Summarize option turns the deepest Organize First By box into one box by summarizing all of the individual boxes it contains.

To gain more information on the potentially problematic targets, you can hover over the target's box and click it. The Properties region, which appears on the right, provides additional details on the target and its members and enables you to drill-down further.

Properties Region

When you click on a box, properties relevant to the selected target are displayed in the Properties region. This may include a breakdown of the member statuses or the number of WebLogic Servers it contains, depending on the current heat map view.

The Properties region displays target properties such as Type and Target Version. It also displays any user-defined properties such as Contact, Location, or Department and so on, if they have been defined.

Incident information about this target and its descendants is also shown. Click on the counts to navigate to the Incidents Manager page where you can search, view, manage exceptions and issues, and track outstanding incidents and problems.

Importance of Color

The color of the boxes is also meaningful. For example, for Status, red indicates that the target is down and green indicates that the target is up. Using the Options region, you can customize the color range, that is, the meaning of red versus the meaning of green. By default, if 60% or less of the members in the target are up, then the box on the Heat Map view will be red; whereas, if at least 95% of the members are up, then the box on the Heat Map will be green. In the case of the WebLogic Servers Only: CPU Usage view, the color represents a range of CPU Usage for the WebLogic Server targets – where the more red the box, the higher the CPU usage.

You can adjust the slider to change the color range.

2.2.1.2 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 **Search** list, located on the left, is used to specify target types, as well as target properties, for example Cost Center. Target types only appear in the list if you have access to at least one target in that area.

Use the **View** menu located at the top left to select the properties you want displayed in chart format. For example, select Lifecycle status to see the distribution of lifecycle statuses across your targets.

The search results display as a hierarchy where all displayed targets match all search fields. The leaf nodes are shown in context with their parents. To show the results as a flat list without this hierarchical information, uncheck the "Show Hierarchy" box in the table toolbar.

To clear the filter, click the x next to the property name. Note that when multiple options for a property are selected in the Search list, that information is displayed at the top of the charts, for example Multiple Target Types.

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

Additional highlights of the Search feature include:

- When options in the Search tree are collapsed, all the hidden search options still apply.
- If you change a search option, the page content is automatically refreshed. Your search criteria is automatically saved 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' target type for targets with contact Smith, only targets 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.
- The table is populated only if the search query results are less than the maximum target.

For example, if the site has 5000 Middleware targets and the threshold is set to display 2000 targets, the table will be empty with a statement explaining that there are too many targets and that you should filter the results. If after filtering there are now 1500 targets that match the criteria, all the targets will appear in the table, since the total number is under the 2000 limit. If the threshold had been set to 6000, you would have seen all the targets on the page.

 **Note:**

If the threshold limit is very large, the page will run slower.

By default the threshold is 2000.

To change the threshold, update the `oracle.sysman.emas.MWTableTargetLimit` property using the following Oracle Management Service `emctl` command:

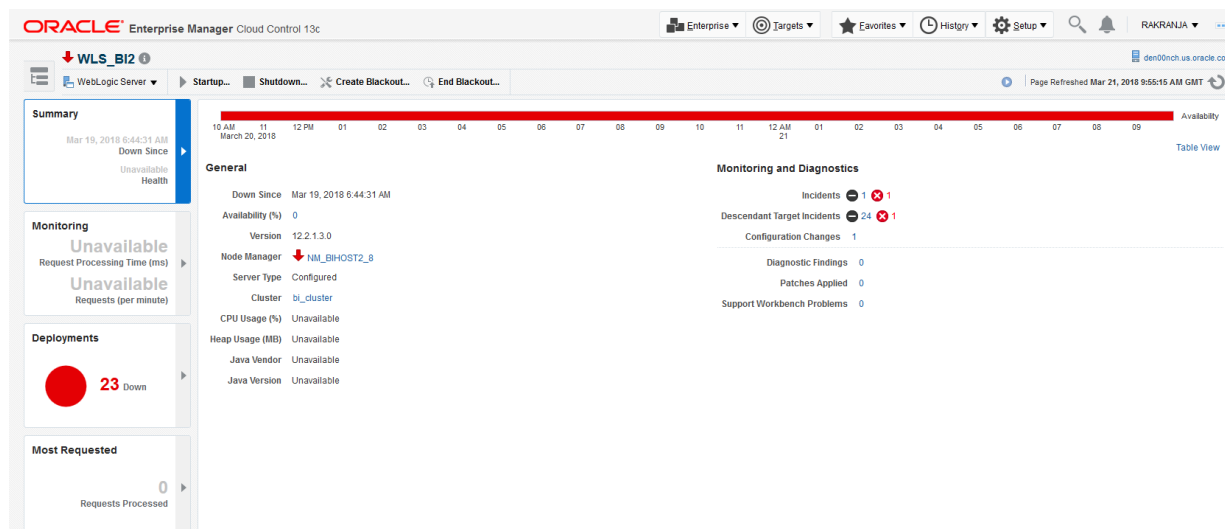
```
emctl set property -name oracle.sysman.emas.MWTableTargetLimit -value 2000
```

2.2.2 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 log in to 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 2-1 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. By default, target navigation is closed. To open the navigation pane, click the Navigation Drawer icon located at the top left.
- **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.
- **Target Name** is the name of the currently selected target.
- **Context Pane** provides the host name.
- **View** lets you 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.)

2.2.3 Predefined Performance Metrics

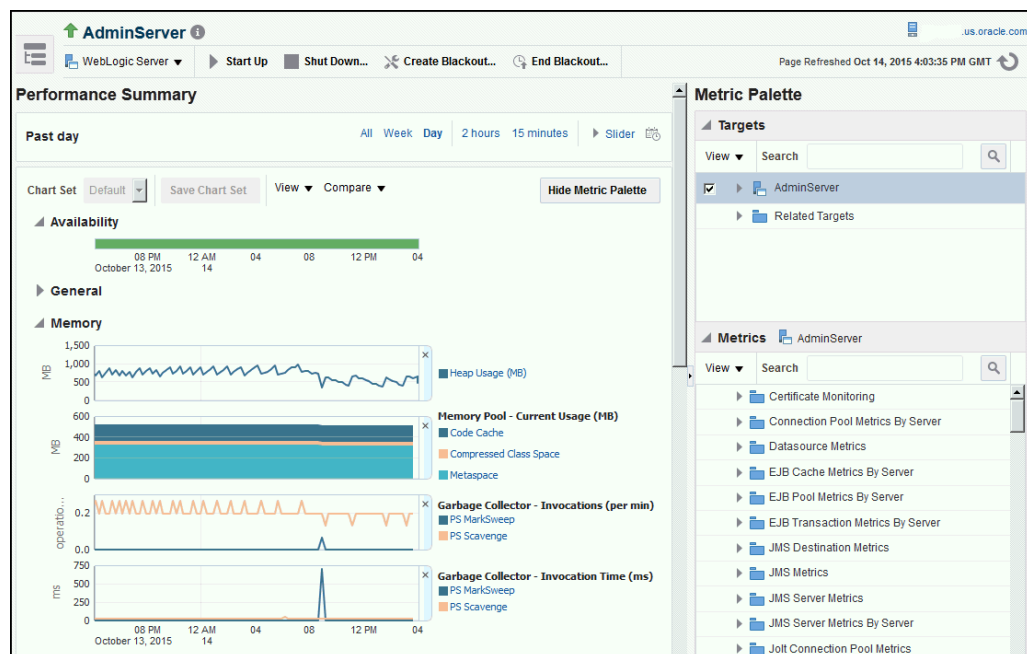
Enterprise Manager provides a set of pre-defined performance metrics for each Middleware target. The metric data is collected and stored in the Management Repository. For more details on the pre-defined metrics, see the *Oracle Fusion Middleware Metric Reference Guide*. For more information, see the Management Repository Data Retention Policies in the *Oracle Enterprise Manager Cloud Control Administrator's 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
- 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 predefined metrics can be viewed from the performance summary pages as shown below:

Figure 2-2 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.

2.2.4 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 Cloud Control, 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)

For more information, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

2.2.5 Setting Metric Thresholds for Alert Notifications

When editing metric settings, use the Threshold Suggestion feature to calculate thresholds based on deviations from 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.

To access the Threshold Suggestion feature from a target's home page:

1. Select **Monitoring** from the target's menu located at the top-left of the page, then select **Metric and Collection Settings**.
2. On the Metric and Collection Settings page, locate the metric in which you are interested and click the pencil icon associated with the metric.
3. On the Edit Advanced Settings page, locate the Threshold Suggestion region and change the thresholds as needed.

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.

2.2.6 Monitoring Templates

You can also use monitoring templates to simplify the task of standardizing monitoring settings across your enterprise. You can specify the settings for performance metrics as well as configuration collections, and apply them across multiple targets of a specific target type.

A Monitoring template defines all the parameters you would normally set to monitor a 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.

2.2.7 Managing and Creating Blackouts and Notification 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.

Blackouts

However, there are occasions when you want to perform maintenance work on your Middleware targets, but do 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.

Notification Blackouts

Notification Blackouts are used for suppressing the notifications on targets during the notification blackout duration. The Oracle Management Agent continues to monitor the target under notification blackout and the Oracle Management Service shows the actual target status along with an indication that the target is currently under

notification blackout. Events are generated as usual during a notification blackout and only their notifications are suppressed.

There are two types of notification blackouts:

- Notification blackout for maintenance (default): The target is under a planned maintenance and you do not want to receive any notifications during this period. Since the target is brought down deliberately for maintenance purposes, the notification blackout duration will not be considered while calculating the availability percentage and service level agreement. In this scenario, you should create a maintenance notification blackout.
- Notification-only notification blackout: The target is having an unexpected down time, for example, a server crash. While you are fixing the server, you do not want to receive alerts as you already know about the issue. The availability percentage computation considers the actual target status during the notification blackout and the service level agreement is computed accordingly. In this scenario, you should create a Notification-only notification blackout.

2.2.8 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 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.

For information on creating management plug-ins, see the *Oracle Enterprise Manager Cloud Control Extensibility Programmer's Guide*. For information on creating metric extensions, see the *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

2.2.9 Using Multi-Tenancy

Multi-tenancy, as it relates to Oracle WebLogic Server (WLS), refers to domain partitions that provide dedicated servers and domains to multiple applications.

A domain partition is an administrative portion of a domain that can be managed independently and can share the runtime capacity in a domain, that is, the managed servers and clusters.

By using domain partitions, you use fewer servers and domains. This enables you to simplify the management of Software as a Service (SaaS) and Platform as a Service (PaaS) applications.

PaaS

Using multi-tenancy with PaaS, encourages increased density by enabling domain sharing, in other words, you can consolidate at the domain level. This makes it easy to:

- Deploy applications from many groups into the same WebLogic infrastructure.
- Share WebLogic infrastructure and underlying resources, for example, domain, clusters, managed servers, hardware, and network.
- Isolate management tasks.
 - WebLogic Administrators manage the infrastructure.
 - Partition administrators manage deployments and related resources.
- Isolate runtime specifics.
 - Security realm per "tenant".
 - Virtual Target (addresses), Database (pluggable database), JNDI (internal traffic), Other runtime resources, for example, JMS.
 - Work managers/resource consumption management.

SaaS

Multi-tenancy encourages increased density by enabling multiple SaaS application instances in a consolidated domain. This makes it easy to:

- Deploy additional instances of an application
- Share WebLogic infrastructure and underlying resources, for example domain, clusters, managed servers, hardware, network.
- Tailor application instance to a tenant, for example, virtual target, pluggable database, runtime resources (JMS).
- Isolate runtime.
 - Security realm, virtual target, database, work managers and resource consumption management.
 - Known and trusted applications.

Enterprise Manager discovers new targets related to WebLogic Server Multi-tenancy (WLS MT) and tracks the performance metrics for the new target types that are related to WLS MT. This includes domain partition, partition application deployment. Enterprise Manager also provides the ability to create, edit and delete resource

groups, resource group templates, virtual targets and partitions in order to provide a sharable infrastructure for use by multiple organizations, and to export/import partitions across domains.

2.3 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.
- Use Java Diagnostics to diagnose performance problems in production. To take advantage of this feature, ensure that JVMD has been deployed.

2.3.1 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.

2.3.2 Diagnostic Snapshots

A diagnostic snapshot consists of 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 that 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.

Video Demonstration

To view a visual demonstration on how you can capture diagnostics snapshots, access the following URL and click **Begin Video**:

https://apex.oracle.com/pls/apex/f?p=44785:24:0::NO:24:P24_CONTENT_ID,P24_PREV_PAGE:5465,1

2.3.3 Log File Viewer

You can centrally search logs generated by WebLogic and Oracle Fusion Middleware across all Oracle Fusion Middleware components and across multiple domains. 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.

2.6 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 and the Fusion Middleware Plug-in 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 administration servers, managed servers, clusters, domains or other Fusion Middleware components (for example, managed and standalone Oracle HTTP Server, Oracle Data Integrator Agents, and so on) immediately or scheduling the operation to occur at a future point in time. For more information, see [Shutting Down, Starting Up, or Restarting a Middleware Target](#) .
- Viewing and editing settings for the Oracle WebLogic Domain, Oracle WebLogic Cluster, Oracle WebLogic Server, Server Template (applicable to only WebLogic release 12 and later), and Machine configurations. Changes made to these configurations are managed by the Change Center feature of the Cloud Control console.
- Creating, editing and deleting resource groups, resource group templates, virtual targets and partitions related to Oracle WebLogic Server Multi-tenancy (applicable only to Oracle WebLogic Server 12.2.1.x and later releases).

2.6.1 Shutting Down, Starting Up, or Restarting a Middleware Target

You can shut down, start up, or restart administration servers, managed servers, clusters, domains, node manager targets, WebLogic domain partitions and partition application deployments, or other Fusion Middleware components (for example, managed and standalone Oracle HTTP Server, Oracle Data Integrator Agents, and so on). To do so, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, select either an administration server, a managed server, a cluster, a domain, or any other Fusion Middleware component (for example, managed or standalone Oracle HTTP Server, Oracle Data Integrator Agents, and so on).
3. On the Home page, from the context menu, select **Control**, then select either **Start Up**, **Shut Down**, or **Restart** depending on your requirement.

 **Note:**

For Oracle WebLogic Domain, only start and stop operations are supported. Restart operation is not supported.

4. On the Start Up, Shut Down, or Restart page, provide the following details, and click **OK**.

Element	Default Value	Description
Create Blackout Before Shutting Down <i>(Appears only for shutdown operation)</i>	Selected	<p>When issuing a shut down operation from Cloud Control, you have an option to put the target(s) in a blackout state. Two different blackout states are available. The first and default option is a traditional blackout where monitoring of the target(s) is suspended in order to perform maintenance (the agent does not perform metric data collection on the target(s) and no notifications will be raised for the target(s)).</p> <p>The second option is a notification blackout where only event notifications on the target(s) is suspended (the agent continues to monitor the target(s)). If you do not want a blackout created, deselect this option.</p> <p>Note:</p> <ul style="list-style-type: none"> • This option does not appear for restart operation. • If the selected target is a composite target, then Enterprise Manager creates blackouts for all its member targets. • If the option is selected, then the blackouts are created on targets even if the start up or shut down operation fails.
End Blackout After Starting Up <i>(Appears only for start up operation)</i>	Selected	<p>Ends blackouts on targets after they are started. By default, the option is selected. Deselect it if you do not want Enterprise Manager to automatically end blackouts on the targets.</p> <p>Note:</p> <ul style="list-style-type: none"> • This option does not appear for restart operation. • If the selected target is a composite target, then Enterprise Manager ends blackouts for all its member targets. • If the option is selected, then the blackouts are ended on targets even if the start up or shut down operation fails.
Include Administration Server <i>(Appears only for Oracle WebLogic Domains)</i>	Not Selected	<p>Select this if you want to start or stop even the Administration Server when the Oracle WebLogic Domain to which the Administration Server belongs, is started or stopped.</p> <p>Note: The Administration Server can be stopped only if the Management Agent that is monitoring it is running on the same host as the Administration Server.</p>
Time Out After (in minutes)	5 Minutes Per Target	<p>Set the time limit (in minutes) for the job to wait while it is trying to start, stop, or restart a target before terminating the attempt and generating an error.</p> <p>By default, it is set to 5 minutes, and it applies to each target. If a composite target is selected, then the timeout is per member target.</p>

Element	Default Value	Description
Process Control Method <i>(Appears only for Oracle WebLogic Domains, Oracle WebLogic Clusters, Oracle WebLogic Servers)</i>	Administration Server	<p>Select one of the following ways in which the shutdown, start-up, or restart operation can be performed:</p> <p>Note: Options not applicable to a particular target type are disabled.</p> <ul style="list-style-type: none"> Administration Server Uses the Administration Server to start up, shut down, or restart a target. For this option, as a prerequisite, ensure that the Administration Server is up and is accessible by the Oracle Management Agent monitoring the server. Node Manager Uses Node Manager for Oracle WebLogic Servers, Oracle WebLogic Clusters, and Oracle WebLogic Domains to start up, shut down, or restart the target. As a prerequisite, ensure that the Node Manager are up and are accessible by the Oracle Management Agent monitoring the target. Monitoring agent should be local to target. Default Script Uses the <code>startManagedWeblogic</code> script and the <code>stopManagedWeblogic</code> script located in the <code><DOMAIN_HOME>/bin</code> directory to start up, shut down, or restart a target. For this option, as a prerequisite, ensure that the Administration Server is up and is accessible by the Oracle Management Agent monitoring the server. Also, configure the <code>boot.properties</code> file for the server. For information on boot identity files and instructions to configure them, see <i>Oracle Fusion Middleware Administering Server Startup and Shutdown for Oracle WebLogic Server</i>. Custom Script Uses a custom script you specify to start up, shut down, or restart a target. For this option, as a prerequisite, ensure that the Administration Server is up and is accessible by the Oracle Management Agent monitoring the server. Also, configure the <code>boot.properties</code> file for the server. For information on boot identity files and instructions to configure them, see <i>Oracle Fusion Middleware Administering Server Startup and Shutdown for Oracle WebLogic Server</i>.
Credentials	Preferred	<p>If default script or custom script is selected then Administration Server Credentials are not required, only agent host credentials are required.</p> <p>You can use preferred or named credentials if you have already registered the credentials with Enterprise Manager Cloud Control, or you can enter a new set of credentials to override the preferred or named credentials.</p>
Targets	Not Selected	<p>You can perform the start/stop operation in context of the selected target (not in context of the job UI). You can pick and choose a the member targets of a domain that you want to start or stop.</p>

 **Note:**

If a remote Management Agent is monitoring a Java EE application target, such as Oracle Data Integrator Agent, then while starting up, shutting down, or restarting that Java EE application target, you might see errors. A remote Management Agent is a Management Agent that is not installed on the host where the target is running.

To circumvent this error, follow these steps:

1. On the host where the Java EE application target is running, navigate to the following location in the middleware home:

```
cd $<MIDDLEWARE_HOME>/wlserver_10.3/server/lib
```

For example,

```
cd /u01/software/middleware/wlserver_10.3/server/lib/
```

2. Generate the `wlfullclient.jar` file:

```
java -jar wljarbuilder.jar
```

3. On the remote host where the Management Agent is running, copy the generated `wlfullclient.jar` file to the following location in the Management Agent home:

```
<AGENT_HOME>/sysman/jlib
```

For example,

```
cp /u01/software/middleware/wlserver_10.3/server/lib/  
wljarbuilder.jar /u01/software/agent/core/12.1.0.3.0/sysman/jlib/
```

 **Note:**

If a job fails at the *Start/Stop/Restart* step with the following error, then follow the workaround steps outlined in this note to resolve the issue.

```
Remote operation finished but process did not close its stdout/stderr
```

1. Open the user-defined custom script file.
2. Identify the line where command, which caused the error, was invoked.

For example,

```
my $startStopScript = "/scratch/aime/wl_home/user_projects/domains/  
base_domain/bin/startManagedWebLogic.sh";
```

3. Add the following code snippet after the above line:

```
if($isWindows){  
    $startStopScript= "cmd /c start /b $startStopScript";  
    # redirecting to NUL  
    close STDOUT;  
    close STDERR;  
    open(STDOUT, ">", "NUL");  
    open(STDERR, ">", "NUL");  
} else{  
    $startStopScript= "$startStopScript > /dev/null 2>&1 &";  
}
```

2.7 Auditing WebLogic-specific Operations

Auditing is the process whereby information about operating requests and the outcome of those requests are collected, stored and analyzed for the purposes of non-repudiation. Auditing produces an electronic trail of computer activity. Enterprise Manager users can enable the auditing of WebLogic-specific operations - including the following:

- Logging in to a domain
- Updating a domain
- Logging out of a domain

By default, these three types of operations are not enabled for auditing. An administrator would have to enable them via the Enterprise Manager Command Line Interface (EMCLI). To audit these events, enter the following EMCLI command:

```
emcli update_audit_settings -  
operations_to_enable="WEBLOGIC_DOMAIN_UPDATE_INVOKE;WEBLOGIC_DOMAIN_LOGIN;WEB_LOGIC_D  
OMAIN_LOGOUT"
```

 **Note:**

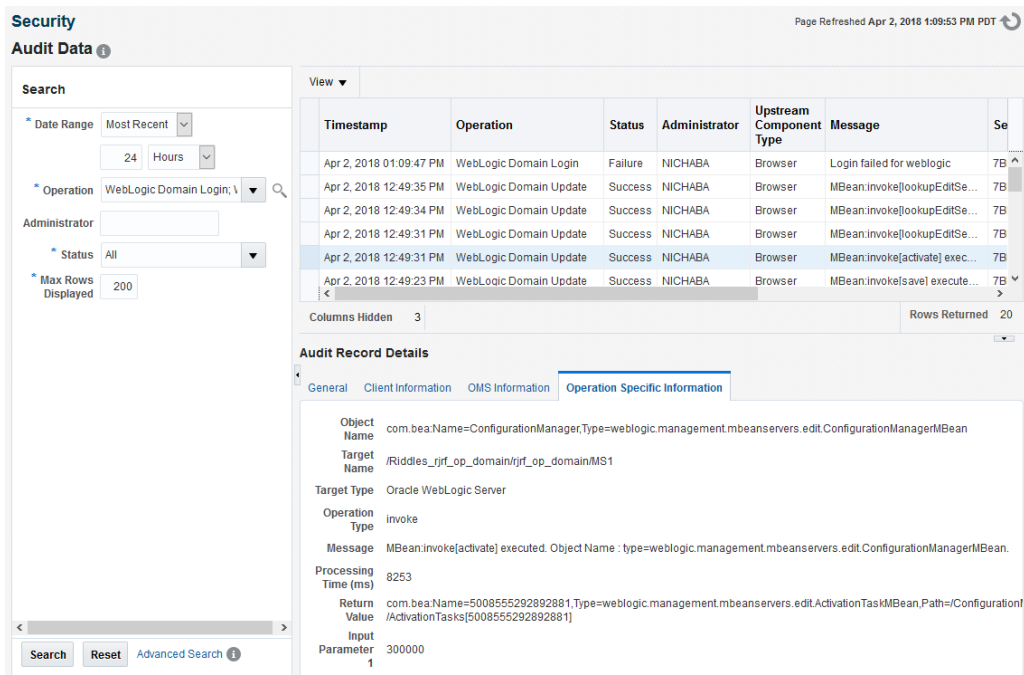
These operations are audited by Enterprise Manager Cloud Control when they are performed from either the Enterprise Manager Cloud Control console or from the EMCLI. If the operations are performed from the Oracle Enterprise Manager Fusion Middleware Control console or from the Oracle WebLogic Server Administration Console or from the WebLogic Scripting Tool (WLST), Enterprise Manager Cloud Control will not audit the operations.

After enabling these events, a super administrator is able to view and analyze the audited data. A super administrator can search for audit data that has been generated over a specified time period, and can also search on the following:

- Audit details of a specific WebLogic user or all WebLogic users.
- Audit details of WebLogic-specific operations with a Success or Failure status.

To view the audit data, a super administrator can navigate from the **Setup** menu, select **Security** and then **Audit Data**. The **Audit Data** page is displayed. Specify the search criteria in the fields and click **Search**. The results are displayed in the summary table.

To drill down to the full audit record details, click on the Timestamp for a row in the summary table.



Security Page Refreshed Apr 2, 2018 1:09:53 PM PDT ↻

Audit Data

Search

* Date Range: Most Recent
24 Hours

* Operation: WebLogic Domain Login; \

Administrator:

* Status: All

* Max Rows Displayed: 200

Timestamp	Operation	Status	Administrator	Upstream Component Type	Message	Se
Apr 2, 2018 01:09:47 PM	WebLogic Domain Login	Failure	NICHABA	Browser	Login failed for weblogic	7B
Apr 2, 2018 12:49:35 PM	WebLogic Domain Update	Success	NICHABA	Browser	MBean:invoke[lookupEditSe...	7B
Apr 2, 2018 12:49:34 PM	WebLogic Domain Update	Success	NICHABA	Browser	MBean:invoke[lookupEditSe...	7B
Apr 2, 2018 12:49:31 PM	WebLogic Domain Update	Success	NICHABA	Browser	MBean:invoke[lookupEditSe...	7B
Apr 2, 2018 12:49:31 PM	WebLogic Domain Update	Success	NICHABA	Browser	MBean:invoke[activate] exec...	7B
Apr 2, 2018 12:49:23 PM	WebLogic Domain Update	Success	NICHABA	Browser	MBean:invoke[save] execute...	7B

Columns Hidden: 3 Rows Returned: 20

Audit Record Details

General Client Information OMS Information **Operation Specific Information**

Object Name: com.bea.Name=ConfigurationManager.Type=weblogic.management.mbeanservers.edit.ConfigurationManagerMBean

Target Name: /Riddles_rjrf_op_domain/rjrf_op_domainMS1

Target Type: Oracle WebLogic Server

Operation Type: invoke

Message: MBean:invoke[activate] executed. Object Name : type=weblogic.management.mbeanservers.edit.ConfigurationManagerMBean.

Processing Time (ms): 6253

Return Value: com.bea.Name=5008555292892881.Type=weblogic.management.mbeanservers.edit.ActivationTaskMBean_Path=/Configuration/ActivationTasks[5008555292892881]

Input Parameter: 300000

1

2.8 About Lifecycle Management

Enterprise Manager Cloud Control offers lifecycle management solutions that help you meet all lifecycle management challenges by automating time-consuming tasks related

to cloning, patching, configuration management, ongoing change management, compliance management, and disaster recovery operations.

2.8.1 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, 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. Configurations 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 configurations. You can also compare one configuration with multiple configurations or one configuration in the Management Repository with a saved configuration. When a comparison operation results in differences that you do not require, you can synchronize the configurations so that one of the configurations replaces the other one. This synchronization can be performed on demand based on the configurations being compared.

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 to parse comparison data and fine-tune comparisons. For example, you can specify rules that indicate which differences must initiate email notifications 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. Enterprise Manager also provides the drift and consistency configurations for Fusion Middleware components. Configuration drift ensures consistency/uniformity across a large number of targets, whereas configuration consistency replicates the changes of target members within a system or group. For example, configuration consistency is used to ensure all of the WebLogic servers within a WebLogic cluster have the same configuration. For more information on configurations, see Oracle Enterprise Manager Lifecycle Management Administrator's Guide.

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. This allows you to identify the cause for improved or deteriorated performance due to a change in the BPEL Process suitcase file.

2.8.2 Compliance Management

Enterprise Manager Cloud Control offers the following compliance management features:

- 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 on 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.

- Compliance standard for the DISA published Security Technical Implementation Guide (STIG Version 1.1 and STIG Version 1.2) for Oracle WebLogic Server 12c, is provided out-of-box with Enterprise Manager Cloud Control 13c Release 1 and later. This compliance ensures that Oracle WebLogic servers installed and configured from Oracle Fusion Middleware Infrastructure Release 12.1.3 are compliant with Oracle WebLogic Server 12c STIG - Version 1, Release 1.

For additional information about compliance management, refer to the Managing Compliance chapter in the *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

2.8.3 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.

You can get automated patch recommendations from My Oracle Support on what patches to apply and then use patch plans to apply them. Patch Plans enable 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. This gives you the ability to apply patches in a rolling fashion to minimize downtime or in parallel fashion, thus implementing the best possible patch rollout for your organization.

Fusion Middleware best uses patch management for:

- Applying one or more patches to WebLogic Servers spanning one or more domains
- Applying patches to SOA Infrastructure targets

- Using validation checking to identify patch conflicts or other potential problems before the patches are actually applied.

For more information about patching, see *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

2.8.4 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 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, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning**.

From the Middleware Provisioning page, you can:

- Gain access to all Fusion Middleware related operations.
- Create profiles in the software library that can be used for future cloning operations. A WebLogic Domain Provisioning Profile consists of the Middleware Home, binaries and the domain configuration. You can create a profile, save it in the Software Library, and then use the saved profile as the source for creating new WebLogic domains. This will ensure that future WebLogic installations follow a standard, consistent configuration.
- Access the deployment procedures, both pre-defined and user-defined, to provision software and configurations.
- Automate the cloning of WebLogic Domains and/or Middleware Homes from a profile present in the software library.
- 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.
- Automate the scaling down of a cluster by removing a managed server from the cluster.

For more information on using provisioning, see Middleware Provisioning section in the *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

2.8.4.1 Deploying / Undeploying Java EE Applications

You can deploy, undeploy, and redeploy Java EE applications (for example, .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.

For more information, see Middleware Provisioning section in the *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

2.9 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.

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.

2.9.1 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.

2.10 Job System

Enterprise Manager has a job system that automates WebLogic administrator tasks. Enterprise Manager offers two types of predefined job types for WLS related targets namely the Fusion Middleware Process Control job type and the WLST Script job type.

In addition to executing and scheduling WLST scripts and Fusion Middleware process control operations from the job system context, you can also execute these administrative tasks as Corrective Actions. That is, you can associate a WLST Script or Fusion Middleware Process Control Corrective Action to automatically run in response to a threshold being crossed. For example, when an Oracle WebLogic Server target goes down, you could have a Fusion Middleware Process Control corrective action to automatically start it again. For more information, see *Utilizing the Job System and Corrective Actions in the Enterprise Manager Cloud Control Administrator's Guide*.

2.11 Routing Topology Viewer

Enterprise Manager provides a Routing 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 to Oracle HTTP Server, to a Managed Server, to a data source, to a database.

The Routing 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. Using filters you can 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 including performance metric data. You can see which entities are up and which are down. You can also print the topology using the Print to File feature on your printer's settings/options. For more details, see the *Enterprise Manager Online Help*.

2.4 Analyzing Middleware Problems Using Problem Analysis

Most of the information presented in Enterprise Manager concerns one or more of the targets managed or monitored by Cloud Control. When a problem is encountered in any managed entity, you must navigate to multiple screens to gather information to triage the problem. The Problem Analysis functionality allows you to see all the related information in one place. You can choose a problematic spike in a metric chart and do root cause analysis based on system knowledge enabling you to narrow the scope of the problem quickly.

You can use the Problem Analysis and Analyze Log pages in Cloud Control to help you inspect metrics, related metrics, target status information, incidents, and logs during troubleshooting.

Accessing Problem Analysis and Logs

There are several navigation methods to access Problem Analysis and log pages:

- **Middleware access method**
 1. From the targets menu of the Cloud Control console, select **Middleware**.

2. Select and click on an **Oracle HTTP Server** or **Oracle WebLogic Server** from the Details Table.
 3. In the Home page that appears, click on a metric legend that appears below the Response and Load chart.
 4. In the pop-up that appears, click on **Problem Analysis** or **Log Messages**.
- **Incident manager access method**
 1. From the targets menu of the Cloud Control console, select **Hosts**.
 2. Click a numbered link in the **Incidents** column of the summary table.
 3. In the Incident Manager page that appears, select an incident in the table, then click on the **Problem Analysis** link located in the Diagnostics section in the lower right portion of the page.
 - **Correlation charts method**

Correlation charts are the pages in which the charts are shown as a stack of charts.

 1. From any correlation chart, click on the chart legend.
 2. In the pop-up that appears, click on **Problem Analysis** or **Log Messages**.
 - **Chart regions method**

Chart regions are charts displayed on the home page, or a single chart shown on some pages.

 1. Click on the chart legend or chart line.
 2. In the pop-up that appears, click on **Problem Analysis** or **Log Messages**.

Using the Tabs on the Problem Analysis Page

You can use the five tabs on the Problem Analysis page to perform the following tasks or view the following data:

- **Related Metrics**

Displays the related metrics which are affected or could affect the source metric. Optionally, you can choose to customize the affected metric.

You can choose **Export Chart Set** from the Chart Sets drop-down to create a Problem Analysis xml metadata file for the particular source metric, including default and custom related metrics defined for the chosen chart set. Similarly, exported chart set data can be imported to a custom chart set using the **Import Chart Set** from the Chart Sets drop-down.
- **Related Targets**

Displays the related targets to the target instance of the source metric. Related Targets information provides key information such as Status, Status change time, Incidents, Configuration changes, Important Key metrics and Patches applied at one place rather than your viewing this information in different locations within the Enterprise Manager console.
- **Related Config**

Displays the related configuration metrics which are affected or could affect the source metric.
- **Related Logs**

Displays charts for each target and depicts the number of messages for each severity level. Data is collected for the related targets irrespective of the search filter criteria. The page includes one graph for each target with different columns for each severity. Data is gathered based on the time range and other values defined in the filter.

- **Topology**
Provides the topology view of any related targets.

Viewing and Analyzing Problems

You can inspect metrics, status information, and logs using Cloud Control by following these steps:

1. Specify the time period for which you want the charts to display data. Near the top of the default Related Metrics tab, adjust the left and right slider to specify the time period, or click and drag within a metric chart to indicate the time period you want to inspect.
2. Inspect the charts for unusual increases in recorded metrics.
 - Out of the box, Enterprise Manager provides two charts: Source Metric and Enterprise Manager Identified Related Metrics. You can add more chart displays to suit your needs by using the Metric Palette. See "Customizing the Display" below for more information.
 - Increased request processing time due to a high number of requests per minute may indicate a need to increase the capacity of your system.
3. If the metric charts do not indicate the cause of the problem, select the **Related Targets** tab and inspect the table for information about target health (status) and recent configuration changes.

If you want to see a reminder of the topology of the components for which data is being displayed, click the **Topology** tab.
4. If the table does not indicate the cause of the problem, return to the Related Metrics tab and click the **View Related Log Messages** link near the top of the tab. This action displays log messages for the selected target and its members during the selected time period.
5. Inspect any log messages that are displayed for possible causes of problems.

Using Log Analysis

You can use Log Analysis in one of two ways:

- **Target Log Analysis** -- Click the **Log Analysis** link on the chart pop-up to view the log for the target on which the metric chart is displayed. The log viewer is launched with the start time and end time as the same time duration of the chart with filters applied.
- **Related Logs** -- Click **Related Logs** on the Problem Analysis page to view all the related log messages for all the related targets during the viewed metric chart duration with all the filters applied.

Customizing the Display

You can create your own metric charts and then recall them at a later time when needed.

1. From the Metric Palette on the Related Metrics tab, select a target from the Targets pane, then select the desired metrics associated with the target from the Metrics pane.

A region named User Identified Related Metrics appears in the lower portion of the page and displays a chart for each metric you have selected in the Metric Palette.

2. *Optional:* Save any modifications to the current chart by clicking **Save**.

You can also save your modified chart to Enterprise manager and have it appear as a choice in the Charts Sets menu for recall at a later time. To do so, select **Save Charts As...** from the Chart Sets menu, then name the chart and click **OK**. To set the saved chart set as the default chart set, select the saved chart set listed on the Chart Set menu and then click **Set as Default Chart Set** from the Chart Set menu.

You can also set this chart as the default chart that appears when you access this page by selecting **Set as Default Chart Set** from the Chart Sets menu.

 **Tip:**

If you prefer seeing the chart data in a tabular format, you can click the **Table View** link below the last chart.

2.5 Managing Problems with 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. Support Workbench also recommends and provides easy access to Oracle advisors that help you repair data corruption problems, and more.

Support Workbench Compatibility with Fusion Middleware Components

You can use Support Workbench with:

- Oracle WebLogic Server
- SOA Infrastructure

 **Note:**

Support Workbench is available only for WebLogic Server Domains with the Java Required Files (JRF) template applied.

Basic Support Workbench Work Flows

You can use Support Workbench to manage problems in two basic ways:

- Respond to alert notifications by packaging associated problems and uploading them to Oracle Support for resolution.
- Proactively package observed problems and upload them to Oracle Support for resolution.

The process by which you receive alerts and use Support Workbench is as follows:

1. The Enterprise Manager Agent has collected one or more metrics that have exceeded the thresholds that have been set.
2. The alert log generates an incident and you are notified of a pending alert.
3. You search for and view problems within Support Workbench.
4. You access My Oracle Support to search for this problem or a similar problem, and to determine a proper course of action to resolve the problem. If the search is unsatisfactory, you continue to the next step.
5. You create a package for My Oracle Support that includes supporting material, such as external files, executed dumps, and so forth.
6. You create a service request.
7. You upload the package to My Oracle Support.

The process by which you proactively observe problems and upload them to Oracle Support is the same as steps 3 through 7 above, but you initiate a user-reported problem before proceeding to step 5.

The following sections provide procedures to perform these tasks.

2.5.1 Accessing and Logging In To Support Workbench

The following sections explain how to access and log in to Support Workbench.

2.5.1.1 Accessing Support Workbench

To access Support Workbench:

1. From the Middleware home page, click on either an **Oracle WebLogic Domain**, **Oracle WebLogic Cluster**, or **Oracle WebLogic Server** in the Details Table.
2. From the Oracle WebLogic Domain or Oracle WebLogic Server menu, select **Diagnostics**, then **Support Workbench**.

2.5.1.2 Logging In

You can log in using either preferred credentials or named credentials you have previously set up. Otherwise, you can choose the New Credentials option to override the other two login options.

- **Prerequisites**

- The host credentials should have write privileges on the AdrHome location of the target.

AdrHome is the location where WebLogic server stores its incidents. Refer to the 'Diagnostic Framework Components' section located in the Oracle Fusion Middleware Administrator's Guide for information.

Logging in to Software Workbench requires you to have only read permissions to the AdrHome. With read-only permissions, you can still log in to Software Workbench and view problems and incidents, create a user reported problem, and execute additional diagnostic dumps. Creating a package requires you to have write permissions.

- The WebLogic credentials should have Monitor privileges on the WebLogic server.
- **Preferred Credentials Choice**

Select this choice if you want to use the credentials that you have already registered as preferred credentials on the Preferred Credentials page.

Preferred credentials simplify access to managed targets by storing target login credentials in the Management Repository. With preferred credentials set, you can access an Enterprise Manager target that recognizes these credentials without being prompted to log into the target. Preferred credentials are set on a per user basis, thereby ensuring the security of the managed enterprise environment.
- **Named Credentials Choice**

Select this choice if you want to use the credentials of a named profile you created on the Named Credentials page.

You can override host or WebLogic Server preferred credentials with this option. A named credential specifies a user's authentication information on a system. A named credential can be a username/password, a public key-private key pair, or an X509v3 certificate.
- **New Credentials Choice**

You can override previously defined preferred credentials or named credentials by using the New Credentials option. When you enter new credentials, you can save the credentials and give them a name, which consequently becomes Named Credentials.

 **Note:**

Support Workbench requires you to save the credentials when you choose the New Credentials option. By saving the credentials, you can submit Enterprise Manager jobs for long running Software Workbench operations (for example, packaging a problem) and Oracle requires access to the saved credential to perform these operations.

2.5.2 Using Fusion Middleware Support Workbench

You can use Support Workbench within Fusion Middleware to:

- [View an aggregated diagnostic summary](#)
- [Execute tests to diagnose a problem](#)
- [Create a problem, package it, and upload it to Oracle Support](#)

The following sections provide procedures for these diagnostic tasks.

2.5.2.1 Viewing Diagnostics

This procedure assumes that an incident occurred on a WebLogic Server, and you received an alert notification. You now need to determine the appropriate action to resolve the problem.

1. From the domain home page drop-down, select **Monitoring**, then **Incident Manager**.
2. Click the link in the **Target** column for the incident you want to investigate.
3. In the Monitoring and Diagnostics section of the page that appears, click the **Support Workbench Problems** numbered link.

2.5.2.2 Viewing an Aggregated Diagnostic Summary

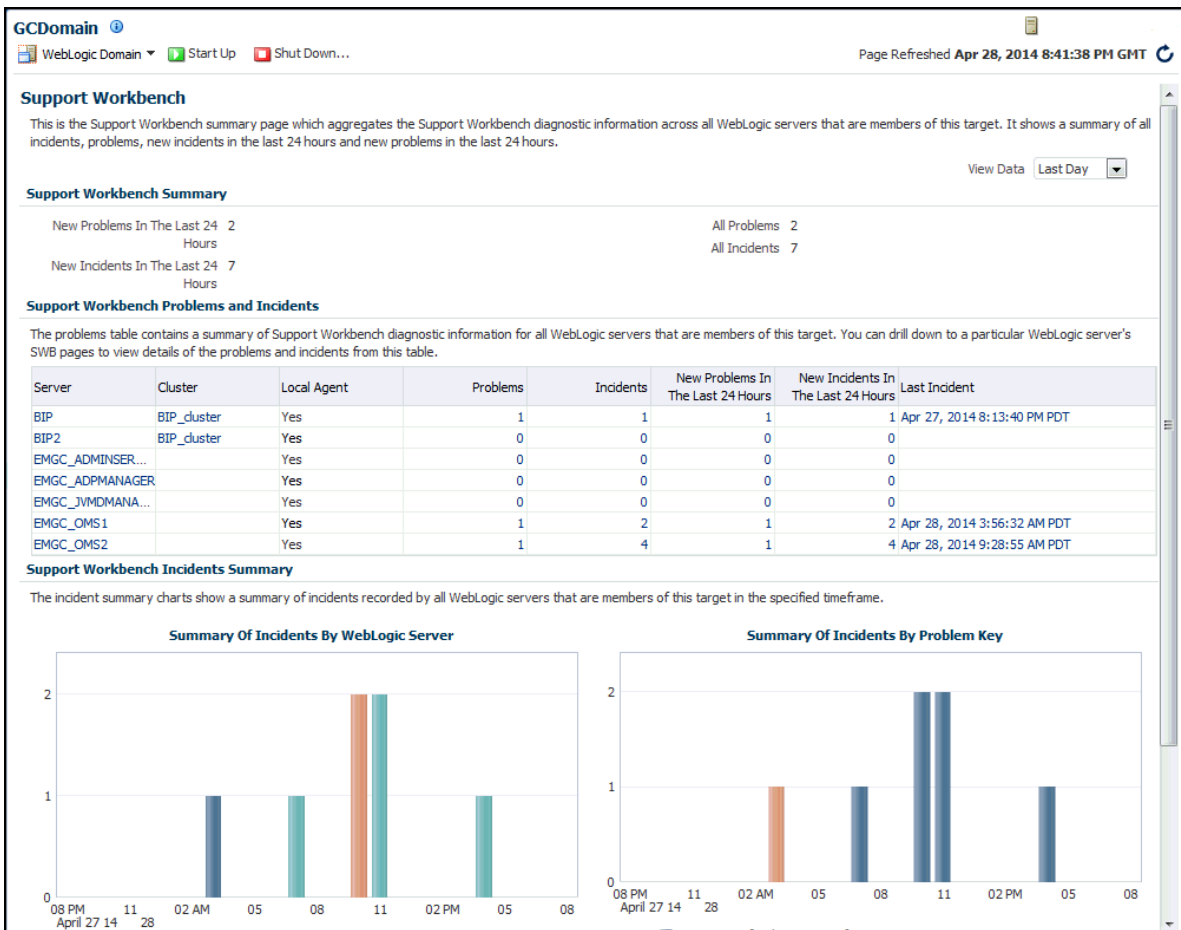
Fusion Middleware is deployed across multiple systems, and incidents are therefore recorded in multiple Automatic Diagnostic Repository homes. The following procedure describes how to get a quick summary of diagnostic data across all targets and Automatic Diagnostic Repository homes aggregated by the instance, product family, or cluster application.

This procedure is applicable to a WebLogic Domain and WebLogic Cluster in the context of Fusion Middleware. The procedure assumes that multiple Fusion Middleware incidents occurred on the servers deployed in a WebLogic domain, and you received multiple alerts from related servers.

1. After receiving alerts from related targets, access the Fusion Middleware instance, product family, or cluster application home page.
2. From the drop-down menu, select **Diagnostics**, then **Support Workbench**.

The Support Workbench home page appears, and displays a summary table with the problems and incidents aggregated by the application.

Figure 2-3 Aggregated Diagnostic Summary



- Sort the tables to see which WebLogic Server(s) have had the highest number of problems and incidents.
- Drill down through an individual server's Support Workbench pages to view detailed diagnostics information for the server, such as problems and incidents.

2.5.2.3 Searching for Problems

The following procedure assumes that problems are already recorded in Enterprise Manager.

- From the Support Workbench home page, enter search criteria in the **Filter by problem key field**, then click **Go**.

Search criteria includes keywords to use in the search, such as date range, problem key, SR number, and bug number.

You can also alternatively click the **Advanced Search** link, provide search criteria, then click **Search**.

2.5.2.4 Annotating a Problem

You may want to add short notes to a problem and then communicate this to other administrators.

1. From the domain home page drop-down, select **Monitoring**, then **Incident Manager**.

The Incident Manager page appears, and displays all open incidents in the table.

2. From the lower right side of the Incident Manager page, click the **Add Comment** link.
3. Add your comment in the pop-up that appears, then click **OK**.

Enterprise Manager records the comment and then redisplay it if this administrator or a different one looks at this problem.

2.5.2.5 Adding More Files

You may want to add more diagnosability information, such as diagnostic dumps, to an incident.

1. From the Support Workbench home page, select the ID link for the problem for which you want to add diagnostics.
2. From the Incident Details page that appears, click the ID for the associated incident.
3. Select the **Additional Diagnostics** tab.
4. Select a diagnostic from the list in the table, then click **Run**.
5. Enter values for required parameters on the Run User Action page, schedule the run, then click **Submit**.
6. When the confirmation message appears, click **OK**.

The diagnostic dump executes, and the results are attached to the incident.

2.5.2.6 Creating a Package

You have two options for creating a package. You can:

- Create a package initiated from alert notifications
- Proactively create a package from observed problems

To create a package initiated from alert notifications:

1. From the Support Workbench home page, select the ID link for the problem that you want to package.
2. From the Problem Details page that now appears, click **Quick Package**.

The Quick Packaging wizard appears.

3. Provide the requisite input in the wizard, then click **Submit**.

Most of the wizard is self-explanatory. Your input is required for the following wizard steps:

- Create New Package
 - Package Name — Accept the default system-supplied name, or provide your own descriptive name.
 - Package Description — Provide a description of any length as a reminder what this package consists of.

- Send to Oracle Support — If you enable this option, a confirmation message appears when processing has completed stating that the upload file for the package has been successfully generated, and also provides the location of the file.

If you decide not to send the package to Oracle support now, you can do so later From the Package Details page. The upload file is generated but not sent to Oracle if you choose No.

- Service Request Number — Enter the SR associated with this package. This is only required if you are uploading.
- Schedule
 - Immediately/Later — If you want to generate the upload files later rather than now, you do not need to change the time zone unless you want to specify a time in another time zone, such as the database time zone or the OMS time zone.
 - Host Credentials — The required host credentials should be the same as the credentials used to start up the target database.

To proactively create a package from observed problems:

1. From the Support Workbench home page, click **Create User-Reported Problem** in the Related Links section.
2. In the page that appears, select the issue type, then click **Continue with Creation of Problem**.
3. Follow the instructions in steps 2 and 3 above.

2.5.2.7 Providing Additional Files

You may want to add more information, such as external files, to a package. This procedure assumes that a package has been created and additional diagnostics have been generated for the problem.

1. From the Support Workbench home page, click the **Yes** link in the Packaged column for the package you want to modify.
2. From the Packages page, click the package name link.
3. From the Package Details page, click **Customize Package**.

The Customize Package page appears, where you can edit the package contents, generate and include additional diagnostic data, or scrub user data.

2.5.2.8 Uploading a Package to Oracle Support

1. From the Package Details page, described in the previous section, click **Generate Upload File**.
2. Indicate the package file type, select the schedule, then click **Submit**.
3. After the confirmation message appears, click **OK**.
4. Click **Send to Oracle**.
5. Choose an existing SR or create a new SR to upload the package to.

2.5.2.9 Creating a Service Request

Following packaging and uploading the problem to Oracle support, you may want to create service request to address a problem through Oracle supp6ort.

1. From the Cloud Control console Enterprise menu, select **My Oracle Support**, then **Service Requests**.

After providing your Single Sign-on credentials, the Service Requests tab of the My Oracle Support site opens.

2. Click **Create "Contact Us" SR**.
3. Provide the necessary input in the wizard that appears, then click **Submit**.

2.5.2.10 Managing Problem Resolution

After the problem is resolved, close it so that Automatic Diagnostic Repository (ADR) can purge the required memory for the problem.

1. From the Support Workbench home page, select the ID link for the problem you want to manage.
2. From the Problem Details page, click the **Manage problem resolution** link in the Investigate and Resolve section of the page.

Several management options are available on the Incident Manager page that appears.

For more information about managing incidents in Enterprise Manager, see [Using Incident Management](#) in the *Enterprise Manager Cloud Control Administrator's Guide*.

3

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](#)
- [Prerequisites and Considerations When Using Application Replay](#)
- [Understanding the Application Capture and Replay Process](#)
- [Creating Application Workload Captures](#)
- [Monitoring the Application Capture Process](#)
- [Replaying Application Workload Captures](#)
- [Importing Replay Session Divergences into OpenScript](#)
- [Troubleshooting Application Replay](#)

3.1 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 a lower risk.

3.2 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 activity 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.

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

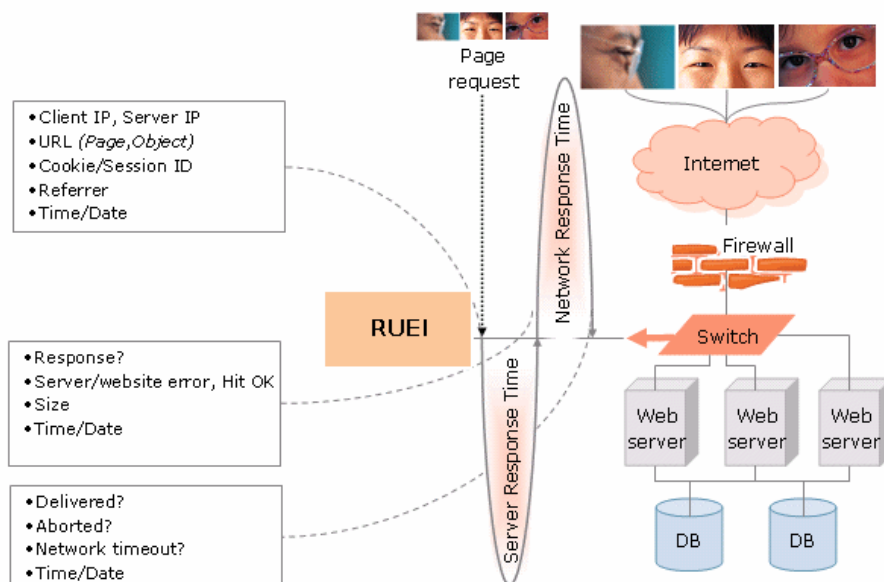
- It provides a system-wide perspective starting from the user's activity. This is in contrast to the traditional piecemeal testing of individual components that provides little information on their combined behavior and performance under a realistic workload.
- Rather than relying on pre-determined 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.

3.3 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 3-1](#).

Figure 3-1 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>

3.4 Prerequisites and Considerations When Using Application Replay

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.

 **Note:**

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](#)
- [Setting up the Test System Database for Application Replay](#)
- [Setting up the Capture Directory for Application Replay](#)

3.4.1 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 [Managing Security-Related Information](#) in *Oracle Real User Experience Insight User's Guide*.

3.4.2 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.

To assign these privileges, find the resource type named Application Replay Entities, then click the pencil icon to add the above privileges.

In addition to the above, users must also be assigned the `PERFORM_OPERATION_ANYWHERE` (Execute Command 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`.

3.4.3 Setting up the Test System Database for Application Replay

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

3.4.4 Setting up the Capture Directory for Application Replay

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

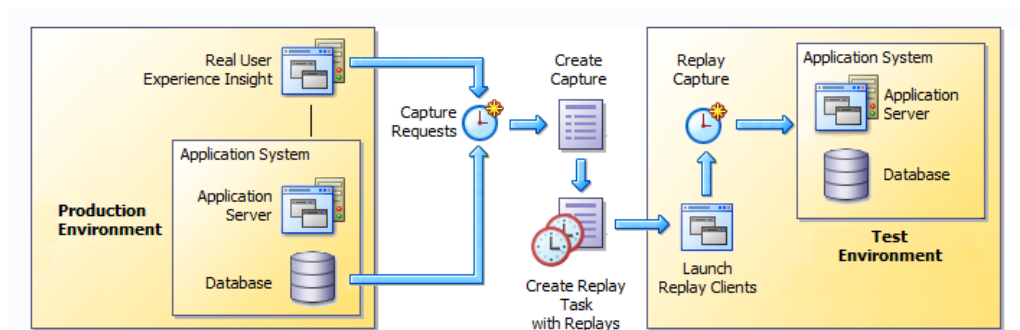
To estimate the required disk space, it is recommended that you run a test capture on your workload for a short duration (typically, a few minutes), and then use this to

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.

3.5 Understanding the Application Capture and Replay Process

Figure 3-2 shows the architecture of the Application Replay facility.

Figure 3-2 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.

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.

3.6 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 3-3.

Figure 3-3 Application Replay Page

The screenshot shows the 'Application Replay' page with the 'Overview' tab selected. Below the tab, there are buttons for 'Create...', 'Create Like...', 'Edit...', 'Delete', and 'Query By Example'. A pagination bar shows '1-16 of 16' items. Below this is a table with the following data:

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_ats_ip	Completed	EBS6170_system	1	PJONES	02-Sep-2011 16:40:55	
<input type="radio"/>	ruei12_form_https_cookiei	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 3-4 appears.

Figure 3-4 Create Capture (Overview) Page

The screenshot shows the 'Create Capture : Overview' page. At the top, there is a progress bar with steps: Overview (selected), System, RUEI Application, Database, Storage, Schedule, and Review. Below the progress bar, there are 'Back', 'Step 1 of 7', 'Next', and 'Cancel' buttons. The main form area contains:


- * Name: Siebel CRM
- Description: North America and Europe Siebel CRM application.

Below the form is a section titled 'Capture Prerequisites' with the text: 'These prerequisites should be completed before performing a capture.' There are two check boxes, both of which are checked:

- 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 3-5 appears.

Figure 3-5 Create Capture (System) Page



5. 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 3-7) is not available, and is skipped. When ready, click **Next**. The page shown in Figure 3-6 appears.

Figure 3-6 Create Capture (RUEI Application) Page



6. 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 3-7](#) appears.

Figure 3-7 Create Capture (Database) Page

7. 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 3-8](#) appears.

Figure 3-8 Create Capture (Storage) Page

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.

 **Note:**

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 3-9](#) appears.

Figure 3-9 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 3-10](#) appears.

 **Note:**

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

Figure 3-10 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|Siebel|CRM|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

- 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**.

3.7 Monitoring the Application Capture Process

After 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 3-11](#).

Figure 3-11 Capture Page



Be aware that there is a lead time of approximately 10 minutes after the start of a capture before progress information about it becomes available. This is available from the Application Replay page (Figure 3-3). 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.

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

3.8.1 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. See [Setting up the Test System Database for Application Replay](#).
- All references to external systems have been resolved. See [Resolving References to External Systems for Application Replays](#).

3.8.2 Understanding Application 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 be grouped into the same replay task in order to facilitate comparison. For example, replays that relate to the testing of the same database upgrade patch should be grouped into the same replay task.

3.8.3 Resolving References to External Systems for Application Replays

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 3-1](#).

Table 3-1 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.

Table 3-1 (Cont.) References to External Systems

Type	Description
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.

 **Note:**

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.

3.8.4 Remapping URLs for Application Replays

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.

3.8.5 Substituting Sensitive Data for Application Replays

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.

3.8.6 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 3-3](#) appears.
2. From the **Replay Tasks** section, click **Create** or **Create Like**. The page shown in [Figure 3-12](#) appears.

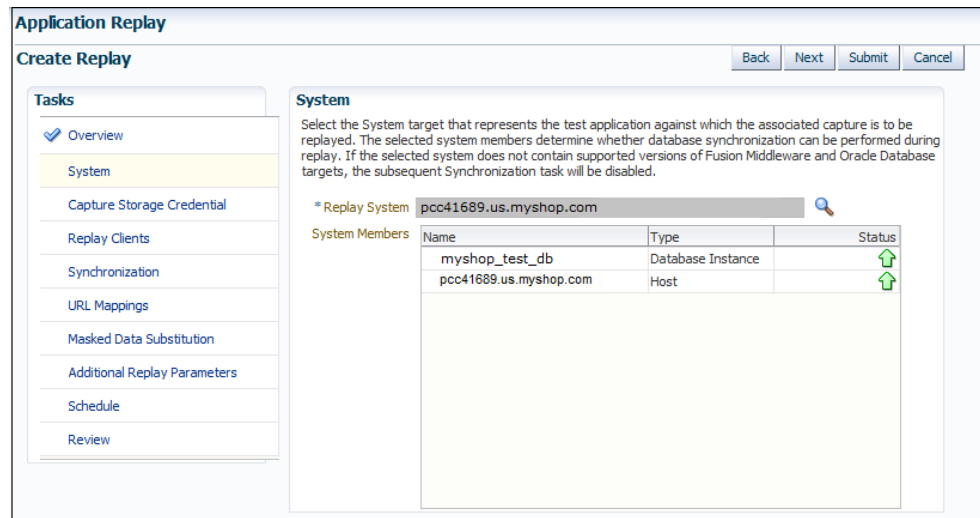
Figure 3-12 Create Replay Task Page

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 3-3](#).
4. Click the newly created replay task. The page shown in [Figure 3-13](#) appears.

Figure 3-13 Create Replay (Overview) Page

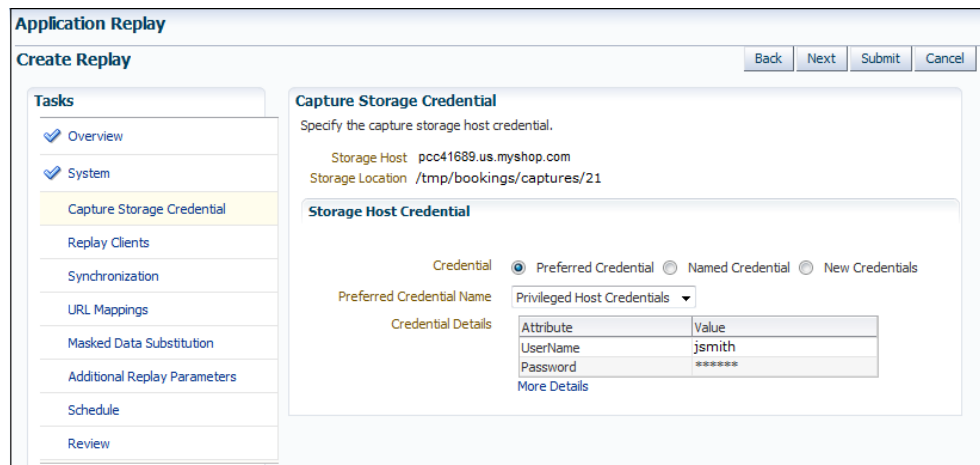
5. Specify a name for the replay. It must be unique among the selected replay task. Optionally, specify a brief description for the traffic to be replayed. It is recommended that you include an indication of the purpose and scope of the replay. Carefully review the required information, and click the acknowledgement check boxes to indicate that they have been met. When ready, click **System**. The page shown in [Figure 3-14](#) appears.

Figure 3-14 Create Replay (System) Page



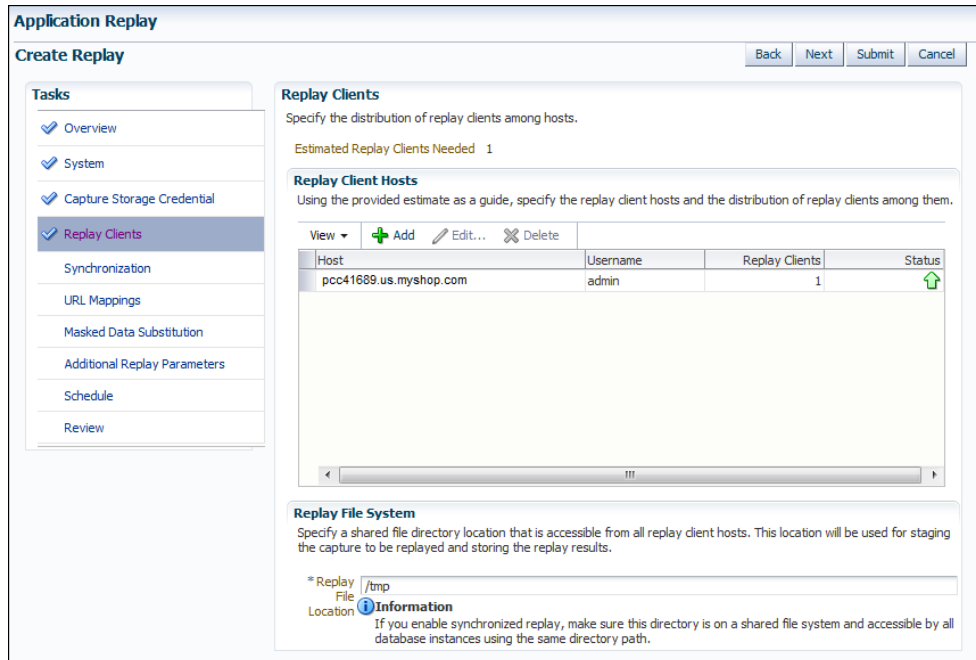
6. Click the **Select System** icon, and select the targets that represent the test environment against which the capture should be replayed. It is recommended that you review the status of the selected components, and ensure that they will be available throughout the planned replay. When ready, click **Capture Storage Credential**. The page shown in [Figure 3-15](#) appears.

Figure 3-15 Create Replay (Capture Storage Credential) Page



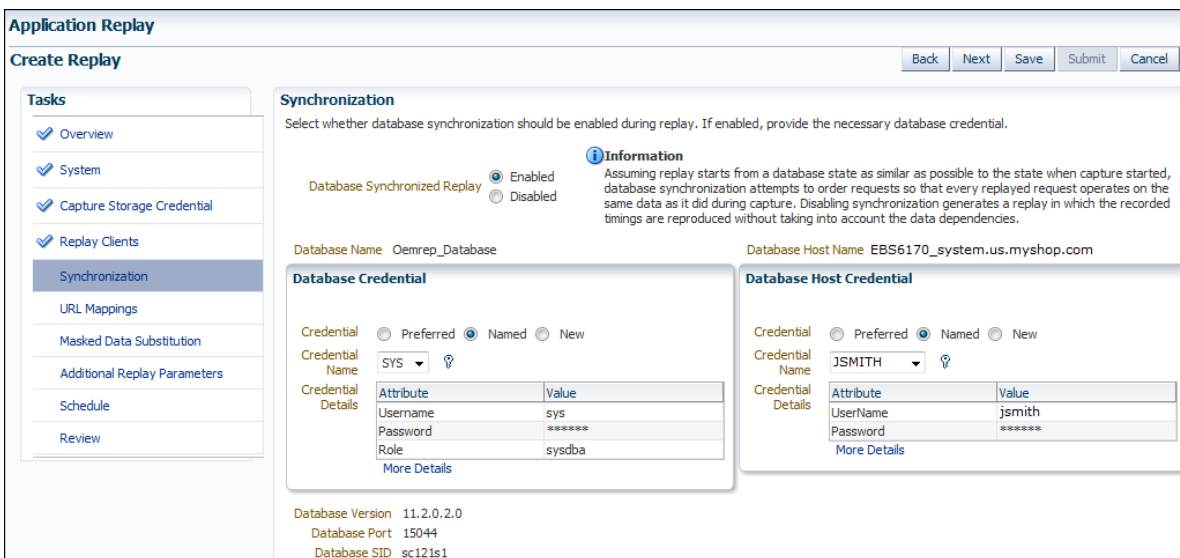
7. Specify the credentials of the selected storage host. When ready, click **Replay Clients**. The page shown in [Figure 3-16](#) appears.

Figure 3-16 Create Replay (Replay Clients) Page



- 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 3-17 appears.


Figure 3-17 Create Replay (Synchronization) Page



- 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 3-18](#) appears.

Figure 3-18 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 3-19](#) appears.

Figure 3-19 Create Replay (Masked Data Substitution) Page



11. 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 3-20](#) appears.

Figure 3-20 Create Replay (Additional Replay Parameters) Page

The screenshot shows the 'Application Replay' interface. On the left is a 'Tasks' sidebar with options: Overview, System, Capture Storage Credential, Replay Clients, Synchronization, URL Mappings, Masked Data Substitution, **Additional Replay Parameters** (highlighted), Schedule, and Review. The main area is titled 'Additional Replay Parameters' and contains the following sections:

- 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

Navigation buttons: Back, Next, Submit, Cancel.

- 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 3-21](#) appears.

Figure 3-21 Create Replay (Schedule) Page

The screenshot shows the 'Application Replay' interface. On the left is a 'Tasks' sidebar with options: Overview, System, Capture Storage Credential, Replay Clients, Synchronization, URL Mappings, Masked Data Substitution, Additional Replay Parameters, **Schedule** (highlighted), and Review. The main area is titled 'Schedule' and contains the following sections:

- Schedule:** Schedule a time for replay to start.
 - Replay Start:**
 - Start: Immediately, Later
 - Date/Time: (UTC-08:00) US Pacific Time

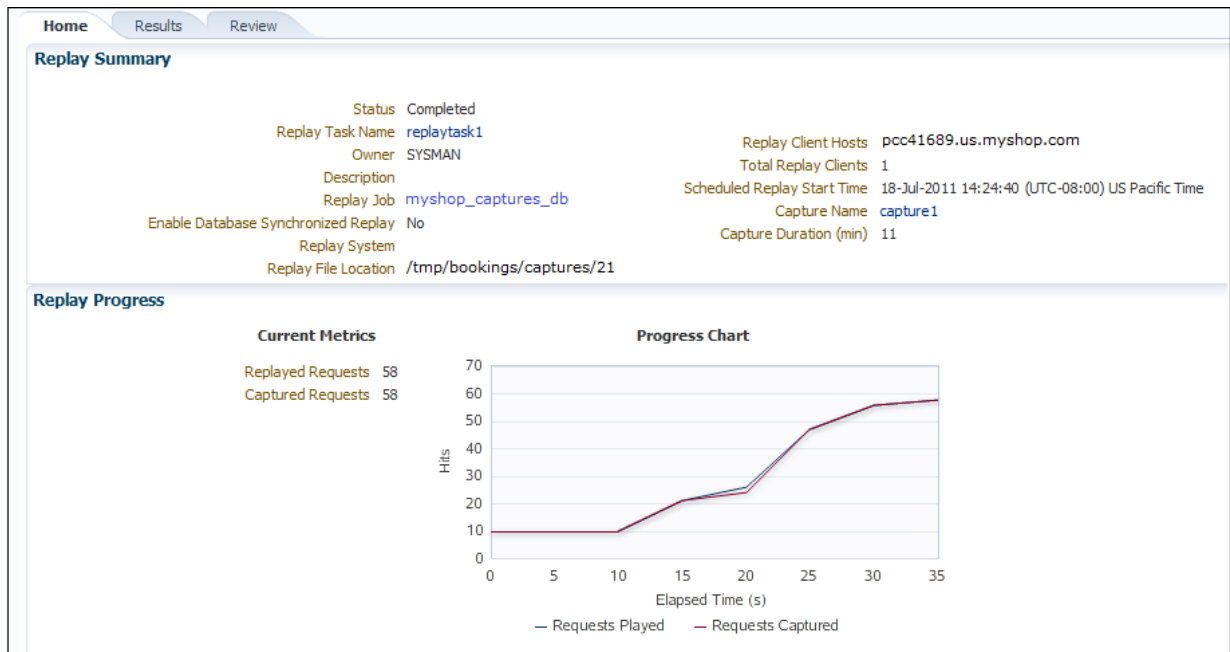
Navigation buttons: Back, Next, Submit, Cancel.

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.

3.8.7 Analyzing Application 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 3-22](#).

Figure 3-22 Example Replay Summary



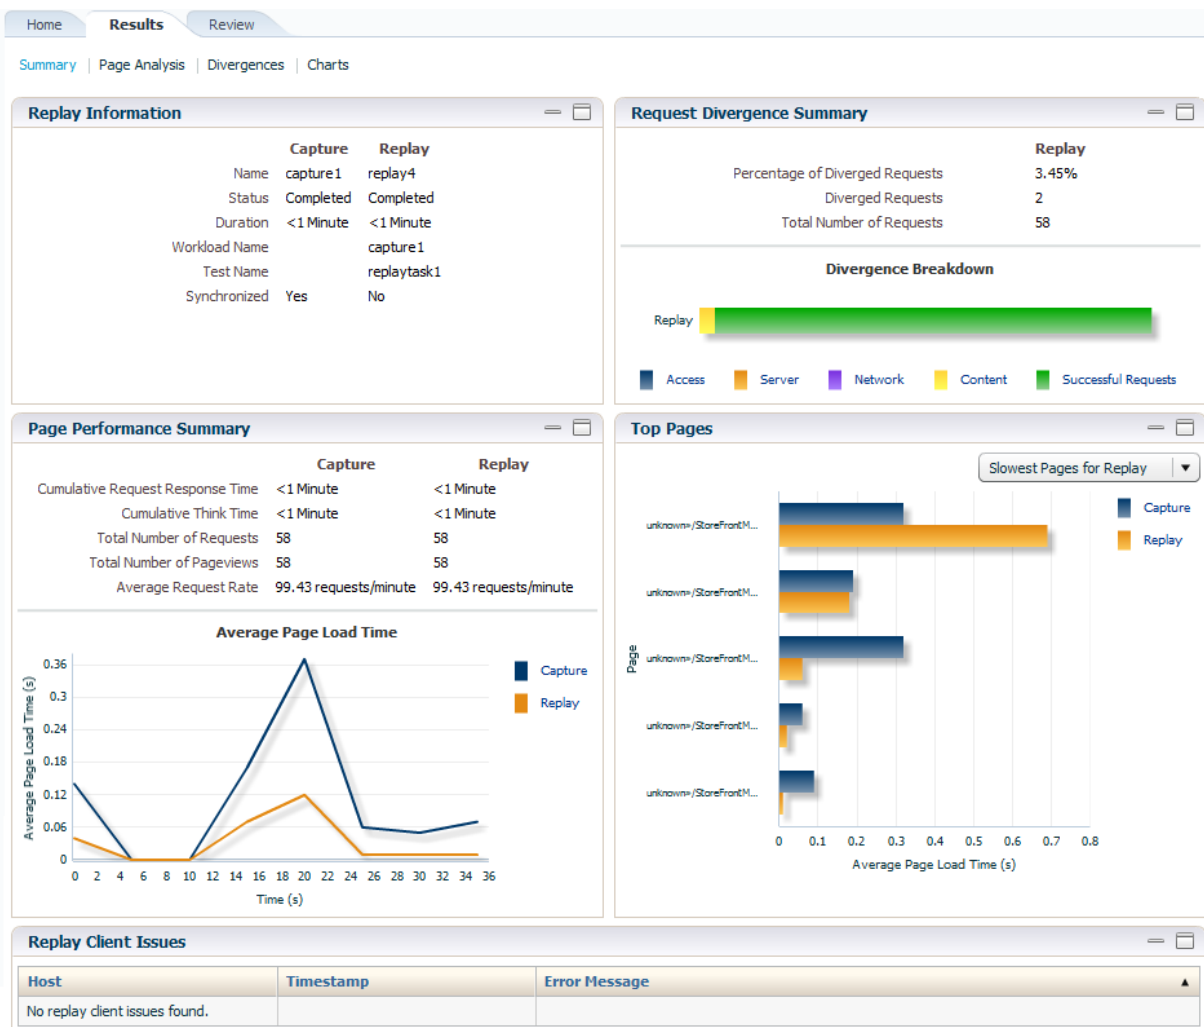
It consists of three parts:

- **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 3-23](#).

Within this section, The **Page Analysis** section allows you to perform 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 3-23 Example Replay Results Summary



3.9 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**.
After the export operation is complete, a link to the session data .zip file appears.
6. Click 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**.

3.10 Troubleshooting Application Replay

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.
- Make sure 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*.
- Make sure 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*.

- Ensure 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
```

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 [Remapping URLs for Application Replays](#). 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 [Substituting Sensitive Data for Application Replays](#).

4

Composite Applications

This section describes how to use composite applications in Enterprise Manager. While individual Java EE applications can be managed using Enterprise Manager, your business needs may require that these applications be managed as a group. This logical application group is called a composite application.

To access composite applications in Enterprise Manager, select **Composite Applications** from the **Targets** menu. From the Composite Applications page you can view existing composite applications or create new ones. For a demonstration of Composite Applications, see [Composite Application Management](#).

This chapter covers the following:

- [Viewing the Composite Application Dashboard](#)
- [Creating a Composite Application](#)
- [Editing a Composite Application](#)
- [Editing a Composite Application Home Page](#)
- [Using Composite Applications](#)

4.1 Viewing the Composite Application Dashboard

Each instance of Composite Application can have a different home page.

You can modify this page by adding and dropping regions using the Personalize Page icon located at the top-right of the page. In turn, you can customize each region by adding content and changing the configurable properties of the region.

The Target Navigation tree, located at the left, shows the direct members of the composite application. These are the members you selected when creating the composite application. The navigation tree also includes the related members that were selected during the creation process.

Each direct target lists its related members. The following lists the possible direct targets and their related members:

- **Application Deployments**
Contains Application Deployments, as well as clustered Application Deployments
- **Databases**
Contains only databases related to the targets in this composite application
- **Fusion Applications**
Contains Fusion Applications, as well as Fusion Application clusters
- **Hosts**
Contains associated host targets
- **Others**

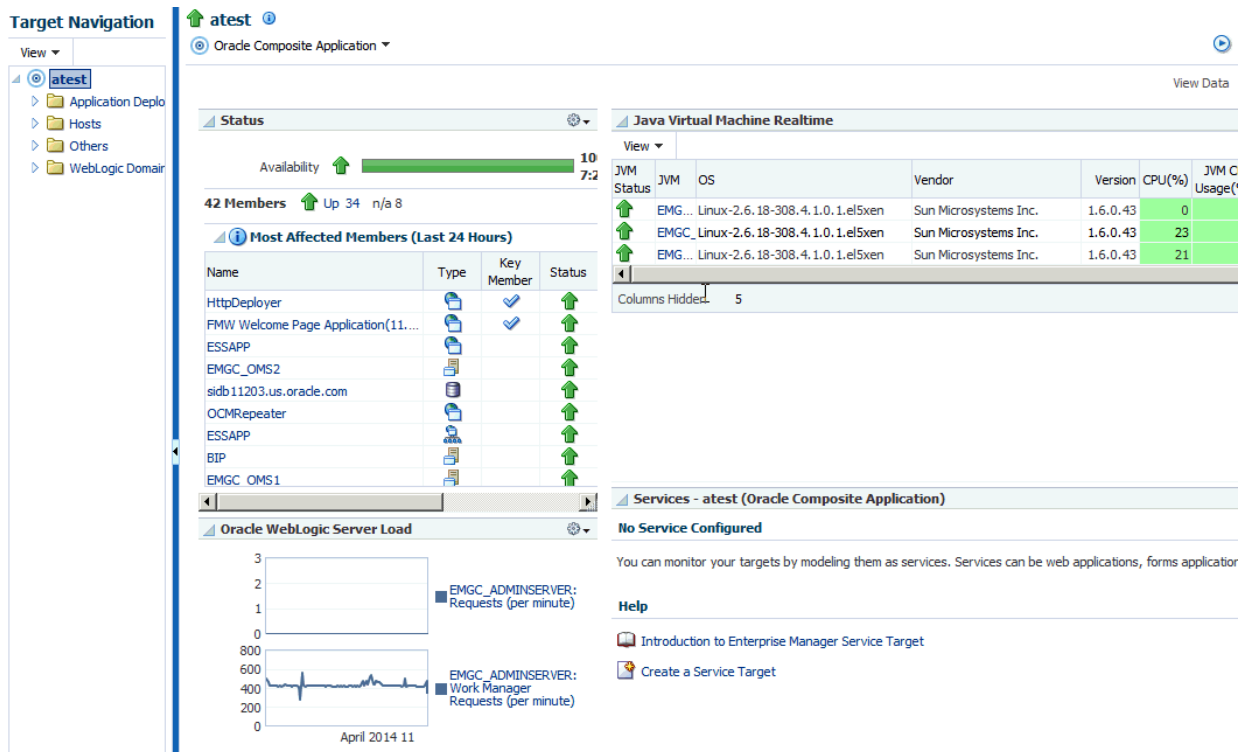
Contains other targets that do not have their own folder, for example, Oracle Homes, Oracle Management Agent, and so on

- Service Oriented Architecture (SOA)
Contains OSBs, SOA Composites, and SOA Infrastructure
- WebLogic Domains
Contains all participating domains

The overall summary page provides additional details for the composite application (see [Figure 4-1](#)):

- Status
Availability of all the members
- Oracle WebLogic Server Load
Includes Requests (per minute) and Work Manager Requests (per minute)
- Request Processing Time and Cache Statements Used (%)
- Overview of Incidents and Problems
Click the number of incidents to view a table listing the reported incidents.
- Java Virtual Machine Realtime
- Services
- SLA Status

Figure 4-1 Composite Application Dashboard



4.2 Creating a Composite Application

To create composite applications, perform the following steps.

1. From the **Targets** menu, select **Composite Applications**.
2. On the Composite Applications page, click **Create**. The first page of the Create Composite Application wizard appears.

Figure 4-2 First Page of the Create Composite Application Wizard

Create Composite Application

Select Applications Create System Identify Signature Services Summary

Create Composite Application: Select Applications

This is the starting step of the wizard to build and define a composite application. Specify the name of the composite application, the timezone and the availability criteria. The Add button must be used to select the

* Composite Application Name

* System TimeZone

Availability Criteria All of The Selected Applications
 Any of The Selected Applications

View

Application Display Name	Application Name	Application Type	Domain Name
No Applications Selected			

3. On the Select Applications page:
 - a. Enter the composite application name.
 - b. Specify a time zone.
 - c. Specify the Availability Criteria. Select either **All of the Selected Applications** or **Any of the Selected Applications**. When you select 'All of the Selected Applications' option, the availability of the composite application is shown as Up when *all* the members of the composite application are up.

When you select 'Any of the Selected Applications' option, the availability of the composite application is shown as Up if *any* of the members of the composite application are up.
 - d. Click **Add**. The **Search and Select: Targets** dialog appears.
 - e. Filter the application list (optional). You can filter the list by Target Type, Target Name, or Host on which the application(s) reside. Specify the desired filter parameters and click **Go**.
 - f. Select the applications to be added. Use the Shift or Control key to select multiple applications.
 - g. Click **Select**. The selected applications now appear in the Select Applications page table.
 - h. Click **Next**.

Note: You can remove applications that you do not want to be part of the composite application. Select the target and click **Remove**.

4. On the Create System page:

Applications previously selected on the **Select Applications** page are displayed in the **Selected Members** table. Based on these applications, related targets are displayed in the **Related Members** table. In the Related Members table, you can

edit the system membership to add additional components or remove existing components.

- a. From the **Related Members** table, click **Add**. The **Select Targets** dialog appears.
- b. Filter the application list (optional). You can filter the list by Target Type, Target Name, or Host on which the applications reside. Specify the desired filter parameters and click **Search**.
- c. Select the applications to be added. Use the Shift or Control key to select multiple targets.
- d. Click **Select**. The **Related Members** table automatically refreshes with the newly added targets.
- e. Click **Next**.

Note: You can remove targets that you do not want to be part of the composite application. Select the target and click **Remove**.

5. On the Identify Signature Services page:

Optionally, you can model the key entry points of the composite application by defining them as Enterprise Manager services, thus identifying them as signature services for the composite application. The **Services List** table lists all Web services exposed by the applications you selected in the first step of the wizard.

To identify signature services:

- a. Select the desired Web services from the **Services List** table.
 - b. Click **Edit**. The **Configure Service** dialog appears.
 - c. Enter the **EM System Name**. Note that you cannot change the name of the system if it is already defined within Enterprise Manager.
 - d. Click **OK** on the Configure Service dialog.
 - e. Click **Next**.
6. On the Summary page:

Review all selections you have made for the composite application. You can go back to any previous step of the wizard to make modifications. When ready, click **Submit** to create the composite application.

4.3 Editing a Composite Application

You can edit a composite application in two ways.

- From the **Targets** menu, select **Composite Applications**. Highlight a composite application and click **Edit**.
- If you are on the Composite Application home page, select **Target Setup** from the Oracle Composite Application menu, then select **Edit Composite Application**.

On the Edit Composite Application page, follow the steps. For more information on how to fill out each step, [Creating a Composite Application](#).

 **Note:**

Ensure to click **Save and Exit** when you have finished making changes. Any changes made will be applied for only *this* target.

4.4 Editing a Composite Application Home Page

You can change the layout of the composite application home page, add content, and edit and remove regions.

You can edit a composite application home page in two ways:

- On the Composite Applications home page, select a composite application and click **Edit Homepage**.
- On the Composite Application home page, click the **Personalize** button located adjacent to the Page Refreshed field.

When you choose to Add Content, you are adding regions to the page. When you edit a region, you change the properties of the region.

Note: Ensure to click **Close** when you have finished making changes. Any changes made will be applied for only *this* target.

4.5 Using Composite Applications

Using the Composite Applications page, you can view the status and statistics of the various components. In addition, using the target navigation tree enables you to access all related targets in the composite application.

Study the following regions to determine if the applications are running at optimal performance and if not, resolve the issues.

- **Status**

Provides the availability of the applications.

- Up (green) arrow means that at least one application is up or all applications are up.
- Down (red) arrow means that at least one application is down.
- n/a (not applicable) means that the target does not have a status.

If an application is down, determine if that is a scheduled down time.

- **Oracle WebLogic Server Load**

Analyze the Requests (per minute) and Work Manager Requests (per minute) metrics. By clicking the metric associated with an application, you can see problem analysis for the metric.

- **Request Processing Time (ms) and Cached Statements Used (%)**

By clicking the metric for the application, analyze the statistics for that metric to provide request processing time and percentage of cached statements used. For more information see the *Enterprise Manager Middleware Plug-in Metric Reference Manual*.

- **Overview of Incidents and Problems**
Click the number associated with either an incident or a problem. For example, if a problem is reported for a particular application, the Incident Manager page summarizes the severity, what target is exhibiting the problem, and so on. The detailed information provides you the opportunity to acknowledge the problem, see other notifications that have been sent regarding the problem, resolve the problem using the guided resolutions, and so on.
- **Java Virtual Machine Realtime**
Provides up-to-date data on JVM.
- **Services**
Provides the overall health of the services, how long the service has been up, and so on.
- **SLA Status**
Provides data regarding service level objectives. This section reports when a service has been breached.

Part II

Monitoring Exalytics Target and Traffic Director

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

The chapters are:

- [Monitoring an Exalytics Target](#)
- [Oracle Traffic Director](#)

5

Monitoring an Exalytics Target

The Oracle Fusion Middleware Management plug-in provides a consolidated view of the Exalytics In-Memory System and Machine within Oracle Enterprise Manager, including a consolidated view of all the hardware components and their physical location with indications of status.

See the *Managing Oracle Exalytics In-Memory Machine with Oracle Enterprise Manager* manual available from the Management page of the Enterprise Manager documentation library:

http://docs.oracle.com/cd/E63000_01/nav/management.htm

In particular, see:

- Features and enhancements provided by the Oracle Fusion Middleware Management plug-in for the Exalytics In-Memory System.
- Instructions for discovering the Exalytics In-Memory System by Oracle Enterprise Manager Cloud Control.
- Instructions for configuring the Exalytics In-Memory System within Oracle Enterprise Manager Cloud Control.

6

Oracle Traffic Director

This section describes how to use Oracle Enterprise Manager to monitor Oracle Traffic Director functions.

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: In a High Availability configuration (Active-Passive or Active-Active), Oracle Traffic Director supports only 2 OTD instances for a given configuration. Hence, Enterprise Manager Cloud Control 12.1.0.3 and higher monitoring capability is limited to at most 2 OTD instances for a given OTD configuration.

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.

- [Before Discovering Traffic Director 11g](#)
- [Adding a Traffic Director to an Exalogic Target](#)
- [About Traffic Director Configuration](#)

- [About Traffic Director Instance](#)
- [About Traffic Director Refresh Flow](#)

For additional information, see *Oracle Traffic Director Administrator's Guide*.

6.1 Before Discovering Traffic Director 11g

Before you discover Traffic Director, you need to configure Traffic Director instances for SNMP monitoring and start the SNMP subagent.

See the Monitoring Using SNMP chapter in the Oracle Traffic Director Administrator's guide. In particular, perform the steps described in the following sections:

1. [Configuring Oracle Traffic Director Instances for SNMP Support](#)
2. [Configuring the SNMP Subagent](#)
3. [Starting and Stopping the SNMP Subagent](#)

Note:

Discovering Oracle Traffic Director 12c is automatically performed as part of Oracle WebLogic Domain discovery. For more information on adding a domain to Cloud Control, see *Discovering and Adding WebLogic Domains in the Oracle Enterprise Manager Cloud Control Administrator's Guide*.

6.2 Adding a Traffic Director to an 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).

6.3 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:

- 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.

After the entities are created, this configuration is deployed on the hosts.

6.3.1 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**.

6.3.2 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 Targets Using Guided Process**. In the Target Types menu, select **Traffic Director**.
3. Click **Add Using Guided Discovery**.

Fill out the pages as described in:

- [Finding Configurations and Instances](#)
- [Discovered Targets](#)
- [Viewing Results](#)

6.3.2.1 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.

6.3.2.2 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.

6.3.2.3 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 more information, see *Oracle Enterprise Manager SNMP Support Reference Guide*.

6.4 About Traffic Director Instance

After 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**.

6.5 About Traffic Director Refresh Flow

After 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.

6.5.1 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.

6.5.2 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.

6.5.3 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.

6.5.4 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 Domains and Oracle GlassFish Domains

The chapters in this part describe how you can monitor Oracle WebLogic Domains and Oracle GlassFish Domains.

The chapters are:

- [Monitoring WebLogic Domains](#)
- [Overview of Oracle GlassFish Server Management](#)

7

Monitoring WebLogic Domains

This section describes how to monitor WebLogic domains.

When using Enterprise Manager version 12.1 and a Secure Socket Layer (SSL) protocol or Transport Layer Security (TLS) protocol to discover and monitor WebLogic servers, the Management 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) that 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.

7.1 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/iops, 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

7.1.1 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.

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.

Example 7-1 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
```

7.1.2 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.

7.1.3 Prerequisites for Domain Discovery When in TLS Mode

If the Oracle Management Service is running in TLS mode only, set the following parameters on the Management Agent of the target. This is the Management Agent which is going to run the discovery of the WebLogic Server Domain.

```
emctl secure agent -protocol TLS  
  
emctl setproperty agent -name  
allowTLSOnly -value true
```

7.2 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.

7.3 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.

7.3.1 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.

7.3.2 Activating Platform MBeans on WebLogic Server 9.x to 10.3.2 versions

The following WebLogic Scripting Tool session shown in [Example 7-2](#) demonstrates how to use, check, and set the PlatformMBeanServerUsed attribute.

User actions are shown in bold.

Example 7-2 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/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/features/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/org.apache.ant
1.6.5/lib/ant-all.jar:/net/mymachine/scratch/shiphomes/wl/wl10/modules/net.sf.antcontrib_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/jrocket_150
15/jre/bin:/net/mymachine/scratch/shiphomes/wl/wl10/jrocket_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/packages/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/adminsw
easvr/oracle/opmn/bin:/usr/X11R6/bin:/home/adminsw/products/valgrind/bin:/home/adminsw/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/features/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/org.apache.ant
1.6.5/lib/ant-all.jar:/net/mymachine/scratch/shiphomes/wl/wl10/modules/net.sf.antcontrib
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/wl/wl10/wlserver
10.0/common/eval/pointbase/lib/pbtools51.jar:/net/mymachine/scratch/shiphomes/wl/wl10/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.

8

Overview of Oracle GlassFish Server Management

This section describes how to use Oracle Enterprise Manager to monitor Oracle GlassFish Domains, Servers and Clusters.

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.

This section 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](#)

8.1 Before Getting Started

Before you start using Oracle GlassFish, ensure that 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 that you have been granted the roles and privileges needed to use Oracle GlassFish.

8.1.1 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

8.1.2 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 from the location where the GlassFish domain's configuration is installed, then copy this server.cer file over to the specified location under the Agent monitoring that domain and servers.

```
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.

```
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(certificcate.cer) file to import
<alias> - alias for the certificate to import

OR

```
keytool -import -alias <alias> -file <file>  
-keystore <AGENT_HOME>/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
<AGENT_HOME> - path to agent_inst directory
<password> - password to the AgentTrust.jks file

8.2 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 that 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, a 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:

- **Diagnostics** - These tools enable you to detect and resolve availability and performance issues on your Java applications. In addition, you can monitor Java applications, configure JVM pools, analyze live threads, as well as view snapshots for threads, heaps, and JFRs. For more information, see [Using JVM Diagnostics](#) .
- **Control** - Allows you to start, restart, and shut down all servers, as well as create and end notification blackouts and blackouts for all servers. Other than notification blackout and 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.

8.2.1 How to Add an Oracle GlassFish Domain To Be Monitored

There are two ways 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.

To watch a video about how to discover an Oracle GlassFish Domain, click [here](#).

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**, and then click the **Add Using Guided Process** button.
4. On the **Add GlassFish Domain: Find Targets** page, provide the required information denoted by an asterisk. Click **Continue**.
5. Reassign the agents.

8.2.2 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/iops, 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 -Dkerborosekey=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.

After 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.

8.2.3 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.

8.2.4 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** 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.
- Notification blackout or blackout the domain when the domain must go down for maintenance. By notification blackout, alert notifications are not sent but the data collection activity continues. By blacking out the domain, the data collection activity is suspended and 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...** or **Create Notification 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.

To view a video about monitoring an Oracle GlassFish Domain, click [here](#).

8.2.5 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**.

8.3 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:

- **Diagnostics** - These tools enable you to detect and resolve availability and performance issues on your Java applications. In addition, you can monitor Java applications, configure JVM pools, analyze live threads, as well as view snapshots for threads, heaps, and JFRs. See [Using JVM Diagnostics](#) for a detailed description of the JVM diagnostics tools.
- **Control** - Allows you to start, restart, and shut down the server, as well as create and end notification blackouts and 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.

8.3.1 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** in the Area field and click **Search**.
3. In the table, expand the Oracle GlassFish Domain node that contains the GlassFish Server of interest.
4. Click the GlassFish Server. 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 and notification 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.

Notification blackouts allow Enterprise Manager users to continue management data collection activity but 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.

8.4 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:

- **Diagnostics** - These tools enable you to detect and resolve availability and performance issues on your Java applications. In addition, you can monitor Java applications, configure JVM pools, analyze live threads, as well as view snapshots for threads, heaps, and JFRs. See [Using JVM Diagnostics](#) for a detailed description of the JVM diagnostics tools.
- **Control** - Allows you to start, restart, and shut down all servers, as well as create and end notification blackouts and blackouts for all servers.

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.

8.4.1 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** in the Area field and click **Search**.
3. 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](#).

8.5 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.

8.6 Creating an Oracle GlassFish Server Configuration Comparison Template

Enterprise Manager provides a monitoring template for GlassFish that 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 **Help** 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:

- [Overview of Oracle SOA Management](#)
- [Discovering and Monitoring Service Bus](#)
- [Discovering and Monitoring the SOA Suite](#)

9

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.

9.1 About Oracle SOA Management Pack Enterprise Edition

The following table highlights the main features of Oracle SOA Management Pack Enterprise Edition.

Table 9-1 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.• 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 list 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 Service Bus) are recorded for troubleshooting and root cause analysis.
Configuration management	Collects configuration information for the BPEL Process Manager server/domains/processes and 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• Service Bus resources from a source domain to a target domain. For detailed information on the provisioning procedures, see <i>Enterprise Manager Administrator's Guide for Software and Server Provisioning and Patching</i> .

Table 9-1 (Cont.) Highlights of Oracle SOA Management Pack Enterprise Edition

Feature	Benefit
Business-IT alignment	Enables you to consolidate the 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.
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.
Historical analysis and reporting	Stores 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.
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.
Error Hospital	Enables you to view an aggregate count of errors that have occurred in all SOA Composites deployed in the SOA Infrastructure. SOA Administrator can use this report to perform bulk recovery on a selected group of similar faults.

10

Discovering and Monitoring Service Bus

This chapter describes how you can discover and monitor Service Bus using Enterprise Manager Cloud Control.

In particular, this document covers the following:

- [New Features in This Release](#)
- [Supported Versions](#)
- [Understanding the Discovery Mechanism](#)
- [Understanding the Discovery Process](#)
- [Discovering Service Bus](#)
- [Enabling Management Packs](#)
- [Monitoring Service Bus in Cloud Control](#)
- [Generating Service Bus Reports Using BI Publisher](#)
- [Troubleshooting Service Bus](#)

10.1 New Features in This Release

The new features that have been introduced in the 13c version of the SOA Suite are:

- Cluster-Level Target pages
- Heat Maps

10.2 Supported Versions

The following are the versions of Service Bus that are supported for monitoring in Enterprise Manager Cloud Control Release 13c.

- Service Bus 12.2.1.x
- Service Bus 12.1.3.x
- Service Bus 11.1.1.7.x
- Service Bus 10.3.2.0.x

10.3 Understanding the Discovery Mechanism

The Service Bus 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 Service Bus 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 Service Bus 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 Domain to which the Oracle WebLogic Managed Server belongs. This will automatically discover the Service Bus that is deployed to it.

For instructions to discover Service Bus, see [Discovering Service Bus](#).

10.4 Understanding the Discovery Process

The following table describes the overall process involved in discovering and monitoring Service Bus in Enterprise Manager Cloud Control. Follow the instructions outlined for each step in this process to successfully discover and monitor your Service Bus.

Table 10-1 Discovery Process

Step	Requirement	Description
1	Service Bus	<p>Install the Service Bus software.</p> <p>Note: Before you launch the Service Bus Deployment Procedure, ensure that Sun JDK has been installed.</p>
2	Enterprise Manager Cloud Control	<p>Install Enterprise Manager 12c.</p> <p>For information about installing the base release of Enterprise Manager Cloud Control, see the Enterprise Manager Cloud Control Basic Installation and Configuration Guide .</p> <p>Oracle recommends that you install the Enterprise Manager Cloud Control components on a host that is different from the host where Service Bus is installed. For example, if Service Bus is installed on host1.xyz.com, then install and configure Oracle Management Service (OMS) and the Management Repository on host2.xyz.com.</p>
3	Oracle Management Agent (Management Agent)	<p>Install Oracle Management Agent 12c on the host where Service Bus is installed.</p> <p>If Service Bus 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 Service Bus is installed. Alternatively, the Management Agent can also be installed on a different host and made to remotely monitor the Service Bus 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 Enterprise Manager Cloud Control Basic Installation and Configuration Guide.</p>
4	Discovery in Enterprise Manager Cloud Control	<p>Service Bus is automatically discovered when the Oracle WebLogic Domain to which it is deployed is discovered and added to Enterprise Manager Cloud Control.</p>

10.5 Discovering Service Bus

The Service Bus 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 Service Bus, 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 Service Bus 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 Domain to which the Oracle WebLogic Managed Server belongs. This will automatically discover the Service Bus that is deployed to it.

This section outlines the instructions for discovering Service Bus for the cases described above. In particular, this section covers the following:

- [Discovering Service Bus Deployed to WLS Not Monitored by Enterprise Manager](#)
- [Discovering Service Bus Deployed to WLS Monitored by Enterprise Manager](#)

10.5.1 Discovering Service Bus Deployed to WLS Not Monitored by Enterprise Manager

To discover Service Bus 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 Service Bus 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 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 Service Bus you just added. Note that, at this point, Service Bus 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 Service Bus to Enterprise Manager Cloud Control, you can monitor its status from the Service Bus 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 Service Bus Console, and turn on the service monitoring and set the level to "pipeline" or "action".

- In the case of clustered Service Bus 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 *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

10.5.2 Discovering Service Bus Deployed to WLS Monitored by Enterprise Manager

To discover Service Bus deployed to Oracle WebLogic Managed Server that is already being monitored in Cloud Control, refresh the membership of the Oracle WebLogic Domain to which the Oracle WebLogic Managed Server belongs. This will automatically discover the Service Bus that is deployed to it.

To refresh the membership of the Oracle WebLogic 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 Domain** target from the list of Middleware targets being monitored.
3. On the Oracle WebLogic Domain Home page, in the General section, click **Refresh Domain**. Enterprise Manager Cloud Control displays the membership page that lists the Service Bus that is currently not being monitored. Click **OK**.

Enterprise Manager Cloud Control refreshes the membership and returns to the Oracle WebLogic Domain Home page.

 **Note:**

On the Oracle WebLogic 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 Service Bus 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 Service Bus has been added.

10.6 Enabling Management Packs

Besides monitoring the status of Service Bus, 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 **Service Bus**.
Enterprise Manager Cloud Control lists all the Service Bus targets being monitored.
3. From the table, for the Service Bus target you are interested in, enable the SOA Management Pack Enterprise Edition and click **Apply**.

10.7 Monitoring Service Bus in Cloud Control

Enterprise Manager Cloud Control helps you monitor the health of Service Bus targets deployed to Oracle WebLogic Managed Servers. When you discover Oracle WebLogic Managed Servers, Cloud Control automatically discovers the Service Bus targets deployed to them and adds them for central monitoring and management.

For each Service Bus target being monitored, Cloud Control provides information about its status, availability, performance, services, alerts, business services, proxy services, pipeline services, and split-join services. It also allows you to view the latest configuration details, save them at a particular time, and compare them with other Service Bus instances. 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 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 Service Bus target depends.

10.7.1 Enabling Monitoring for Service Bus Services

If you are not able to view Service Bus data on Enterprise Manager pages, it may be because monitoring is disabled for Service Bus Services. Before you can view Service

Bus data in Enterprise Manager, check to see if monitoring is enabled for Service Bus Services. You can do that by following these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, select a Service Bus target. The Service Bus home page is displayed.
3. On the Service Bus home page, click **Fusion Middleware Control**.
4. Log in to the Service Bus target.
5. To enable monitoring, on the Global Settings tab, select **Monitoring Enabled** option.

 **Note:**

For a more information about monitoring Service Bus, see Oracle Fusion Middleware Administrator's Guide for Service Bus.

6. Click **Apply**.

10.8 Generating Service Bus Reports Using BI Publisher

You can use Enterprise Manager to print 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 Service Bus Services Report that describes the way Service Bus services have been performing over a period of time. The report provides charts to list the top 5 Service Bus Services and a table with critical metric details for all the services.

The following table describes the Service Bus-related reports you can choose.

Table 10-2 Service Bus Reports

Service Bus Report	Description
Service Bus Service Summary Report	The Service Bus 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 Service Bus Service Summary Report displays a chart with the top 5 Service Bus services based on Average Response Time or Throughput across the selected Service Bus 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 Service Bus Service by Projects or by selecting individual services.

Table 10-2 (Cont.) Service Bus Reports

Service Bus Report	Description
Service Bus Service Operations Summary Report	The Service Bus 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 Service Bus Service by Projects or by selecting individual services.

To print Service Bus 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 Service Bus 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 Service Bus 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.

10.9 Troubleshooting Service Bus

This section describes the errors you might encounter while discovering Service Bus, and the workaround steps you can follow to resolve each of them.

10.9.1 System and Service

The following error occurs if configuration information has not been collected for the selected Application Server.

Table 10-3 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. Click Configuration , and then select Last Collected from the Application Server menu.

10.9.2 SOAP Test

The following error occurs when the Management Agent is upgraded to Enterprise Manager 13c with OMS 10.2.0.5.

Table 10-4 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 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 Service Bus Homepage and from the Related Links section, select Monitoring Configuration.	If the Management Agent has been upgraded to 12c, the following workaround must be applied to support the SOAP test. In the Monitoring Configuration page for the Service Bus 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.example.com:7001/</code>

11

Discovering and Monitoring the SOA Suite

This chapter describes how you can discover and configure the components of the SOA Suite 11g and 12c using Enterprise Manager Cloud Control.

In particular, this document covers the following:

- [List of Supported Versions](#)
- [Monitoring Templates](#)
- [Discovering the SOA Suite](#)
- [Configuring the SOA Suite with Target Verification](#)
- [Metric and Collection Settings](#)
- [Integration Workload Statistics \(IWS\)](#)
- [Setting Up and Using SOA Instance Tracing](#)
- [Viewing Composite Heat Map](#)
- [Monitoring Dehydration Store](#)
- [Publishing a Service to UDDI](#)
- [Generating SOA Reports](#)
- [Exporting a Composite .jar File](#)
- [Provisioning SOA Artifacts and Composites](#)
- [Diagnosing Issues and Incidents](#)
- [Searching Faults in the SOA Infrastructure](#)
- [Recovering Faults in Bulk](#)
- [Generating Error Hospital Reports](#)
- [Recovering BPMN Messages](#)
- [Troubleshooting](#)

11.1 List of Supported Versions

The following versions of the SOA Suite and the SOA Cloud Services (SOACS) are supported in Enterprise Manager Cloud Control:

- 11.1.1.6.0 (PS5)
- 11.1.1.7.0 (PS6)
- 11.1.1.9.0 (PS7)
- 12.1.3 (SOA 12c)
- 12.2.1 (SOA12c)
- 12.2.1.1(SOA12c)

11.2 Monitoring Templates

The following Oracle-certified default templates are being shipped for Enterprise Manager Cloud Control 12c Release 2 and Enterprise Manager Cloud Control 12c Release 3 agents. [Table 11-1](#) describes the available templates, and the agents to which they apply:

Table 11-1 Monitoring Templates

Target Type	Template Name
SOA Infrastructure	Oracle Certified Fusion Middleware Template for SOA Infrastructure
SOA Infrastructure	Oracle Certified Fusion Applications Template for SOA Infrastructure
SOA Composite	Oracle Certified Fusion Middleware Template for SOA Composite
SOA Composite	Oracle Certified Fusion Applications Template for SOA Composite



Note:

The templates created using older versions of OMS (Enterprise Manager Cloud Control 12c Release 2, Enterprise Manager Cloud Control 12c BP1, and so on) should not be used in Enterprise Manager Cloud Control 12c Release 3.

11.3 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 SOACS Instance Using the Hybrid Cloud Agent](#)

11.3.1 Discovering the SOA Suite Using a Local Agent

If you use 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. On the Middleware page, from the **Add** list, select Oracle Fusion Middleware / WebLogic Domain and click **Go**.
3. On the **Find Targets** page, specify the **Administration Server Host**, **Port**, **Username**, **Password**, and **Agent** (local or remote) details.

Figure 11-1 New Domain Discovery

In the Advanced section, select the **JMX Protocol** from the list. By default, the Discover Application Versions appears checked which enables administrators to discover all versions of deployed SOA Composites. However, if you uncheck this option, then you can discover only the latest default version of SOA composites.



Note:

When the SOA Infrastructure application is down, if you uncheck the **Discover Application Versions** check box, then, only composites with single version is discovered. If there are composites with multiple versions, they are ignored.

Figure 11-2 Upgrade Domain Discovery

 **Note:**

- If you have targets which were discovered with the **Discover Application Versions** box checked (which is the default, see [Figure 11-1](#)), but now want to disable this option, perform the following steps:
 - Go to the WebLogic Domain target page.
 - On the Monitoring Configuration page, update the value of Discover application versions to false. (See [Figure 11-2](#).)
 - Perform a domain refresh.

Doing this will discover new composite targets (without any version numbers in their names) that will not contain the metric history from the previous targets.

- Once you are in a state where you have composite targets without version numbers in their names, if you add more SOA composite versions, the version specified as the default version in the SOA Suite will be monitored. Historical metrics will be retained on the same target whenever the default version changes.

Click **Continue**.

4. 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 refresh the WebLogic Domain.

To refresh the domain manually, see My Oracle Support note 1586853.1.

To enable Automatic Refresh of the domain, see My Oracle Support note 1531733.1.

 **Note:**

For a successful monitoring of SOA 12.2.1.1 and above targets, you must have Enterprise Manager Cloud Control 13.2 PG Release and Agent with 13.2 PG Plug-in. Monitoring of SOA 12.2.1.1 and above versions is not compatible using older Enterprise Manager or Agent versions. Metrics related to SOA Composite entities will not be collected due to agent side dependency and hence, SOA Composite entities will not be discovered. Features that require SOA Composite as input parameter will not work.

11.3.2 Discovering the SOA Suite Using a Remote Agent

You can discover the SOA Suite 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 provision SOA Artifacts remotely, or capture the host metrics.

To collect metric data, ensure that you copy the jar files listed in [Table 11-2](#) from the SOA HOME install location to the Agent Home Directory, which is located at: `$AGENT_HOME/plugins/oracle.sysman.emas.agent.plugin_<plugin version>/archives/jlib/extjlib`. 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.

Table 11-2 Metric Data Collection

SOA Target	Files Names
SOA PS5 (11.1.1.6.0) and higher targets	soa-infra-mgmt.jar oracle-soa-client-api.jar jrf-api.jar
SOA 12c targets	soa-infra-mgmt.jar oracle-soa-client-api.jar tracking-api.jar jrf-api.jar To enable Error Hospital and Instance Tracing, you additionally require: wlthint3client.jar
To enable BPMN instance tracing	For SOA 11g targets: oracle.bpm.bpmn-em-tools.jar wsclient_extended.jar For SOA 12c targets: rulesdk2.jar xmlparserv2.jar com.oracle.webservices.fabric-common-api.jar oracle.bpm.bpmn-em-tools.jar

11.3.3 Discovering the SOACS Instance Using the Hybrid Cloud Agent

You can discover a SOA instance on the Oracle Public Cloud in Enterprise Manager Cloud Control using the Hybrid Cloud Agent. To understand the architecture of the hybrid cloud see Enabling Hybrid Cloud Management in the *Oracle Enterprise Manager Cloud Control Administrator's Guide*. The steps required to discover the SOACS instance are also a part of the same chapter. However, a detailed listing of the steps along with the relevant links to the sections is given below.

1. Meet the prerequisites for configuring a Hybrid Cloud Gateway Agent.
See, Prerequisites for Configuring a Hybrid Cloud Gateway Agent.
2. Configure a Management Agent as a Hybrid Cloud Gateway Agent.
See, Configuring an Management Agent as a Hybrid Cloud Gateway Agent.

3. Meet the prerequisites for installing Hybrid Cloud Agents.
See, Prerequisites for Installing Hybrid Cloud Agents.
4. Install a Hybrid Cloud Agent.
See, Installing a Hybrid Cloud Agent.
5. Discover the SOACS instance.

 **Note:**

Once the Hybrid Cloud Gateway Agent is deployed in the on-premise environment and the Hybrid Cloud Agent is deployed in the Oracle Cloud environment, the Oracle Cloud virtual hosts become manageable targets in Enterprise Manager Cloud Control. The procedure to discover and promote the targets running on an Oracle Cloud virtual host is the same as the procedure to discover and promote targets running on any normal host in the on-premise environment. However, for discovering SOA instances running on Oracle Cloud virtual hosts, you should use the public IP address and port **9001** (representing the custom t3 channel that is configured by default on these Admin Servers).

Follow the steps provided in [Discovering the SOA Suite Using a Local Agent](#).

11.4 Configuring the SOA Suite with Target Verification

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:**

You will not be able to click on the torch icon available next to the database system field in Dehydration Store repair pop-up if an association exists between the database system and SOA infrastructure. It is enabled only when the association is missing. When the association is missing, you can select appropriate database system target from target selector popup. Pop-up can be launched by clicking on the torch icon.

11.4.1 Running Functionality-Level Diagnostic Checks

To run diagnostic scans on the functionalities associated with an Enterprise Manager target and to 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 [Repairing Target Monitoring Setup Issues](#).

11.4.2 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.

11.4.3 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 [Running Functionality-Level Diagnostic Checks](#).
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 pre-populated 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 pre-populated by default is an editable field, and appears in multiple rows if it is 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.

11.5 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 11-3 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, select a SOA Infrastructure home. 3. On the SOA Infrastructure home page, from SOA Infrastructure menu, select Monitoring, and click Metric and Collection Settings. 4. 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 11-3 (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. Update the collection interval for the metric SOA Composite Status. <p>Note: This change is applicable to all the SOA composites which are deployed in that SOA Infrastructure</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. Update the collection interval for the metric SOA Infrastructure - Recoverable Faults.

Table 11-3 (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. Update the collection interval for the metric SOA Infrastructure - Recoverable And Rejected Messages.

11.6 Integration Workload Statistics (IWS)

Integration Workload Statistics (IWS) reports provide SOA system-wide reports that can help you analyze utilizations, identify potential bottlenecks and backlogs, and perform top-down analysis of your integration system.

So, for instance, if there are stressed components or endpoints in your SOA system that are slowing down the system, IWS reports can help you narrow down on these. For example, a slow FTP or database adapter reference endpoint can be identified in the reports. Likewise, a BPEL process running slower than usual can also be identified. You can look at internal queue backlogs, like BPEL queues and EDN queues. SOA composite-wise summaries are also available.

IWS reports can include metrics like system resource usage, composite statistics, statistics for internal system queues, statistics for synchronous and asynchronous business processes, and endpoint statistics. The components supported in this release include BPEL Service Engine, EDN, Web Service Binding, File Adapter, JMS Adapter, FTP Adapter, DB Adapter, AQ Adapter, and MQ adapter.

IWS takes periodic snapshots of performance statistics to generate these reports. You can enable or disable IWS data collection. You can also set the collection frequency and the granularity of data collected for your IWS reports. The following table illustrates the data collection levels, or statistics levels, and the data collected for each level.

Data Collection Level/Statistics Level	Data Collected
MINIMUM	System-wide resource usage data.
BASIC	MINIMUM + Service and reference endpoint statistics, BPEL and EDN backup queue statistics, BPEL instance statistics.
NORMAL	BASIC + Data on BPEL activities like Receive, Invoke, Pick, and onMessage.
FINEST	NORMAL + Data on all BPEL activities.

The sections contains the following topics:

- [Statistics in an IWS Report](#)
- [Enabling, and Configuring, or Disabling IWS](#)
- [Generating an IWS Report](#)

11.6.1 Statistics in an IWS Report

An Integration Workload Statistics (IWS) report contains various statistics, depending on the data collection level that you have set. In addition to system-wide resource usage data, the report can include service and reference endpoint statistics, BPEL and EDN backup queue statistics, and BPEL instance statistics. Statistics on BPEL activities may also be included.

The IWS report contains the following broad sections when the data collection level is set to finest:

Parameter	Description
System Resource Usage	Statistics include Java Virtual Machine (JVM) statistics like CPU utilization and memory utilization (for JVM heap and non-heap memory), SOA Data Source statistics that show active connections and connection pool details, and SOA Work Manager statistics that include details on threads.
Composite (Rollup) Statistics	Aggregate composite-wise statistics that indicate flow rate (throughput/transactions per second) and latency (in milliseconds) for the composite endpoints and internal backup queues (EDN and BPEL queue).
Slowest Composite Endpoints	Aggregate composite-wise statistics that indicate the latency (in milliseconds) and flow rate (throughput) for the slowest endpoints.
Backups in Internal Queues	Aggregate statistics for the backups in internal system queues (BPEL queue and EDN queue).
Longest Running Business Processes	Aggregate statistics for top asynchronous and synchronous business (BPEL) process instances based on execution time.
Most Time-Consuming Business Process Activities	Aggregate statistics for top business process activities (BPEL activities like Receive, Invoke, etc) based on execution time.

11.6.2 Enabling, and Configuring, or Disabling IWS

Integration Workload Statistics (IWS) snapshot data is collected at periodic intervals. You can enable snapshot data collection, configure snapshot interval, and the granularity of data collected.

To enable, and configure, or disable Integration Workload Statistics (IWS) follow the steps below:

Note:

The IWS Configuration feature is active only if the Preferred Credential for the Weblogic domain is set to an user having SOA Administrator role. The credentials can be updated using the SOA Infrastructure menu option - **Target Setup Verification**.

1. From the SOA Infrastructure menu, select **Diagnostics** and then click **Generate IWS Report**.
2. Click **IWS Settings**.
3. Set the IWS Collection to **ON**.
OFF disables IWS data collection until it is manually set to **ON** again.
4. Select a **Snapshot Interval** in minutes.
The snapshot interval is the periodic interval at which data snapshots are collected.
5. Select a **Data Collection Level**. The level selected determines the metrics that are collected.
Use the **Minimum** level to collect only system-wide resource usage data. The **Basic** level additionally includes service and reference endpoint statistics, BPEL and EDN backup queue statistics, and BPEL instance statistics. If you choose **Normal**, it includes additional statistics on BPEL activities like Receive, Invoke, Pick, and onMessage. The **Finest** level additionally includes data on all BPEL activities.
6. Click **Save Changes** to save your configuration changes.

11.6.3 Generating an IWS Report

The Integration Workload Statistics (IWS) reports help you identify bottlenecks and backlogs in the system. IWS include metrics like system resource usage, composite statistics, statistics for internal system queues, statistics for synchronous and asynchronous business processes, and endpoint statistics.

To generate an IWS report, follow the steps below:



Note:

You must have already configured IWS data collection and set a snapshot interval before generating an IWS report. See [Enabling, Disabling and Configuring IWS](#) for more information.

1. From the SOA Infrastructure menu, select **Diagnostics**, and then click **Generate IWS Report**.
2. Select the period for which you wish to generate a report. Select timestamps for **Start Date** and **End Date**.
Ensure that the time period does not span server restarts, or periods where you have disabled IWS by setting Data Collection Level to **OFF**.
3. Click the appropriate **Report Format** to generate and download the report.
You can choose between HTML, XML formats, and CSV (comma-separated values).
4. Optionally, choose a partition name if you are using composite partitions and wish to limit your report to a particular partition.

The **Select Composites** field is displayed. This option enables you to select from all composites in the selected partition.

5. Under **Select Composites**, optionally choose one or more composite names to restrict your report to the specified composite applications.
6. Click **Generate Report**.

11.7 Setting Up and Using SOA Instance Tracing

Instance Tracing allows you to trace the message flow across SOA Composites and SOA Infrastructures monitored by Oracle Enterprise Manager Cloud Control. The flow of message can be traced across servers, clusters, and WebLogic domains.

The section contains the following topics:

- [Configuring Instance Tracing \(SOA 11g Targets Only\)](#)
- [Setting Search Criteria for Tracing an Instance](#)
- [Tracing an Instance Within a SOA Infrastructure](#)
- [Tracing Instance Across SOA Infrastructures](#)

11.7.1 Configuring Instance Tracing (SOA 11g Targets Only)

Before enabling Instance Tracing, ensure that the SOA infrastructure is monitored by an Oracle Management Agent.

To enable Instance Tracing for any SOA Infrastructure 11g instances involved in executing composite instances:

1. Set the host and the WebLogic Domain preferred credentials using Target Setup Verification. For details, see [Configuring the SOA Suite with Target Verification](#).
2. To view the state of the listed SOA instances, enable the Capture Composite State flag on the instance tracing page as follows:
 - a. On the SOA Infrastructure home page, from **SOA Infrastructure** menu, select **Fusion Middleware Control**.
 - b. Navigate to the home page of the SOA Infrastructure target.
 - c. From **SOA Infrastructure** menu, select **SOA Administration**, and then click **Common Properties**.
 - d. On the SOA Infrastructure Common Properties page, select the **Capture Composite Instance State** check box.

11.7.2 Setting Search Criteria for Tracing an Instance

Select the appropriate search link based on the version of your SOA target:

- [Instance Tracing for SOA 11g Targets](#)
- [Instance Tracing for SOA 12c Targets](#)

11.7.2.1 Instance Tracing for SOA 11g Targets

To search for faults and messages, enter details as described in the following table, and click **Search**.

Table 11-4 Setting Search Criteria

Field	Description
Instance ID	Specify the ID of the instance that is to be traced. The flow trace is a runtime trail of a message flow identified by an Instance ID. It enables you to track a message flow that crosses instances of different composites.
Start Time From - To	The time period the instances were initiated.
Name	The name of the instance.
Conversation ID	The conversation ID of the instance.
Instance Count	The number of instances that should be retrieved by the Search.
ECID	The ECID enables you to track the message flow across different SOA Composite instances that span across SOA Infrastructure.
Composite Name	The name of the composite. Use this to restrict your search for business flows to a specific composite. Note that wild-card search is supported. For example, (%<part_of_composite_name>%).

Click **Search** after you have specified the required criteria. A list of Instance IDs that meet the criteria are displayed. Click Trace to generate trace data for the specified instance and period.

**Note:**

To trace an instance, credentials must be set for the WebLogic domain of each SOA Infrastructure monitored by Oracle Enterprise Manager Cloud Control and for the host on which the Management Agent monitoring each SOA Infrastructure application is present.

Click the Instance ID link to see the flow trace which includes the list of SOA Infrastructure instances involved in the flow, faults, the domain, and the list of faults.

11.7.2.2 Instance Tracing for SOA 12c Targets

To search for faults and messages, enter details as described in the following table, and click **Search**.

Table 11-5 Setting Search Criteria

Field	Description
Time	<p>Use this filter to restrict your query to a specific time in the past. A time filter is required to search for faults. Ensure that you enter appropriate values in Instance Created From and Instance Created To fields. By default, all the instances created in the last one day are displayed.</p> <p>Additionally, you can add the following filters:</p> <ul style="list-style-type: none"> Instance Updated If you set this value to None, then it means that instance updated filter is not set at all. Fault Occurred
Composite	<p>Use to restrict your search for business flows to a specific composite.</p> <p>If you trace an instance at the composite level, then the Composite value is pre-populated. However, if you trace an instance at SOA infrastructure level, then select any of the following:</p> <ul style="list-style-type: none"> Initiating limits your search to only the business flows that started in the selected composite. Participating allows you to search for all business flows in that composite. Click the torch icon. In the Search and Select Targets wizard, select the target name from the table and click Select. A faults search is performed on the selected composite.
Sensor	Ensure that you select a composite to view the sensors associated with it.
Flow Instance	<p>Flow ID: Use this to search for the flow ID of the business flow instance.</p> <p>Flow Correlation ID: Use this to search for the flow correlation ID of the business flow instance.</p> <p>Initiating ECID: Use this to search for the ECID of the business flow instance.</p> <p>Flow Instance Name: Use this to search for unique system and business identifiers that help you isolate a specific flow instance</p> <p>Composite Instance Name: Use this to specify the name or title of the composite instance name.</p>
State	<p>Select one of the following states:</p> <p>Select Active to search active instances. If you select a blank, then the filtering is ignored.</p> <ul style="list-style-type: none"> All active: Finds all business flows in active states. Running: A business flow is currently running. The flow may include a human task component that is currently awaiting approval. Suspended: A business flow that is typically related to migration of one version of the SOA composite application to another. Recovery: A business flow with a recoverable fault. <p>Select Inactive to search inactive instances. If you select a blank, then the filtering is ignored.</p> <ul style="list-style-type: none"> All inactive: Finds all terminated business flows. Completed: A business flow has completed successfully. There are no faults awaiting recovery. Failed: Finds completed business flows with non-recoverable faults. Aborted: Finds business flows explicitly terminated by the user or for which there was a system error.

Table 11-5 (Cont.) Setting Search Criteria

Field	Description
Fault	<p>Use to limit your search for business flows to only those with faults. If you leave this field blank, the Fault filter is ignored.</p> <p>To search for faults in any state, select All.</p> <p>To search for faults in a particular state, select one of the following:</p> <ul style="list-style-type: none"> • Recovery Required indicates business faults and some specific system faults. For example, 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. • Not Recoverable, indicates rejected messages, most system faults, non-existent references, service invocation failures, and policy faults. • Recovered, indicates flows that contain at least one recovered fault. • System Auto Retries, indicates the faulted flows in which system auto retries occurred.
Fault Type	<p>To search for all types of faults, select All</p> <p>To search for a particular type of fault, select one of the following:</p> <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.
Fault Owner	<p>Use the Name field to enter a fault owner name. Ensure that the name entered is in the following format:</p> <p><code><partition>/<composite name>!<composite version>/<component name></code></p> <p>Use this to further filter your search for faulted business flows to stuck flows awaiting a particular type of recovery action from the administrator. To search for faults belonging to all the owners, select All.</p> <p>To drill down to a particular fault owner, select one of the following:</p> <ul style="list-style-type: none"> • BPEL • BPMN • Mediator • Human Workflow • Decision • Spring • Case Management
Fault Details	<p>You can fine grain your search parameters to drill down to granular result by providing all or some of the following details:</p> <ul style="list-style-type: none"> • Error Message Contains: Use to find only faulted business flows with the same error message text. You can enter any part of the message. This search is case sensitive. • Fault Name: Use to find only faulted business flows with a specific descriptive fault name such as Negative Credit. You must enter the exact name (the entire string). This search is case sensitive. <p>Expand Other to display additional fields for filtering:</p> <ul style="list-style-type: none"> • HTTP Host • JNDI Name

Table 11-5 (Cont.) Setting Search Criteria

Field	Description
Restrict Search Rows	<p>By default, the search results are restricted to 10 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>

11.7.3 Tracing an Instance Within a SOA Infrastructure

To trace an instance within the context of a SOA Infrastructure, follow these steps:

1. In Cloud Control, from the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the target you are interested in. For example, SOA Infrastructure.
3. From the SOA Infrastructure menu, select **Trace Instance**.
4. On the Instance Tracing page, perform instance search. To do so, see [Table 11-5](#).
5. To trace an instance across composites, do the following:
 - For a SOA 12c target, click **Flow Instance ID**.
 - For a SOA 11g target, click **Composite Instance ID**.

You can further drill down to the component audit trail by clicking the component instance available in the trace table.

6. Click **OK**.

11.7.4 Tracing Instance Across SOA Infrastructures

To trace an instance across SOA Domains, follow these steps:

1. In Cloud Control, from the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the target you are interested in. For example, SOA Infrastructure.
3. From the SOA Infrastructure menu, select **Trace Instance**.
4. On the Instance Tracing page, perform instance search. To do so, see [Table 11-5](#).
5. Select an instance, and Click **Trace**.
6. In the Trace Instance dialog box, click **Add** to add the other SOA Infrastructure targets where this SOA instance has been executed.

7. In the Search and Add targets dialog box, select the other SOA Infrastructure targets, and click **Select**.
8. Click **Set** to set the WebLogic Domain Credentials, and Host Credentials if you haven't already set them.
9. Click **OK**. A flow trace job is scheduled to run immediately to collect the instance trace data across domains. On completion, a status is displayed in Trace Job Status column. Click the status link to drill down to the Flow Trace page.
10. Click **OK**.

11.8 Viewing Composite Heat Map

Composite heat map is a graphical representation of a set of metrics depicted as colored boxes.

To view the composite heat map follow the steps below:

1. In the SOA Infrastructure home page, click the **Deployed Composites** tab.
2. Click the **Composite Heat Map** link on the top-right.
A graphical representation of a set of metrics is displayed.
3. From the options section, select the metric size and the metric color. They can be grouped, and you can select the time period for which you want to analyze the data.
4. Enter the composite count and click **Refresh**.

Metric size represents the size of the block. The bigger the block the higher the numeric value of the metric size. Metric color represents the color code of the metric defined by you. In the color scroll bar, the left slider with a number indicates the number below which the metric is displayed as green and similarly the right slider with the number indicates the number above which the metric is displayed in red. The number range (metric range) between the left slider and the right slider is denoted by the range of colors in between green and red.



Note:

Heat map displays the blocks or nodes only for the composites that have values above 0.

In the heat map display, if you want to view the details of a box, click the box. A complete summary of the service is displayed in a dialog box. You can view the details on a chart that helps you analyze how the service is being used over a period of time. The same analysis is available in a table format as well.

11.9 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 Performance 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.

11.9.1 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. Navigate to the Target Verification page to run the functionality check for dehydration store. For more details, see [Configuring the SOA Suite with Target Verification](#).

If you do not see data after these configuration details have been specified, you must wait for the next collection interval.

11.9.2 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 Performance** 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.
 - Database diagnostics by clicking the SQL ID in the Top SOA SQL table.
 - JVM diagnostics by clicking the JVM Diagnostics link for the respective SQL ID.
 - Tablespace utilization for the SOA schema.

- Performance data recorded by the ASH.
- Key SOA tables and tablespace details related to the SOA schema.

11.10 Publishing a Service 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.

11.11 Generating SOA Reports

This section describes the steps to use Enterprise Manager to print SOA reports using BI Publisher Enterprise Reports, or using Information Publisher.

- [Generating SOA Reports Using BI Publisher](#)
- [Generating SOA Reports Using Information Publisher](#)
- [Generating SOA Diagnostic Reports](#)
- [Viewing SOA Diagnostics Jobs](#)

11.11.1 Generating SOA Reports Using BI Publisher

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 11-6 SOA Reports

SOA Report	Description
SOA Infrastructure Performance Report	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. 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.
SOA Composite Detailed Performance Report	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. 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). 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.
Top 5 SOA Composites (From Dehydration Store)	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.

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.

11.11.2 Generating SOA Reports Using Information Publisher

This section describes the procedure to create SOA Reports.

Note:

These reports can be generated only for SOA 11g targets. Information Publisher reports are not supported for SOA 12c targets.

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.

11.11.3 Generating SOA Diagnostic Reports

To collect the SOA diagnostics data from SOA Dehydration Store, and generate report, follow these steps:

1. Ensure that you set the SOA Database Host Credentials and SOA Database user Credentials before scheduling a SOA diagnostics job.
2. From the **Targets** menu, select **Middleware**.

3. On the Middleware page, select a SOA Infrastructure target. The SOA Infrastructure home page is displayed.
4. From the SOA Infrastructure target menu, select **Diagnostics**, then click **Schedule SOA Diagnostics Job**.
5. In the General section, enter a name and description for the job.
6. In the Target section, select a database instance from the table. To add an instance, click **Add**. From the target selector dialog box, select a database instance, and click **Select**.
7. In the Parameters section, enter the following details:
 - **Report Time Period** is the period for which you want to collect the diagnostic data. This is a mandatory field. By default, data for last one week is collected.
 - Optionally, you can select a desired value for System Backlog Report.
 - To get details about open instances, completed instances, or rolled back instances for a product, you must choose the Instance Growth Report.
 - To get a report on invoke process delays, callback delays, callback processing delays, select BPEL Execution Report, and BPEL Performance Report
 - To understand invoke delays, and engine time better, select Mediator reports like Mediator Execution Report, and Mediator Performance Report.
 - To understand pending events in an event queue, select **EDN Report**.
 - To get a summary of all the faults, select Fault Summary Report and Detailed Fault Report.
 - To view the human workflow tasks, select Human Workflow Report.
 - To receive a SOA Diagnostic report through an email, select **Email Notification**.
 - Subject, enter a subject for your email.
 - E-mail To, add contacts to whom this report must be sent.
 - E-mail Cc, add contacts who must be copied on the diagnostics report email.
8. In the Credentials section, provide the SOA Infra Dehydration Store user Credentials, and host credentials for the SOA Dehydration Store.
9. In the Schedule section, you can choose to either run job once or repeatedly. You can additionally schedule to run the job immediately or at a later point.
10. The Access table gives a summary of all the users and roles who have access to this job.
11. Click **Submit**.

11.11.4 Viewing SOA Diagnostics Jobs

To view all the SOA diagnostics jobs, follow these steps:

1. Ensure that you set the SOA Database Host Credentials and SOA Database user Credentials before scheduling a SOA diagnostics job.
2. From the **Targets** menu, select **Middleware**.
3. On the Middleware page, select a SOA Infrastructure target. The SOA Infrastructure home page is displayed.

4. From the SOA Infrastructure target menu, select **Diagnostics**, then click **All SOA Diagnostics Job**.

This page displays all the diagnostics jobs that have run already, and are scheduled to run.

11.12 Exporting a Composite .jar File

Exporting a Composite from a SOA instance provides you the option of deploying it on another SOA instance. The export feature allows administrators to:

- Export a Composite from the on-premise SOA instance
- Export a Composite from the SOACS on the cloud

To export a Composite .jar file perform the following steps:

1. In the SOA Infrastructure home page, click the **Deployed Composites** tab.
2. Select a Composite from the table.
3. Click **Export Composite**.

The Composite Name, Partition and Revision fields are auto-populated. These fields can be edited if required. If you did not select a composite before clicking on Export Composite, edit these fields manually.

4. Select one of the following export options:
 - Export with all post deployment changes - to generate a composite archive file containing the original, design-time definitions of the composite as well as the post-deployment information including the metadata and property update.
 - Export with runtime/metadata changes only - to generate a composite archive file containing the original composite and post-deployment changes such as task definitions, rule changes, and so on.
 - Export with property changes only - to generate a composite archive file containing the original composite and any post-deployment property changes, such as binding properties or policy settings.
 - Export with no post deployment changes - to generate a composite archive file containing only the pre-deployment, design-time definitions of the composite. Any property settings that you may have made on a running composite, or any other runtime metadata, will NOT be included.
5. Optionally, change the default staging location only if required. Click **Advanced** and provide a staging directory location on the server side host.
Utilize this option if you have issues accessing the default staging location.
6. Click **OK**.

11.13 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.

 **Note:**

The features listed above are also supported on the SOA instances running on the cloud.

For more details on the SOA Artifacts Deployment Procedure, see the *Enterprise Manager Administrator's Guide for Software and Server Provisioning and Patching*.

11.14 Diagnosing Issues and Incidents

To access the 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.

11.15 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 \(Simple Recovery\)](#)

11.15.1 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.

11.15.2 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 11-7 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

11.15.3 Prerequisites for Searching, Viewing, and Recovering Faults

Meet the following prerequisites before searching, viewing, and recovering SOA composite application faults:

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.

11.15.4 Searching and Viewing Faults

To search and view SOA composite application faults, follow these steps:

1. Meet the prerequisites. See [Prerequisites for Searching, Viewing, and Recovering Faults](#).
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 [Setting Search Criteria](#).
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 [Filtering Displayed Search Results](#)

For example, to filter and list all faults related to the BPEL engine type, in the **Engine Type** column, type `bpel`.

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.

11.15.4.1 Setting Search Criteria

To search for faults and messages, enter details as described in the following table, and click **Search**.

Table 11-8 Setting Search Criteria

Field	Description
Time	<p>Use this filter to restrict your query to a specific time in the past. A time filter is required to search for faults. Ensure that you provide values in the Fault Time From and Fault Time To fields.</p> <p>For example, enter 1/13/14 5:33:25 AM and 2/13/14 5:33:25 AM in the respective fields to query for all the faults that have occurred in this one month time window.</p>
Composite	<p>Use to restrict your search for business flows to a specific composite.</p> <p>Click the torch icon. In the Search and Select Targets wizard, select the target name from the table and click Select.</p> <p>A faults search is performed on the selected composite.</p>
Flow Instance	<p>Enter the Flow ID to isolate a specific flow instance. For each workflow involving different composites a unique flow ID gets generated. When there is an error in any component in a particular flow, the ID gets listed on the Faults and Rejected Messages tab. This ID is useful in assessing the error trend.</p>
Fault	<p>Use to limit the search for business flows to only those with faults. If you leave this field blank, the Fault filter is ignored.</p> <p>To search for faults of any type, select All or blank.</p> <p>To search for faults in a particular type, select one of the following:</p> <ul style="list-style-type: none"> • Recovery Required indicates business faults and some specific system faults. For example, 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. • Not Recoverable, indicates rejected messages, most system faults, non-existent references, service invocation failures, and policy faults. • Recovered, indicates flows that contain at least one recovered fault. • System Auto Retries, indicates the faulted flows in which system auto retries occurred.
Fault Type	<p>To search for all types of faults, select All.</p> <p>To search for a particular type of fault, select one of the following:</p> <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.

Table 11-8 (Cont.) Setting Search Criteria

Field	Description
Fault Owner	<p>Use this to further filter your search for faulted business flows to stuck flows awaiting a particular type of recovery action from the administrator. To search for faults belonging to all the owners, select All.</p> <p>To drill down to a particular fault owner, select one of the following:</p> <ul style="list-style-type: none"> • BPEL • BPMN • Mediator • Human Workflow • Decision • Spring • Case Management
Fault Details	<p>You can fine grain your search parameters to drill down to granular result by providing all or some of the following details:</p> <ul style="list-style-type: none"> • Error Message Contains: Use to find only faulted business flows with the same error message text. You can enter any part of the message. This search is case sensitive. • Fault Name: Use to find only faulted business flows with a specific descriptive fault name such as Negative Credit. You must enter the exact name (the entire string). This search is case sensitive. <p>Expand Other to display additional fields for filtering:</p> <ul style="list-style-type: none"> • HTTP Host • JNDI Name
Restrict Search Rows	<p>By default, the search results are restricted to 10 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>

11.15.4.2 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 [Searching and Viewing Faults](#).
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.

11.15.4.3 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 [Setting Search Criteria](#)). 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 it takes to retrieve the faults. Entering a higher value than the default in Restrict Search Result (rows) field can lead to longer time to get the faults, and therefor result in a longer load time.

11.15.4.4 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 may be 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, 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.

For more information, see [Searching and Viewing Faults](#).

11.15.4.5 Searching Faults in a Particular Service Engine

There might be faults across various service engines such as 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. Then search.

For more information, see [Setting Search Criteria](#).

11.15.4.6 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. 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`.

For more information, see [Setting Search Criteria](#).

11.15.4.7 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 [Searching and Viewing Faults](#).
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`.

11.15.5 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 [Prerequisites for Searching, Viewing, and Recovering Faults](#).
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 [Setting Search Criteria](#).
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 [Overview of the Recovery Actions for Resolving Faults](#).
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 [Recovering Faults in Bulk](#).
 - 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 [Tracking Bulk Recovery Jobs, and Viewing Their Results and Errors](#).
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.

11.16 Recovering Faults in Bulk

The process of recovering similar type of faults in a single operation is called Bulk Recovery. In case of SOA 11g targets all the *Recoverable* faults can be recovered through bulk recovery option, and similarly for SOA 12c targets, all the *Recovery required* faults can be recovered through bulk recovery.

Note:

For SOA 12c targets, you can supply either the composite details or the fault details to recover faults. It is mandatory that you supply at least one of these parameters, if not, bulk recovery cannot be performed. For SOA 11g targets, you must supply the composite details.

Bulk recovery can be performed when the following criteria are met:

- All faults to be recovered are in the same partition.
- The recovery required count is greater than zero.
- The **Fault Owner** type of the selected row is bpmn, mediator or bpel.

- A state for the fault is specified.

You can perform bulk recovery from Faults and Rejected Messaged tab, or Error Hospital tab available on the SOA Infrastructure home page. This way, the context of the fault is maintained, and is accordingly pre-populated on the Create Bulk Recovery Page. However, if you access it from the Bulk Recovery Jobs page, you will need to enter all the details afresh.

In particular, this section covers the following:

- [Performing Bulk Recovery from the Bulk Recovery Jobs Page](#)
- [Performing Bulk Recovery from Faults and Rejected Messages Tab](#)
- [Performing Bulk Recovery from the Error Hospital Tab](#)
- [Tracking Bulk Recovery Jobs](#)
- [WorkFlow Examples for Bulk Recovery](#)

11.16.1 Performing Bulk Recovery from the Bulk Recovery Jobs Page

To directly recover a large number of faults from the SOA database, follow these steps to perform a bulk recovery:

1. Meet the prerequisites.
2. From the **Targets** menu, select **Middleware**.
3. On the Middleware page, select a SOA Infrastructure target.
4. On the SOA Infrastructure target page, from the **SOA Infrastructure** menu, select **Fault Management**, then select **Bulk Recovery**.
5. On the Bulk Recovery Jobs page, click **Create Job**.
6. On the Create Bulk Recovery Job, in the Composite section, enter the following details:
 - Select **Initiating** or **Participating** composite type from the menu.
 - Click **Add** to add additional composites for which faults must be searched. In the Search and Select dialog box, select all the targets that you want to add to the list, and click **Select**.
 - Click **Remove** to delete a composite.

Note:

You can add only up to 10 composites.

7. In the Time section, enter the suitable values in the following fields to filter out the faults that you want to recover: **Instance Created From**, **Instance Created To**, **Instance Updated**, **Fault Time To**, and **Fault Time From**.
8. In the Fault Details section, set the details of the faults you want to recover. To do so, see [Setting Fault Details for Recovering Faults in Bulk](#).
9. In the Recovery Options section, set the recovery and batch parameters. To do so, see [Setting Recovery and Batch Details for Recovering Faults in Bulk](#).

10. In the Job Parameters section, schedule the bulk recovery job. To do so, see [Scheduling Bulk Recovery Jobs to Run Once or Repeatedly](#).
11. 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.
12. Click **Submit**.
13. Track the status of the bulk recovery job. For more information, see [Tracking Bulk Recovery Jobs, and Viewing Their Results and Errors](#).

 **Note:**

For a SOA 12c target, faults with following recovery states are recovered:

- Admin Recovery
- Mediator Recovery
- BPEL Invoke Message Recovery
- BPEL Callback Message Recovery
- BPEL Activity Message Recovery

However, faults with recovery states EDN Recovery, Rejected Messages Recovery, and Human Workflow Recovery, cannot be recovered.

For a SOA 11g target, all faults with state **Recoverable** are recovered. However, faults with recovery states BPEL messages, rejected messages, and human workflow faults cannot be recovered.

11.16.1.1 Setting Fault Details for Recovering Faults in Bulk

To set the fault details while recovering faults in bulk, follow these steps:

1. In the Fault Details section, from the Engine Type menu, select an engine, so that fault search could be restricted to the selected type.
2. From the Fault Type menu, select the type of fault you want to recover. This could be, **System Faults, Business Faults, or OWSM Faults**.
3. In the Error Message Contains field, enter a keyword you are looking for in the error messages so that only faults with such error messages are recovered.
4. In addition to this, you can refine your fault search by providing details like Fault Name, Fault Code, HTTP Host, and JNDI Name.

11.16.1.2 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. In the Recovery Options section, from the **Recovery Action** list, select a recovery action.

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 60 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 (sec)** field. The minimum delay time is 5 seconds and the maximum delay time is 900 seconds.

Batch Recovery ensures that all the faults that occurred in the specified fault time period are recovered in a phased manner. For example, lets assume:

```
Fault time period: 1 Mar 2013 2.00am to 1 Mar 2013 3.00am
Batch time period: 10mins
Batch Delay: 300secs (i.e 5mins)
```

This means, there are $60\text{mins}/10\text{mins} = 6$ batches in all. The first batch recovers faults between 2.00am to 2.10am. The second batch recovers faults between 2.10am to 2.20am, and so on. After each batch runs, there is a delay of 300secs (5mins), after which the next batch execution begins.

11.16.1.3 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, select an appropriate value from the **Repeat** menu, and set the corresponding frequency.

Note: For a repeating job, ensure that you do not set a custom time period. If you do so, the job cannot track the faults properly, and in-turn recovers the same faults again and again. Instead, you can set a relative time period. For example, select **Last 1Day** from the **Fault Occurred** menu.

- 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.

11.16.2 Performing Bulk Recovery from Faults and Rejected Messages Tab

To recover a large number of faults from the SOA database, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, select a SOA Infrastructure target.
3. On the SOA Infrastructure target page, click **Faults and Rejected Messages** tab.
4. In the Faults and Rejected Messages tab, set the search criteria. To do so, see [Table 11-8](#).
5. Click **Search**.
6. In the table, select one or more faults, and click **Bulk Recover**.
7. In the Navigate to Bulk Recovery wizard, select the details that you want to carry forward from the selected faults in the table to the Create Bulk Recovery page. Select one or more from the following list: **Composite**, **Fault start time**, **Fault end time**, and **Error message** for the fault, and click **OK**.
8. In the Composites section, the composite name and partition field is pre-populated with the values passed from the Faults and Rejected Messages tab. If you want to add additional composites that need recovery, then click **Add**. You can add only up to 10 composites.
9. In the Time section, if you have passed custom values for **Faults Start Time** and **Fault End Time**, then the **Instance Created From** and **Instance Created To** fields are also updated with the same values. You can change these values if required. However, if you select **Last 1 Day**, then all the faults that have occurred across instances in the last one day since the previous bulk recovery job was submitted are displayed.
10. In the Fault Details section, the Error Message field may appear pre-populated if you have passed error message attribute using the Navigate to Bulk Recovery dialog box. If not, you can update this section. For more information, see [Setting Fault Details for Recovering Faults in Bulk](#).
11. In the Recovery Options section, set the recovery and batch parameters. To do so, see [Setting Recovery and Batch Details for Recovering Faults in Bulk](#).
12. In the Job Parameters section, schedule the bulk recovery job. To do so, see [Scheduling Bulk Recovery Jobs to Run Once or Repeatedly](#).
13. 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.
14. Click **Submit**.
15. Track the status of the bulk recovery job. For more information, see [Tracking Bulk Recovery Jobs, and Viewing Their Results and Errors](#).
16. Search for faults again (How?) 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 successful.

 **Note:**

For a SOA 12c target, faults with following recovery states are recovered:

- Admin Recovery
- Mediator Recovery
- BPEL Invoke Message Recovery
- BPEL Callback Message Recovery
- BPEL Activity Message Recovery

However, faults with recovery states EDN Recovery, Rejected Messages Recovery, and Human Workflow Recovery, cannot be recovered.

For a SOA 11g target, all faults with state **Recoverable** are recovered. However, faults with recovery states BPEL messages, rejected messages, and human workflow faults cannot be recovered.

11.16.3 Performing Bulk Recovery from the Error Hospital Tab

To recover a large number of faults from the SOA database, follow these steps:

1. Meet the prerequisites. [Prerequisites for Searching, Viewing, and Recovering Faults](#).
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 **Error Hospital**.
5. In the Error Hospital tab, set the search criteria. To do so, see [Table 11-10](#).
6. Click **Search**.
7. In the table, select one or more faults, and click **Bulk Recover**.
8. The composite section appears pre-populated with **Composite**, **Composite type**, and **Fault Owner** details. You cannot add more composites or edit this section.
9. In the Time section, details like **Instance Created From** and **Instance Created to** are picked up from the Error Hospital page. Additionally, if you had provided **Fault Created From**, **Fault Created To**, **Instance Updated From** and **Instance Updated To** values, then these values will also appear pre-populated on this page. If not, you can enter these values to refine your search.
10. In the Fault Details section, usually, one of the fault parameters appear pre-populated, by default, it is fault name. However, if you have grouped your Error Hospital Report by other categories, then those values are populated accordingly. To refine your search, you may update the other fields in this section. For more information, see [Setting Fault Details for Recovering Faults in Bulk](#).
11. In the Recovery Options section, set the recovery and batch parameters. To do so, see [Setting Recovery and Batch Details for Recovering Faults in Bulk](#).

12. In the Job Parameters section, schedule the bulk recovery job. To do so, see [Scheduling Bulk Recovery Jobs to Run Once or Repeatedly](#).
13. 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.
14. Click **Submit**.
15. Track the status of the bulk recovery job. For more information, see [Tracking Bulk Recovery Jobs, and Viewing Their Results and Errors](#).
16. Search for errors again to verify if the errors you selected for recovery still appear in the search results.
If they do not appear, then the recovery operation for those errors has been successful.



Note:

For a SOA 12c target, faults with following recovery states are recovered:

- Admin Recovery
- Mediator Recovery
- BPEL Invoke Message Recovery
- BPEL Callback Message Recovery
- BPEL Activity Message Recovery

However, faults with recovery states EDN Recovery, Rejected Messages Recovery, and Human Workflow Recovery, cannot be recovered.

For a SOA 11g target, all faults with state **Recoverable** are recovered. However, faults with recovery states BPEL messages, rejected messages, and human workflow faults cannot be recovered.

11.16.4 Tracking Bulk Recovery Jobs

This section describes the following:

- [Tracking Bulk Recovery Jobs, and Viewing Their Results and Errors](#)
- [Creating Bulk Recovery Jobs Using EMCLI and Web Services](#)

11.16.4.1 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	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 . 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.

Section Name	Description
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.</p> <p>To run the recovery job again, reduce the batch time period to a value lower than the value you entered earlier, and submit the bulk recovery job again. For more information, see Recovering Faults in Bulk.</p>

11.16.4.2 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](#)

11.16.4.2.1 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 is a sample Property File:

```
target_list=<soa-infra target name>:oracle_soainfra
variable.CompositeList=<compositel target name>, <composite 2 target name>
variable.BatchDelay=300
variable.BatchSize=10
variable.EnableBatching=1
variable.EngineType=BPEL
variable.ErrorMessage=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=WLCredsNormal -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=HostCredsNormal -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.example.com>:<host>=NAMED:xxxx
cred.SOADomainCreds.<target_name>:<target_type>=NAMED:xxxx
```

11.16.4.2.2 Viewing the Submitted Jobs and Outputs Using EMCLI

The following table describes certain other operations that can be performed using EMCLI commands.

Table 11-9 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>

11.16.4.2.3 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.

11.16.5 WorkFlow Examples for Bulk Recovery

This section covers the following examples:

- [Running Bulk Recovery Job Every Night](#)
- [One Time Job with Specific Time Interval to Recover Faults](#)

11.16.5.1 Running Bulk Recovery Job Every Night

To schedule a bulk recovery job that runs at 12.00am every night, to recovers faults that have occurred through the day:

1. In the composites section, add the desired composites.
2. In the Time section, enter the following values:
 - a. For a SOA 12c target, from Instance Created menu, select **Custom**, and provide the custom values. To recover instances created during the day alone, select **Last 1 Day**.
 - b. From the Fault Occurred menu, select **Last One Day**.
 - c. Click **Estimate Faults** to view the number of faults that will be recovered.
3. In the Fault Details section, enter appropriate values.
4. In the Recovery Option section, enter the following values:
 - a. Select **Batch by Fault Time**.
 - b. In Batch Time Period, enter **10 mins**. This would mean that, every batch would recover faults in 10mins time window. Since you have already selected Last One day (Fault Time From value), there will be $24 \times 60 / 10 = 124$ batches in all.
 - c. In Delay Between Batches, enter **200 secs**. This will be the delay between each batch. The main intention behind a delay is to allow the SOA System time to stabilize after each recovery.
5. In the Job Parameters section, enter the following values:
 - a. Select **Immediately** to start the job as soon as it is submitted.
 - b. From repeat menu, select **Every N Days**.
 - c. Enter Frequency as **1 day**.
6. Click **Submit**.

11.16.5.2 One Time Job with Specific Time Interval to Recover Faults

To schedule a bulk recovery job that runs one time, and recover faults in a specific time interval, follow these steps:

1. In the composites section, add the desired composites.
2. In the Time section, enter the following values:
 - a. For a SOA 12c target, from Instance Created menu, select **Custom**, and provide the custom values. To recover instances created during the day alone, select **Last 1 Day**.
 - b. From the Fault Occurred menu, select **Custom**. Enter **3:00 am** in Fault Time From field, and **4:00 am** in Fault Time To fields.
 - c. Click **Estimate Faults** to view the number of faults that will be recovered.
3. In the Fault Details section, enter appropriate values.
4. In the Recovery Option section, enter the following values:
 - a. Select **Batch by Fault Time**.

1. Error Hospital Report acts as a quick view of fault count for administrators to determine the error trends.
2. A consolidated report with all an aggregate error count is available on a single page.
3. You can also perform bulk recovery on a selected group of similar faults in a single operation.
4. Autoretries feature allows system to continuously retry a recoverable fault. When a fault is in recovery required state and an autoretry is setup, then a automated system call is generated at a certain interval to try and recover the error. This feature greatly benefits the Administrator as they have lesser faults to manually track.
5. Trace Instance option allows you to navigate to the Instance Tracing page based on the fault group selected.

To set the search criteria for Error Hospital, enter details as described in the following table, and click **Search**.

Table 11-10 Setting Search Criteria for Error Hospital

Field	Description
Time	<p>Use this filter to restrict your query to a specific time in the past. A time filter is required to search for faults. Ensure that you enter appropriate values in Instance Created From and Instance Created To fields. By default, all the instances created in the last one day is displayed.</p> <p>Additionally, you can add the following filters:</p> <ul style="list-style-type: none"> • Instance Updated • Fault Occurred
Composite	<p>Use to restrict your search for business flows to a specific composite.</p> <p>You can select the following option:</p> <ul style="list-style-type: none"> • Initiating limits your search to only the business flows that started in the selected composite. • Participating allows you to search for all business flows in that composite. <p>Click the torch icon. In the Search and Select Targets wizard, select the target name from the table and click Select. A faults search is performed on the selected composite.</p>
State	<p>Select one of the following states:</p> <p>Select Active to search active instances. If you select a blank, then the filtering is ignored.</p> <ul style="list-style-type: none"> • All active: Finds all business flows in active states. • Running: A business flow is currently running. The flow may include a human task component that is currently awaiting approval. • Suspended: A business flow that is typically related to migration of one version of the SOA composite application to another. • Recovery: A business flow with a recoverable fault. <p>Select Inactive to search inactive instances. If you select a blank, then the filtering is ignored.</p> <ul style="list-style-type: none"> • All inactive: Finds all terminated business flows. • Completed: A business flow has completed successfully. There are no faults awaiting recovery. • Failed: Finds completed business flows with non-recoverable faults. • Aborted: Finds business flows explicitly terminated by the user or for which there was a system error.

Table 11-10 (Cont.) Setting Search Criteria for Error Hospital

Field	Description
Fault	<p>Use to limit the search for business flows to only those with faults. If you leave this field blank, the Fault filter is ignored.</p> <p>To search for faults of any type, select All or blank.</p> <p>To search for faults in a particular type, select one of the following:</p> <ul style="list-style-type: none"> • Recovery Required indicates business faults and some specific system faults. For example, 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. • Not Recoverable, indicates rejected messages, most system faults, non-existent references, service invocation failures, and policy faults. • Recovered, indicates flows that contain at least one recovered fault. • System Auto Retries, indicates the faulted flows in which system auto retries occurred.
Fault Type	<p>To search for all types of faults, select All</p> <p>To search for a particular type of fault, select one of the following:</p> <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.
Fault Owner	<p>Use the Name field to enter a fault owner name. Ensure that the name entered is in the following format:</p> <pre><partition>/<composite name>!<composite version>/<component name></pre> <p>Use this to further filter your search for faulted business flows to stuck flows awaiting a particular type of recovery action from the administrator. To search for faults belonging to all the owners, select All.</p> <p>To drill down to a particular fault owner, select one of the following:</p> <ul style="list-style-type: none"> • BPEL • BPMN • Mediator • Human Workflow • Decision • Spring • Case Management
Fault Details	<p>You can fine grain your search parameters to drill down to granular result by providing all or some of the following details:</p> <ul style="list-style-type: none"> • Error Message Contains: Use to find only faulted business flows with the same error message text. You can enter any part of the message. This search is case sensitive. • Fault Name: Use to find only faulted business flows with a specific descriptive fault name such as Negative Credit. You must enter the exact name (the entire string). This search is case sensitive. <p>Expand Other to display additional fields for filtering:</p> <ul style="list-style-type: none"> • HTTP Host • JNDI Name

Table 11-10 (Cont.) Setting Search Criteria for Error Hospital

Field	Description
Restrict Search Rows	<p>By default, the search results are restricted to 10 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>

In particular, you can perform the following tasks from this page:

- [Generating an Error Hospital Report](#)
- [Customizing the Error Hospital Report](#)

11.17.1 Generating an Error Hospital Report

To generate and view an error counts that have occurred across all SOA Composites using the search fields, follow these steps:

1. Meet the prerequisites. See [Prerequisites for Searching, Viewing, and Recovering Faults](#).
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 **Error Hospital**.
5. In the Error Hospital tab, set the search criteria. For more information, see [Table 11-10](#).
6. Click **Search**.
7. View the results:
 - To view the aggregate count of errors for each fault, see the **Total Faults** column in 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 text-box placed above the column header. For more information, see [Limiting Faults Searched and Retrieved from the SOA Infrastructure](#).
 - To group the faults by different categories, select the relevant category. For more information, see [Customizing the Error Hospital Report](#).
 - To recover the faults in bulk, click **Bulk Recover**. For more information, see [Performing Bulk Recovery from the Error Hospital Tab](#).

11.17.2 Customizing the Error Hospital Report

After generating the report, if you want to group the results by some other category, then follow these steps:

1. Create an error report. See [Generating an Error Hospital Report](#).
2. In the Error Hospital page, select the fault attribute by which data is aggregated. To do so, from the **Group By** menu select one of the following fault attributes. By default, the faults are aggregated by the **Fault Name**. However, you can select any of the following options:
 - **Fault Code:** Aggregates the fault code.
 - **Fault Name:** Aggregates the fault name. This aggregation option is selected by default.
 - **Fault Type:** Aggregates the fault type:
 - **System:** Network errors or other types of errors such as a database server or a web service being unreachable.
 - **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).
 - **OWSM:** Errors on Oracle Web Service Manager (OWSM) policies attached to SOA composite applications, service components, or binding components. Policies apply security to the delivery of messages.
 - **JNDI Name:** Aggregates the JNDI name (for example, `eis/FileAdapter`).
 - **Composite:** Aggregate faults by the SOA composite application name.
 - **Fault Owner:** Aggregate faults by the name of the service component, service binding component, or reference binding component that handled the fault. In some cases, this can be both the fault owner and fault location.
 - **Fault Owner Type:** Aggregates the type of component, service, or reference that handled the fault (for example, if a BPEL process service component owns the fault, BPEL is displayed).
 - **Partition:** Aggregates the partition of the SOA composite application in which the fault occurred.
 - **HTTP Host:** Aggregates the HTTP host on which the fault occurred.

11.18 Recovering 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**.

11.19 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 Errors](#)
- [Recent Faults](#)
- [Fault Management](#)
- [Information Publisher Reports](#)
- [BI Publisher Reports](#)
- [Systems and Services](#)
- [BPEL Recovery](#)
- [SOA License Issue](#)
- [Dehydration Store Issue](#)

11.19.1 Discovery

The following error occurs when the SOA instances are being discovered.

Table 11-11 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.

11.19.2 Monitoring

The following error occurs when the collection frequency causes a delay in the collection of configuration data.

Table 11-12 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.
Metric Collection	If target setup verification is not run, you may see metric collections errors for a few metrics. Collection is suspended for these metrics till you unsuspend them. To batch unsuspend metrics with collection errors for SOA infrastructure targets, use the <code>unsuspend_soametrics.pl</code> script available in the agent scripts directory under <code>%emd_root%</code> .

11.19.3 Instance Tracing Errors

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 11-13 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.

11.19.4 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 11-14 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.

11.19.5 Fault Management

This section contains the troubleshooting information for fault management:

- [Bulk Recovery](#)
- [Fault Search and Recovery](#)
- [Fault Management and Instance Tracing Errors](#)

11.19.5.1 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 11-15 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.example.com:8001 javax.naming.CommunicationException [Root exception is java.net.ConnectException: t3://slc03dms.us.example.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>

11.19.5.2 Fault Search and Recovery

The following error occurs when you are unable to connect to the SOA Infrastructure target:

Table 11-16 Error Message

Error Message	Workaround Steps
<pre>Error connecting to SOA Infra t3://slc03dms.us.example.com:8001.</pre>	<p>Ensure that the SOA Infrastructure is up and running.</p>

11.19.5.3 Fault Management and Instance Tracing Errors

The following errors occur when the SOA database is not functional:

Table 11-17 Error Message

Error Message	Workaround Steps
<pre>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.</pre>	<p>Ensure the SOA Database is up and running.</p>
<pre>t3://slc03dms.us.example.com:80 01 javax.naming.CommunicationExce ption [Root exception is java.net.ConnectException: t3://slc03dms.us.example.com:80 01/soa-infra: Destination unreachable; nested exception is: java.net.ConnectException: Connection refused; No available router to destination]</pre>	<p>Ensure the SOA Database is up and running.</p>
<pre>Error occurred when getting faults oracle.sysman.emSDK.agent.comm .exception.ConnectException: Unable to connect to the agent at https://slc03dms.us.example.com :3872/emd/main/ [Connection refused]</pre>	<p>Ensure the SOA Database is up and running.</p>

11.19.6 Information Publisher Reports

This section lists report related errors.

Table 11-18 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.

11.19.7 BI Publisher Reports

This section lists BI Publisher report related errors.

Table 11-19 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 12. 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>

11.19.8 Systems and Services

The following error occurs when you try to refresh a service that has not been created.

Table 11-20 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.

11.19.9 BPEL Recovery

The following error occurs when invalid credentials are provided.

Table 11-21 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 be available in the preferred credential store. Set the valid credentials and try again.

11.19.10 SOA License Issue

The following error occurs if the SOA Management Pack EE has not been enabled.

Table 11-22 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.

11.19.11 Dehydration Store Issue

Data is not displayed on the Dehydration Store page.

Table 11-23 Error Message

Error Message	Workaround Steps
Data is not displayed in the Dehydration Store page.	<p data-bbox="769 338 1453 422">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 data-bbox="769 443 1453 653" style="list-style-type: none"><li data-bbox="769 443 1453 527">1. Compare the Database Host and SID value of the database target with the value collected for the WebLogic Server JDBC Datasource configuration.<li data-bbox="769 548 1453 653">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 13c.

This part contains the following chapter:

- [Discovering and Monitoring Oracle Business Intelligence Instance and Oracle Essbase](#)

12

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 13c 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 13c. 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](#)
- [Scaling Out Oracle Business Intelligence Domains](#)
- [Creating Oracle Business Intelligence Instance Provisioning Profiles](#)
- [Cloning Oracle Business Intelligence Instances](#)

12.1 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 13c Release 1 or higher. It contains the following:

- [Oracle Business Intelligence Instance](#)
- [Oracle Essbase](#)

12.1.1 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 12-1](#) describes these components.

Table 12-1 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.

12.1.2 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.

12.2 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, see [Oracle Fusion Middleware Installation Guide for Oracle Business Intelligence](#).

2. Install the Enterprise Manager Cloud Control 13c Release 1 or higher. If you are using an earlier version of Enterprise Manager Cloud Control, upgrade it to 13c Release 1 or higher.

For information on how to install the Enterprise Manager Cloud Control 13c system, see [Oracle Enterprise Manager Cloud Control Basic Installation Guide](#).

For information on how to upgrade to Enterprise Manager Cloud Control 13c Release 1 or higher, see [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 12.1.0.5.0 or higher. If the host has a Management Agent of version earlier than 12.1.0.2.0 installed, upgrade the Management Agent to 12.1.0.5.0 or higher.

For information on how to install a Management Agent, see [Oracle Enterprise Manager Cloud Control Basic Installation Guide](#).

For information on how to upgrade a Management Agent, see [Oracle Enterprise Manager Cloud Control Upgrade Guide](#).

4. 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, see [Using Plug-Ins in 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 [Discovering Targets of an Undiscovered WebLogic Domain](#).

For information on how to discover BI Instance and Essbase targets part of a previously discovered WebLogic domain, see [Discovering New or Modified Targets of a Discovered WebLogic Domain](#).

6. Monitor the BI Instance and Essbase targets.

For information on how to monitor BI Instance and Essbase targets, see [Monitoring Oracle Business Intelligence Instance and Essbase Targets](#).

12.3 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](#)

12.3.1 Discovering Targets of an Undiscovered WebLogic Domain

To discover BI Instance and Essbase targets that are 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. For additional information about Fusion Middleware discovery, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

Enabling Automatic Discovery of Targets

Using 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.

Auto-discovery only discovers the domains which are not visible under the Middleware tab, for monitoring. To make a domain visible, it should first be promoted to Enterprise Manager. You can promote a domain by using the Autodiscovery page.

Discovering Targets Manually

Using 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.

12.3.2 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. For additional information about Fusion Middleware discovery, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

Enabling Automatic Discovery of Targets

Using 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.

Discovering Targets Manually

Using 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.

12.4 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 and 12c targets.

12.4.1 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](#)

12.4.1.1 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 files, and configuration files for viewing or editing purposes.

[Table 12-2](#) describes the elements of the **Summary** section.

Table 12-2 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.
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.

12.4.1.2 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 12-3](#) describes the elements of the **General** section.

Table 12-3 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.
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.

12.4.1.3 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.

12.4.1.4 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.

12.4.1.5 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.

12.4.1.6 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.

12.4.1.7 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 12-4](#) describes the elements of the **Monitoring and Diagnostics** section.

Table 12-4 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.
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.

12.4.1.8 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 12-5 describes the elements of the Alert History page.

Table 12-5 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.

12.4.1.9 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*.

12.4.1.10 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 12-6](#) describes the elements of the Log Messages page.

Table 12-6 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.
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.

Table 12-6 (Cont.) Target Log Messages

Element	Description
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.

12.4.1.11 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.

12.4.1.12 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.

12.4.1.13 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*.

12.4.2 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](#)

12.4.2.1 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**.

You can also perform a SQL drill down of the dashboard by selecting **Top Dashboards by Resource Usage in last 7 days**, and then clicking the SQL details icon.

In order to do a SQL drill down, you need to discover the database from where the dashboards are fetching the data. To do this, from the **Business Intelligence Instance** menu, select **Target setup**, and then select **Monitoring Credentials**. Select the database which has been discovered and add it using the target selector icon.

[Table 12-7](#) describes the elements of the Dashboard Reports page.

Table 12-7 Oracle Business Intelligence Dashboard Reports

Element	Description
User	User who accessed the dashboard.
Total Sessions	Total number of user sessions which accessed the dashboard.
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.

Table 12-7 (Cont.) Oracle Business Intelligence Dashboard Reports

Element	Description
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.

12.4.2.2 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 12-8](#) describes the elements of the Scheduler Reports page.

Table 12-8 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.

12.4.2.3 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 [Viewing Target Metrics](#).

[Table 12-9](#) describes the elements of the **Metrics** section.

Table 12-9 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.
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.

12.4.2.4 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 12-10 describes the elements of this section.

Table 12-10 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.

12.4.2.5 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 12-11 describes the elements of this section.

Table 12-11 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.

Table 12-11 (Cont.) Oracle Essbase Application Data Storage Details

Element	Description
Data Cache Size (KB)	Size, in KB, of the buffer in memory that holds uncompressed data blocks.

12.5 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 [Monitoring Oracle Business Intelligence Instance and Essbase Targets](#).

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](#)

12.5.1 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](#)

12.5.1.1 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.

12.5.1.2 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.

12.5.1.3 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.

12.5.1.4 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 12-12](#) describes the elements of the Monitoring Configuration page.

Table 12-12 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.

12.5.2 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](#)

12.5.2.1 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](#).

12.5.2.2 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 12-13](#) describes the elements of the Monitoring Credentials page.

Table 12-13 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.

12.6 Scaling Out Oracle Business Intelligence Domains

This section describes the procedure to scale out a Business Intelligence Domain.

Prerequisites for Scaling Out a Business Intelligence Domain

Before you run the deployment procedure, meet the following prerequisites:

- The Business Intelligence Domain that is scaled up must be an existing domain that has been discovered with Cloud Control.
- If you are scaling out a domain, ensure that the destination machine contains sufficient space. If the size of the Middleware Home on the source machine is 3

GB, you need approximately 3 GB in the working directory on the source and destination machines. Additionally, the destination machine should also have 3 GB of space for the Middleware Home. The working directory is cleaned up after deployment procedure has been successfully completed.

- The Middleware Home directory you specify on the destination machine must be a new directory or must be empty.
- The Management Agent must be installed on the source (where the Administration Server is running) and the destination machines. The Administration Server for the domain must be up and running.
- The Administration Server and Managed Server (being cloned) must be up and running before you run the deployment procedure.
- The Managed Server and Node Manager ports must be free.
- The Listener Address should be updated in the admin console for Adminserver.
- The oralnst.loc file should be created in the Destination Host before starting the Scale Out deployment procedure.
- Both the Source and Destination Hosts should have the same username and target user.
- The user must have the following permissions:
 - Read permissions on:
 - * Administration Server Host Middleware Directory
 - * Administration Server Host Domain Directory
 - Write permissions on:
 - * Administration Server Host Working Directory
 - * Working Directory of all the destination Managed Server hosts
 - * Middleware Directory of all the destination Managed Server hosts
 - * Domain Directory of all the destination Managed Server hosts

Procedure for Scaling Out a Business Intelligence Domain

To scale out an existing Business Intelligence (BI) domain in online mode, follow these steps:

1. You can access the BI instance scale out deployment procedure page, by using any of the following two ways:
 - On the BI instance home page, from the **Business Intelligence Instance** menu, select **Provisioning**, and then select **Scale Out Business Intelligence**.
 - From the **Enterprise** menu, select **Provisioning and Patching**, and then select **Middleware Provisioning**. On the Middleware Provisioning page, in the Deployment Procedures section, select **Scale Out Business Intelligence**, and then click **Launch**.

The Oracle Business Intelligence Scale Out wizard opens.

2. On the Business Intelligence Details page, select a Business Intelligence instance. In the Working Directory field, specify the directory on the Administration Server machine on which the domain scale out related files are temporarily stored. If this

directory is not present, it will be created. When the scale out operation has been completed, the directory and its contents will be deleted.

 **Note:**

The Working Directory must not be created under the Middleware Home or the BI Domain Home directory.

3. In the Source Information section, the details of the source domain including the Middleware Home, BI Domain Home, and the BI Domain Location are displayed. Verify the details.

4. In the Oracle Instance Details section, to select the host where you want to scale out the Oracle Business Intelligence domain, click **Add Hosts**.

In the Target Selector dialog box, select a host, and then click **Select**. You can select more hosts by clicking **Add Host** again.

Click **Next**.

You can add the name of the instance and select the number of BI servers, Presentation servers, and Java hosts. You can also select the port range.

5. In the Credentials page, specify the Host credentials and the Weblogic Administrator credentials.
 - **Preferred Credentials:** The preferred credentials stored in the Management Repository are used. The Preferred Credentials option will be available only if it has already been defined in Cloud Control.
 - **Named Credentials:** The credentials stored in the Management Repository is used. To use the Named Credentials stored in the Management Repository, you must have already registered each of the preferred credential types with a unique name. Select the desired Named credential from the list available in Credential Name. If you have created all necessary named credentials, you can use them now. If they have not been created, you can create them using this deployment procedure.

Click **Next**.

6. In the Schedule page, specify a Deployment Instance name. If you want to run the procedure immediately, then retain the default selection, that is, **Immediately**. If you want to run the procedure later, then select **Later** and provide time zone, start date, and start time details. You can set the notification preferences according to deployment procedure status. If you want to run only prerequisites, you can select **Pause the procedure after the necessary prerequisite checks have been completed** to pause the procedure execution after all prerequisite checks are performed.

Click **Next**.

7. On the Review page, review the details you have provided for the Deployment Procedure. If you are satisfied with the details, then click **Submit** to run the Deployment Procedure according to the schedule set. If you want to modify the details, click the **Edit** link in the section to be modified or click **Back** repeatedly to reach the page where you want to make the changes. You can also directly go to the required page by clicking the train button corresponding to the desired page, on the top of the page.

8. Navigate to the Procedure Activity page. From the **Enterprise** menu, select **Provisioning and Patching**, and then select **Procedure Activity**.

In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the **Status** link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.

12.7 Creating Oracle Business Intelligence Instance Provisioning Profiles

This section describes the procedure to create a Business Intelligence (BI) instance snapshot.

To create a BI instance snapshot, follow these steps:

Note:

Before you create a snapshot, you need to configure the NFS software library or the Agent SW library.

Provisioning of BI 12c targets are not supported.

1. From the **Enterprise** menu, select **Provisioning and Patching**, and then select **Middleware Provisioning**.
2. On the Middleware Provisioning page, in the Deployment Procedures section, select **Create Business Intelligence Profile**, and then, click Launch.
3. On the Create Snapshot: Business Intelligence Details page, click the search icon to search and select the Oracle Business Intelligence instance for which you want to create a snapshot.
4. On the Create Snapshot: Credentials page, specify the Host credentials and the Weblogic Administrator credentials.
 - **Preferred Credentials:** The preferred credentials stored in the Management Repository are used. The Preferred Credentials option will be available only if it has already been defined in Cloud Control.
 - **Named Credentials:** The credentials stored in the Management Repository is used. To use the Named Credentials stored in the Management Repository, you must have already registered each of the preferred credential types with a unique name. Select the desired Named credential from the list available in Credential Name. If you have created all necessary named credentials, you can use them now. If they have not been created, you can create them using this deployment procedure.

Click **Next**.

5. In the Business Intelligence Details page, verify or specify the BI instance details. Click **Next**.
6. In the Snapshot Details page, verify or specify the details of the snapshot to be created. Click **Next**.

7. On the Review page, review the details you have provided for the Deployment Procedure. If you are satisfied with the details, then click **Submit** to run the Deployment Procedure according to the schedule set. If you want to modify the details, click the **Edit** link in the section to be modified or click **Back** repeatedly to reach the page where you want to make the changes. You can also directly go to the required page by clicking the train button corresponding to the desired page, on the top of the page.
8. Navigate to the Procedure Activity page. From the **Enterprise** menu, select **Provisioning and Patching**, and then select **Procedure Activity**.

In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the **Status** link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.

12.8 Cloning Oracle Business Intelligence Instances

This section describes the procedure to clone a Business Intelligence Instance from a provisioning profile present in the Software Library.

Prerequisites for Cloning a Business Intelligence Instance

Before running this deployment procedure, the following prerequisites must be met:

- The user must have Write permissions on:
 - The Working Directory on all destination hosts.
 - Middleware Home on all destination hosts.
- The ports on the Administration Server, Managed Server, and Node Manager must be free.
- A Business Intelligence Instance Provisioning Profile must be present in the NFS Software Library or in the Agent Software Library. For details on creating this profile, see [Creating a Business Intelligence Provisioning Profile](#).

Else, the root credentials have to be used in the credentials page.

Procedure for Cloning a Business Intelligence Instance

To clone a Business Intelligence (BI) instance from a profile, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning**.
2. On the Middleware Provisioning page, in the Procedure Deployments section, select **Business Intelligence Cloning**, and then, click **Launch**.
3. On the Clone BI Instance: Snapshot page, click the search icon. In the Select Snapshot dialog box, search for or select a snapshot from the list, and then click **Select**.

Click **Next**.
4. On the Clone BI Instance: Destination Host page, click **Add** to select the destination host where you want the business intelligence instance to be cloned.

 **Note:**

The destination base directory is read-only and automatically gets configured to match the source location.

The Temporary Work Directory gets deleted when the content movement deployment procedure has been run.

5. On the Clone BI Instance: Credentials page, specify the following:

- Normal Host Credentials
- Root Host Credentials
- Weblogic Administration Server Password
- Global Node Manager Password
- Keystore Passphrase

Preferred Credentials: The preferred credentials stored in the Management Repository are used. The Preferred Credentials option will be available only if it has already been defined in Cloud Control.

Named Credentials: The credentials stored in the Management Repository is used. To use the Named Credentials stored in the Management Repository, you must have already registered each of the preferred credential types with a unique name. Select the desired Named credential from the list available in Credential Name. If you have created all necessary named credentials, you can use them now. If they have not been created, you can create them using this deployment procedure.

6. On the Clone BI Instance: Domains page, to configure the BI domain information, do the following:

- In the Destination Hosts section, search for the database hosts that were configured for the destination environment, enter the password that was assigned for each database, and then click **Test Connection** to confirm that the entry is valid.
- In the RPD Configuration section, specify the data source and the password.
- In the Business Intelligence Publisher Configuration section, specify the XMLP Data Source URL.
- In the Email Options section, specify the SMTP server name, the port, the From display name, and the email address of the sender.
- Click **Next**.

7. On the Review page, review the details you have provided for the Deployment Procedure. If you are satisfied with the details, then click **Submit** to run the Deployment Procedure according to the schedule set. If you want to modify the details, click the **Edit** link in the section to be modified or click **Back** repeatedly to reach the page where you want to make the changes. You can also directly go to the required page by clicking the train button corresponding to the desired page, on the top of the page.

8. Navigate to the Procedure Activity page. From the **Enterprise** menu, select **Provisioning and Patching**, and then select **Procedure Activity**.

In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the **Status** link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.

Part VI

Monitoring Application Performance

Monitoring distributed applications requires the use of several products, each of which examines a different aspect of application performance. The chapters in this part explain how you can use these products singly and together to monitor your application. It also provides a summary of the workflow required to install, configure, and work with these products. It includes the following chapters:

- [Monitoring Performance](#) introduces the process of monitoring distributed applications. It describes RUEI, JVMD, and EM, which you use to monitor performance, it explains how you set up end-to-end monitoring, and it looks at how security schemes translate across different monitoring contexts.
- [Understanding the User Experience](#) explains how you use RUEI to understand how users are interacting with your product. Using the measurements that RUEI collects, you can assess the effectiveness of user interface design, the responsiveness of web servers and the internet, and the success of user operations.
- [Getting Detailed Execution Information](#) explains how you use Java Virtual Machine Diagnostics to look at the finest details of code execution and to identify problems like race conditions, blocked threads, and memory leaks.
- [Monitoring Business Applications](#) describes how you create a Business Application, and how you use the Enterprise Manager (EM) console to get summary and detail information about the user experience and transaction performance related to that Business Application.
- [Monitoring End-to-end Performance](#) provides an example that illustrates how you use RUEI, and JVMD together to troubleshoot an issue from the user experience to the finest machine-level details.

The chapters in this part are meant to be read sequentially, from beginning to end. If you are familiar with any of the individual components described, Oracle still recommends that you read those subsections that describe how you navigate from one component to others.

The information in this part is not exhaustive. It is a map rather than a compendium. The bulk of material describing how monitoring components work, is found in other documents. Cross references to additional material are provided for your convenience.

13

Monitoring Performance

This section describes the issues and tasks involved in monitoring the performance of distributed applications.

Service-oriented, distributed applications, which are characterized by modular development and dynamic binding, have a critical need for a single point of management from where one can monitor the behavior of the application as a whole, identify actual or potential problems, and take corrective action.

This chapter includes the following sections:

- [Monitoring Views and Dimensions](#)
- [Using ECIDs to Track Requests](#)
- [Setting up End-to-end Monitoring](#)
- [User Roles and Privileges](#)

To monitor the performance of distributed applications, you must be able to do the following.

- Examine the user experience to assess the quality of service rendered and to understand use patterns.
- Discover the components that make up the application, identify request flows of interest, and determine where performance issues or errors occur in the flow.
- Find the root cause of poor performance and failure by looking at the infrastructure supporting the logical application, or by obtaining more detailed information.

Used together, the products described in this guide offer the functionality described above. You do not need to use all these to learn about your application's performance. For example, you could start by monitoring the end-user experience and then later, add transaction monitoring. The next section describes the different monitoring options that are available to you.

13.1 Monitoring Views and Dimensions

End-to-end performance monitoring requires multiple views and dimensions:

- A complete view of the topology of the logical application, including routing schemes and database access.
- A complete view of the underlying infrastructure.
- Varying detail about the distributed application components used.
- For web-based applications, the ability to access html source for the web pages visited by users.
- Access to machine-level execution detail for application components running in a Java Virtual Machine.
- The ability to go from the logical to the physical view of the application.

RUEI, JVMD, and Enterprise Manager provide the functionality required for end-to-end performance monitoring. As mentioned before, you do not need to install and configure all of these. You can use the piece that addresses your most immediate concerns and add more later.

- **Real User Experience Insight (RUEI)**
Helps you identify problems with user interfaces, evaluate the quality of service offered, and understand and anticipate use patterns. With Enterprise Manager 13c registering a RUEI system automatically creates End User Services (EUS) corresponding to each RUEI Application/Suite. The EUS has the same name as the RUEI Application/Suite, and can be used to create a business application.
- **Java Virtual Machine Diagnostics (JVMD), Enterprise Manager.**
Provides a server-level view of the request flow and of the internal workings of the application execution environment for those services that execute in a Java Virtual Machine. JVMD uses an execution context ID (ECID) to allow you correlate events. Fusion Middleware injects an ECID that can be used to track transaction operations (messages). With Enterprise Manager 13c, JVMD can set an ECID in the datastream if an ECID is not already present. This makes JVMD more useful in environments that are use middleware from suppliers other than Oracle.
- **Selenium beacons, Enterprise Manager.**
Selenium is a framework for testing web applications. You can use it to create tests that are integrated into Enterprise Manager as a **Generic Service - Test Based**. For more information on creating service tests and the online help for details about the Selenium Transaction test type, see the *Enterprise Manager Cloud Control Administrator's Guide* .
- **Business Application, Enterprise Manager**
Allows you to define Business Applications, in which context you can view and analyze RUEI information, and to access more detailed monitoring information.
With Enterprise Manager 13c, there is a new service level target, called an **End User Service** (EUS) automatically created, corresponding to each RUEI Application/Suite. If your goal is to monitor a single End User Service (RUEI Application Suite), then you may not need to create a Business Application. However Business Applications allow you to combine various targets (for example, multiple End User Services or a combination of EUS targets).

Figure 13-1 illustrates a simple application configuration where:

- **RUEI** monitors the actions of web users and can create reports, segmented in a variety of ways, that tell you who has requested a page, what pages were requested, which servers were affected, what the response time was, and what the throughput rate is for a given session or user flow.
- **JVMD** provides insight into code running in each JVM, allowing for real-time and historical diagnostics on your Java applications.

 **Note:**

This release of Enterprise Manager automatically adds an ECID header for Java applications that do not automatically provide an ECID. You can now perform tracing and diagnostics that were previously only available to Oracle Fusion Middleware components as described in [ECIDs for Components Other Than Oracle Fusion Middleware Components](#).

Figure 13-1 RUEI and JVMD Monitoring

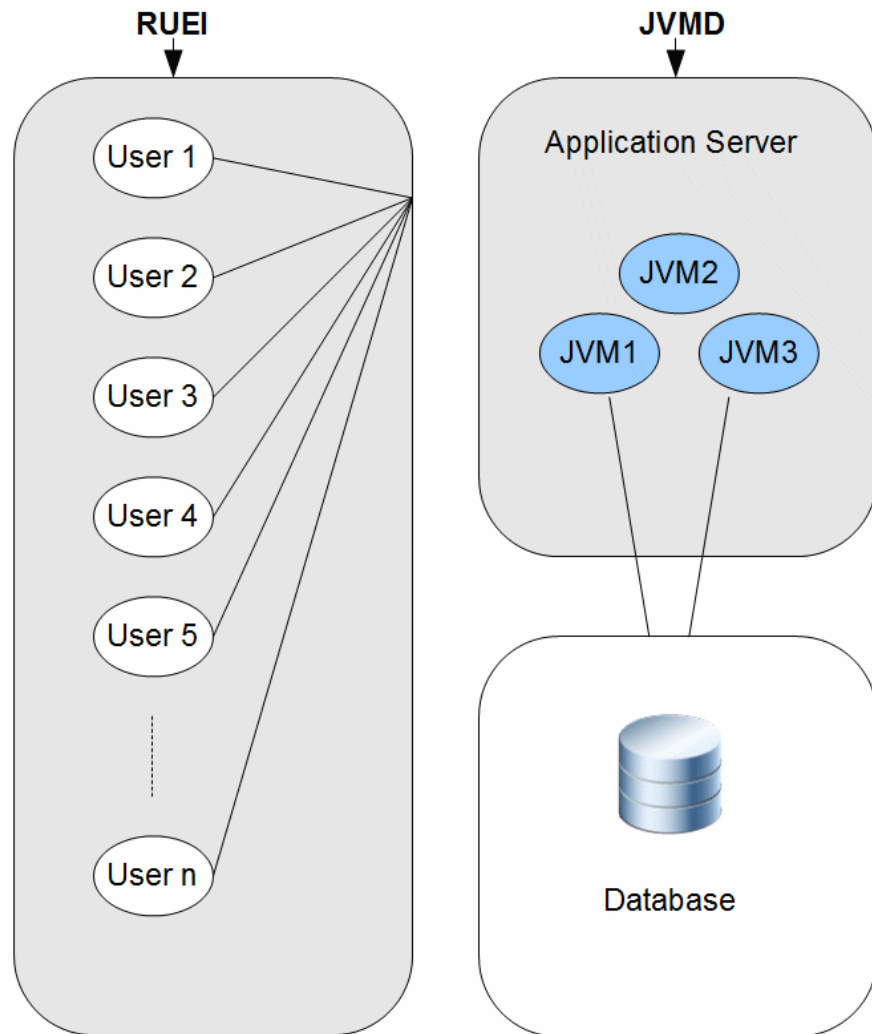
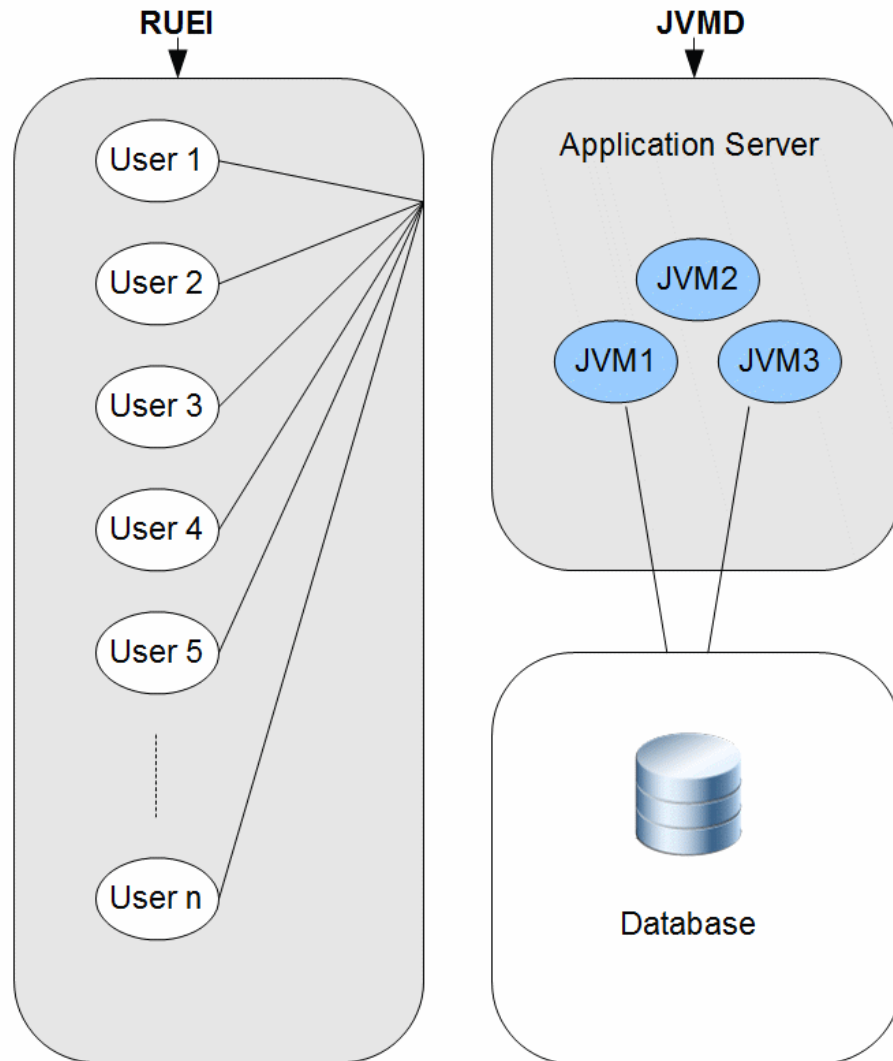


Figure 13-2 a more complex application configuration and to define and monitor the request flows (transactions) that are critical to our understanding of application performance. For a given time period, we can determine the number of started and completed transactions, the throughput, and the average and maximum response times.

Figure 13-2 RUEI and JVMD Monitoring



In addition to using RUEI, and JVMD to monitor end-to-end application performance, you can also use the Enterprise Manager (EM) console to monitor Business Applications that include RUEI applications transactions. For more information, see [Monitoring Business Applications](#) .

13.2 Using ECIDs to Track Requests

Because RUEI and JVMD have a different focus and level of granularity, it helps to have some shared identifier to help us realize that we are looking at a shared process or element.

An Execution Context ID (ECID) is an identifier for tracking a request for components in the Oracle technology stack. An ECID is usually generated by the outer-most Oracle component handling the request and may be propagated to the Oracle components handling that request, potentially crossing server boundaries.

 **Note:**

This release of Enterprise Manager automatically adds an ECID header for Java applications that do not automatically provide an ECID. You can now perform tracing and diagnostics that were previously only available to Oracle Fusion Middleware components as described in [ECIDs for Components Other Than Oracle Fusion Middleware Components](#).

The creation and propagation of ECIDs enable the sharing of context and of diagnostic data between components. Although ECIDs are not universally used, where they are used, they provide good support for end-to-end diagnostic work.

Several technologies generate ECIDs for message traffic; these include RMI, JAX-RPC, JAX-WS, EJB, JMS, JDBC, Servlets, and SOA. (In some cases, ECIDs are supported only when communication occurs between WebLogic servers.)

Where ECIDs are used, they can help the user determine whether they are indeed looking at the same object across execution contexts. For example, you can correlate error messages from different target components if they share the same ECID.

The components used in end-to-end performance monitoring all support the use of ECIDs.

- RUEI displays ECIDs assigned to page objects in the history shown for a particular page.
- ECIDs are also used at the lowest level to further identify threads running in the Java Virtual Machine.
- ECIDs can also be used in the correlation of log entries for Oracle Fusion Middleware components that use the Oracle Diagnostic Logging (ODL) framework.

To have ECIDs generated by default by an HTTP server or Web Logic server, follow the instructions given in My Oracle Support Knowledge Document 1527091.1.

13.2.1 ECIDs for Components Other Than Oracle Fusion Middleware Components

This release of Enterprise Manager automatically adds an ECID header for Java applications that do not automatically provide an ECID. You can now perform tracing and diagnostics that were previously only available to Oracle Fusion Middleware components.

The Oracle Dynamic Monitoring Service (DMS) enables Oracle Fusion Middleware components to put an ECID in the X-ORACLE-DMS-ECID header of a response. If you are not using an Oracle Fusion Middleware component, JVMMD creates a dummy ECID in the X-ORACLE-DMS-ECID header. This ECID serves the same purpose as the DMS-generated ECID allowing Enterprise Manager identify the request.

 **Note:**

Generation of dummy ECIDs is performed only in environments where:

- DMS is not active
- JVMD is installed on the application server
- Instrumentation is enabled on the Enterprise Manager JVM target configuration page

13.3 Setting up End-to-end Monitoring

To obtain end-to-end monitoring, you must install, configure, and connect the products described in [Monitoring Views and Dimensions](#). You might not need to deploy all these pieces at once. You can start with the piece that gives you the functionality you need and add other pieces later.

This section describes the steps required to set up end-to-end application performance monitoring for each dimension of performance monitoring. The purpose of each step is explained, and references are given to the relevant documentation. This section includes the following:

- [Set up Enterprise Manager](#)
- [Set up Java Virtual Machine Diagnostics](#)
- [Set up Real User Experience Insight](#)
- [Create the Business Application](#)

Before looking at the set-up instructions, take a moment to look over the following illustrations, which provide a topological view of the pieces that you can deploy to enable end-to-end monitoring.

[Figure 13-3](#) shows how the RUEI collector, EM agents, and JVMD agents are deployed in a monitored environment.

- The RUEI collector must be deployed in front of the web server.
- The EM agent must be deployed on the machine hosting the application servers and database servers used by the distributed application.
- The JVMD agents must be deployed in the application servers where application components are deployed.

Of course, which of these you deploy, depends on the views you need. For example, if you are not interested in machine-level runtime information, you do not need to deploy the JVMD agent.

Figure 13-3 Agents and Observers in the Monitored Environment

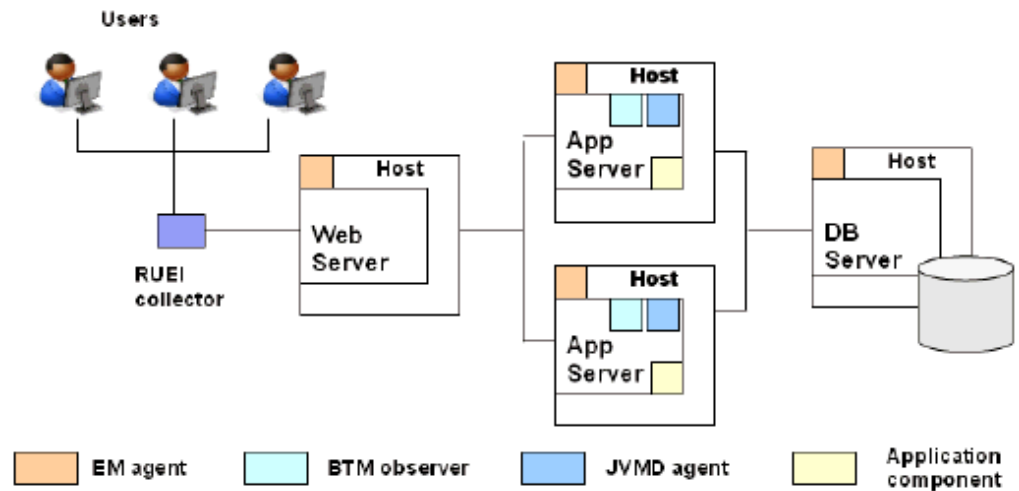
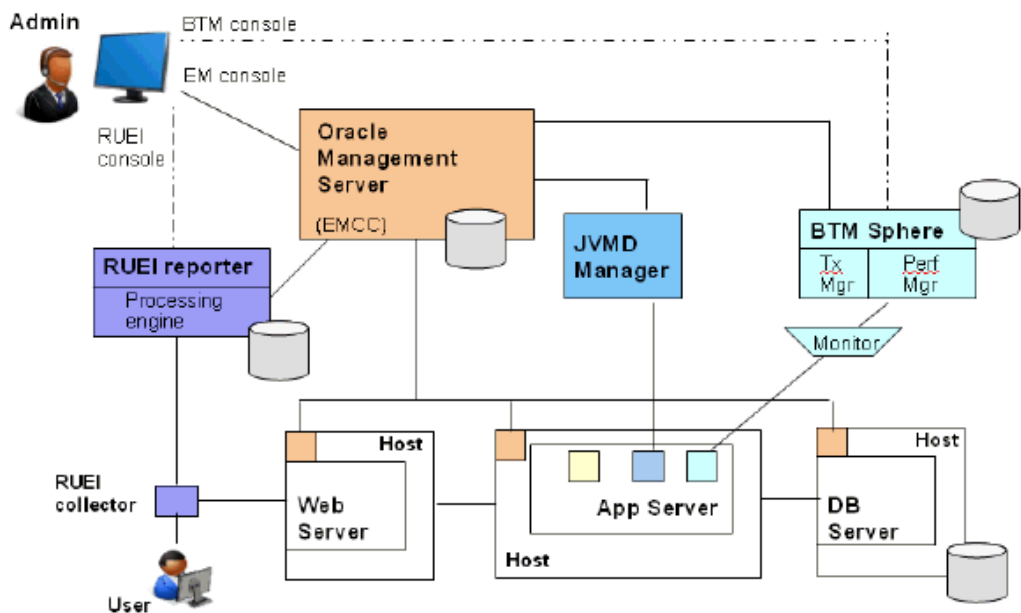


Figure 13-4 shows how RUEI, OMS, and JVM are connected to one another and to their corresponding data collection points.

- It shows that each processing engine connects to its respective console, which allows the administrator to create and update monitored objects.
- It shows how RUEI, and JVM are connected to the Oracle Management Server, which allows the sharing of data that enables the creation and monitoring of Business Applications.

A minimal user environment is shown below the administrative layer.

Figure 13-4 Processing Engines in the Monitored Environment



13.3.1 Set up Enterprise Manager

You need Enterprise Manager to create and monitor Business Applications. Enterprise Manager is also required if you want to do deep-dive diagnostics by looking at machine-level performance data. For more info on this option, see [Set up Java Virtual Machine Diagnostics](#).

To set up Enterprise Manager:

1. **Install and configure Enterprise Manager.** For more information, see [Oracle Enterprise Manager Cloud Control Basic Installation Guide](#).
2. **Install an Oracle management agent** on the hosts where targets and application components monitored by RUEI is running. For more information, see [Oracle Enterprise Manager Cloud Control Basic Installation Guide](#).
3. **Launch Enterprise Manager** and use the **Enterprise Manager** console to create and monitor Business Applications.

13.3.2 Set up Java Virtual Machine Diagnostics

To access machine-level performance data using the JVMD or RID views, you must install the JVMD manager and JVMD agents. JVMD is an integral part of Enterprise Manager, so the latter must be installed before you install JVMD.

To set up Java Virtual Machine Diagnostics:

Install a JVMD agent on all nodes where targets and services monitored by RUEI is running. You need this step to collect JVM data for a given server. For information, see "Installing JVM Diagnostics" in *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.

13.3.3 Set up Real User Experience Insight

To obtain information about the user experience, you must install and configure RUEI, and then register it with Enterprise Manager.

To set up RUEI:

1. **Install and configure RUEI.** This step includes the following:
 - Install collectors, processor, and reporter in the monitored environment.
 - Install the reporter database.
 - Configure the RUEI reporter.

Configuration teaches RUEI to identify users, to specify the collection of pages that make up an application, to specify the scope of monitoring, to configure mail notification, and to provide security options. It is also at this time that you set up a connection to the Oracle Enterprise Manager. For information, see *Oracle Real User Experience Insight Installation Guide* from the appropriate RUEI documentation library:

<http://www.oracle.com/technetwork/documentation/realuserei-091455.html>

2. **Register RUEI with Enterprise Manager.** Specify the port where the Reporter system can be accessed and provide access credentials. You need this step to

establish communication between RUEI and EM. For information, see [Monitoring Business Applications](#) .

13.3.4 Create the Business Application

For end-to-end monitoring, you want to create a Business Application that includes your RUEI applications (End User Services), and the system that supports these. You can build up your Business Application as you go along. You can start by including only the RUEI application, and then add any related transactions. You can even start by looking at the system that supports your distributed applications without including either a RUEI application.

To create a Business Application:

1. **Create a system in Enterprise Manager** that specifies the hosts and containers where monitored application components are running. These hosts and containers are the infrastructure of your distributed application. You need this step for Enterprise Manager to collect and return information about the health of the underlying infrastructure. For more information, see the online help for the Enterprise Manager Console.
2. **Create a Business Application** using the Enterprise Manager console. This step specifies the RUEI applications (End User Services) to be included in a Business Application, and it specifies which system (Step 1) supports the Business Application. For information, see [Monitoring Business Applications](#) .

Note:

With Enterprise Manager 13c, there is a new service level target, called an End User Service (EUS) corresponding to each RUEI Application/Suite. The EUS has the same name as the RUEI Application/Suite, and is used to create a business application.

3. **Monitor the Business Application.** Use the Enterprise Manager Console to monitor the performance of your business application. For more information, see [Monitoring Business Applications](#) .
4. **Edit the RUEI application** if needed.

If you have defined a RUEI application and monitoring results in Enterprise Manager show that you need to change its definition to segment data differently or to re-set key performance indicators, you will need to use the RUEI console to change the application definition. For information, see the appropriate *Oracle Real User Experience Insight User's Guide* at the following URL:

<http://www.oracle.com/technetwork/documentation/realuserei-091455.html>

Enterprise Manager is automatically updated with the new definitions.

13.4 User Roles and Privileges

User roles and privileges define accessibility to component functions. The following guidelines apply as you work with components to monitor application performance.

- Overall, higher privileges are required to create entities in RUEI than to monitor them in EM.

- Clicking through from one component to another exposes you to each console's native authentication system. Make sure that you have the privileges required for each component to perform your work.
- With the exception of the `admin` and `superAdmin` rules, in EM roles are always associated with targets. What is visible to you in the EM console depends on your role with regard to a particular target.
- You need the `super admin` role to register a RUEI system with EM. Once the RUEI system is registered with EM, you don't need the `super admin` role to create a business application.
- You need `Create Any Target` privilege and `View Target` privilege on the RUEI system target to access the credentials used by EM to talk to RUEI.
- You need `Manage Business Application` and `Business Application Menu Item Application Performance Management` resource privileges.
- To view JVM Diagnostics data, you must have `JVM Diagnostics User` privileges.
- To manage JVM Diagnostics operations such as creating and analyzing heap and thread snapshots, tracing threads, and so on, you must have `JVM Diagnostics Administrator` privileges.

14

Understanding the User Experience

This section describes the Real User Experience Insight (RUEI) stand-alone product. For information on using RUEI monitoring functions from the Enterprise Manager console, see [Monitoring Business Applications](#). RUEI allows you to monitor application performance. In particular, RUEI monitors the user's interaction with a web browser, usually the first step (application component) in your distributed application. This first step is a crucial one because it identifies those problems that are most visible to users and because it discovers use patterns that can help you improve the design and effectiveness of your user-facing services.

This section introduces the concepts and tasks involved in working with RUEI to understand the user experience. It includes the following topics:

- [What Does RUEI Discover?](#)
- [Viewing and Analyzing RUEI Data](#)
- [What Questions Can RUEI Answer?](#)
- [What Aspects of RUEI Can You Access from the EM Console?](#)
- [How Does RUEI Work with JVM Diagnostics?](#)

RUEI offers a rich set of features, for complete information about its use, see *Oracle Real User Experience Insight User Guide*.

RUEI can be configured to allow you to monitor performance without requiring access to the network infrastructure, however this chapter assumes that network data collection is used, but the features described are available for the other non-network data collection configurations. Specifically, [What Does RUEI Discover?](#) in this chapter mentions requirements, for example port configuration and network data collection, that are only required if you configure network data collection. For further information on non-network data collection see the RUEI documentation. To view a visual demonstration on how you can use RUEI, navigate to the following URL and click Begin Video:

https://apex.oracle.com/pls/apex/f?p=44785:24:0::NO:24:P24_CONTENT_ID,P24_PREV_PAGE:5783,1

14.1 What Does RUEI Discover?

Users work with your application by interacting with a web page that contains one or more objects. Interacting with an object, for example clicking on a link, the user sets in train a sequence of calls that invoke the services that make up your distributed application.

Typically, a single RUEI instance is installed to collect network data before the Web servers, behind a firewall in the DMZ. RUEI can monitor all users accessing a web page, and it does so without affecting server or network response time.

When you install and configure RUEI, you specify the following information:

- The ports that it should watch for traffic (scope of monitoring)

- How to identify users (using cookie information or log-in information)
- How to deal with security issues and how to monitor encrypted data
- How to identify pages that are associated with a RUEI application

A RUEI *application* is a collection of pages. In the configuration process, you teach RUEI which pages are associated with a given application.

After RUEI begins to monitor traffic on the ports you have specified, it can identify and organize the information it discovers according to the scheme you have defined when you configured RUEI.

Figure 14-1 shows how RUEI collects data associated with a page request.

1. When the user performs an action on a monitored page, RUEI sees the request and starts measuring network timings and the time it takes the Web Server 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).

2. When the Web server responds and sends the object to the user, RUEI sees that response and stops timing the server response time.

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

Figure 14-1 How RUEI Monitors User Requests for Network Data Collection



Every time an object on a page associated with a RUEI application is accessed, RUEI gathers the following information:

- Who requested the page and what object they requested
- Which server hosted the page
- The response time and the correctness of the response
- The size of the object
- Whether the object was completely received or aborted
- The internet throughput for this request/response sequence

The next section explains the various ways in which you can view and analyze this data using RUEI.

14.2 Viewing and Analyzing RUEI Data

Using the information it collects while the user is interacting with your application, RUEI can present a number of views to help you understand performance issues and use patterns relating to the user experience.

In addition to monitoring data on an ongoing basis, you have the option of creating Service Level Agreements that specify the expected level of service. This agreement is expressed in terms of a number of Key Performance Indicators (KPI) that define benchmark values. For more information, see [KPIs and Service Level Agreements](#).

Another aspect of evaluating performance is the monitoring of use patterns. You can define a *user flow* as a sequence of pages, and monitor whether the steps of the flow are completed. For more information, see [User Flows](#).

Data reported is scoped either to active sessions (5 minute duration) or closed sessions which might stretch for several days.

This section introduces some of the most commonly used RUEI views and also describes some additional ways of analyzing the information it gathers. The screenshots show the RUEI user interface, but much of the information is available in Enterprise Manager. The following sections are included:

- [Dashboards](#)
- [Reports](#)
- [Session Diagnostics](#)
- [User Flows](#)
- [KPIs and Service Level Agreements](#)


14.2.1 Dashboards

The RUEI Dashboard offers the most comprehensive view of user activity. It provides the following views of activity over the last twenty four hours (much of this information is available from the Business Application or EUS dashboard in Enterprise Manager):

- A regional, map-based view of the current session activity
- The five most active applications by page view
- The five top problem pages

- The most recent alerts across all monitored applications
- The status of defined KPIs across all monitored applications, showing how much they have changed from the previously recorded value
- A chart showing the proportion of errors due to network errors, client aborts, server errors, website errors, and content errors
- Charts showing average page-load time, and the relationship of page views to sessions
- You can view data from more than one RUEI instance or view the data aggregated from all connected RUEI instances

Figure 14-2 RUEI Dashboard Tab



14.2.2 Reports

RUEI provides an extensive library of pre-defined reports that allow you to display collected user information in a standard way. You use controls in the **Reports** tab to generate and view reports.

You begin by using controls in the **Reports** tab to specify a time period and to select the report you want to generate. Reports are grouped by category, for example **Applications** or **Clients**. Each category offers a variety of reporting options. For example, the **Clients** category allows you to generate reports for Performance per country, Sessions per browser, Sessions per language, Sessions per OS, and so on.

Reports are displayed in table or graphic form and they can be saved as PDF files or exported to other tools.

You can customize reports, you can create new reports, you can create shortcuts to your favorite reports, and you can define filters to constrain reported findings.

14.2.3 Session Diagnostics

The session diagnostics facility allows you to perform root cause analysis of operational problems that have occurred in a given time period.

Diagnostics information is available in a variety of categories; for example, All sessions, failed URLs, slow URLs, Failed pages, and so on. The specific search criteria varies with each group. For example, in the Failed pages category, you can narrow the search by application name, Client IP address, and User ID. You can also use additional filters to limit results.

For some diagnostics categories, you can also specify a search order. For example you can search the most active sessions first.

To use the facility you specify a time period, search criteria (including filters), and search order. RUEI returns all user records that match your search criteria in the order you specified. You can then search further within the currently displayed user records to isolate specific sessions.

The user record that is returned to you includes the complete session page history for a five minute period. You can inspect each page to see its loading satisfaction level, whether it is a key page, and whether it contains an error. You can also select a page to display full page content and the underlying html code received by the server and the client.

In some cases, you can click the **Replay** icon beside a viewed page to replay the complete user session. This allows you to review each page viewed by the visitor during a session, together with any reported error messages.

You can also click out to external tools from the Session diagnostics facility from selected functional areas. For more information, see [How Does RUEI Work with JVM Diagnostics?](#).

You can export complete session contents to external utilities for further analysis, to integrate with other data, or to create the basis for generating test scripts.

14.2.4 User Flows

You create a user flow to define a logical task. A user flow is a collection of web pages and actions. It contains a number of steps that need to be performed to complete the task. In addition to viewing user flows in the RUEI user interface, you can also view user flows as an option on the Business Application or EUS dashboard in Enterprise Manager. For example, a Purchase user flow might have the following defined steps:

- Item selection
- Shipping information
- Billing information
- Confirmation

Each step can consist of multiple pages. For example, the Item selection step might include a number of pages from which items are selected.


User flow steps are defined in terms of conditions specifying the requirements that must be met for the step to be considered complete. For example, if the Billing information includes conditions relating to alternate methods of payment, only one of

these conditions need be satisfied for the step to complete. Steps can be labeled as required or optional. Steps can also have an associated time period against which time-outs and the user experience can be evaluated.

User flows can be associated with a specific application or they can stand on their own.

User flow activity is reported at the most generic level using a funnel shape that illustrates the transition of the visitor through the flow steps for a given time period. The narrowing of the funnel represents visitors lost due to time-outs or visitor aborts. [Figure 14-3](#) shows a sample illustration of a user flow.

Figure 14-3 User Flow Illustration




The flow starts at the top and narrows as users drop off. Each step of the flow is shown in a different color. To the right of the figure are numbers showing how many users aborts and user time outs made up the loss of users for a given step. Following the funnel illustration is more detailed information (not shown in [Figure 14-3](#)) about the activity for each step.

RUEI provides further insight into user flow activity with a view that compares user active time with idle time for each flow. This kind of analysis might suggest which of your pages are most difficult for the user to complete. An example of this view is shown in [Figure 14-4](#):

Note the difference between the Ferry Booking and Flight Booking average idle time. Greater idle time might reflect poor web page design.

User flows provide an excellent means of finding trouble spots, identifying patterns of use, and improving the overall user experience.

Figure 14-4 Active Time vs Idle Time in User Flow Steps



14.2.5 KPIs and Service Level Agreements

In addition to the continuous, passive monitoring provided by RUEI, you can set up active monitoring using Key Performance Indicators (KPIs) to monitor specific aspects of performance, and you can define Service Level Agreements that alert you when the specified benchmarks are breached. You can review this data using dashboards and reports. In addition to viewing KPI and SLA information in the RUEI user interface, you can also view an SLA dashboard or SLA summary region as an option on the Business Application or EUS dashboard in Enterprise Manager.

An SLA defines an expected level of service, typically expressed in terms of one or more Key Performance Indicators. For example a KPI might test whether a service is available 99% of the time, and an SLA might be defined to report when availability falls below this value.

KPIs are grouped into categories such as load times, sessions, throughput, and so on. You can define your own category; for example, user flow completion or website availability.

When you define a KPI you specify the following information:

- Whether to associate it with data from a specific application or whether it will be generic

- What metric to apply
- Whether filters are needed to further define the scope of the KPI. For example, if you selected the user-flow-load-time or Ended user flows metric, you need to specify the user flow to which it refers.
- Whether the KPI has a minimum or maximum target range. Targets can be fixed or relative to historical performance.
- Whether and how the KPI should be incorporated into an SLA
- Whether an alert should be associated with the KPI

RUEI gives you very fine control over active monitoring. You can create service-level and alert schedules that are sensitive to normal periodic variation in target values, and you can define alert profiles and escalation procedures to specify who should be notified when an alert is triggered.

14.3 What Questions Can RUEI Answer?

RUEI can answer questions such as the following about the user experience:

- *What time of the day are the greatest number of page hits?*
Look at the chart that relates page views to sessions on the Dashboard tab.
- *What regions in Europe are experiencing the greatest user activity.*
Look at the Session origin map for Europe, in the Dashboard tab.
- *What percentage of total errors is due to client aborts?*
Look at the Functional errors chart in the Dashboard tab.
- *What are my most problematic pages?*
Look at the Problem Pages listing in the Dashboard tab.
- *Which browser is most heavily used by clients in France?*
Select the **Sessions per browser** report from the **Clients** category in the **Reports** tab, and filter by client-location/country.
- *Show me user records for the Bookings application that have a specific ECID.*
Select the **Session diagnostics** group, and then specify the application name and the ECID of interest. For information about ECID, see [Using ECIDs to Track Requests](#).
- *In what step of my Booking user flow am I losing the most customers?*
Look at the user flow funnel and status details.
- *How many users returned to a previous step in my user flow?*
Look at the Status Details for a user flow to see the number of users returning for each step. A high number of returning users might indicate the need to carry some status information forward into the following screen.
- *When has the availability of my creditCheck service fallen below 95%?*
Define a KPI for that metric, and define a Service Level Agreement that alerts you when the desired value is breached.

14.4 What Aspects of RUEI Can You Access from the EM Console?

You can access monitoring information about the user experience from the Enterprise Manager console. However you cannot define or edit user flows, KPIs, SLAs, or custom Reports in the Enterprise Manager console. All that needs to be done using the RUEI console.

What information is provided in the Enterprise Manager console depends on how you have defined your application and monitoring features in RUEI. Should you find that you need different information, you can use the RUEI console to edit the appropriate elements. Enterprise Manager will be automatically updated with the new definition, and it will display the information you need after you have run additional traffic.

Overall, the information you can access from the Enterprise Manager console includes the following for each RUEI application associated with the current business application:

- On the **Business Application Home** page, you can view the Key Performance Indicators (KPIs) defined for your application, their status, and their defined thresholds. You can also view an overview of incidents and problems associated with the business application. Some of these might have been generated by RUEI.

 **Note:**

With Enterprise Manager 13c, there is a new service level target, called an End User Service (EUS) corresponding to each RUEI Application/Suite. The EUS has the same name as the RUEI Application/Suite, and is used to create a business application.

- The alerts generated by KPIs defined for RUEI applications are reported as events in **Incident Manager**. To view these events select **Monitoring** and then **Incident Manager** from the **Enterprise** menu. Then open the **Events Without Incidents** predefined view. Click the event of interest to view more information.

To reach more detailed monitoring information for RUEI applications, select **Real User Experience (RUEI)** and then **Real User Experience (RUEI) Data** from the **Business Application** drop down. You will be able to see the following regions:

- **RUEI Key Performance Indicators** region, which gives more detailed information for defined KPIs
- **Top User and Application Violations** region, which allows you to examine the application pages with the highest number of violations
- **Top executed User Requests** region, where you can view the most frequent user requests and actions, and assess their impact on the business application
- **Top Users** region, where you can monitor the most active users of the targets associated with the business application

To perform root cause analysis of operational problems, you can use the **RUEI Session Diagnostics** facility. You access this facility by selecting **Real User**

Experience (RUEI) and then **RUEI Session Diagnostics** from the **Business Application** drop down.

To view the **RUEI Metrics** page, select **Real User Experience (RUEI)** and then **RUEI Metrics** from the **Business Application** drop down.

For complete information about working with RUEI in the Enterprise Manager Console, see [Monitoring Business Applications](#) .

14.5 How Does RUEI Work with JVM Diagnostics?

RUEI can work seamlessly with JVMD if you install and configure these as described in [Setting up End-to-end Monitoring](#). Options include the following:

- You can click out to JVMD to get activity information for the selected request based on its ECID. You can access the Request Instance Diagnostics page by a right-click on a record in a RUEI Session Diagnostics view.

For more information about how RUEI works with external tools, see [Configuring Clickouts to External Tools](#) in *Oracle RUEI User's Guide*.

Getting Detailed Execution Information

This section describes how to use Enterprise Manager JVM views to get detailed execution information about failing or problematic operations.

There are times when the views offered by RUEI is not sufficient to understand performance issues. If the suspect services are executing in a Java Virtual Machine, it is possible to go deeper and get detailed execution information that helps you diagnose the root cause of such problems.

JVM Diagnostics is a tool that allows you to view the details of an executing JVM process. These details include the stack frames for executing threads, thread state information, aggregate information about the frequency and cost of method execution, information regarding the holding of Java and database locks, and details about the objects in the Java heap. Using this tool you can also access historical data for each JVM monitored.

This chapter includes the following sections:

- [Using JVM Diagnostics](#)
- [Using Request Instance Diagnostics](#)

When you invoke one of these views from RUEI to further analyze performance, Enterprise Manager selects and displays data generated in the time interval for the selected RUEI page object operation instance. One additional piece of information that might be shown for the data displayed is its execution context ID (ECID).

An ECID is an identifier used to track a request, for components in the Oracle technology stack. The creation and propagation of ECIDs enable the sharing of context and of diagnostic data between components. ECIDs are also used to identify threads running in the Java Virtual Machine. Where ECIDs are available, they can help you correlate data shown in RUEI with data shown in the JVM Diagnostics view or Request Instance Diagnostics view. For additional information, see [Using ECIDs to Track Requests](#).

To access JVM views from RUEI, you must do some preliminary set-up work. For more information, see [Setting up End-to-end Monitoring](#).

JVMD offers a rich set of features that we cannot hope to describe in a single chapter. For complete information about its use, see the chapters describing JVMD in this book.

15.1 Using JVM Diagnostics

Java Virtual Machine Diagnostics (JVMD) information is accessed in one of the following ways:

- In the **Message Log** tab for a service, endpoint, logical operation, physical operation, or transaction. Right-click on a row and select **Drilldown to JVMD** from the context menu.
- In the **Transaction Instance Inspector**, right click on an operation (in either the graph or grid view), and select **Drilldown to JVMD** from the context menu.

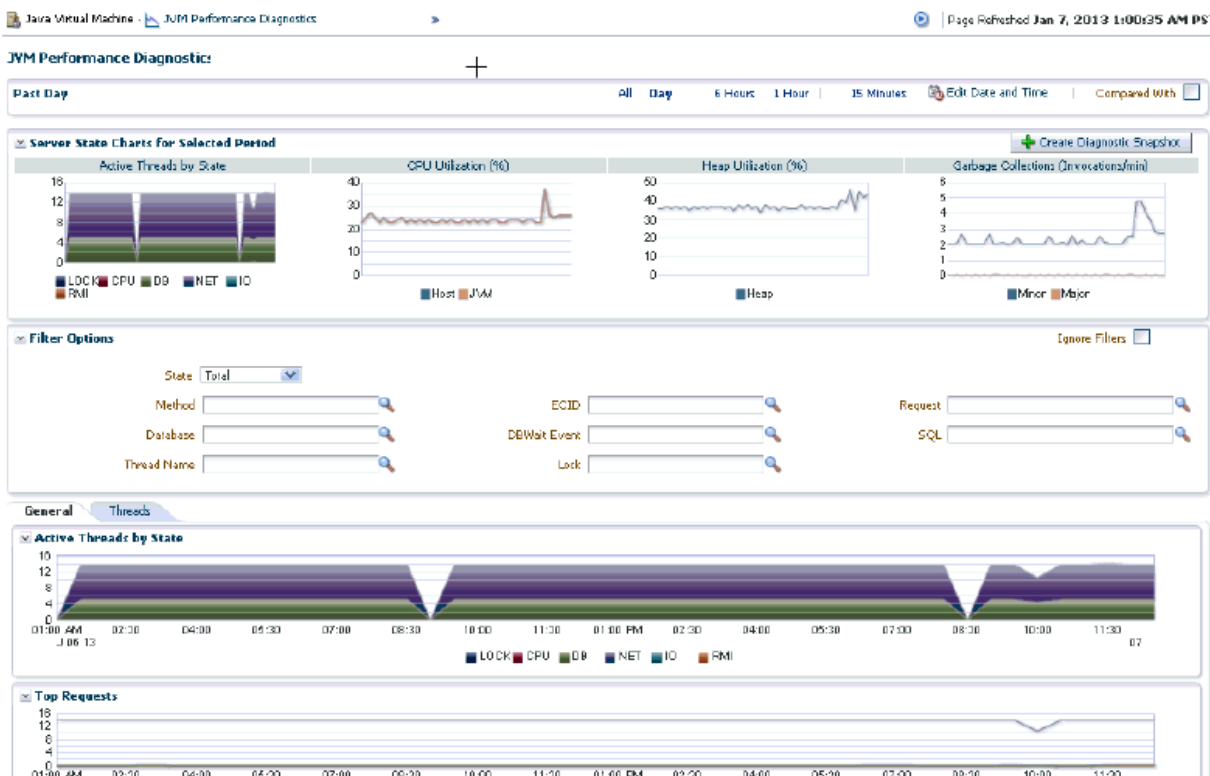
- In the **Message Search Log** tool, right click on a message row, and select **Drilldown to JVMD** from the context menu.

In Enterprise Manager, you can access JVMD information for a transaction operation by selecting a transaction in the Business Application and opening the transaction summary page. Then do one of the following:

- Right click one of the operation nodes in the topology diagram and select JVMD diagnostics from the context menu.
- Right click one of the operation rows in the operations table and select JVMD diagnostics from the context menu.

In each case, a new window is displayed showing the JVM Performance Diagnostic view. In the multi-VM case, JVMD shows a VM group target and aggregate information for the group. [Figure 15-1](#) shows the JVM Performance Diagnostic view.

Figure 15-1 JVM Performance Diagnostic View



This view shows the summary details of the JVM in which the selected operation is running. It shows Server state charts, Active Threads by State, Top Methods, Top Requests, Top DBWait Events, TopSQLs, and Top Databases. You can filter the data that is displayed by specifying various criteria.

Click on the **Threads** tab to view the Thread State transition chart. This chart shows how the threads have transitioned from one state to another in the selected period. Click on a bar graph in the Thread State Transition chart to view the Sample Analyzer, which provides a detailed analysis on the thread of the thread.

Click the **Live Thread Analysis** control to see all threads running in the JVM. Click on a thread to view additional information about that thread.

15.2 Using Request Instance Diagnostics

You can access the Request Instance Diagnostics (RID) view either from RUEI.


- From the RUEI stand-alone application, the ECID is used to correlate the data shown. You can access RID by a right-click on a record in a RUEI session diagnostics view.
- From a RUEI Session Diagnostics, object view in EM, you can access RID by a right-click on the Oracle logo icon. (The icon is displayed only if there's an ECID)

Note:

Java Virtual Machine Diagnostics (JVMD) must be installed and active before you can access the Request Instance Diagnostics (RID) view.

Figure 15-2 shows part of the Request Instance Diagnostic view for a given ECID.

Figure 15-2 The Request Instance Diagnostics View



The JVMs panel lists all the JVMs through which the request was executed. Select a JVM to display the following information:


- RID: The Relationship ID, an ordered set of numbers describing the location of each task in the tree of tasks.
- The start time and the duration of the request.
- Step Name: The individual steps in the request. For example, the first step could be jsp, the second could be EJB, and the third could be DB.
- CPU utilization by the JVM
- GC Major/Minor indicates the number of objects added to the major and minor garbage collections.

If you select a JVM from the list, a bar graph is displayed in the **Instance Samples** tab of the JVM Instance Analysis panel. This graphs shows the thread state in each JVM snapshot taken within the duration of the request. A color key, to the right of the display, indicates a different thread state; Runnable, Lock, IO wait, DB Wait, NW wait, and RMI Wait. Hover over the graph to get an in-depth view of the thread.

To view aggregate metrics collected for the selected JVM during the specified period, click on the **Request Aggregate** tab. [Figure 15-3](#) shows a sample tab.

To view measurements for a given metric type, select the desired type from the drop down Metric menu, as shown in the figure.

Figure 15-3 RID: Request Aggregate Tab.



16

Monitoring Business Applications

This section describes how you use Oracle Enterprise Manager to monitor Business Application performance.

A *Business Application* is an Enterprise Manager target that represents a logical application; for the user, it defines a unit of management. A Business Application is composed of RUEI applications. Using the Enterprise Manager Console, you view a Business Application to access RUEI and information about the application's supporting infrastructure: the hosts and servers where the application services are executing.

You cannot use Enterprise Manager to create RUEI applications. That work must be done using the RUEI products, which were introduced in previous chapters. You must complete the steps described in [Setting up End-to-end Monitoring](#), to be able to set up Business Application monitoring. You must also complete the tasks described in [Prerequisites and Considerations](#).

Note:

With Enterprise Manager 13c, there is a new service level target, called an End User Service (EUS) corresponding to each RUEI Application, Suite or web service. The EUS has the same name as the RUEI Application, Suite or web service, and is used to create a business application. See [Monitoring an End User Service](#) for more information.

This chapter covers the following:

- [Introduction to Business Applications](#)
- [Prerequisites and Considerations](#)
- [Registering RUEI Systems](#)
- [Creating Business Applications](#)
- [Monitoring Business Applications](#)
- [Monitoring End User Experience](#)
- [Monitoring an End User Service](#)
- [Upgrading Business Applications](#)

16.1 Introduction to Business Applications

By using Oracle Enterprise Manager to monitor your Business Applications, you can make sure that your applications are performing at their peak and that end users are satisfied with their performance.

The use of Business Applications offers a number of significant advantages over traditional IT-centric approaches that only focus on system health issues. In particular, Business Applications:

- Allow you to manage your applications in their business context, measuring, and alerting on the basis of the end-users' experience.
- Provide customizable dashboards with complete visibility across multi-tier composite applications.
- Provide a visualization of all target relationships within a business service.

16.1.1 Systems, Services, and Business Applications

Within Oracle Enterprise Manager, there are two types of targets: systems and services. A Business Application is an aggregate service target that can be associated with a minimum of either a system target or a sub-service.

System Targets

Consider an example business application that contains an order entry application implemented by a collection of physical (system) resources. The application is deployed in a Web Logic domain modeled as a system target whose members are the individual managed servers. The Business Application could include transactions deployed in containers. Each of these containers is an application server, possibly within a single Web Logic domain. In this case, the Web Logic domain is a system target. (In the case that the transaction spans multiple domains, it is recommended that you create a composite application within Oracle Enterprise Manager.)

System monitoring provides insights into the behavior of the monitored application infrastructure. It collects metrics and reports on the health of all components from the hosts to the application servers and the deployed Java EE applications. It also provides deep-dive diagnostics tools for the application servers and the databases.

Sub-services

You can define a service by creating one or more tests that simulate common end-user functionality. You can also define services based on system targets, or on both system and service tests. You can then use these services as sub-services to define a Business Application. For more information on services, see the *Enterprise Manager Cloud Control Administrator's Guide*.

Business Applications

A business application can be associated with a mix of system targets and sub-services, monitoring them to determine the business application's availability.

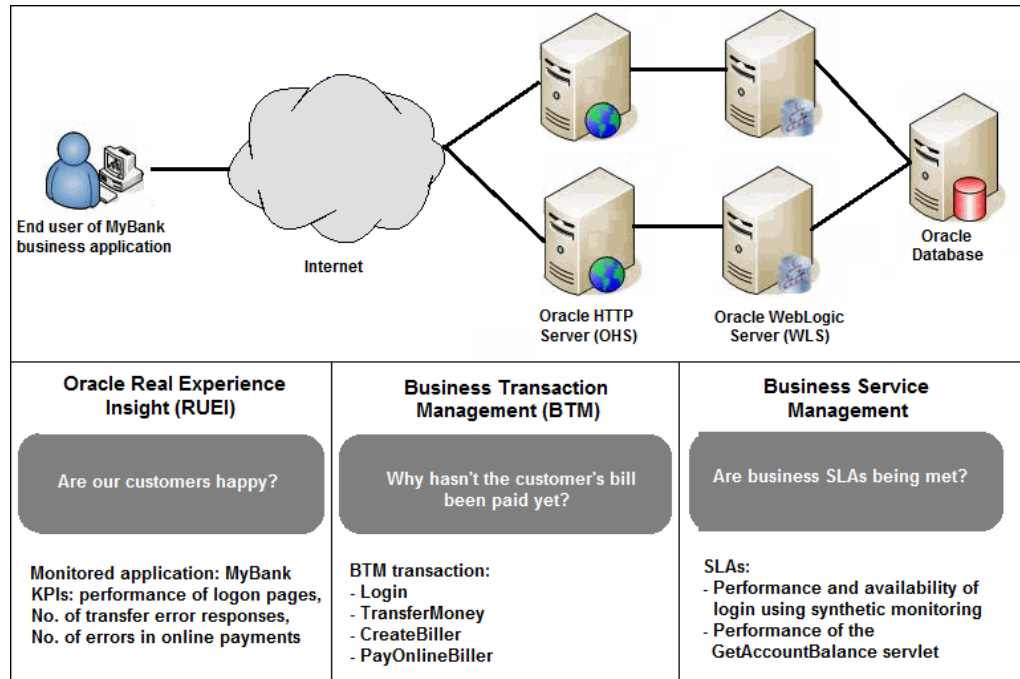
For system targets, you can specify which key components within a system target should be monitored to determine the business application's availability. For instance, for a transaction, the key components will be the servers where the services that comprise the transaction are running.

16.1.2 MyBank: An Example Business Application

To illustrate the nature of a Business Application, consider the situation in which end users access a banking application (MyBank) that allows them to perform such tasks as the payment of bills. This business application is delivered through the infrastructure shown in [Figure 16-1](#).

The end-user experience of the MyBank business application is monitored through RUEI, while Key Performance Indicators (KPIs) are used to monitor its key aspects, such as the availability and performance of the logon page, and the number of errors in transfer responses and online payments.

Figure 16-1 The MyBank Business Application



Proactive application monitoring is achieved by defining business objectives that set acceptable levels of performance and availability. Within Oracle Enterprise Manager, these business objectives are referred to as *Service Level Agreements* (SLAs) and are composed of *Service Level Objectives* (SLOs) that measure specific metrics.

Insight into each of these key aspects of a business application's operation and delivery is available through a number of dedicated regions of the Oracle Enterprise Manager console.

16.2 Prerequisites and Considerations

This section describes the requirements that must be met and the issues that should be considered to use the Business Applications facility. It is strongly recommended that you carefully review this information before proceeding with the creation of business applications.

 **Note:**

It is recommended that you review the My Oracle Support website to obtain up-to-date information about the supported RUEI as well as patches, configurations, known issues, and workarounds.

This section covers the following:

- [Requirements for Using RUEI](#)

16.2.1 Requirements for Using RUEI

To use RUEI to monitor the performance of your Business Applications, you must ensure that the following requirements have been met:

- RUEI version 12.1.0.7 (or higher) has been installed and configured to monitor the required applications, suites, and services. Information about deployment options and requirements is available from the *Oracle Real User Experience Insight Installation Guide*.
- The Enterprise Manager for Oracle Fusion Middleware plug-in must be deployed to both Oracle Management Service (OMS) and to each Management Agent monitoring the business application targets.

For details on deploying the plug-in to OMS, see the "Deploying Plug-Ins to Oracle Management Service" chapter in the *Enterprise Manager Cloud Control Administrator's Guide*.

For details on deploying the plug-in to a Management Agent, see the "Deploying Plug-Ins on Oracle Management Agent" chapter in the *Enterprise Manager Cloud Control Administrator's Guide*.

- The Reporter system must be accessible to Oracle Enterprise Manager via an HTTPS connection on port 443. Other component host systems (such as Collector, Processing Engine, and database servers) do not need to be accessible to Oracle Enterprise Manager unless you intend to make them managed targets. For more information, see [Registering RUEI Systems](#).
- The statistics data retention setting (which governs the availability of statistical information such as violation counters) has been configured to be consistent with your business application reporting requirements. The procedure to do this is described in the *Oracle Real User Experience Insight User's Guide*.
- If you intend to export session information from the Session Diagnostics facility, you should ensure that the exported session is not older than the period specified for the Full Session Replay (FSR) data retention setting. In addition, the URL prefix masking setting should be specified as "Complete logging". For more information, see the *Oracle Real User Experience Insight User's Guide*.

16.2.1.1 Registering RUEI Installations with Self-Signed Certificates

A RUEI installation can use a self-signed certificate. This is explained in the *Oracle Real User Experience Insight Installation Guide*. However, Oracle Enterprise Manager only accepts SSL certificates issued by a trusted Certificate Authority (CA), and that contain a valid Common Name (CN). Therefore, in order to be able to register a RUEI installation with Oracle Enterprise Manager, you need to do the following:

 **Note:**

All instructions on the Oracle Enterprise Manager system need to be carried out as the user running the oms and agent.

1. Verify that the certificate is valid. One way to do this is to attempt to access the Oracle Real User Experience Insight system through a browser via HTTPS and view the certificate details. You should ensure the certificate's date validity. If the certificate's date range does not include the period your Oracle Real User Experience system is running, you will not be able to use it.
2. Download the certificate to your Oracle Enterprise Manager system. Many browsers provide an option when creating a security exception for a self-signed certificate to also save the certificate to a file. If you have already approved the security exception in your browser, the following example works in Mozilla Firefox:
 - a. Click the security icon to the left of the hostname.
 - b. Click **More information**, then click **View certificate**.
 - c. Select the **Details** tab and click **Export**.

The exported file should be copied to your system running Enterprise Manager. The examples below assume that you stored the file containing the certificate in `~/ruei.cert`.

A more direct way to download the certificate to your Oracle Enterprise Manager system can be carried out on the system itself. Issue the following commands on the Oracle Enterprise Manager system:

```
openssl s_client -showcerts -connect <RUEI_REPORTER_HOST>:443 </dev/null \
| openssl x509 -inform PEM > ~/ruei.cert
```

3. Add the certificate to the keystore. Within Oracle Enterprise Manager, two components are used to communicate with a RUEI system via SSL: one for polling the status of RUEI, and one for the communication with RUEI. Both keystores need to contain the same certificate. Issue the following commands on the Oracle Enterprise Manager system:

Agent:

```
cd <agent instance home>/bin
./emctl secure add_trust_cert_to_jks \
[-password <keystore password, default "welcome">] \
-trust_certs_loc ~/ruei.cert -alias <unique alias>
```

OMS

```
<path_to_Oracle_WT>/jdk/bin/keytool -import \
-keystore <path_to_wlserver_10.3>/server/lib/DemoTrust.jks \
-file ~/ruei.cert -alias <unique alias> -storepass DemoTrustKeyStorePassPhrase
```

4. In order for Oracle Enterprise Manager to work with the new certificate, perform a bounce of the OMS and the AGENT. Issue the following commands:

```
<OMS oracle home>/bin/emctl stop oms -all
<OMS oracle home>/bin/emctl start oms
<AGENT oracle home>/bin/emctl stop agent
<AGENT oracle home>/bin/emctl start agent
```

16.3 Registering RUEI Systems

Before you can create Business Applications based on RUEI-monitored applications and services, you must first register the appropriate RUEI with Oracle Enterprise Manager.

 **Note:**

You must have Super Administrator privileges in order to access the Middleware Management Setup page.

With Enterprise Manager 13c, there is a new service level target, called an End User Service (EUS) corresponding to each RUEI Application, Suite or web service. The EUS has the same name as the RUEI Application, Suite or web service, and is used to create a business application. A synchronization job is automatically created to maintain End User Services (EUS) entry details corresponding to each RUEI Application/Suite. For more information on jobs, see the *Enterprise Manager Cloud Control Administrator's Guide*. If you later create a new RUEI Application, Suite or web service, you may need to run the job manually as described in [Troubleshooting an End User Service](#).

 **Note:**

To successfully register a RUEI system, make sure that the Enterprise Management Repository database instance target is properly registered as described below, and that the database instance target has associated monitoring credentials with the SYSDBA role. These credentials are required so that KPIs from RUEI can be monitored by Oracle Enterprise Manager.

During a typical installation of Enterprise Manager, the database instance target for the management repository is not fully registered automatically. The database is auto-discovered by Enterprise Manager but must be promoted as described in the Discovering and Adding Database Targets chapter of the Enterprise Manager Administrator's guide. Use the appropriate section instructions for your database setup, for example, Discovering and Adding Single Instance Database Targets or Discovering and Adding Cluster Database Targets.

1. From the **Setup** menu, select **Middleware Management**, then select **Setup**. The **Setup** page shown in [Figure 16-2](#) appears. The currently registered systems are listed.

Figure 16-2 Middleware Management Setup

The screenshot shows the 'JVMD Agents' section with a count of 2 and a 'Manage JVMD Agents' button. Below it is the 'JVMD Load Balancer(s)' table:

Load Balancer URL	Status	Associated Engine(s)
https://solsparcc...	↑	javmdengineEMGC_OMS2, javmdengineEMGC_OMS1

Below the table is the 'RUEI / BTM / JVMD Engines' section with a toolbar containing 'View', '+ Add', 'Redeploy', 'Remove', 'Configure', and 'Troubleshoot'. The table below shows the following entries:

Name	Host	Port
RUEI Systems (0)		
BTM Systems (0)		
▲ JVM Diagnostics Engines (2)		

2. Select **Real User Experience Insight System** from the **Add** drop down. A page similar to the one shown in [Figure 16-3](#) appears.

Figure 16-3 Discover RUEI System Page

The screenshot shows the 'Discover RUEI System: Find Targets' page. It includes a 'Test Connection' button, a 'Discover' button, and a 'Cancel' button. The page contains the following text and form fields:

Enterprise Manager can be configured to manage a RUEI instance

Enter the host name and port number where system RUEI is running, along with valid credentials to access the target. If the server running system RUEI is managed by Enterprise Manager, you can select it, then select the Management Agent that will monitor the target.

* Host

* Port

SSL

* Username

* Password

* Target Prefix

* EM Agent

TIP A Management Agent with Oracle Fusion Middleware plug-in of version 12.1.0.3 or above is required.

**Note:**

You can register more than one RUEI system.

3. Specify the host system where the Reporter system is located. Click **Select Target**. A new window opens that allows you to view the available systems. You can use the **Target Type** menu to search for specific target types.
4. Specify the port number for communication with the RUEI Reporter, for example 443.
5. Specify a secure connection for communication with the RUEI Reporter. Only use an insecure connection for testing purposes.
6. Specify a valid user name and password combination. For a RUEI system, the specified user must have Oracle Enterprise Manager access permissions. Note that Oracle SSO authentication for this user is not supported. The Security Officer privilege is also recommended to allow downloading of sessions and the showing of replay details in the Enterprise Manager UI. With this in place, Enterprise Manager will be able to retrieve this data. Moreover, additional (per-end-user) Enterprise Manager roles will be applied to reveal session-zip download and content-download buttons.
7. Optionally, specify a string to be attached to the RUEI system name.
8. Specify the host and port of the Management Agent to be used to collect metric information about the system. If it is managed by Oracle Enterprise Manager, you can click **Select** to specify it.
9. Click **Test Connection** to verify whether a working connection to the RUEI system can be made.

 **Note:**

A secure connection to a RUEI installation fails if you have not completed the process described in [Registering RUEI Installations with Self-Signed Certificates](#).

10. Click **Discover**. An overview of the components and End User Services associated with the selected system is displayed. An example is shown in [Figure 16-4](#).

Figure 16-4 Discover RUEI Instance: View Targets Page

Setup Page Refreshed 04-Nov-2015 08:33:32 PST ↻

Discover RUEI System: Add Targets Back Add Targets Cancel

Targets Found 3

Following are the discovered RUEI targets. Click Add Targets to enable these targets to be monitored by Enterprise Manager.

Name	Type	Host
doctest_Oracle Real User Experi...	RUEI System	192.168.1.100
doctest_reporter-adc01jld.us.orac...	RUEI Reporter Engine	192.168.1.100
doctest_RUEI	End User Service	192.168.1.100

Edit Credential

Incident Alert Credentials

Stores the Management Repository connection details on the RUEI target, and adds Enterprise Manager credentials for alert integration. Have the RUEI wallet password available when performing the update. Also verify that the backlink works after completing the update, even if the update succeeds.

▲ **Oracle RUEI Wallet Credentials**
Specify the Oracle wallet password for RUEI system "doctest_Oracle Real User Experience Insight".

* Wallet Password

For RUEI systems, you also need to enter credentials to enable incident communication from RUEI in a section of the screen labelled **Edit Incident Alert Credential**.

In the **Oracle RUEI Wallet Credentials** section, enter the RUEI Wallet password. Click **Test** to make sure you have entered the password correctly.

11. Click **Add Targets** to have each of the system's components become a managed target within Oracle Enterprise Manager. Note that if you do so, each system must be accessible to a Management Agent. Further information about managed targets is available from the *Oracle Enterprise Manager Cloud Control Administrator's Guide*. If you need to edit any of the credentials you specified above, you can select the target and click the **Configure** button to display a screen that will allow you edit the values.

16.3.1 Setting Up a Connection Between RUEI and the Oracle Enterprise Manager Repository

The following procedure describes setting up a connection so that KPIs from RUEI can be monitored by one or more Oracle Enterprise Manager instances. This procedure can be used for the following situations:

- After changing Enterprise Manager hostname
- After changing the sysman user credentials
- After changing the TNS settings of the Enterprise Manager database
- Correcting issues with initial KPI setup

1. From the **Setup** menu, select **Middleware Management**, then select **Setup**. The page shown in [Figure 16-2](#) appears. The currently registered systems are listed.
2. Select the RUEI system you would like to set up and click **Configure**. The RUEI Setup Page appears.
3. Select the **Edit credential** tab and click **Edit**. The **RUEI Setup page** shown in [Figure 16-5](#) appears. Enter the appropriate RUEI wallet password, this is typically specified while setting up the RUEI repository.

Figure 16-5 RUEI Credentials

RUEI Setup page

Edit dimension listing | **Edit credential**

Incident Alert Credentials Save Remove Return

Stores the Management Repository connection details on the RUEI target, and adds Enterprise Manager credentials for alert integration. Have the RUEI wallet password available when performing the update. Also verify that the backlink works after completing the update, even if the update succeeds.

▲ **Oracle RUEI Wallet Credentials**
Specify the Oracle wallet password for RUEI system "Almere_Oracle Real User Experience Insight".

* Wallet Password *

Test

Note:

If you change the RUEI wallet password, you must edit the Enterprise Manager RUEI credential on this screen to maintain the RUEI connection. Do not edit using Enterprise Manager named credential feature for EUS_ENGINE_USER.

4. On the RUEI host, configure RUEI to use the mkstore utility:
 - a. Determine the location of the mkstore utility. This utility is included with the Oracle Database and Oracle Client runtime. In both cases, it is located in `$ORACLE_HOME/bin/mkstore`.
 - b. Edit the `/etc/ruei.conf` file and add the following line, where `mkstore_location` is the path determined in the step above:


```
export MKSTORE_BIN=mkstore_location
```
 - c. Restart RUEI by selecting **System**, then **Maintenance**, and then **System reset**. Select **Reapply latest configuration** option and click **Next** to apply the changes you have made.

16.4 Creating Business Applications

To create a Business Application, you need to specify the RUEI-monitored applications, suites, and services, upon which it is based. You are not required to include both RUEI applications in a Business Application.

Do the following:

1. From the **Targets** menu, select **Business Applications**. The currently defined Business Applications are displayed. The page (partially) shown in [Figure 16-6](#) appears.

Figure 16-6 Business Application Page

Name	Health	Status	Incidents			
			⊖	⊗	⚠	🚩
ba_ruei_btm	100	↑	-	2	-	-

2. Click **Create**. The page shown in [Figure 16-7](#) appears.

Figure 16-7 Create Business Application (Name) Page

Business Application

Progress bar: Name (active), End User Service Associations, BTM Associations, System, Service Associations, Review

Create Business Application: Name [Back] Step 1 of 6 [Next] [Cancel]

Enter a unique name for the new Business Application.

* Business Application Name

3. Specify a unique name for the new business application. It is recommended that you include an indication of the purpose and scope of the business application as part of the name. Do not accept the default name, which is invalid if it contains either of the < or > characters. Note that business applications cannot be renamed later. When ready, click **Next**. The page shown in [Figure 16-8](#) appears.

Figure 16-8 Create Business Application (End User Service Associations) Page

Business Application

Name End User Service Associations BTM Associations System Service Associations Review

Create Business Application: End User Service Associations Back Step 2 of 6 Next Cancel

Notice: There are multi RUEI systems. Select End User Services from the same RUEI system.
Use Add to optionally associate one or more End User Services with the Business Application.

End User Services

+ Add X Remove

Name	Type
No End User Services associated	

4. Click **Add**. A new window opens that allows you to select one or more End User Services upon which the business application should be based. There is an End User Service item corresponding to each RUEI application/suite. You can multi-select from the list. When ready, click **Next**. The page shown in [Figure 16-9](#) appears.

Figure 16-9 Create Business Application Page

Business Application

Name End User Service Associations BTM Associations System Service Associations Review

Create Business Application: BTM Associations Back Step 3 of 6 Next Cancel

No BTM system available. To discovery a BTM system, click [Setup](#).
Use Add to optionally associate one or more BTM Business Transactions with the Business Application.

Business Transactions

+ Add X Remove

Name
No Business Transactions associated

5. Click **Add**. When ready, click **Next**. The page shown in [Figure 16-10](#) appears.

Figure 16-10 Create Business Application (System) Page

Business Application

Name End User Service Associations BTM Associations **System** Service Associations Review

Create Business Application: System Back Step 4 of 6 Next Cancel

Optionally associate a system with the Business Application. Associating with system is optional.

A "system" is the infrastructure used to host one or more services. A system consists of components such as hosts, databases and other targets.

Select a system target on which the service will be based.

System <No system selected> 🔍

Type

6. Click **Select System** and select the system that hosts the business application. This should be a system that encompasses the infrastructure that the business application runs on.

When ready, click **Next**. The page shown in [Figure 16-11](#) appears.

Figure 16-11 Create Business Application (Service Association) Page

Business Application

Name End User Service Associations BTM Associations System **Service Associations** Review

Create Business Application: Service Associations Back Step 5 of 6 Next Cancel

Use Add to optionally associate one or more member services with the Business Application.

Member Services

Name	Type	Status
No service added. Click on 'Add' button to include services.		

7. Click **Add**. A new window opens that allows you to select the service upon which the business application should be based. When ready, click **Next**. The page shown in [Figure 16-12](#) appears.

Figure 16-12 Create Business Application (Review) Page

Business Application

Progress bar: Name, End User Service Associations, BTM Associations, System, Service Associations, Review (selected)

Create Business Application: Review Back Step 6 of 6 Next Create Business Application Cancel

Displays a summary of all of the previous steps. Clicking the Create Business Application button will create a Business Application with the data shown here.

▲ **Name**
Business Application Name Test

▲ **End User Service Associations**

Name	Type
RUEI1_ADF	End User Service

▲ **BTM Associations**

Name
No Business Transactions associated

▲ **System**

Name /bi_mrg_bifoundation_domain/bifoundation_domain

▲ **Service Associations**

Name	Type
No Service associated	

- Review the new business application's properties before creating it. If necessary, use the **Back** and **Next** buttons to amend its properties. When ready, click **Create Business Application**. The newly created Business Application home page (Figure 16-13) appears in personalization mode, allowing you to rearrange the regions or add new regions.

16.5 Monitoring Business Applications

Once a Business Application has been created, you can use the **Business Application** home page to monitor its performance and availability, as well as the status of the systems (hosts, databases, and middleware components) that support it.

It is also from the **Business Application** home page that you can access more detailed information about RUEI components:

- To get more information about RUEI components, select one of the End User Service related views from the **Business Application** drop down menu. See [Monitoring End User Experience](#).

 **Note:**

If there are timeout issues associated with monitoring the Business Application, you can set the `APM_WEBSERVICE_CREATE_TIMEOUT` system property in the Enterprise Manager WebLogic configuration to a value appropriate to your network configuration, for example 60 seconds.

To view the **Business Application** home page, do the following:


1. From the **Targets** menu, select **Business Applications**. The currently defined business applications are listed. An example is shown in [Figure 16-6](#).

In addition to the names of existing Business Applications, this page also provides summary information about status, system, RUEI metrics, and RUEI KPIs. Use the **View > Columns** menu option to add or delete columns to this display.

You can also click across to view information about system targets by clicking on one of the names listed in the System pane.

2. Click the Business Application of interest. The home page for the selected Business Application is displayed. An example is shown in [Figure 16-13](#).

Figure 16-13 Business Application Home Page



Each region provides specific information on the various operational aspects of the selected business application. By default, the following regions are available:

- **General:** indicates its status, availability and health.
- **Incidents and Problems Overview:** Count of incidents and problems across different categories such as severities or target types.
- **Component Status:** indicates the availability of the components that deliver the business application and the history of the current status. Expand the displayed tree to show status for systems, services and service tests. Click an item for further detail.
- **SLA Status:** indicates the status of the Service Level Agreement if an SLA is configured.
- **RUEI - Geographic Location Overview:** displays RUEI metrics on a geographic map. You can select which metric you would like to view and restrict the data to a region.
- **RUEI - Key Performance Indicators (KPI):** indicates the status of the KPIs defined for the applications, suites, and services associated with the business application.
- **RUEI - Usage and Violations Overview:** displays a chart of issues to help you isolate issues by page.
- **Weblog - Most Requested:** displays the most requested services over the last 24 hours. Click a tab to view the services of a particular type, for example RESTful Services.

In addition to the default regions, more regions can be added by clicking the Personalize Page icon, including regions relating to JVM, RUEI and SLA activities. In this mode, you can also rearrange the regions.

**Note:**

A Business Application is an aggregate service as explained in the Configuring and Using Services chapter of the *Enterprise Manager Cloud Control Administrator's Guide*.

16.6 Monitoring End User Experience

The **Business Application** drop down menu, accessible from the Business Application home page, includes the following options for the End User Experience item:

- **End User Experience data**, whose contents are described in [Monitoring End User Experience Data](#).
- **Session Diagnostics**, whose contents are described in [Monitoring End User Experience Data](#).
- **Metrics**, whose contents are described in [Monitoring End User Experience Metrics](#).

As discussed in [Creating Business Applications](#), a business application can consist of an End User Service which corresponds to the application or suite being monitored by

RUEI. If you monitor a End User Service directly, as described in [Monitoring an End User Service](#) the **End User Service** drop down menu, accessible from the End User Service home page, includes the following options for the End User Experience item:

- **Session Diagnostics**, whose contents are described in [Working With Session Diagnostics](#).
- **Metrics**, whose contents are described in [Monitoring End User Experience Metrics](#).
- **User Flows**, whose contents are described in [Monitoring User Flows](#).

16.6.1 Monitoring End User Experience Data

Selecting the End User Experience Data option from the **Business Application > End User Experience** menu, displays a page that includes **Key Performance Indicators**, **Usage Data** and **Violations Data** tabs.

16.6.1.1 Key Performance Indicators

When displayed, the page shows the KPIs and Users Flows data from RUEI.

 **Note:**

In order to view KPI alerts within Incident Manager, you will need to set up a connection between RUEI and the Oracle Enterprise Manager Repository. The procedure to do this is described in [Setting Up a Connection Between RUEI and the Oracle Enterprise Manager Repository](#).

16.6.1.2 Usage Data

When displayed, the page shows the most active users and most executed user requests data from RUEI.

The **Top Users** region enables you to monitor the most active users of the targets associated with the business application. This includes session and page view information, as well as user and application violation indicators. An example is shown in [Figure 16-14](#).

Use this region to verify the performance of the most popular user requests associated with a business application (such as downloads or payment handlings).

Figure 16-14 Top Users

The screenshot shows a web interface with three tabs: 'Key Performance Indicators', 'Usage Data', and 'Violations Data'. The 'Usage Data' tab is active, displaying a table titled 'RUEI - Top Users'. The table has a 'View' dropdown and a 'Row Limit' set to 10. The table columns are: User ID, Sessions, Page (Views, Load Time (sec)), Total, User (%), and Application (%). The data shows 'blah' as the top user with 25 sessions, 2989 views, and a load time of 0.99 seconds. Other users have 1 session each and 408 views.

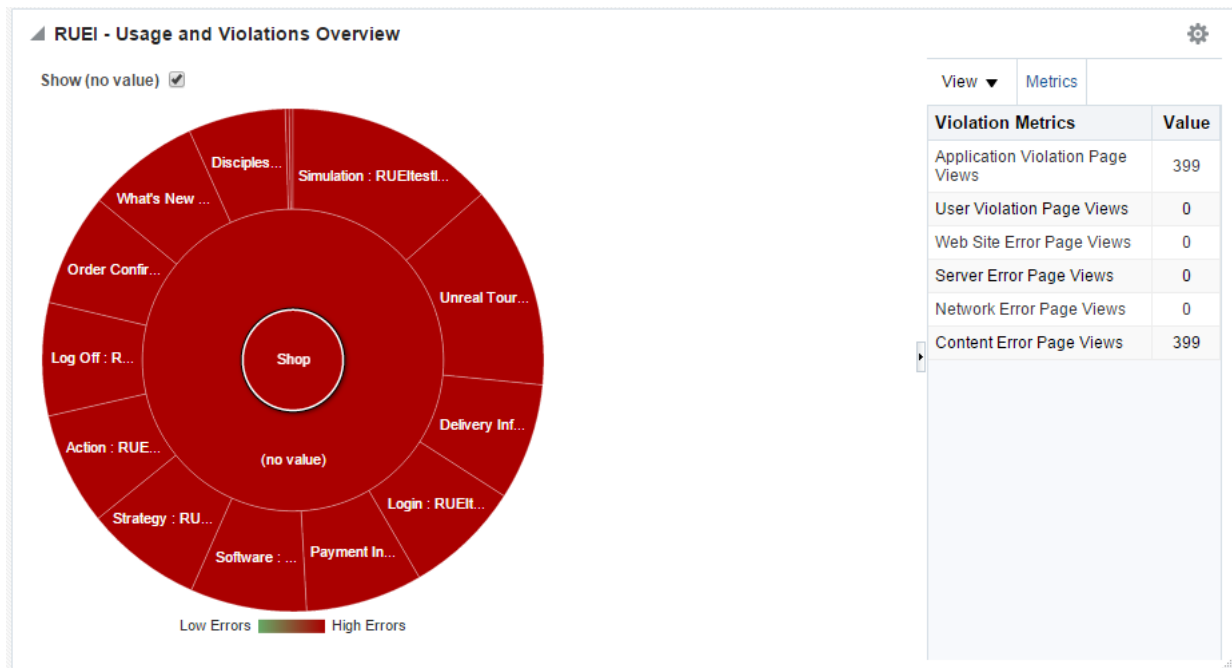
User ID	Sessions	Page		Total	User (%)	Application (%)
		Views	Load Time (sec)			
blah	25	2989	0.99	0	0.00	0.00
elliopleach	1	408	1.63	0	0.00	0.00
connorlgough	1	408	0.61	0	0.00	0.00
ameliaflane	1	408	1.73	0	0.00	0.00
alicejnorton	1	408	1.90	0	0.00	0.00
danieladark	1	408	1.67	0	0.00	0.00
leahmryan	1	408	1.71	0	0.00	0.00
yasminzmann	1	408	1.81	0	0.00	0.00
bradleyhwebster	1	408	0.63	0	0.00	0.00
kieranmwest	1	408	0.73	0	0.00	0.00

Selecting a user opens the RUEI Session Diagnostics facility and displays detailed information about the selected user. For more information, see [Working With Session Diagnostics](#).

16.6.1.3 Violations Data

When displayed, the page shows a sunburst chart to illustrate the violations data from RUEI. An example is shown in [Figure 16-15](#).

Figure 16-15 Usage and Violations Overview



The color of a segment indicates the number of violations that are associated with that page. Hover your mouse pointer over a segment to display the exact number or violations associated with the page. Click on the page to use the session diagnostics facility, see [Working With Session Diagnostics](#), Double click on a page to explore the next level of detail available.

16.6.2 Working With Session Diagnostics

The Session Diagnostics facility allows you to perform root-cause analysis of operational problems. It supports session performance breakdown, including the impact of failing pages and hits on sessions, the full content of each failed page, and the relationship between objects, page views, and sessions. Moreover, it offers the opportunity to track exactly what error messages visitors to the monitored website receive, and when. With this ability to recreate application failures, you can identify and eliminate annoying or problematic parts of your web pages.

This section explains the use of the Sessions Diagnostics facility. It covers the following topics:

- [Creating an Enterprise Manager User for Session Diagnostics](#)
- [Getting Started with Session Diagnostics](#)
- [Customizing Session Diagnostics Reporting](#)
- [Exporting Full Session Information](#)
- [Exporting Session Pages to Microsoft Excel](#)

16.6.2.1 Creating an Enterprise Manager User for Session Diagnostics

With Oracle Enterprise Manager 13c, a new role called EM_EUS_DIAGNOSTICS exists to allow users access session diagnostics data. To create a user with this role:

1. From the **Setup** menu, select **Security**, then select **Administrators**.
2. Click **Create** to create a new user.
3. On the **Properties** page, enter a username and password.
4. On the **Roles** page, add the EM_EUS_DIAGNOSTICS role to the new user definition.
5. On the **Target privileges**, first add the **RUEI Reporter engine** target, then grant the **configure_target** privilege.
6. Review and complete creating the new user.
7. Log out of Enterprise Manager and log in as the new user you have just created.
8. To test the user access, use session diagnostics as described below and review the payload message for session activity (camera icon).

16.6.2.2 Getting Started with Session Diagnostics

To locate the diagnostics information you require, do the following:

1. This item is available on both the Business Application home page and the End User Service home page. Select **End User Experience** and then **Session Diagnostics**.
2. Use the **View Data** menu to select the required period. Note that the availability of session diagnostics information is determined by the Statistics and Session Diagnostics data retention policy settings specified for the associated RUEI instance. For more information, see the *Oracle Real User Experience Insight User's Guide*.

Figure 16-16 Session Diagnostics

The screenshot shows the 'Session Diagnostics' interface. At the top, there are search filters for 'ADF Component Type', 'ADF Taskflow Code', and 'ADF View ID', each with a 'Starts With' dropdown and a search input field. There is also a 'User ID' search field. A 'Saved Search' dropdown is set to 'ADF Framework'. Below the filters are 'Search', 'Reset', 'Save...', and 'Add Fields' buttons. The main data table has the following structure:

Period	User ID	Client Location	Info
04/11/15 06:30 - 04/11/15 06:31	(no value)	Client Country (no value) Client City (no value) Client IP 226.164.159.221	Page Load Time per Page (sec) 1.1 Application Violation Page Views 0 Client Aborts 0 Content Errors 0 Content Notifications 0 Frustrated Page Views 0 Network Errors 0 Page Views 7 Server Errors 0 User Violation Page Views 0 Violation Page Views 0 Web Site Errors 0
		Client Country (no value)	Page Load Time per Page (sec) 3.4 Application Violation Page Views 0 Client Aborts 0 Network Errors 0 Page Views 12 Server Errors 0

3. Specify the appropriate search criteria to locate the required user record(s). The available default search criteria are controlled by the RUEI instance configuration

(described in [Customizing Session Diagnostics Reporting](#)). You can click **Add Fields** to make additional search criteria available. Be aware that while the use of wildcard characters (*) is supported, all other search characters are treated as literals. Also, *all* criteria specified for the search must be met for matched user records to be reported.

You can specify multiple values for a single dimension by clicking **Add Fields**, and selecting the required dimension. In this case, only *one* of the specified values needs to be found in order for a match to be made.


After updating the appropriate search filters, you can save the search combination by clicking **Save**. Note that changes to saved searches can influence the available fields within the **Add Fields** facility. In addition, the predefined list of available dimensions is based on the business application definition. For example, only oracle Fusion-specific dimensions are available if the business application is defined as a Oracle Fusion suite.

When ready, click **Search**. The results of the search are shown in the lower part of the page.

If you specify an ECID as the search criteria and the ECID value refers to a spurious hit then no results are displayed if you also specify a filter relating to a page property, because there is not a page associated with the spurious hit.

4. Click the user record of interest from the displayed list. Information like the one shown in [Figure 16-17](#) is displayed.

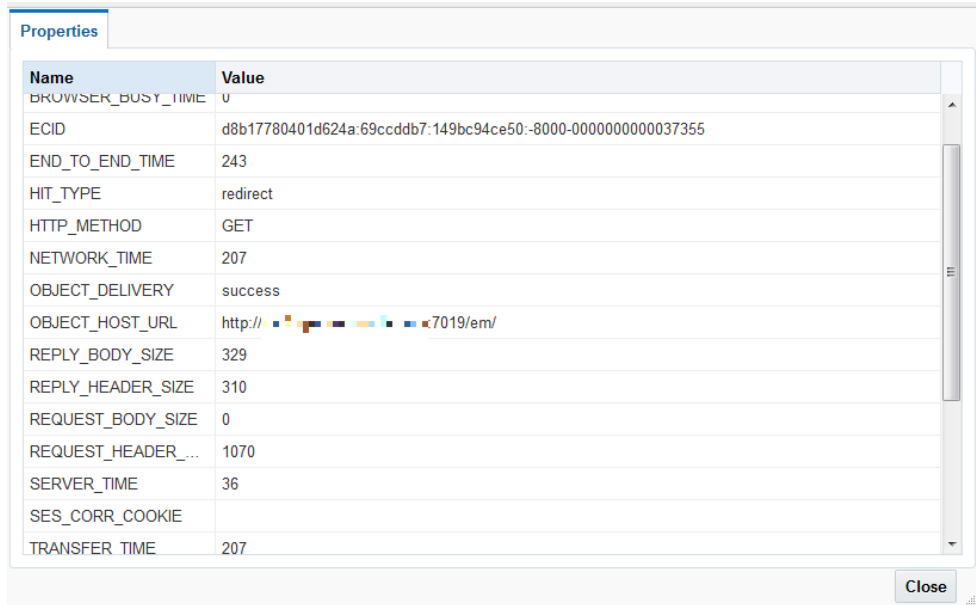
Figure 16-17 Example Session Activity Listing




5. The overview shows the pages and actions recorded within the selected user record. You can hover your mouse over an icon to see a tooltip. Icons indicate slow or failed objects, page-loading satisfaction, whether replay content is available, and whether clickout is available to JVM Diagnostics to provide activity information. (Clickout capability is shown by the Oracle icon.) The camera replay can show a pop-up with full contents of the hit. This allows access to the original html in full detail. The icons shown in the info column are described under [Figure 16-23](#), for example icon 1 provides a link to associated logs to help debug an issue.

- You can click a page or object within the selected user session to open a window with detailed technical information about it. An example is shown in [Figure 16-18](#).

Figure 16-18 Page Properties Window



The screenshot shows a 'Properties' window with a table of technical details. The table has two columns: 'Name' and 'Value'. The data is as follows:

Name	Value
BROWSER_BUSY_TIME	0
ECID	d8b17780401d624a69ccddb7-149bc94ce50-8000-0000000000037355
END_TO_END_TIME	243
HIT_TYPE	redirect
HTTP_METHOD	GET
NETWORK_TIME	207
OBJECT_DELIVERY	success
OBJECT_HOST_URL	http://  7019/em/
REPLY_BODY_SIZE	329
REPLY_HEADER_SIZE	310
REQUEST_BODY_SIZE	0
REQUEST_HEADER_...	1070
SERVER_TIME	36
SES_CORR_COOKIE	
TRANSFER_TIME	207

A 'Close' button is located at the bottom right of the window.

When replay content is available an icon is displayed in the Session Activity Listing as shown in [Figure 16-17](#). You can click on the icon to open a pop up showing the Replay Content that was recorded. An example is shown [Figure 16-19](#).

Figure 16-19 Full Session Replay

The screenshot displays the 'Full Session Replay' window in 'Source View'. It is divided into three main sections:

- Request Headers (9):** A table listing headers such as Host, Connection, TE, User-Agent, Accept, Accept-Language, x-oracle-slm-message-id, Accept-Encoding, and ECID-Context.
- Response Headers (11):** A table listing headers such as Date, Server, X-Powered-By, Set-Cookie, Expires, Cache-Control, Pragma, Vary, Content-Encoding, Content-Length, and Content-Type.
- Response Content:** A text area showing the raw HTML of the response, including DOCTYPE, head, meta, and title tags.

Click the **Replay View** link in the upper right corner of the screen to be redirected to the RUEI server, where you see the browser view of the Session Replay Content. You might be required to log in to the RUEI server.

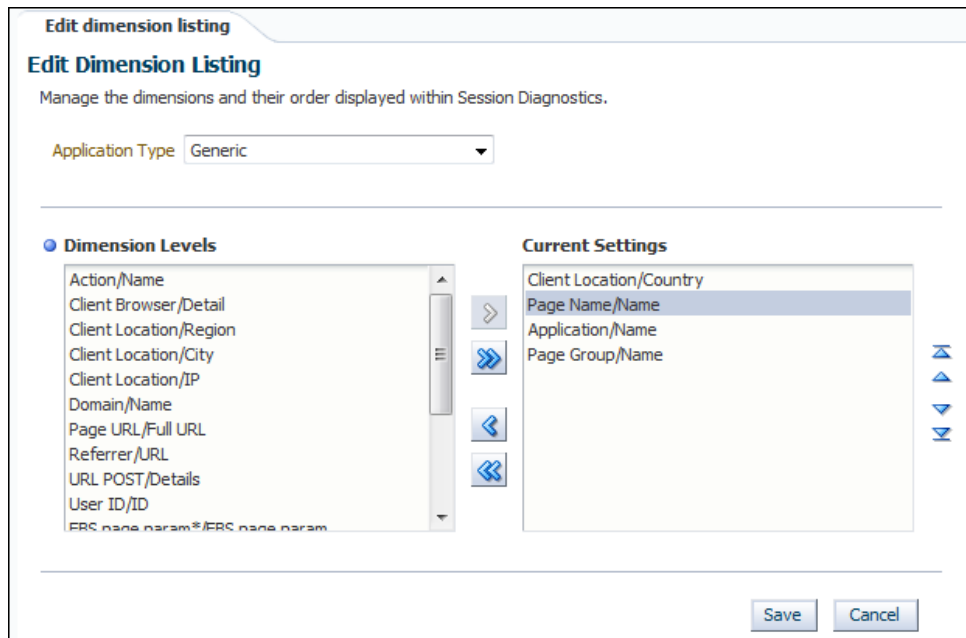
7. The list of matched user sessions shown in [Figure 16-16](#) is based upon the period selected in the **View** menu. For example, if the period "Last hour" is selected, the list of matched user sessions is based on sessions that were active during that period. However, they may have started or finished outside this period. For this reason, you can use the slider at the bottom of [Figure 16-17](#) to restrict the displayed page views and actions to a more specific period.
8. Optionally, click **Export as Zip** to export the session's complete contents to external utilities for further analysis (described in [Exporting Full Session Information](#)) or **Export as XLS** to export a summary of the pages within the session (described in [Exporting Session Pages to Microsoft Excel](#)).

16.6.2.3 Customizing Session Diagnostics Reporting

You can control the specific dimensions reported in Session Activity part of the Session Diagnostics for applications, suites and services. To do so:

1. From the **Setup** menu, select **Middleware Management**, then select **Setup**. The currently registered RUEI instance is shown in the **RUEI Systems** row of the page shown in [Figure 16-2](#).
2. Select the required RUEI system. Click **Configure**. The **Edit Dimension Listing** page shown in [Figure 16-20](#) is displayed.

Figure 16-20 Edit Dimension Listing Page



3. Use the **Application Type** menu to select whether you want to modify the dimension listings for generic applications (that is, applications that are not suite-based), services, or suites. If the latter, you will need to specify the suite type.
4. Use **Move** and **Remove** to select the dimensions that should be listed. Once selected, you can control the order in which the item appears in the list. When ready, click **Save**.

16.6.2.4 Exporting Full Session Information

In addition to viewing session information, you can also export complete session contents to external utilities for further analysis or integration with other data. For example, you could use complete real-user sessions as the basis for test script generation. Test platforms, such as Oracle Application Testing Suite (ATS), can easily be configured to generate automated test scripts for an application's most common usage scenarios.

In addition, this facility can also be used to support root-cause analysis. Complete user session information can be provided to application or operations specialists to help identify unusual or difficult to isolate issues. Sensitive information within the exported data is masked according to the actions defined in the HTTP protocol item masking facility. This is described in the *Oracle Real User Experience Insight User's Guide*.

To export session information:

1. Locate the required session, and click **Export as Zip**.
2. Depending on how your browser is configured, you are either prompted to specify the location to which the zip file should be saved, or the session is immediately saved to the defined default location.

Important

In order for the session export files to be created correctly, you should do the following:

- Ensure that the requirements for exporting session information described in [Prerequisites and Considerations](#) have been met.
- Verify the exported content files (described in the following section) are present before attempting to import an exported RUEI session into an external utility.

Understanding the Structure of the Exported Data

The exported session zip file contains the following files:

- `data.tab`: contains the direct (raw) hit information for the selected session extracted from the Collector log file.
- `page.tab`: contains the direct (raw) page information for the selected session extracted from the Collector log file.
- `content_hitno.tab`: contains the complete (raw) content information for the indicated hit. There is a file for each hit within the `data.tab` file that has content. For example, if the third and sixth hits had content available for them, two files would be created: `content_3.tab` and `content_6.tab`.

Viewable versions of the files cited in the hit file are also available under the `content_viewer` directory. This means that data transferred with chunked encoding can be immediately viewed. Note that the same `hitno` as in the `data.tab` file is used in their file naming.

- `index.html`: allows developers and other interested parties outside RUEI to view and analyze session details as they would appear within the Session Diagnostics facility, with access to source, page and object details, and element identification.

Note:

The log files used as the basis for creating exported session files are also used internally by RUEI. The format and contents of these files is subject to change without notice.

16.6.2.5 Exporting Session Pages to Microsoft Excel

You can export a summary of the pages within the currently selected session to Microsoft Excel.

1. Locate the required session, and click **Export as XLS**. Depending on how your browser is configured, you are either prompted to specify the tool with which to open the file directly (by default, Microsoft Excel), or the session is immediately saved to the defined default location.
2. Within Microsoft Excel, you can view and edit the generated file. The exported page view history and session summary can be used to compile sets of real-user sessions that could be used as the basis for testing or performance analysis.

Controlling Row Creation and Ordering

Be aware that the rows that appear in the Microsoft Excel export are based on the currently specified RUEI configuration. This is described in [Customizing Session Diagnostics Reporting](#).

16.6.3 Monitoring End User Experience Metrics

As part of Business Application and End User Service monitoring, the End User Experience Metrics page presents a useful overview of user-selectable metrics within a given timespan. These metrics can be counts (for example, page views) or aggregate values (such as, median page load time).

To view the Metrics page, select **End User Experience** and then **Metrics** from the **Business Application** or **End User Service** menu.


This page allows you to select performance metrics and view their associated average and median data graphically.

The time-period for the data can be set and search filters similar to those on the Session Diagnostics Page allow you to further refine the data returned. Filter settings can be saved for subsequent use. The metrics are displayed in two tabs, **Aggregation** shows the metric aggregated over time and **Instances** shows individual events.

Aggregation

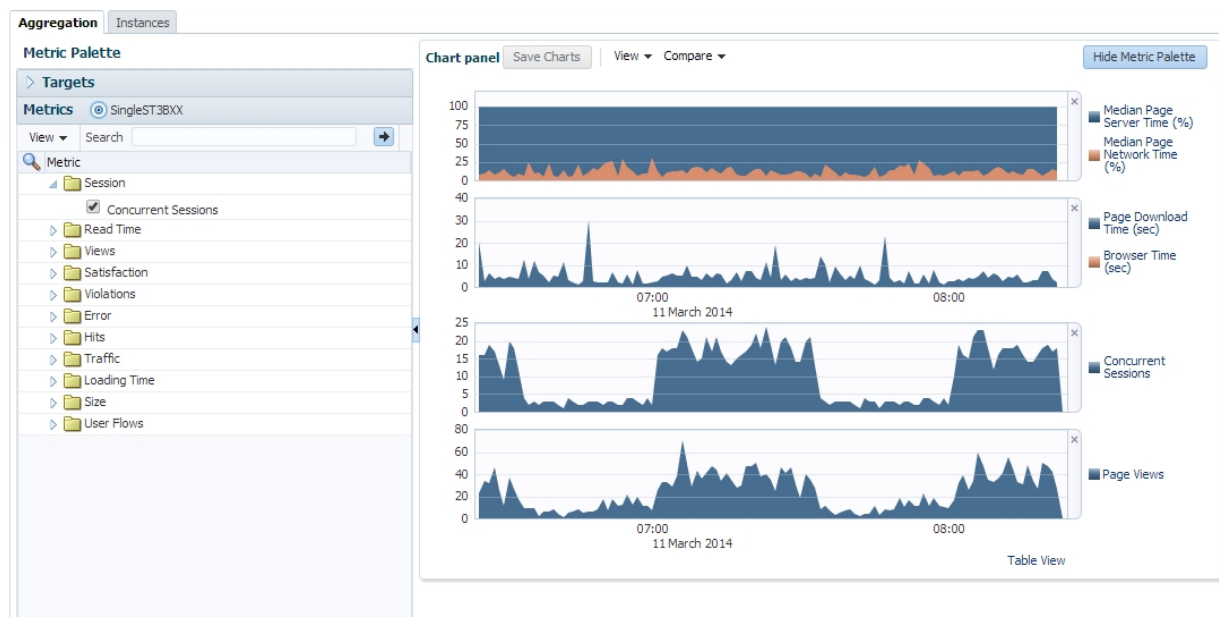
On the **Aggregation** tab the selectable metrics are arranged in an hierarchical tree palette. The list displays a set of items that is appropriate for the configured business application type. This set can expand to include system metrics. For example, if the application associates to a WebLogic server, JVM metrics become available in addition to RUEI metrics. An example of a metrics palette is shown in [Figure 16-21](#).

Figure 16-21 Metric Palette



Note that the individual metrics graphs can be combined into one chart using the graph toolbar. Also some of the listed graphs show data from different parts of the metric palette combined into one chart. In [Figure 16-22](#), the top chart shows graphs for both **Median Page Server Time** and **Median Page Network Time**.

Figure 16-22 Sample Metrics Graph



Incident Manager allows you to navigate directly from a KPI event to the Metrics page. In this case, the Metrics page is populated with time and filter settings relevant to the context of the KPI event. Therefore, you can inspect relevant metrics around the offending event, be it in time-period or in filters broader than those that correspond to the original KPI event.

From the **Metrics** page there is a direct link to the session diagnostics facility. When you click this link, the search properties and time-span will be re-used to find sessions that match the active criteria.

Note:

Selecting median values over a large timespan can impact performance.

Instances

The instances tab displays events filtered using the criteria you set, for example Client Browser = Chrome. If you specify an ECID as the criteria and the ECID value refers to a spurious hit then no results are displayed, because there is not a page associated with the spurious hit. For each item in the results, you can:

- Review the Session Activity, timing and user information (if available). Where applicable, you can expand and collapse items to display further detail, for example specific URLs associated with the session activity, or page attributes.
- Display session information for the event, see [Working With Session Diagnostics](#).
- Review the icons (numbered 1 to 5) for each item where:

Icon 1 provides a link to the log viewer, showing logs associated with the current item.

Icon 2 provides a link to Request Instance Diagnostics, that is, JVM Diagnostics for the current item.

Icon 3 indicates page loading satisfaction, a tooltip appears displaying the satisfaction level, for example **Satisfied**.

Icon 4 provides a link to replay content (for example, the original html if available).

Icon 5 provides a link to further session diagnostics, see [Working With Session Diagnostics](#).

Some of the icons shown are the same as those for Session Diagnostics, see [Figure 16-16](#).

Figure 16-23 Incident Icons



16.6.4 Monitoring User Flows

To view the User Flows page, select **End User Experience** and then **User Flows** from the **End User Service** menu.

This page allows you to view user flows defined for the associated RUEI suite/application. User flows indicate the time components when users were active, average page-load time within user flows, idle time, and when users were outside the user flow. For further information, see the *Oracle Real User Experience Insight User's Guide*.

16.6.5 Monitoring Logs

To view detailed logs, navigate to a session as described in "[Working With Session Diagnostics](#)." Then click the icon labelled 1 in [Figure 16-23](#). This displays a screen similar to [Figure 16-24](#).

This page allows you to view log details and search using criteria including ECID, time and session diagnosis ID.

Figure 16-24 Log Viewer

The screenshot shows the Log Viewer interface with the following elements:

- Search Section:**
 - Search mode: Selected Fields, All Fields
 - Match mode: All, Any
 - Message filter: contains [] Add Fields
 - ECID filter: contains [d8b17780401d624a:69ccddb7:149bc94ce50--8000-0] ✖
 - Search button
- View Section:**
 - View: [] Show: Messages [] View Related Messages: [] Export Messages to File: []
- Table:**

Time	Message Type	Message ID	Message
Nov 19, 2014 6:13:50 PM	PS Trace	BEA-000...	ServletContext@2147159989[app.emgc module:/em path:/em spec-version:2.5] ChainedSecurityM...
Nov 19, 2014 6:13:50 PM	PS Error	ADFC-56...	ADFC: Required task flow input parameter 'rueiVersion' not passed into task flow '/WEB-INF/apm/b...
Nov 19, 2014 6:13:50 PM	PS Error	ADFC-56...	ADFC: Required task flow input parameter 'btmVersion' not passed into task flow '/WEB-INF/apm/b...
- Summary:**
 - Rows Selected: 1
 - Columns Hidden: 21
- Diagnostic Information:**
 - message Level: 1
 - Diagnostic Session ID: 0000KbAyQ122jKP5fk3yf1KQPvh0000E5
 - Relationship ID: 0
 - Component: EMGC_OMS1
 - Address: 10.245.29.1
 - User: SYSMAN
 - Thread ID: [ACTIVE] ExecuteThread: '24' fc
 - ECID: d8b17780401d624a:69ccddb7:1

16.7 Monitoring an End User Service

With Enterprise Manager 13c, there is a new service level target, called an End User Service (EUS) corresponding to each RUEI Application/Suite. The EUS has the same name as the RUEI Application/Suite, and can be used to create a business application as described in [Creating Business Applications](#). It is also possible to monitor an End User Service directly:

1. From the **Targets** menu, select **Services**. The currently defined services are listed. Filter for end user services. Alternatively, you can navigate to the End User Service from the Business Application home page by clicking on the **Sub Services** tab. The resulting list may include End User Services if the Business Application includes one.
2. If Fatal/Critical KPI is associated with an application defined in RUEI system, EUS status is also affected by status of associated Critical/Fatal KPI.

If KPI status is changed to down, EUS status will also be changed to down. If Status is up in Key Components, but EUS status is down, refer:

Name	Type	Status	Availability	Service Level Agreement Status	Incidents	System	Key Components				Key Tests		
							Status	Incidents	Status	Monitors			
vux535_KMS_rest_ap4	End User Service	↑	-	-	-	vux535_Oracle Real User Experience Insight	↑2	0	0	0	0	n/a	0
vux535_ose-rue	End User Service	↑	-	-	-	vux535_Oracle Real User Experience Insight	↑2	0	0	0	0	n/a	0
xru_Creater	End User Service	↑	System	-	-	xru_Oracle Real User Experience Insight	↑1	0	0	0	0	n/a	0
vux535_testapp0998	End User Service	↑	-	-	-	vux535_Oracle Real User Experience Insight	↑2	0	0	0	0	n/a	0
vux535_sabel	End User Service	↑	-	-	-	vux535_Oracle Real User Experience Insight	↑2	0	0	0	0	n/a	0
xru_Toyco	End User Service	↑	System	-	-	xru_Oracle Real User Experience Insight	↑1	0	0	0	0	n/a	0
vux535_ZiekenOmroep	End User Service	↓	-	-	-	vux535_Oracle Real User Experience Insight	↑2	0	0	0	0	n/a	0
vux535_WebService_dummy-webst1000	End User Service	↑	-	-	-	vux535_Oracle Real User Experience Insight	↑2	0	0	0	0	n/a	0
vux535_Toyco	End User Service	↑	-	-	-	vux535_Oracle Real User Experience Insight	↑2	0	0	0	0	n/a	0
vux535_KMS_rest_ap89	End User Service	↑	-	-	-	vux535_Oracle Real User Experience Insight	↑2	0	0	0	0	n/a	0
vux535_KMS_rest_ap1	End User Service	↑	-	-	-	vux535_Oracle Real User Experience Insight	↑2	0	0	0	0	n/a	0
vux535_RUE899	End User Service	↑	-	-	-	vux535_Oracle Real User Experience Insight	↑2	0	0	0	0	n/a	0
vux535_ADPStoreFrontModule	End User Service	↑	-	-	-	vux535_Oracle Real User Experience Insight	↑2	0	0	0	0	n/a	0

You can check the Fatcat/Critical KPI in **Incident Manager** page,

Severity	Summary	Target	Priority	Status	Age	Time Since Last Update	Owner	Ackn	Escal	Type	Category
Critical	RUEI KPI ZiekenOmroep 1 has failed	vux535_ZiekenOmroep	None	New	19 hours...	19 hours 2 min...	-	No	No	Incident	Load
Critical	The End User Service is down due to critical or fatal KPI alerts happened	vux535_ZiekenOmroep	None	New	19 hours...	19 hours 15 min...	-	No	No	Incident	Availability

3. Click the End User Service of interest. The home page for the selected end user service is displayed. This is very similar to the Business Application home page shown in [Figure 16-13](#).

The **End User Service** drop down menu, accessible from the End User Service home page, includes the following options for the **End User Experience** item:

- **Session Diagnostics**, whose contents are described in [Working With Session Diagnostics](#).
- **Metrics**, whose contents are described in [Monitoring End User Experience Metrics](#).
- **User Flows**, whose contents are describe in [Monitoring User Flows](#).

16.7.1 Troubleshooting an End User Service

If you encounter any issues with data from an end user service, consider the following list of checks to troubleshoot the issue:

1. Refresh the page. If the metrics displayed do not correspond to the metrics configured in RUEI, click refresh the refresh icon on the Enterprise Manager page.
2. Synchronize the RUEI system target. Navigate to the appropriate RUEI system home page and choose **Synchronization** from the **Configuration** menu item.

16.8 Monitoring KPI and SLA Alert Reporting

This section explains the KPI-related information that is available for both RUEI applications.

 **Note:**

To monitor KPI and SLA alerts you must first complete all the steps described in "[Setting Up a Connection Between RUEI and the Oracle Enterprise Manager Repository.](#)"

The alerts generated by KPIs defined for the applications, suites, and services, as well as for the SLAs for the transactions that comprise your business applications are reported as events in **Incident Manager**. To view these events:

1. From the **Enterprise** menu, select **Monitoring**, and then **Incident Manager**.
2. Open the **Events Without Incidents** predefined view.
3. Click the event of interest to view more information about it.

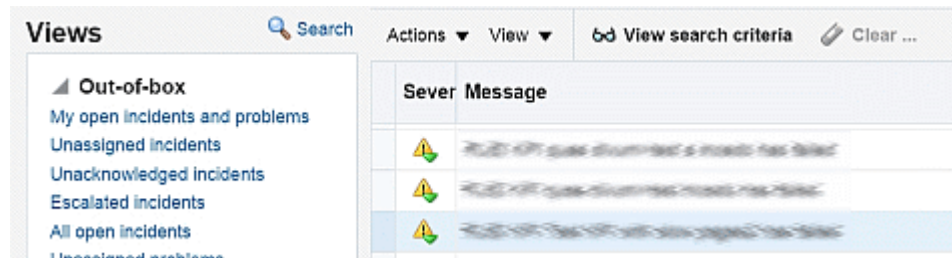
Event detail information varies depending on whether the event is based on a RUEI KPI. This is described in the following sections.

For RUEI related events, you can also access the **Events Without Incidents** view from the home page of a business application using the **Business Application** menu. Select **Monitoring**, then **Incident Manager**, and **Events without Incident**. This option shows events in the context of the selected Business Application.

RUEI Event Detail


After accessing the **Events Without Incidents** view, you will see events listed.

Figure 16-25 Incident Manager



When you click on the RUEI event, you'll see event details at the bottom of the screen in the **Guided Resolution** region.

Figure 16-26 Accessing RUEI Metrics from Incident Manager



Select RUEI metrics to drill down to the **RUEI Metrics** page in context of the filters that were set up for the KPI as well as in the time frame of the KPI violation. The **RUEI Metrics** page will also show the metric that is the basis for the KPI definition.

The status of the KPIs defined for the applications, suites, and services that comprise your business applications are reported in the RUEI - Key Performance Indicators (KPIs) tab.

This provides information about the Business Application associated with the KPI, as well as the metric upon which the KPI is based. Note that for ease of management, KPIs within RUEI are grouped into categories that can be customized to contain related performance indicators. For example, separate categories could be defined for business and IT-related issues, such as user flow completion, visitor traffic, website availability, and so on.

 **Note:**

In order to view KPI alerts within Incident Manager, you will need to set up a connection between RUEI and the Oracle Enterprise Manager Repository. The procedure to do this is described in [Setting Up a Connection Between RUEI and the Oracle Enterprise Manager Repository](#).

16.9 Upgrading End User Service

When upgrading from Enterprise Manager Cloud Control 13.2 to Enterprise Manager Cloud Control 13.3, you must perform the following steps:

1. Remove the RUEI target from Enterprise Manager 13.3c.
2. Re-register the RUEI target with Enterprise Manager 13.3c.
3. Verify same functionality as original End User Services in [Monitoring an End User Service](#).

 **Note:**

If you did not create Business Applications in your EM system before upgrade, you can Re-register RUEI target on EM side. But, if you have created Business applications before upgrade, refer to [Upgrading Business Applications](#).

16.10 Upgrading Business Applications

Use the following procedure if you are upgrading Business Applications from Enterprise Manager 12c to Enterprise Manager 13c or Enterprise Manager 13.2c to Enterprise Manager 13.3c , and that Business Application has a RUEI connection:

1. Document the details of Enterprise Manager association information with RUEI (the Business Application) before upgrading to Enterprise Manager.
2. Remove the RUEI target from Enterprise Manager.

3. Re-register the RUEI target with Enterprise Manager to ensure the discovery of End User Services targets.
4. Create a new Business Application using the details from step 1.
5. Verify that the Business Application provides the same functionality as the original Business Application.
6. Delete the original Business Application before upgrade.

17

Monitoring End-to-end Performance

This section describes how you can use the application monitoring components to identify the underlying cause of poor user experience. It then poses a series of questions to test your understanding of end-to-end monitoring.

This chapter includes the following sections:

- [Troubleshooting: A Case Study](#)

The demonstration uses the stand-alone versions of RUEI.

You can view a live demonstration of the case study described in this chapter by navigating to the following site:

http://apex.oracle.com/pls/apex/f?p=44785:24:0::NO:24:P24_CONTENT_ID,P24_PREV_PAGE:5781,1#prettyPhoto/0/

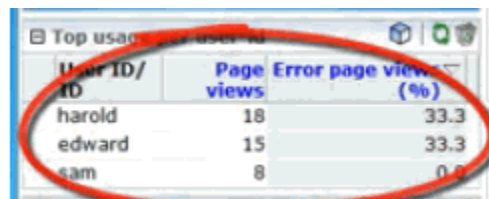
17.1 Troubleshooting: A Case Study

This demonstration aims to traverse all the functional layers of a distributed application. Only partial views of screens are shown.

Looking at the User Experience

Our investigation begins with the RUEI dashboard, the first place to review the overall user experience.

Looking at the **Top usage by User ID** panel, we note a very high percentage of Error page views for the users Harold and Edward:



User ID/ ID	Page views	Error page views (%)
harold	18	33.3
edward	15	33.3
sam	8	0.0

To get more details about this situation, we select to display browser data by clicking on the cube icon in the upper-right hand corner.

From the **Browser data** display, we select a user and select user diagnostics to get session diagnostics for the user, filtering on a specific application, in this case the Toyco application.

Filter on	Value
▼ User ID/ID	harold
▼ Application/Name	Toyco

Session diagnostics

Search user records for the specified period using the available criteria. All strings are regarded as literal strings.

Search filters

Application/Name:

User ID/ID:

Client location/IP:

Add more filters

Dimension level:

Value:

Dimension level	Value
No filters	

Search result order

- Session start time
- Most active sessions
- Fastest sessions
- Slowest sessions
- Shortest sessions
- Longest sessions
- Most failure sessions

Next, we retrieve session information for the user for a given time period. The results are displayed in the **Session diagnostics** pane:

Session diagnostics

Search user records for the specified period using the available criteria. All strings are regarded as literal strings.

Order: Session start time ▼ Dimension level: « Select » ▼ Value: « Select »

Period/Hour	User ID/ID	Client network/IP
<input type="button" value="🔍"/> 15:00 - 16:00	harold	144.25.146.189
<input type="button" value="🔍"/> 15:00 - 16:00	harold	144.25.146.189
<input type="button" value="🔍"/> 15:00 - 16:00	harold	144.25.146.189
<input type="button" value="🔍"/> 15:00 - 16:00	harold	144.25.146.189
<input type="button" value="🔍"/> 15:00 - 16:00	harold	144.25.146.189
<input type="button" value="🔍"/> 15:00 - 16:00	harold	144.25.146.189
<input type="button" value="🔍"/> 16:00 - 17:00	harold	144.25.146.189
<input type="button" value="🔍"/> 16:00 - 17:00	harold	144.25.146.189
<input type="button" value="🔍"/> 16:00 - 17:00	harold	144.25.146.189
<input type="button" value="🔍"/> 16:00 - 17:00	harold	144.25.146.189

We select one of the session listed in the grid view to find out more about the session. Information is displayed in the **Session activity** pane.

Session activity		Page load time (s)	Info
Toyco	Toyco > Toyco - Purchasing Client	2.7	
Toyco	Toyco > Toyco Order Form Content error > error string: Purchase Failed	120.2	
Toyco	Toyco > Toyco - Purchasing Client	1.0	

We see that one of the load times is excessive and that there's an error listed as well.

We click on the page icon in the **Info** column to view the page as the user saw it. It is shown next. Indeed, at the bottom of the page is the error message "Purchase failed."

Customer:	Hardware Hotel		
Address:	11 Houston St		
City:	Honolulu	State:	HI
Country:	USA		

ID	Product	Item Price	Quantity
1001	Call of Duty: Black Ops for Xbox 360	59.99	
1002	Star Wars: The Force Unleashed II for Xbox 360	58.50	
1003	Kinect Adventures for Xbox 360	69.99	
1004	FIFA Soccer 11 for PlayStation 3	59.95	
1005	Rubix Cube 2011	9.99	

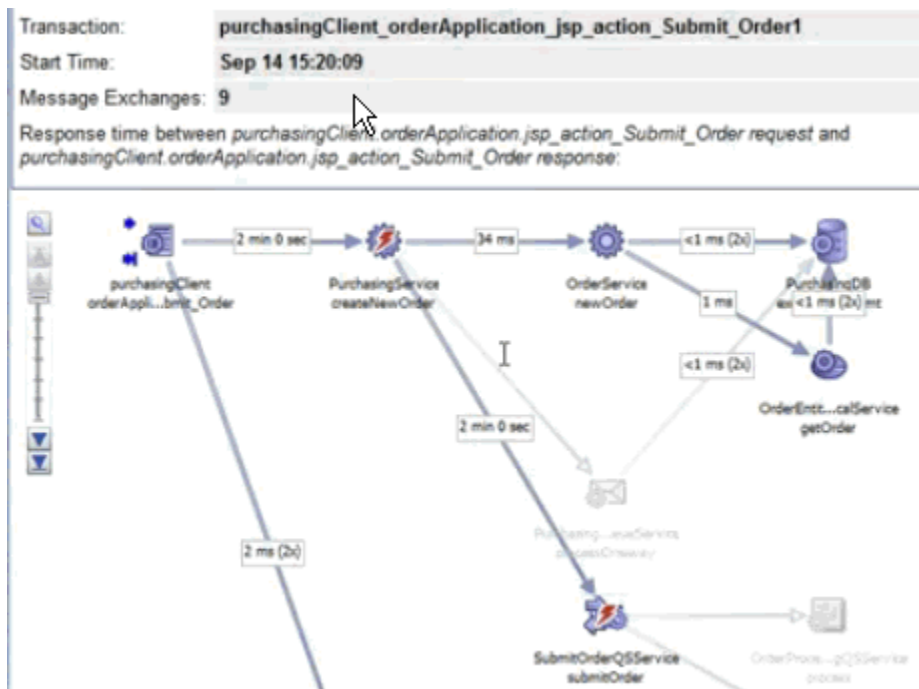
Ship Using:		
Select	Carrier	JMS API Used
<input checked="" type="radio"/>	FedEx	TextMessage (SOAP)
<input type="radio"/>	UPS	TextMessage (non-SOAP)
<input type="radio"/>	DHL	ObjectMessage
<input type="radio"/>	USPS	MapMessage
<input type="radio"/>	SpeedPost	ByteMessage
<input type="radio"/>	Cargo	StreamMessage
<input type="radio"/>	Air	Message

Purchase Failed

The error message in the user view suggests that this is a functional error.

Looking at Diagnose transaction

Selecting the problematic application and selecting **Diagnose transaction** from the context menu displays the **Instance inspector** view.



The red thunderbolt icons identify the failing services.

Suspecting that the call to the database is the culprit, we take a look at the message content, which suggests that the trouble is in the message response.



Choosing to view the XML, we find ourselves in a Java stack trace, and we see a fault string:

```

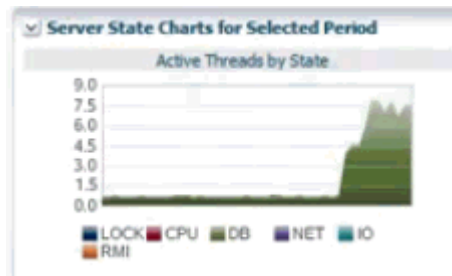
- <soap:Envelope>
- <soap:Body>
- <soap:Fault>
  <faultcode>tns:BEA-382515</faultcode>
  - <faultstring>
    Callout to java method *public static void com.oracle.callouts.CustomSocket.commitSocketOpen
    java.net.SocketTimeoutException: Read timed out at java.net.SocketInputStream.socketRead0(
    java.net.SocketInputStream.read(SocketInputStream.java:182) at com.oracle.callouts.CustomSo
    com.oracle.callouts.CustomSocket.callSocket(CustomSocket.java:42) at com.oracle.callouts.Cu
    sun.reflect.GeneratedMethodAccessor1185.invoke(Unknown Source) at sun.reflect.Delegating
    java.lang.reflect.Method.invoke(Method.java:597) at stages.transform.runtime.JavaCalloutRuntin
    weblogic.security.acl.internal.AuthenticatedSubject.doAs(AuthenticatedSubject.java:363) at web
    weblogic.security.Security.runAs(Security.java:61) at stages.transform.runtime.JavaCalloutRuntin
    com.bea.wli.sb.pipeline.StatisticUpdaterRuntimeStep.processMessage(StatisticUpdaterRuntimeS
    com.bea.wli.sb.pipeline.debug.DebugggerRuntimeStep.processMessage(DebuggerRuntimeStep.ja
    com.bea.wli.sb.stages.StageMetadataImpl$WrapperRuntimeStep.processMessage(StageMetad
    
```

We'll need to drill down to the **Java Virtual Machine Diagnostics** page.

We return to the transaction graph and right-click on the offending operation to get the JVMD view.

Looking at Machine-Level Information

Looking at the **Active Threads by State** graph in the JVMD view, it looks like we have a database problem.



Looking at the **Threads State Transition** display on the same page, we note that a number of threads are stuck.

Thread Name
[STUCK] ExecuteThread: '0' for queue: 'weblogic.kernel.Default (self-tuning)'
[STUCK] ExecuteThread: '2' for queue: 'weblogic.kernel.Default (self-tuning)'
[STUCK] ExecuteThread: '4' for queue: 'weblogic.kernel.Default (self-tuning)'
[STUCK] ExecuteThread: '8' for queue: 'weblogic.kernel.Default (self-tuning)'
[STUCK] ExecuteThread: '5' for queue: 'weblogic.kernel.Default (self-tuning)'
[STUCK] ExecuteThread: '7' for queue: 'weblogic.kernel.Default (self-tuning)'
[STUCK] ExecuteThread: '3' for queue: 'weblogic.kernel.Default (self-tuning)'
[STUCK] ExecuteThread: '6' for queue: 'weblogic.kernel.Default (self-tuning)'
[STUCK] ExecuteThread: '1' for queue: 'weblogic.kernel.Default (self-tuning)'
[STUCK] ExecuteThread: '9' for queue: 'weblogic.kernel.Default (self-tuning)'

Noting that this is a current problem, we select the **Live Thread Analysis** button to get more information.

In the **Live Thread Analysis** display, we see ten threads waiting for the database, with three of them locked.

CPU(%)	JVM CPU Usage(%)	OSR	Memory(%)	Runnable	DB Wait	Lock	Network Wait	IO Wait	RMI Wait	Object Wait	Sleep
0	0	1	34	1	10	3	0	0	0	26	2

We can drill down to the database by selecting the **State (DB Wait)** link. This page shows us the SQL details.



The display confirms our suspicion that the trouble lies with database access.

This ends the troubleshooting session, which traversed all the layers of distributed application performance: from the user layer, to back-end supporting services, to the underlying infrastructure.

17.2 Finding Solutions

See if you can guess the answer to the following questions, which test your understanding of end-to-end performance monitoring.

Is the problem with my application?

The following problems relate either to the user experience or to back-end services.

- *Are users unable to complete a task?*
Look at statistics for user flows in RUEI.
- *Do I have a memory leak?*
Look at heap analysis information in JVMD for a given time period.
- *Am I getting out-of-bounds values?*
Check SLA-based alerts defined for RUEI.

Is the problem with deployment architecture?

- *Do I need to replicate and load-balance services?*
Check high throughput values for transaction links. These might indicate bottlenecks.
- *Do I need a failover scheme?*
Use the Enterprise Manager **Business Applications** page or the **Business Application** home page to check for servers that are often unavailable.

Is the problem with supporting infrastructure?

- *Is a server down or slow?*

Use the Enterprise Manager **Business Applications** page or the **Business Application** home page to check for servers that are often unavailable.

- *Is thread-lock causing services to fail?*

Use the JVM Diagnostics page in Enterprise Manager to get information about executing threads.

- *Is the network slow?*

Look at NetworkWait information in the JVM Diagnostics page in Enterprise Manager.

- *Are any of my routers down?*

If you have included your routers in the definition of your System for Enterprise Manager, you can get information about these in Enterprise Manager.

18

Troubleshooting Middleware Applications Using Enterprise Manager

This section describes various methods to troubleshoot middleware applications using Enterprise Manager.

As the enterprise software world becomes more complex, troubleshooting issues can become more difficult. A quick glance at the documentation for Oracle middleware illustrates the complexity that is possible:

<https://docs.oracle.com/en/middleware/>

This section introduces various methods to troubleshoot applications using Enterprise Manager and associated products that do not require the troubleshooter to have expertise relating to the failed middleware component.

This chapter includes the following sections:

- [Introduction to Troubleshooting Middleware Applications](#)
- [Preparing the Environment to Troubleshoot Applications](#)
- [Configure the Environment to Help Troubleshoot Applications](#)
- [Analyzing Issues Using Enterprise Manager and RUEI](#)
- [Resolving Issues Using Enterprise Manager](#)

18.1 Introduction to Troubleshooting Middleware Applications

This chapter provides an introduction to some of the features of Enterprise Manager that you may not be aware of and documents some techniques that experienced people in Oracle Support have learnt and would like to pass on to you, our customers.

To summarize the methodology used in this guide:


1. Prepare your environment by understanding and preparing your systems to be monitored, as describe in [Preparing the Environment to Troubleshoot Applications](#) and [Configure the Environment to Help Troubleshoot Applications](#).
2. After you have isolated the issue to a single middleware tier, use Enterprise Manager to further isolate and resolve the issue .
3. If necessary, resolve any issues as described in [Resolving Issues Using Enterprise Manager](#).

 **Note:**

Troubleshooting requires that you have an issue that needs to be resolved, this might require that you artificially load the system to produce monitoring results in Enterprise Manager. This is especially true for Java loads, however, the techniques required for producing these loads is outside of the scope of this guide.

For example, the environment might consist of a web application, which is using information from various servers.

Figure 18-1 Example Environment



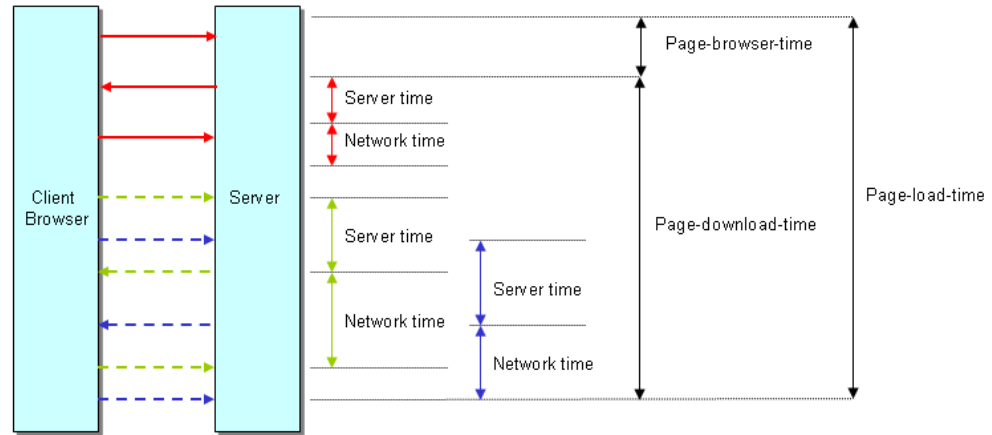
A typical issue might be expressed by a user encountering a difficulty with a specific aspect of the web application. Occasionally, it is obvious which middleware tier is responsible and how to resolve the issue. However, with the complexity of modern applications, and the specialization of personnel, it is more productive to analyze the system using monitoring software that allows you to discover the root cause of an issue and to put the correct fix in place. For example, monitoring a production environment might reveal an issue that is unique to that environment and that could not be revealed by analysis of individual middleware tiers or that could not be reproduced in a test environment.

The diagram is a simple view of a system and does not include Oracle Coherence. Coherence as well as Managed Coherence (Coherence running in a WebLogic

domain) includes a new feature, Coherence Heat Map, which together with Coherence Log Viewer/Search can help isolate middleware issues.

Also, note that a performance issue can be related to items outside of the control of Enterprise Manager, therefore it is important to understand the timing involved in delivering your application before attempting to resolve performance issues.

Figure 18-2 Typical Timing for a Web Application



18.2 Preparing the Environment to Troubleshoot Applications

The amount of time and effort required to troubleshoot applications can be reduced by performing certain steps to prepare the environment. This section outlines additions to the environment that can help the troubleshooting process. You may have already performed the steps in this section, or some of the steps might not be possible in your environment. In this case, skip to [Configure the Environment to Help Troubleshoot Applications](#).

The following table outlines the situations when you can install extra software to prepare the environment:

Table 18-1 Installation Options

Environment Configuration	Installation Option
All	Install Management Agents on all nodes.
Web Application	Install Oracle Real User Experience Insight (RUEI).
Java Applications	Deploy JVM Diagnostics (JVMD) agents.

18.2.1 Document the Topology of the Systems in the Environment

The scope of this task varies depending on the systems that you need to monitor. It is important to monitor all components of the system that can affect performance, including any external services that your system relies on.

Enterprise Manager can automatically provide complete topology documentation for a typical modern environment. If all the systems involved are provided by Oracle, you can be sure that Enterprise Manager is capable of discovering and monitoring all the components. Enterprise Manager can also be used to monitor many non-Oracle systems, for example, using the Oracle System Monitoring Plug-in for Microsoft Active Directory.

However, many systems include components that are unique, for example a system that still relies on a call to a legacy mainframe system. Before you can troubleshoot a complex system, it is important to understand each component and dependency. In some organizations, dependencies can be forgotten as staff changes, so it is important to produce a document relating to that system topology to enable troubleshooting to occur in the future. Some examples of Deployment Topologies are:

- MySOACompany Topology with OAM available from:
http://docs.oracle.com/cd/E23943_01/core.1111/e12036/intro.htm
- WebCenter Topology with OAM available from:
http://docs.oracle.com/cd/E23943_01/core.1111/e12037/intro.htm

18.2.2 Install Management Agents on All Systems in the Environment

Oracle Management Agent (Management Agent) is one of the core components of Enterprise Manager Cloud Control that enables you to convert an unmanaged host to a managed host in the Enterprise Manager system. The Management Agent works in conjunction with the plug-ins to monitor the targets running on that managed host.

Therefore, at any point in time, if you want to monitor a target running on a host, ensure that you first convert that unmanaged host to a managed host by installing a Management Agent, and then manually discover the targets running on it to start monitoring them.

To install a Management Agent, use the Add Host Targets Wizard that is accessible from within the Enterprise Manager Cloud Control console, or use EM CLI. Oracle recommends that you use this wizard, or EM CLI, for the mass-deployment of Management Agents.

For more information on installing Management Agents, see the *Installing Oracle Management Agents* chapter of the *Enterprise Manager Cloud Control Basic Installation Guide*.

18.2.3 Install RUEI to Help Troubleshoot Web Applications

The usage of web applications and services continues to grow. This includes not only the use of the Internet as a marketing channel, but also Extranet-based supply chain and back-office integration, and Intranet deployment of internal applications. Increasingly, it also includes the utilization of web services which implement clearly defined business functions. RUEI is designed for measuring, analyzing, and improving the availability and performance of all of these deployment scenarios, allowing you to isolate issues relating to the end user experience. To achieve this, RUEI is capable of performing end user data collection from network traffic, ADF servers and/or data collection using Javascript browser instrumentation, that is, RUEI provides an end-user data monitoring solution.

To install RUEI, use the appropriate installation guide for your version of RUEI from:

<http://www.oracle.com/technetwork/apps-tech/realuserei-091455.html>

 **Note:**

After installing RUEI, most actions relating to monitoring end users can be performed from Enterprise Manager.

18.3 Configure the Environment to Help Troubleshoot Applications

This section describes configuration options that can help improve troubleshooting. The possible configuration options depend on the software installed in the environment. For example, you cannot create a RUEI application if you have not installed RUEI. See [Table 18-1](#) for a list of installation options.

18.3.1 Discover All Targets in the Environment

Discovery refers to the process of identifying unmanaged hosts and targets in your environment. Targets are entities such as host machines, databases, Fusion Middleware components, that can be managed and monitored in Enterprise Manager Cloud Control. Make sure that all hosts and targets are monitored in Enterprise Manager so that troubleshooting issues is as easy as possible.

Use Enterprise Manager to discover all targets in your environment as described in the Enterprise Manager Cloud Control Administrator's Guide.

 **Note:**

If a target (for example, a host) fails and it has not been discovered, as described above, there will be no indication of the source of the issue in the Enterprise Manager console.

18.3.2 Deploy JVM Agents in the Environment

Java Virtual Machine Diagnostics (JVMD) is one of the critical functionalities in Enterprise Manager Cloud Control that enables administrators to diagnose performance problems in Java applications in the production environment. By eliminating the need to reproduce problems, it reduces the time required to resolve these problems.

The JVMD Agent is deployed on the targeted JVM (the one running a production WebLogic Server or other app 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 JVMD Engine (a built-in component of the Oracle Management Services) and the JVMD Agent can be a secure (SSL) or non-secure connection.

To Deploy JVMD Agent on a WebLogic domain, select the "Setup JVMD Agent" option from the Diagnostics sub menu in the Domain menu. To Deploy JVMD agent on other JVMs, download the agent from the Application Performance Management page that is accessible from within the Enterprise Manager Cloud Control console.

For more details on deploying JVMD agents, see Enterprise Manager Cloud Control Basic Installation Guide.

18.3.3 Define Composite Applications to Help Troubleshoot Multiple Tier Applications

Composite applications allow you to group a set of targets according to tier, so that you can distinguish between issues occurring in different tiers.

 **Note:**

A composite application can be defined to represent multiple tiers (for example, WebLogic servers, SOA Suite and Coherence servers), or it could be defined to represent a web application (WebLogic servers and database). This allows different Enterprise Manager users to create different definitions to represent their unique perspective of the environment, for example application administrators might define a different composite application compared to DBAs.

Defining a composite application is described in the *Composite Applications* chapter of this guide.

Using the environment in [Figure 18-1](#) as an example, the composite application might be defined as:

1. Add all WebLogic servers from Service 2 to a composite application in Enterprise Manager. (The managed server becomes a key member of the composite application).
2. Add a service target to the composite application.
3. Define a SLA rule that determines when the composite application is considered to be down.

18.3.4 Define Synthetic Monitoring Beacons in Enterprise Manager

Synthetic Monitoring Beacons are targets that are used to monitor service tests, primarily to measure performance of the service or business function from a different geographic location.

For more details on Beacons, see the *Configuring and Using Services* chapter of the Enterprise Manager Cloud Control Administrator's Guide.

18.3.5 Define Thresholds in Enterprise Manager

There are monitoring situations in which different workloads for a target occur at regular (expected) intervals. Under these conditions, a static alert threshold would prove to be inaccurate. For example, the accurate alert thresholds for a database

performing Online Transaction Process (OLTP) during the day and batch processing at night would be different. Similarly, database workloads can change based purely on different time periods, such as weekday versus weekend. In both these situations, fixed, static values for thresholds might result in false alert reporting.

Advanced Thresholds allow you to define and manage alert thresholds that are either adaptive (self-adjusting) or time-based (static).

Adaptive Thresholds are thresholds based on statistical calculations from the target's observed behavior (metrics).

Time-based Thresholds are user-defined threshold values to be used at different times of the day/week to account for changing target workloads.

For more details on Beacons, see the *Advanced Threshold Management* chapter of the Enterprise Manager Cloud Control Administrator's Guide.

18.3.6 Set up Compliance Management in Enterprise Manager

Compliance Management provides the ability to evaluate the compliance of targets and systems as they relate to business best practices for configuration, security, and storage. This is accomplished by defining, customizing, and managing compliance frameworks, compliance standards, and compliance standard rules. In addition, Compliance Management provides advice of how to change configuration to bring your targets and systems into compliance. Compliance Management can help you maintain security and performance across all tiers with automated policy violation alerts for settings, for example, ensuring automated policy violation alerts for log levels, JVM versions, and patch levels. If you require legal compliance relating to security, you can set up WebLogic Security Technical Implementation Guidelines (STIG) rules as required.

For more details on compliance, see the *Managing Compliance* chapter of the Enterprise Manager Lifecycle Management Administrator's Guide. Also see the *Managing Configuration Information of Oracle Enterprise Manager Lifecycle Management Administrator's Guide* for information on a new feature, configuration drift that ensures consistency (uniformity) across a large number of targets.

18.3.7 Create a RUEI Application to Help Troubleshoot Web Applications

RUEI can monitor various types of application, and separate the data from each application for reporting. In RUEI, you create either an application, suite or service to correspond with the set of services that you want to monitor. The term, RUEI suite, is used if these application is based on certain Oracle Enterprise architectures (such as Oracle E-Business Suite, Siebel, and WebLogic Portal).

Before attempting this task, make sure that you have installed RUEI as outlined in [Install RUEI to Help Troubleshoot Web Applications](#).

After installing the software, perform the following tasks to create a RUEI application or suite.

1. If you want to monitor network timing, you must set up a network tap as described in the *RUEI Installation Guide*. However, you can skip this step if you want to create a tag-based application, where you insert javascript (Browser JS Library) into your web templates, and that javascript reports events from the client browser.

2. Perform initial configuration of RUEI as described in *Configuring RUEI* chapter of the *RUEI Installation Guide*.
3. If you want to monitor ADF applications, you can use RUEI's ADF Monitoring Service, where ADF-specific timings are reported. See the *Configuring RUEI for ADF Monitoring* chapter of the *RUEI Installation Guide* for more information.
4. Check the RUEI documentation for any specific information relating to a component you intend to monitor. For example, if you intend to monitor an environment that uses Oracle Access Manager, see the *Configuring the Oracle Access Manager* chapter of the *RUEI Installation Guide*.
5. Create a RUEI Application or Suite using the instructions in *Identifying and Reporting Web Pages* and *Working With Suites and Web Services* chapters of the *RUEI User's Guide*.

 **Note:**

With Enterprise Manager 13c, there is a new service level target, called an End User Service (EUS) corresponding to each RUEI Application/Suite. The EUS has the same name as the RUEI Application/Suite, and is used to create a business application. See [Monitoring an End User Service](#) for more information.

18.3.8 Define RUEI Service Level Agreements

Requirements Initially, RUEI reports are only available from the RUEI user interface. For RUEI data to be available in EM, you must connect RUEI to EM as described in [Registering RUEI Systems](#).

Procedure Key Performance Indicators and Service Level Agreements provide a method for business users to monitor processes and to be alerted if an issue arises. However, they can also be used to gather data that can be used to help troubleshoot problems in a complex environment.

18.3.9 Create a Business Application in Enterprise Manager

A Business Application is an Enterprise Manager target that represents a logical application; for the user, it defines a unit of management. A Business Application is composed of RUEI applications, Services, and System. Using the Enterprise Manager Console, you view a Business Application to access RUEI and information about the application's supporting infrastructure: the hosts and servers where the application services are executing.

Creating a business application is described in the [Monitoring Business Applications](#) chapter of this guide.

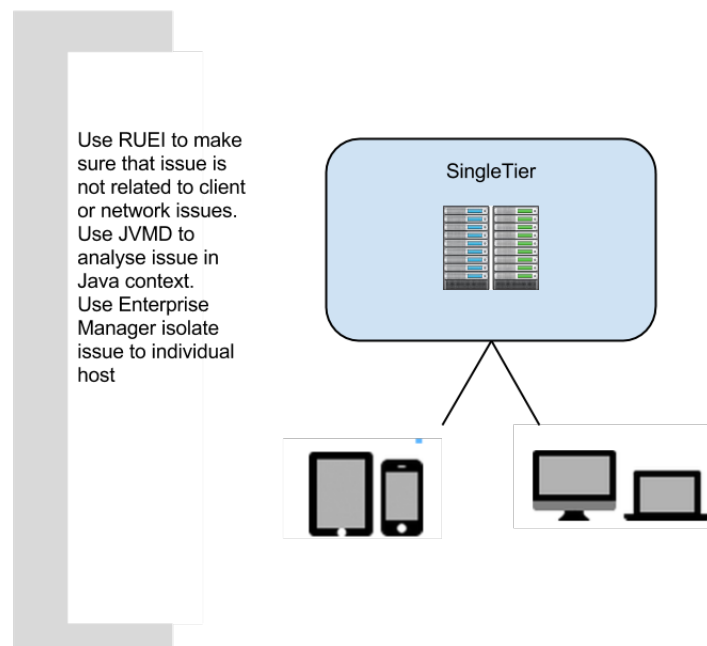
18.4 Analyzing Issues Using Enterprise Manager and RUEI

Enterprise Manager and RUEI provide tools to analyze incidents, transactions and user experience issues, and these methods are documented in the appropriate documentation set. However, this section outlines how to use all of the tools together to troubleshoot an issue, and that method can be summarized as:

1. If your system consists of many middleware tiers, try to isolate an issue to a single middleware tier.
2. Once you have isolated the issue to a single middleware tier, use Enterprise Manager to further isolate and resolve the issue.

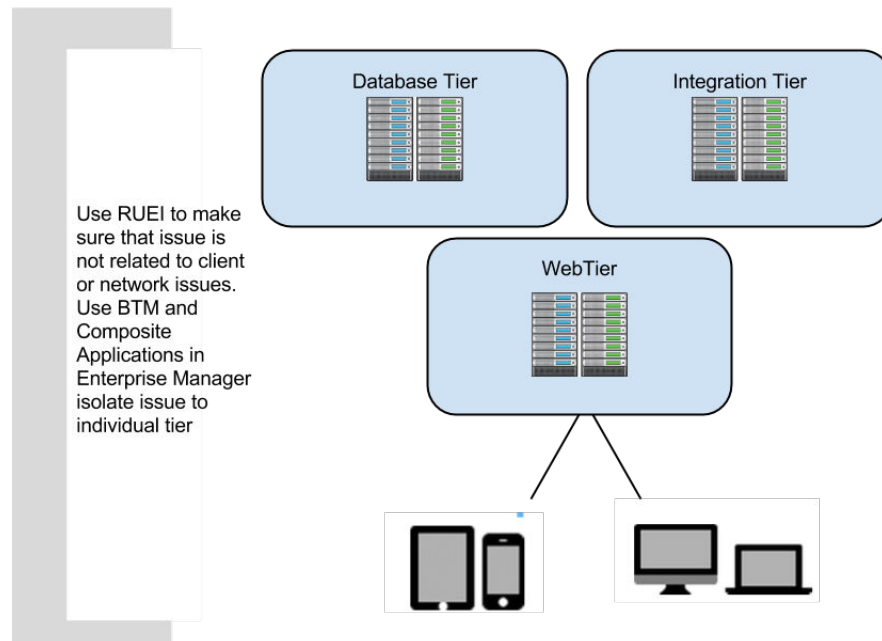
For example, if the application consists of a single tier, it might be represented as follows:

Figure 18-3 Troubleshooting a Single Tier



If the application consists of multiple tiers, it might be represented as follows:

Figure 18-4 Troubleshooting Multiple Tiers



This section describes the following approaches to analyzing issues:

- [Analyzing Incidents using Log Files](#)
- [Analyzing Incidents using Business Applications](#)
- [Analyzing Incidents](#)

18.4.1 Analyzing Incidents using Log Files

Regardless of the complexity of the environment, you can use Enterprise Manager to quickly analyze an incident using the data in log files. For most targets:

1. Select **Monitoring** from the **Enterprise** menu, then select **Logs**.
2. Click **Search** and select the targets you want to analyze.
3. Specify the criteria for the search, for example select **Trace** for very detailed logs.
4. Click **Search** to view the results.

If you are using RUEI, see [Monitoring Logs](#).

Note:

With this release, you do not need to deploy JRF to use the log viewer feature.

18.4.2 Analyzing Incidents using Business Applications

If you created a business application as described in [Create a Business Application in Enterprise Manager](#), you can use Enterprise Manager to analyze all the monitoring information from RUEI and Enterprise Manager using one dashboard, which provides you with the capability to drill down into the root cause of the issue. For more information, see [Monitoring Business Applications](#).

18.4.3 Analyzing Incidents

This section describes how to analyze incidents. Examples include:

- A set of WebLogic servers (a typical scenario for a website)
- An Oracle Access Manager installation that provides authentication and authorization services
- A WebCenter Portal

18.4.3.1 Check EM Dashboards to Analyze Incidents

If an issue is reported to you, the simplest troubleshooting method is to:

1. Log into Enterprise Manager.
2. Check the dashboard for issues relating to the hardware, operating system, database or middleware tier.
3. If an item is highlighted, click through to view the detail.
4. Restart any stopped items or note any suggestions made by Enterprise Manager to resolve the issue.

If Enterprise Manager is monitoring a large set of targets, the method above may not be practical. To help troubleshoot a subset of targets, consider creating a Business Application or Composite Application to make the task more manageable.

18.4.3.2 Use RUEI to Check Pages Affected by an Incident

With Enterprise Manager 13c and RUEI, you can immediately explore web application issues by:

1. Log into Enterprise Manager.
2. Navigate to the End User Service associated with the web application you want to troubleshoot. From the **Targets** menu, select **Services**. The currently defined services are listed. Filter for end user services. Alternatively, you can navigate to the End User Service from the Business Application home page by clicking on the **Sub Services** tab. The resulting list may include End User Services if the Business Application includes one.
3. Click the End User Service of interest. The home page for the selected end user service is displayed. From this page you should be able to determine whether the issue affects all pages or individual pages.
4. If the issue only affects an individual page, you can use other techniques (for example, session diagnostics) to determine the root cause.

18.4.3.3 Use JVMD to Isolate Issue

In Enterprise Manager, you can access JVMD information for a transaction operation by selecting a transaction in the Business Application and opening the transaction summary page. Then do one of the following:

- Right click one of the operation nodes in the topology diagram and select JVMD diagnostics from the context menu.
- Right click one of the operation rows in the operations table and select JVMD diagnostics from the context menu.

For more information, see [Getting Detailed Execution Information](#) .

JVMD also allows you to view related logs as described in [Using JVM Diagnostics](#) .

18.4.3.4 Use Thresholds and Compliance to Analyze Incidents

Thresholds and compliance provide a method to perform incident management, provided you have performed the set up in advance. For example, if an issue is not corresponding to an incident on the dashboard, but you suspect a memory issue:

1. Set a threshold for Component Memory Usage using the instructions provided in the *Enterprise Manager Cloud Control Administrator's Guide*.
2. If the timing of the threshold incident correlates with the reported issue, try increasing the memory allocated to the component.
3. If the timings do not correlate, you have established that the issue is probably not memory related and you can develop a different theory and set a new threshold and repeat the process.

18.5 Resolving Issues Using Enterprise Manager

After finding the issue using Enterprise Manager (or RUEI), there are many times you can use Enterprise Manager to perform the tasks required to resolve the issue.

18.5.1 Resolve an Issue Using Configuration Tools

Cloud Control collects configuration information for all managed targets across the enterprise. Collected configuration information is periodically sent to the Management Repository over HTTP or HTTPS, allowing you to access up-to-date configuration information for your entire enterprise through Cloud Control.

Cloud Control enables you to view, save, track, compare, search, and customize collected configuration information for all managed targets known to Enterprise Manager. Additionally, the Configuration Topology Viewer provides a visual layout of a target's relationships with other targets; for example, you can determine a system's structure by viewing the members of a system and their interrelationships.

For more information on a new feature, configuration drift that ensures consistency (uniformity) across a large number of targets, see *Managing Configuration in the Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

18.5.2 Work with Application Developers or DBAs to Resolve Application Issues

If you have identified an issue caused by code written within your organization, you must work with the developers to provide a solution. For example, the following steps outline a typical scenario:

1. An issue in the checkout code has been identified by operations personnel as causing an average 40 second load time for an application web page (using RUEI).
2. The developer cannot replicate the issue in the development environment, but has accepted that the issue exists after seeing RUEI.
3. Operations and developers meet and use JVMD to isolate the source of the issue.
4. The developer changes the code and operations sets KPIs to monitor the application so that the issue will not reoccur unnoticed.

Similarly, if you have identified an issue with a database, you can work with the appropriate DBA to resolve the issue.

18.5.3 Resolve a Capacity Issue Using Provisioning Tools

Some issues can be resolved by provisioning extra resources. Oracle Enterprise Manager can help with that task, providing automation of provisioning for database, and middleware. See the *Enterprise Manager Lifecycle Management Administrator's Guide* for details on how to perform provisioning using Enterprise Manager.

Two typical scenarios are provisioning extra nodes or provisioning extra memory. Extra nodes can be added to services like WebLogic servers and Oracle Real Application Clusters. Some issues can be resolved by provisioning extra memory to provide extra capacity. Extra memory can be provided in the following ways:

- Extra memory can be allocated to the JVM.
- Extra memory can be allocated to a Virtual Machine.
- Extra physical memory can be added to a server.

Part VII

Using JVM Diagnostics and MDA Advisor

The chapters in this part provide information regarding JVM Diagnostics and MDA Advisor.

The chapters are:

- [Introduction to JVM Diagnostics](#)
- [Using JVM Diagnostics](#)
- [Troubleshooting JVM Diagnostics](#)
- [Using Middleware Diagnostics Advisor](#)

For more information about JVMD Hybrid Cloud, see Deploying JVMD for Hybrid Cloud in the *Enterprise Manager Cloud Control Administrator's Guide*.

19

Introduction to JVM Diagnostics

This section provides an overview of JVM Diagnostics. It contains the following:

- [Overview](#)
- [New Features in this Release](#)
- [Supported Platforms and JVMs](#)
- [User Roles](#)

19.1 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 13c'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](#)

19.1.1 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.

19.1.2 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.

19.1.3 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 frequency is once every 200ms) for a short period of time. This allows you to identify interdependency of threads, bottleneck resources (DB, I/O, CPU, Locks, Network, RMI) and top methods.

19.1.4 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.

19.1.5 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 stopping the application. Additional heap analysis is provided with the Memory Leak Report, and the Anti-Pattern Report. Administrators 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.

19.1.6 JVM Pooling

JVM Diagnostics allows administrators to group sets of JVMs together into JVM pools. This grouping 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 feature allows administrators to visualize problems naturally and intuitively.

19.1.7 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.

19.2 New Features in this Release

This section lists some of the new JVM Diagnostics features in Oracle Enterprise Manager Cloud Control 13c Release 3:

- New option added to configure JVM Settings like Heap Dump Location for the JVM Pool. User can also mention the `libdir` at the time of deployment.
- Support to deploy/upgrade/remove JVMD agent using cluster target.
- Request metrics based on all executed instances (vs. on sampled instances only).
- Deploying agent in OFF mode.
- Deploying the agent on the cluster level. The agent will be targeted to all managed server in the cluster, including those that will be added at a later time.
- Allow to specify LIBDIR at the time of deployment.
- EMCLI support (for JVM and JVM Pools):
 - Configure Heap/JFR Dump Directory
 - Enable/disable monitoring
- Improvements to DB requests data collection.
- Call tree view enhancements:
 - Optional view in method level (in addition to the code line level)
 - Auto expand option to reveal the significant code path

19.3 Supported Platforms and JVMs

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>

19.4 User Roles

To use JVM Diagnostics, you must have either of the following JVM Diagnostics resource 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 information, see *Enterprise Manager Cloud Control Administrator's Guide*.

20

Using JVM Diagnostics

This section describes the tasks you can perform by using JVM Diagnostics. In particular, it contains the following:

- [Setting Up JVM Diagnostics](#)
- [Accessing the JVM Diagnostics Pages](#)
- [Managing JVM Pools](#)
- [Managing JVMs](#)
- [Managing Thread Snapshots](#)
- [Analyzing Trace Diagnostic Images](#)
- [Viewing Heap Snapshots and Class Histograms](#)
- [JVM Offline Diagnostics](#)
- [Viewing JVM Diagnostics Threshold Violations](#)
- [Using Java Workload Explorer](#)
- [Troubleshooting JVM Diagnostics](#)
- [Using Middleware Diagnostics Advisor to View and Diagnose Performance Issues](#)
- [Enable or Disable Monitoring of JVM Targets using EMCLI](#)

20.1 Setting Up JVM Diagnostics

Follow these steps to set up and configure JVM Diagnostics:

1. From the **Setup** menu, select **Middleware Management**, then select **Engines and Agents**. A list of RUEI/JVMD Engines are listed. Under the JVM Diagnostics Engines row, the following details are displayed when all the columns are selected:
 - **Name:** The name assigned to the JVM Diagnostics Engine. This ID identifies the JVM Diagnostics Engine in all the processes.
 - **Type:** JVM Diagnostics Engine (By default, this column is not listed).
 - **Host:** Machine on which the JVM Diagnostics Engine has been deployed.
 - **Port:** HTTP port of the server on which the JVM Diagnostics Engine has been deployed.
 - **SSL Port:** SSL Port of the server on which the JVM Diagnostics Engine has been deployed.
 - **Availability:** Percentage of time when the engine has been available.
 - **Status:** Status of the JVM Diagnostics Engine. Options are Active, Inactive, or n/a (not available).
 - **Server:** Server on which the engine is located.

- **Version:** Build version of this JVM Diagnostics Engine.
2. 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:
 - JVMD Configuration (See [Configuring the JVM Diagnostics Engine](#))
 - JVMs and Pools (See [Configuring JVMs and JVM Pools](#))
 - Register Databases (See [Registering Databases](#))
 - Heap Analysis Hosts (See [Configuring the Heap Analysis Hosts](#))

**Note:**

JVMD Load Balancers information is also displayed on the Setup page. The table includes Load Balancer URL, its status, and a list of engines associated with it.

20.1.1 Configuring the JVM Diagnostics Engine

You can configure the JVM Diagnostics Engine by defining engine parameters and advanced settings. You can then create new idle thread rules and system call rules. Operations on existing rules include importing and exporting a rules, as well as deleting a rule.

1. From the **Setup** menu, select **Middleware Management**, then select **Engines and Agents**. Select the JVM Diagnostics Engine row in the RUEI/JVMD Engines table and click **Configure**.
2. Click the **JVMD Configuration** tab.
3. You can modify the following details in the JVMD Engine Parameters region.
 - **JVMD Engine Log Level:** The log level for console diagnostics messages. Log levels 1 to 5 are supported where:
 - ERROR–1
 - WARN–2 (warning)
 - INFO–3 (information)
 - DEBUG–4
 - TRACE–5The default log level is INFO–3.
 - **Cross Tier Log Level:** The log level for cross-tier diagnostic messages. Log levels 1 to 5 are supported where:
 - ERROR–1
 - WARN–2 (warning)
 - INFO–3 (information)
 - DEBUG–4
 - TRACE–5

The default log level is INFO–3.

- **Agent Request Timeout (secs):** 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.
- **Enable Monitoring:** Select this check box to start or stop monitoring.

4. Advanced Settings

- **Purge Detail Data Older Than (hours):** The period for which the detailed monitoring samples should be retained.
- **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.
- **System Sample Interval (secs):** The frequency at which system details (cumulative CPU counters, heap size, and number of garbage collections (GCs)) should be collected in monitoring.
- **Purge Aggregated Data Older Than (days):** The period for which the aggregated monitoring samples should be retained.

Note:

This field is not applicable to the JVMTI (level 0) optimization.

Click **Save** to save the parameters.

5. In the Thread Rules region, you can define the following:

- **Idle Thread Rules:** 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 and specify the idle thread rule information.
 - **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 (class + method). You can specify a wild card, for example, if you specify `weblogic.servlet.*`, all the threads that meet this criteria will be ignored.
 - Note:** The Current call of the stack is the first frame of the stack, traversing from top to bottom, such that the function (class + method) is not a System call. System calls are assumed to be the calls which are not relevant to the user application like `java.*`, and so on.
 - Thread Name:** Select this type if you want to ignore threads with a particular name.
 - **Rule Value:** The Rule Value should contain the fully qualified class name, method, followed by the class+method.

An example of a Current Call is `weblogic.socket.PosixSocketMuxer->processSockets`

An example of a Monitor (Waiting on Lock) is

```
weblogic.socket.PosiSocketMuxer$1
```

All threads that meet the criteria specified in the Idle Thread Rule will not appear in the View Active Threads screen.

- **Global Rule:** Select this check box to apply the idle thread rule to all JVM Pool targets. If this box is unchecked, you must select one or more JVM Pools for which this rule will be applicable.

All threads that meet the criteria specified in the Idle Thread Rule will not appear in the View Active Threads screen.

- **System Call Rules:** Mark a method as a system call by adding it to the System Call Rules. System calls are assumed to be the calls which are not relevant to the user application like java.*, and so on. Click **New Rule** to create a new system call and specify the matching pattern such as sun.*, java.*, and so on.

All methods that match the rules listed in the System Call Rules table are identified as system calls.

20.1.2 Configuring JVMs and 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 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 **Engines and Agents**. Select the JVM Diagnostics Engine row in the RUEI/JVMD 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 **Yes**, JVMs belonging to this pool will be polled for active requests periodically based on the Poll Interval.
3. Click **Create Pool** to create a new pool.
 - a. In the Create New JVM Pool dialog box, 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 **Create** to save the JVM Pool information.
4. To delete a pool, highlight the pool click **Remove**.
5. Select a JVM Pool or a JVM and click **Details** to view additional details about the JVM Pool or JVM.
6. 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.

- **Load JFR:** Contains scripts to upload JFR snapshots to the repository. Use JMC (Java Mission Control) to download and analyze the JFR snapshot.
7. Select a JVM Pool or a JVM and click **Configure**.
For JVM Pools, see [Configuring a JVM Pool](#).
For JVMs, see [Configuring a JVM](#).

20.1.3 Registering Databases

Follow these steps to register databases:

1. From the **Setup** menu, select **Middleware Management**, then select **Engines and Agents**. Select the JVM Diagnostics Engine row in the RUEI/JVMD Engines table and click **Configure**.
2. Click the **Register Databases** tab. The list of registered databases is displayed. The database name, host, Oracle SID for the monitored database, and listener port number is displayed.
3. You can do the following:
 - **Add a Database Instance:** From the **Add** menu, select **Database Instance** to register a single instance or Oracle Real Application Cluster (RAC) 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, Service, User name, Password, and Confirm 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.
 - **Edit:** You cannot edit a Database Instance. Only custom databases can be edited. Select a custom database from the list and click **Edit**.
 - **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 Database URL(s) to the Database**, select a Database URL and click **Add** and specify the URL of the database to be associated.

Note:

The DB User must be the same as the user running the application that is being monitored and 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 a user such as `jvmsadmin`, enter a command as follows:

```
SQL> grant select on SYS.GV_$LATCHNAME to jvmsadmin
```

Multiple registrations may be necessary for a single database agent if different database users are running multiple applications.

- **Export:** This option provides the information in a spreadsheet format. You can either open the spreadsheet or save it for future use.
- 4. 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.
- 5. Click **Downloads** to manually download the various binaries such as JVM Diagnostics Agent, Load Heap, Load JFR zip, 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.
 - **Load Heap:** The `loadheap.zip` is saved to a specified location.
 - **Load JFR:** Contains scripts to upload JFR snapshots to the repository. Use JMC (Java Mission Control) to download and analyze the JFR snapshot.

20.1.4 Configuring the Heap Analysis Hosts

Note:

The analysis and load heap steps have significant memory requirement on the Heap Analysis Host. Ensure you have a 64-bit JVM and sufficient free memory on the Heap Analysis Host.

To configure the heap analysis hosts, follow these steps:

1. From the **Setup** menu, select **Middleware Management**, then select **Engines and Agents**. Select the JVM Diagnostics Engine row in the RUEI/JVMD Engines table and click **Configure**.
2. Click the **Heap Analysis Hosts** tab. The Configure Heap Analysis Hosts page appears.
3. 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.
4. Click **Save**.

20.1.5 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 **Engines And Agents**.
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.
 - **Type:** The type of engine, in this case, JVM Diagnostics Engine. (By default, this column is not listed.)
 - **Host:** The machine on which the JVM Diagnostics Engine has been deployed.
 - **Port:** HTTP port of the server on which the JVM Diagnostics Engine has been deployed.
 - **SSL Port:** SSL Port of the server on which the JVM Diagnostics Engine has been deployed.
 - **Availability (%):** Percentage of time when the engine has been available,
 - **Status:** Status of the JVM Diagnostics Engine. Options are Active, Inactive, or n/a (not available).
 - **Server:** Server on which the engine is located.
 - **Version:** Build version of this JVM Diagnostics Engine.

Highlight the JVM Diagnostics Engines row and click **Configure** to configure the JVM Diagnostics Engine parameters, JVMs and Pools, databases, and Heap Analysis hosts.

- JVMD Configuration
- JVMs and Pools
- Register Databases
- Heap Analysis Hosts
- Hybrid Cloud Gateways Configuration

20.2 Accessing the JVM Diagnostics Pages

From the **Targets** menu, select **Middleware**, and click on a Java Virtual Machine Pool or Java Virtual Machine target. The Home page for the target is displayed.

To start using JVM Diagnostics, select the appropriate option from the Java Virtual Machine Pool menu.

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.

20.3 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 Java Virtual Machine Pool Home Page](#)
- [Viewing the JVM Pool Live Thread Analysis Page](#)
- [Managing Thread Snapshots](#)
- [Analyzing Heap Snapshots](#)
- [Configuring a JVM Pool](#)
- [Removing a JVM Pool](#)
- [Adding a JVM Pool to a Group](#)
- [Using Java Workload Explorer](#)

20.3.1 Viewing the Java Virtual Machine Pool Home Page

The Java Virtual Machine Pool Home page shows the summary and configuration information of all the JVMs in the JVM pool.

It shows the following details:

- **Summary:** Shows whether polling is enabled and the Polling Interval. It also shows the number of incidents and the number of configuration changes. Click the incident number to drill down to the Incident Manager page
- **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.
- **JVM Activity (Last 15 Minutes):** This region shows the active threads for each JVM in the pool. You can also select the following JVM Activity graphs: Active Thread States, Memory Utilization, CPU Utilization, GC Overhead, and Response and Load.
- **Overview (Last 15 Minutes):** This region shows the status for the last 15 minutes for each JVM in the pool. The current activity of the JVM including CPU usage, memory, average number of threads waiting for a database response, network response, or average number of threads waiting for synchronization lock, idle threads, and configuration changes are displayed. Additional information includes: Target Status, Diagnostics Findings, GC Overhead %, Host CPU Usage, MAX JVM Heap Used %, Major GC Count, Minor GC Count, Major GC Time (ms), Minor GC Time (ms), Host, OS, Vendor, JVM Version, Min Heap Size (MB), Max Heap Size (MB), Open File Descriptor (%), Swap Space (%), Host Memory (%), Context Switch (per sec), and OSR.

If JVMs displayed are present in different WebLogic domains, you can view the WebLogic Domain and the host on which the JVM is running. Click on the JVM link to drill down to the JVM Home page.

- **Top Requests (Last 15 Minutes):** This region shows the top requests for each JVM in the pool. Statistics include: JVM time, JVM CPU, thread allocation, count,

maximum duration of each request, the average duration of each request, throughput (per minute) and minimum duration (ms).

20.3.1.1 Promoting JVM Diagnostics Events to Incidents

An event is a notable occurrence detected by Enterprise Manager that is related to target, job, monitoring template at a particular point in time, which may indicate normal or problematic behavior. Example for events – database target going down, performance threshold violation based on metrics, unauthorized change in the application configuration file changes, failure in job execution, and so on.

An incident is an event or set of closely correlated events that represent an observed issue requiring resolution, through (manual or automated) immediate action or root-cause problem resolution.

By default JVM Diagnostics events are not promoted to incidents and will not appear in the JVM Pool or JVM Home page. To promote events to incidents, follow these steps:

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. Click **Create Rule Set**. In the Create Rule Set page, in the Targets region, select the **All Targets of types** option. Select Java Virtual Machine and Java Virtual Machine Pool target types.
3. Click **Create** in the Rules region and in the Select Type of Rule to Create window, choose **Incoming events and updates to events** option and click **Continue**. The Create New Rule: Select Events page appears. In the Type drop down list, select **JVM Diagnostics Threshold Violation**.
4. Then select Specific events of type JVM Diagnostics Threshold Violation.
5. Click **Add**. The JVM Diagnostics Threshold Violation Rule window appears. Select the JVM Diagnostics metrics that will trigger threshold violation events. These events will be promoted to incidents. Click **OK**. Click **Next**, review the rules, and click **Continue** to save the rule. All events that match the criteria will be promoted to incidents and will appear in the JVM Diagnostics Pool Home page.

20.3.2 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.

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 selected 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 or DMS is configured, the Request Name and Request Age values 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 Waiting Time column.

Click the link to view the database diagnostics details. See [Performing 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.

- Additional columns include: App Name, Module Name, Work Manager, Frames Count, Is Stuck, Is Hogger, Read Characters, Write Characters, Blocked Count, Blocked Time (ms), Wait Count, Waited Time (sec), IO File Name, OS Thread ID, Age (ms), Waiting Time (sec), Lock Held, ECID, RID, User, Session, and Is Idle.

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.
- **Search/Filter:** To minimize the number of reported rows, in the Search field select the column name and then provide a filter. For example, select Thread Name then type the word *job*. The search reports on threads that include the word *job*.
- **Show Idle Threads:** Select this check box to list all the Idle Threads in the JVM Threads table.
- **Detach:** Select a thread and click Detach to view the table in a separate window.
- **Take Snapshot of Selected Thread:** Select a thread, 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 selected time interval is displayed on the Java Workload Explorer page.
- **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 and will no longer be collected in the monitoring data. Marking a thread idle in JVMD will not affect the OS or JVM thread management.
- **Mark Active:** Select an Idle thread and from the **Action** menu, select **Mark Active** to change the status to Active.
- **Mark System Call:** Apart from the threads defined as System Calls in the JVMD Configuration page (see [Configuring the JVM Diagnostics Engine](#)), you can mark specific threads as system calls so that JVMD will not consider the marked method as a user call method.

Select a thread from the JVM Threads table. From the **Action** menu, select **Mark System Call** to mark this thread as a **System Call**. All user calls that are marked in this manner will now be treated as System Calls. If you selected a marked call and click **Unmark System Call**, the thread will now be treated as a User Call.

- **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 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:
 - **Class Method:** The class and method for the stack frame. Click the link to view the method locals.
 - **File:** The file where the class is defined.
 - **Line:** Current execution point in the method. If a method is inlined or native, the line number might not be available.

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.

You can refresh the data that is displayed by specifying the **Auto Refresh** period.

20.3.3 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 [Updating Pool Thresholds](#).

20.3.3.1 Updating Pool Thresholds

Follow these steps to edit the pool thresholds on the Edit JVM Pool Information page:

1. 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.
2. Click **Add** to add a corrective action. Select a corrective action from the list and click **OK**. You can select:
 - **No Action:** No corrective action 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 (ms):** Interval after which snapshot should be repeated.
 - **Poll Duration: (sec)** Duration for which the snapshot should be taken.
 - **Description:** Describe the purpose of the trace dump. Include information pertinent to other users.
 - **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.
 - **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 [Viewing JVM Diagnostics Threshold Violations](#).

- **JFR Snapshot:** Select this option to generate a dump of the JFR.

JFR snapshot creation is supported for Java JVM and for Oracle JDK 1.7.0_04 onwards. For Oracle JDK 1.7.0_04 onwards but prior to JDK 1.8.0_40, JVM process should be run with the following java options '-XX:+UnlockCommercialFeatures -XX:+FlightRecorder'. These java options are not required for JDK 1.8.0_40 onwards.

- **Class Histogram:** Select this option to generate a dump of the class histogram.
- **Diagnostic Snapshot:** Select this option to generate a diagnostic dump.
Select the Duration in Minutes to be used in this snapshot.

3. Click **Save** to save the threshold values.

20.3.4 Removing a JVM Pool

You can remove a JVM Pool from the following:

- **Middleware page:** Highlight the JVM Pool and click Remove.
- **JVM Pool home page:** From the Java Virtual Machine Pool menu, select Target Setup, then click Remove Target.

You will see a warning message that 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.

20.3.5 Adding a JVM Pool to a Group

From the JVM Pool home page, select the **Java Virtual Machine Pool** menu. Select **Target Setup**, then click **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.

20.4 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 Diagnostics Performance Summary](#)
- [Viewing the JVM Live Thread Analysis Page](#)
- [Viewing Memory Diagnostics](#)
- [Working with Class Histograms](#)
- [Taking a Heap Snapshot](#)
- [Taking a Heap Snapshot and Loading Into the Repository](#)

- [Analyzing Heap Snapshots](#)
- [Managing JFR Snapshots](#)
- [Configuring a JVM](#)
- [Removing a JVM](#)
- [Adding a JVM to a Group](#)

20.4.1 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.
2. The JVM Home page with the following details is displayed.
 - **Summary:** Shows details of the JVM such as the availability status of the JVM, heap size, WebLogic Server it belongs to, and composite applications. The region also includes the number of open incidents that have occurred and the number of configuration changes. Click the number of incidents to drill down to the Incident Manager page.
 - **JVM Activity (Last 15 Minutes):** The number of active threads in the JVM in the last 15 minutes. Click on a thread to see the detail of the thread. You can also select the following JVM Activity graphs: Active Thread States, Memory Utilization, CPU Utilization, GC Overhead, and Response and Load.
 - **Overview (Last 15 Minutes):** Shows the state of the various threads in the JVM in the color-coded columns. This region can be added using page customization.

The current activity of the JVM including Target Status, Threads (CPU, DB Wait, Lock, Network Wait, IO Wait, RMI Wait), JVM Time (sec), Diagnostic Findings, Resource (%), Host CPU (%), JVM DPU(%), JVM Heap (%), Max JVM Heap (%), GC Overhead (%), Major GC Count, Minor GC Count, Major GC Time (ms), Minor GC Time (ms), Host, OS, Vendor, Version, Container Name, Container Type, Min Heap Size (MB), Max Heap Size (MB), Open File Descriptors (%), Swap Space (%), Host Memory (%), Context Switch (per sec), and OSR are displayed.

- **Top Requests (Last 15 Minutes): Shows:** Shows the top requests for the JVM. Statistics include: JVM time, JVM CPU, thread allocation, count, maximum duration of each request, the average duration of each request, throughput (per minute), and minimum duration (ms).

20.4.2 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 is displayed on this page for the JVM target. To view the JVM performance metrics, follow these steps:

1. From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine target.
2. Select the **Performance Summary** option from the **Java Virtual Machine** menu. The following page appears:

3. A set of default charts that show the values of the JVM performance metrics over a period of time is 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. You can then create a snapshot that can be examined for leaks. See [Taking a Heap Snapshot](#).

4. 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 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.

20.4.3 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:

1. From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine target.
2. Click the **Live Thread Analysis** link at the top of the page or select the **Live Thread Analysis** option from the **Java Virtual Machine** menu.

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 Process statistics include: CPU (%), Memory (%), Context Switch (per sec), and Open File Descriptors (%). Threads statistics include: CPU, DB Wait, Lock, Network Wait, IO Wait, RMI Wait, and Idle Threads.
- **JVM Threads:** This table shows a list of all the threads running in the JVM. Click on a thread to view the thread details. To view all the available columns, from the **View** menu, select **Columns**, then select **Show All**.

Thread details include:

- 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 thread is the DB Wait state, click on the link to view the database diagnostics 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.

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:

- **Export:** Select a thread and click **Export** to export the thread details along with thread stacks information to an Excel file.
- **Search/Filter:** To minimize the number of reported rows, in the Search field select the column name and then provide a filter. For example, select Thread Name then type the word *job*. The search reports on threads that include the word *job*.
- **Show Idle Threads:** Select this check box to list all the Idle Threads in the JVM Threads table.
- **Detach:** Select a thread and click Detach to view the table in a separate window.
- **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 Action 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.
- * Description: Description of the snapshot.
- * 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.

A snapshot file with details of the selected thread or active threads (depending on your selection) 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 from the list and from the **Action** menu, select **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 from the **Action** menu, select **Mark Active** to mark it as an active thread.
- **Mark System Call**: Apart from the threads defined as System Calls in the JVMD Configuration page (see [Configuring the JVM Diagnostics Engine](#)), you can mark specific threads as system calls. Select a thread from the JVM Threads table. From the **Action** menu, select **Mark System Call** to mark this thread as a **System Call**. All user calls that are marked in this manner will now be treated as System Calls. If you selected a marked call and click **Unmark System Call**, the thread will now be treated as a User Call.
- **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 [Performing Cross Tier Analysis](#).
- **Thread Stack**: The Thread Stack table shows the details of the selected thread such as:
 - Class Method: The class and method for the stack frame. Click on the link to view the method locals.
 - File: The file where the class is defined.
 - Line: Current execution point in the method. If a method is inlined or native, the line number might not be available.

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 the file in .csv format.
- **Mark / Unmark System Call:** You can mark a selected method as a system call. Select a method from the Thread Stack table and from the **Action** menu, select **Mark System Call**. All methods marked in this manner will be treated as system calls. If you select a marked call and click **Unmark System Call**, the method will now be treated as a user call.
- **Auto Refresh:** You can refresh the data that is displayed by specifying the Auto Refresh period.

20.4.3.1 Performing 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. This enables you to drill-down from JVM Diagnostics pages to Database Diagnostics pages.

Note:

If a database is not registered with JVMD, the Database JDBC details, SQL statement, SQL ID, and schema name will be collected by the JVMD agent for threads in DB Wait state.

In the case where the database is not an Enterprise Manager target, you can still register the database with JVMD as a "Custom database" which will track the database activity to this database by its name. The SQL statement and SQL ID would be fetched real-time from the database but you cannot drill-down from JVM Diagnostics pages to Database Diagnostic pages.

To register the database:

1. From the **Setup** menu, select **Middleware Management**, then select **Setup**. Select the JVM Diagnostics Engine row in the RUEI/JVMD Engines table and click **Configure**.
2. Click the **Register Databases** tab. 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 column indicating "JVMD Supported DB". If this column has a value of **Yes**, you can proceed with the cross tier analysis. If the column has a value of **Unavailable**, you cannot perform cross tier analysis because the JDBC connection to the database cannot be established.

 **Note:**

If cross tier correlation is not established even after registering the database with JVMD, select a database and click **Manage DB URL**. In the **Associate / Disassociate a Registered Database** field, select a Database URL and click **Add** and specify the URL of the database to be associated. After the URL has been associated with the registered database, the JVM Diagnostics Engine will start monitoring the cross-tier calls between JVM targets and the underlying database.

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.
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 to navigate to the Database Diagnostics pages.

 **Note:**

If cross tier is not established, the Database Details popup is displayed which shows the host, port, SID, user, and JDBC URL for the target database. This can happen when the database is not registered with JVMD or if JVMD is unable to find the registered database corresponding to the JDBC URL of the database. For registering the database, click on the link in **To view Register Database page click here** and this will take you to the Register Databases tab. In the case where the database is already registered, associate the JDBC URL with a registered Database by clicking the link in **To associate a Registered database to the above Database URL click here**. This will open a popup that will enable you to associate the JDBC URL for the database with a registered database.

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 made 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 SQLS sorted by the number of samples.

Click a SQL in Top SQLS to view the list of registered databases on which the SQL was executed. "Not Defined" represents SQLS that were executed on databases that are not registered with JVMD. Click a Database in Top Databases opens a popup that shows a link to the database name, clicking on which takes to

Database diagnostics page. In case the database name is "Not Defined", the popup shows the different JDBC URLs for the databases that were not registered and the corresponding number of samples for each Database. It also shows the reason as to why the cross tier was not established.

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.

20.4.3.2 Establishing Cross-Tier Correlation in Oracle RAC Databases

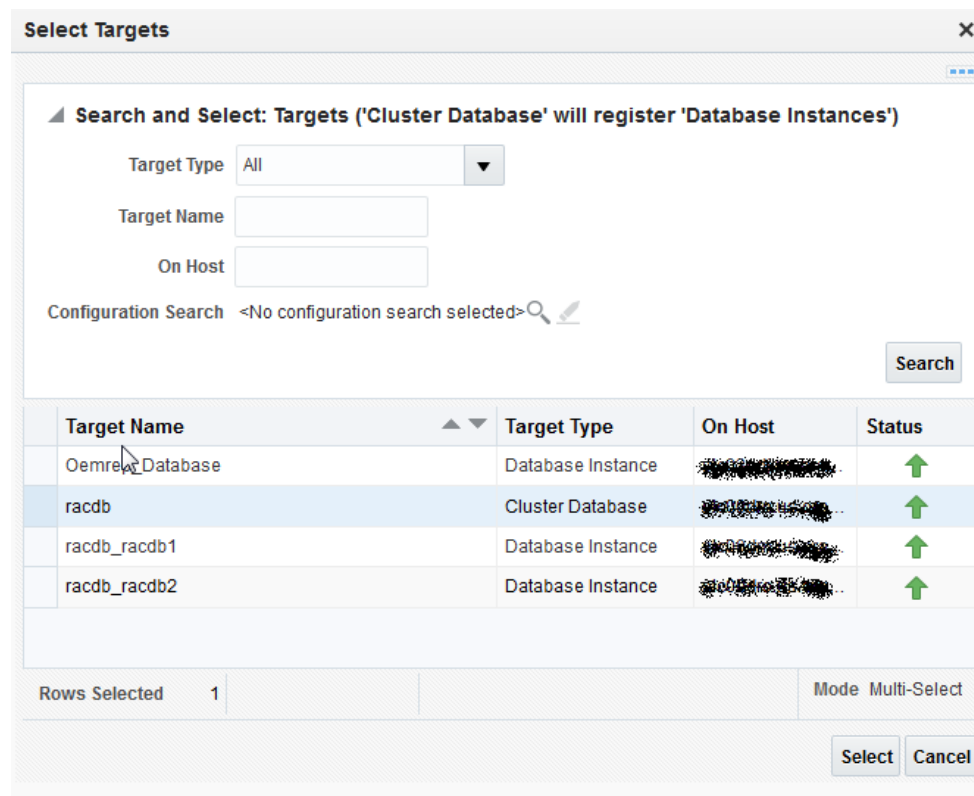
Oracle Real Application Cluster (Oracle RAC) databases have a complex configuration of database instances and listeners. User applications use Oracle RAC services to connect to the database instead of SIDs that are used for single instance databases. User applications can connect to Oracle RAC listeners 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 you have the cluster (RAC) database discovered in your EM and want to do the Xtier correlation with JVMD, then you can go and add the target type of **Cluster Database**. Refer the figure below:

Note:


You should not select the database instance which are part of the cluster. For example, in figure given below **racdb** is cluster database and **racdb_racdb1** and **racdb_racdb2** are part of cluster **racdb**, so if you are using the cluster database in your weblogic data source then you should consider adding cluster database. JVMD will get the associated database instances for you automatically.

Figure 20-1 Adding Cluster Database



In the above figure, the cluster Database **racdb** is added and the below figure shows both database instances **racdb_racdb1** and **racdb_racdb2** being added to registered databases.

Figure 20-2 After registering cluster database




If you want to add only database instances which are part of cluster, you can do that but you may miss some of the mappings if the SQL execution request goes via another db instance. Hence, it is recommended to add cluster database directly if it's discovered in EM since this covers more big spectrum for Xtier mapping.

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 using SID for each instance of RAC DB and add the jdbc url (which is used to connect to application via configured data source) in **Manage DB URL**.

Figure 20-3 Manage Database URL



Consider the following example:

RAC DB interface: host_rac, port:1521, ServiceName:S121, having the three db instances configured as,

DB1--> hos1 SID:S1212 Port:1521

DB2--> host2 SID:S1211 Port:1521

DB3--> host3 SID:S1213 Port:1521

And data source is configured as jdbc:oracle:thin:@//host_rac:1521/S121.

Then in above scenario you must add three custom DB using their SID and for each custom DB add the jdbc url configured in the data source via Manage DB URL.

20.4.4 Viewing Memory Diagnostics

This page provides you the details regarding current memory pool usages and the garbage collections statistics. It also provides statistics related to the class loading and class compilations, and the means to get and save a live histogram, and view all the histograms.

Follow these steps to do a real-time analysis of the JVM heaps that have been loaded into the repository.

1. From the **Targets** menu, select **Middleware**, then click on a JVM target.
2. From the **Java Virtual Machine** menu, select **Memory Diagnostics**.
3. The following tabs are available:

- Heap Memory Usage

These charts depict java memory pool usage.

- JVM Heap Memory Usage

The SunBurst chart shows the java heap structure and sizes. The Heap node in the center of the sunburst represents the heap-memory of the JVM process. When you mouse over the Heap node, it shows the size of the heap in MBs. Surrounding these nodes are other memory pools which are part of the JVM heap memory. Double clicking any pool node expands the sunburst and that pool node becomes the center of the chart and its children "Used" and "Free" are shown.

For HotSpot JVM, there are three children of heap node: Eden Space, Survivor Space, and Old Gen (generation). Mouse over these nodes to view there sizes.

The bar chart shows percent utilization of all the memory pools (heap and non-heap). If the pool usage is less than 75%, the color of the bar is green. If the pool usage is between 75% and 95%, the color of the bar is yellow. If the pool usage is 95% and higher, the color of the bar is red.

- Tenure Distribution Information

Tenure represents the age of the JVM objects that survived the number of minor collections. For example, an object with Tenure Size of 2 represents that these objects have survived two minor collections.

The statistics that display are dependent on the version of the JVM and the configuration of the garbage collection.

- TLAB (Thread Local Allocation Buffer) Statistics

This region shows Thread Local Allocation Buffer (TLAB) related JVM performance counters.

To avoid the pointer contention in the Eden Space while allocating objects, each thread is given a private memory area where it does object allocation. This private memory area is called TLAB.

- JVM Metrics

This region shows different metrics charts (by percentage and by value) related to JVM heap. Monitoring data from the JVM is used to prepare the correlation charts.

Use the time selector to select the desired time period for the charts. By default, chart duration is 15 minutes.

To view the information in table format, click **Table View**.

 **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.

- Garbage Collection

These charts and tables report the number of objects that have been added to the garbage collection (gc). The type of garbage collection, that is minor or major, and the number of garbage collections of a particular type are displayed.

- Last Garbage Collection Information

The Last Garbage Collection Information region provides the information pertaining to the last major and minor garbage collection that occurred in the JVM including: start time, end time, duration of the last collections, garbage collector name, current garbage collection count, and number of garbage collection threads. If provided by the JVM, this tab also shows the cause of the garbage collection.

Also included are bar charts which show the changes in pool usages after the garbage collection.

- Garbage Collections

The Garbage Collections region shows the cumulative statistics of major and minor collections including total collection count, total gc pause time (in milliseconds), rate of collections (invocation per minute), and the gc overhead percentage.

- JVM Metrics

This region shows different metrics charts (by percentage and by value) related to garbage collection. Monitoring data from the JVM is used to prepare the correlation charts.

Use the time selector to select the desired time period for the charts. By default, chart duration is 15 minutes.

To view the information in table format, click **Table View**.

- Class Loading Data

This tab shows data related to class loading and class compilations.

- Class Loading Stats

This region shows the total number of classes loaded, the total number of classes unloaded, and the total class loading time (in milliseconds) since the JVM started. It also shows total compilations (just-in-time [JIT] + on

stack replacement [OSR]) done and total compilation time (in milliseconds) since the JVM started. Finally, it shows the number of objects which are pending finalization.

- JVM Metrics

This region shows different metrics charts related to class loading and compilations. Monitoring data from the JVM is used to prepare the correlation charts.

Use the time selector to select the desired time period for the charts. By default, chart duration is 15 minutes.

To view the information in table format, click **Table View**.

- Class Histograms

Use this tab to generate a live histogram, save it, and see all the saved histograms. When you click **Get Live Histogram**, the JVM heap is traversed and a histogram containing class name, instance count, total size and average size is generated. You can then study the histogram and see the classes whose instances are consuming the most memory.

The JVM Class Details 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.
- Total Size (KB): The size of the JVM heap.
- Average Size (KB):
- Get Live Histograms

Click this option to generate a histogram.

- Schedule

Click **Schedule** to add the JVM Class Histogram data to the repository by scheduling a job. You can 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.

- View Saved Histograms

Click this option to view the saved histograms in the Available Heap Snapshots page.

4. You can do the following from any of the tabs:

- **Take Heap Snapshot**

Click **Take Heap Snapshot** to take a heap snapshot.

- **Export**

Click **Export** to export live heap data to an Excel file.

- **Save**

Click **Save** to save the classes in the JVM Classes table to the repository. The Save Class Histogram popup appears. Enter a name for the snapshot, and a description, and click **OK**.

20.4.5 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](#)

20.4.5.1 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 **Memory Diagnostics** option from the Java Virtual Machine menu.
2. In the Memory Diagnostics page, click the **Class Histograms** tab. Click **Get Live Histogram** to get Class Histogram data. Click **Save**.
3. In the Save Class Histogram window, enter a name for the snapshot and a description and click **OK**.

20.4.5.2 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 **Memory Diagnostics** option from the Java Virtual Machine menu.
2. In the Memory Diagnostics page, click the **Class Histograms** tab. 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.

20.4.5.3 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 **Memory Diagnostics** option from the Java Virtual Machine menu.
2. In the Memory Diagnostics page, click the **Class Histograms** tab. Click **Schedule**. The Schedule Settings page 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.

20.4.5.4 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 **Memory Diagnostics** option from the Java Virtual Machine menu.
2. In the Memory Diagnostics page, click the **Class Histograms** tab. 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.

20.4.5.5 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 **Memory Diagnostics** option from the Java Virtual Machine menu.
2. In the Memory Diagnostics page, click the **Class Histograms** tab. 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.

20.4.6 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 **Memory Diagnostics** from the **Java Virtual Machine** menu, then click **Heap Memory Usage**.

3. Click **Take Heap Snapshot**. The Load Heap Snapshot page appears.
 4. The Load Heap Snapshot page appears. In the Heap Snapshot Formats and Analysis Options region, specify the following:
 - **Heap Snapshot Type:** Select the heap dump file format. This can be:
 - JVMD Format (txt)
 - HPROF Format (binary)
 - Choose this option to take a heap snapshot and load it manually to repository using the loadheap script.
 - **Request GC before taking heap snapshot:** Click the check box if you want to delete unused objects before the heap snapshot is taken.
 - **Heap Analysis Options:** Select the required option from the drop down 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. See [Taking a Heap Snapshot and Loading Into the Repository](#).
 - **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.
 - **Antipattern 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).
-  **Note:**
Leave this field blank if you want to do a heap dump only.
- To take a heap snapshot and load it manually to the repository, select the Heap Snapshot Type and leave the Heap Analysis Options field blank.
- **Heap Snapshot Time:** Schedule the heap snapshot to occur now or schedule it for a later time.
5. Click **Submit**. On the resulting dialog box, click **Yes** to continue the heap snapshot job. You can monitor the progress of the Heap Snapshot job. 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.

20.4.7 Taking a Heap Snapshot and Loading Into the Repository

Select this option to take a heap snapshot and automatically load it into the repository.

Prerequisites

- 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 `sqlldr` & `sqlplus` 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. The script is located inside `loadheap.zip`. You can find `loadheap.zip` at:

```
$MW_HOME/plugins/oracle.sysman.emas.oms.plugin_12.X.X.X.X/archives/loadheap.zip
```

The script is called by `loadheap.sh`. If you execute the script directly, you will be asked to input the required data. There is a `README.txt` file inside the `loadheap.zip` file that provides additional information.

To take a heap snapshot and load it into the repository, follow these steps:

1. Select the **Heap Snapshot and Load Into Repository** option. You can select this option if the Management Agent is running on the JVM Diagnostics Agent and the Heap Loader Host.
2. Specify whether the heap snapshot is taken immediately or at a later date.
3. Specify the credentials for the host on which the JVM Diagnostics Agent is running.
4. 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**.
5. 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.
6. 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.

20.4.8 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:

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 [Working with Class Histograms](#).

20.4.8.1 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 menu.
2. Select **Heap Snapshots and Class Histograms** from the Java Virtual Machine or Java Virtual Machine Pool target menu.

3. The list of available heap snapshots is displayed. Select a heap snapshot and click **Detail** to view the number of objects and memory reachable from each root. Click the **Roots** tab 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.
 - **Objects:** The total number of objects reachable from this root.
 - **Total Memory:** The total amount of memory reachable from this root.
 - **Adjusted Memory:** The adjusted memory reachable from this root. This parameter is useful in tracking the memory leak hot-spots.
4. Click the **Usage** tab to view the [Heap Usage by Objects](#).
5. Click the **Dominator Roots** tab to view the total size of all objects that would be removed when garbage collection is performed on this node.
6. Click the **Memory Leak Report** tab to view the [Memory Leak Report](#).
7. Click the **Anti-Pattern Report** tab to view the [Anti Pattern Report](#).

20.4.8.1.1 Top 40 Objects

This page shows the top 40 objects reachable from a root. The objects are sorted in descending order by the adjustable 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.

1. From the Targets menu, select **Middleware**, then click on a JVM or JVM Pool target.
2. On the JVM or JVM Pool Home page, select **Heap Snapshots and Class Histograms** from the Java Virtual Machine/Java Virtual Machine Pool menu.
3. The list of available heap snapshots is displayed. Select a heap snapshot and click **Detail** to view the number of objects and memory reachable from each root. Click the **Roots** tab to view the objects directly reachable from the root.
4. Click a root to view the top 40 objects.

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.
- **Field:**
- **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 (bytes):** 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.

20.4.8.1.2 Heap Object Information

This page shows information about a specific object in the heap snapshot. The following details are displayed:

- 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.
 - 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.
- 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.
 - 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 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.
- Retained Memory: The total size of all objects that would be removed when garbage collection is performed on this node.
- Depth: How far this parent is from the root.

20.4.8.1.3 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. From the **Java Virtual Machine** or **Java Virtual Machine Pool** menu, select **Heap Snapshots and Class Histograms**.
3. The list of available heaps is displayed. Highlight a heap and click **Detail**. 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 and adjusted reachable memory.
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.

20.4.8.2 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.
- **System:** System details.

Click **Compare with** to compare the heap snapshot with another one. See [Comparing Heap Snapshots](#).

20.4.8.3 Memory Leak Report

Click the **Memory Leak Report** tab to view the 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.

20.4.8.4 Anti-Pattern Report

Click the **Anti-Pattern Report** tab. The Anti-Pattern report is divided into different sections. Each section either shows the summary or one kind of anti-pattern issue. For example, 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.

20.4.9 Managing JFR Snapshots

Java Flight Recorder (JFR) provides a wealth of information on the inner workings of the JVM as well as on the Java program running in the JVM. You can use this information for profiling and for root cause analysis of problems. Furthermore, JFR can be enabled at all times, without causing performance overhead—even in heavily loaded, live production environments.

You can create JFR snapshots that include thread samples, which show where the program spends its time, as well as lock profiles and garbage collection details.

To create a JFR snapshot, follow these steps:

1. From the **Targets** menu, select **Middleware**, then click on a Java Virtual Machine target.
2. From the **Java Virtual Machine** menu, select the **JFR Snapshots** option.
3. Click **Create**. In the Create JFR Snapshot window, enter a description, schedule the snapshot, and click **Create**. The newly created snapshot appears in the JFR Snapshots page. Fields in the table include:
 - **Date:** Date the JFR snapshot was created.
 - **JVM Target:** JVM Target associated with this JFR snapshot.
 - **JFR Snapshot Description:** Description of the JFR snapshot.
 - **Host:** Host containing the JFR snapshot.

- **Path:** Path used to access the JFR snapshot
- **File Name:** File containing the JFR snapshot.

Settings

To start, stop, and remove a JFR recording, highlight a JFR in the table, and click **Settings**.

The JFR Administration page provides the statistics of the recording: Status, Data End Time, Data Start Time, Destination Compressed, Destination File, Continuous Recording, Duration (sec), Maximum Size (KB), Maximum Age (sec), and Start Time.

View Reports

Click **View Reports** to view GC and Latency reports. Latency and GC reports are available when the data is detected.

Downloading a JFR Snapshot

Use JMC (Java Mission Control) to download and analyze the JFR snapshot. Select the JFR snapshot and click **Download**. You are prompted for the host credentials. Enter the credentials and click **Download** and specify the location on which the snapshot is to be saved.

20.4.10 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. You can also modify the Bytecode Instrumentation (BCI). Click **Save** to save the changes.

20.4.11 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.

20.4.12 Adding a JVM to a 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.

20.5 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 that 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.

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 [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**.

20.5.1 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.
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.

20.6 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.

20.7 Viewing Heap Snapshots and Class Histograms

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.

To view and analyze the heap usage, select **Heap Snapshots and Class Histograms** from the Java Virtual Machine Pool or Java Virtual Machine menu. The following regions are displayed:

- **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.
 - 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.

- Used(%): The percentage of heap used.

You can do the following:

- 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.
- 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. The following options are available:
 - **Details:** Click this option to drill down to a detailed view of the heap.
 - **Compare:** Select two rows and click **Compare**. The Class Name, Instance Size (size of each snapshot), and Number of Instances (for each snapshot) are displayed.

20.8 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 Diagnostic Snapshots Page](#)
- [Analyzing a Diagnostic Snapshot](#)
- [Viewing a Diagnostic Snapshot](#)

20.8.1 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 **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.

20.8.2 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.

20.8.3 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.

20.8.4 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.

20.9 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, 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**.
2. In the View panel, click **Events without Incidents**. The JVM Diagnostics events are displayed if there are any outstanding JVMD threshold violations.
3. 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.

20.10 Using Java Workload Explorer

Java Workload Explorer provides a detailed view of all performance statistics associated with the JVM and JVM Pool targets.

20.10.1 Accessing Java Workload Explorer

To use Java Workload Explorer:

1. From the Targets menu, select **Middleware**, then select either a Java Virtual Machine target or a Java Virtual Machine Pool target.

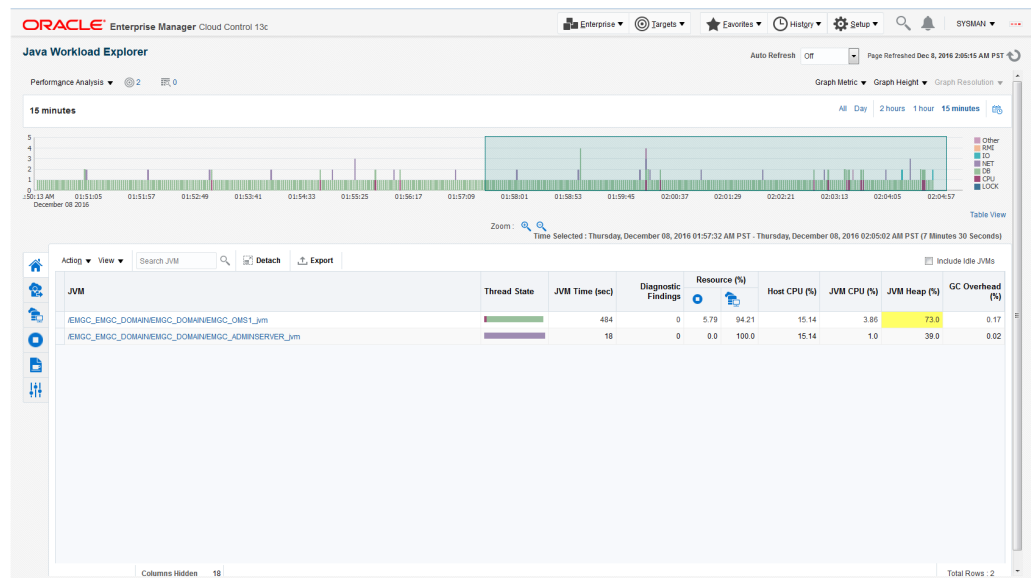
You can also access this page by selecting **Middleware** from the **Targets** menu and then selecting the **Middleware Features** menu. Select **Java Workload Explorer**.

2. On the resulting page, click the **Java Workload Explorer** link at the top of the page.

20.10.2 Performance Analysis and Search Criteria

The Performance Analysis menu provides the following features:

Figure 20-4 Performance Analysis



- **Compare**
Compare available snapshots against current data or sets including data from multiple JVMs across domains. This enables you to compare current activity to a saved baseline snapshot. Use this to proactively spot deviations after new application deployments, upgrades, or configuration changes in the target JVM.
- **Targets**
Select the targets you want to analyze. Remove targets that no longer apply.
- **Sets**
Use this option to create, open, save, and manage sets against which to compare current data sets.
- **Java Workload Report**
Provides insight into the performance of the JVM in a selected time window. The report is available for a maximum of 10 targets with a duration of no more than one hour duration. Creating the report enables you to analyze the data with data reported at different points in time.

The following tables are displayed in the Java Workload Report:

Summary tables for each target include:

- JVM Summary
- Diagnostic Findings
- Threshold Violations
- JVM Statistics
- OS Statistics
- GC Statistics

Multiple tables aggregated by all targets in the context and sorted by important metrics include:

- Requests Statistics
- Request Instances Statistics
- Session Statistics
- User Statistics
- Application Statistics
- Thread Statistics
- Method Statistics
- Class Statistics
- Packages Statistics
- Databases Statistics
- SQLs Statistics
- Database Events Statistics
- Database Schema Statistics
- Database Modules Statistics
- Database Actions Statistics
- Other External Resources Statistics
- Locks Statistics
- Files Statistics
- Supplemental Information
 - Contains JVM startup parameters, full SQLs, and Stacks

Search Criteria

The Search Criteria provided throughout the page enables you to fine tune your search and minimize the reported data.

Figure 20-5 Search Criteria

By default, the following keywords are used: ECID Duration (ms), SQL Duration (ms), and State. Using State enables you to select the Thread State in which you have interest. During any search, you can elect to ignore the field.

Using the Add Field menu, you can select fields for any of the following: request, user session, database, internal resource, code, and threads.

20.10.3 Graph Highlights

The graph at the top of the page provides a visual representation of the workload. Use this graph to quickly narrow down the time selection to the interval of interest. By default, the graph provides statistics on the active threads: RMI Wait, I/O Wait, Network Wait, DB Wait, CPU, and Lock. The data in the graph is available in table format.

Using the Graph Metric menu, you can narrow the graph to report on Memory Utilization, CPU Utilization, GC (garbage collection) Overhead, and Response and Load.

The Graph Height menu enables you to adjust the graph to display more details on the statistic.

The Graph Resolution menu enables you to see more spikes in the chart by increasing the number of data-points on the time axis (x-axis).

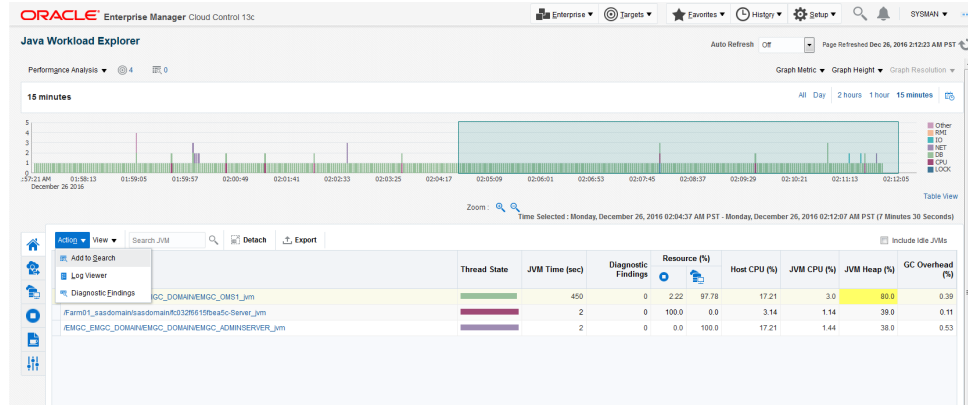
20.10.4 Diagnostics

The statistical data associated with the JVM or JVM Pool is available in the form of tabs. A diagnostic tab corresponds to a user intention based on a region or a set of related regions. A tab can have associated subtabs.

The majority of the tabs have an Action menu and a View menu. The options on the Action menu often replicate the options available on other parts of the screen, most notably Add to Search and Add to Set. Additional menu options are:


- Add to Search: It adds the element to the Search Criteria.

Figure 20-6 Add to Search



- Add to Set: It adds the element to the Set Criteria.

Figure 20-7 Add to set

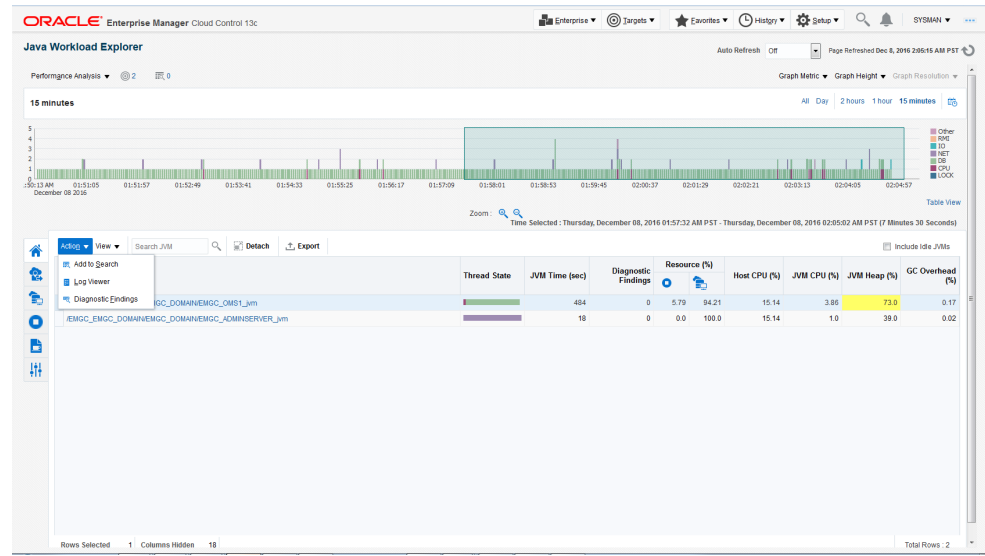


- Log Viewer (Displays Log Messages).
- Diagnostic Findings (Displays Middleware Diagnostics Advisor).
- View Call Tree (Displays the methods and the percentage of time for the call to execute to method).
- View Thread Transition (Displays the graphical view of how threads change over time).
- Session Diagnostics (This is a RUEI (Real User Experience Insight) based analysis and will be enabled only if the JVM target has been enabled on a RUEI system).

The tabs are:

- Overview:

Figure 20-8 Overview Tab

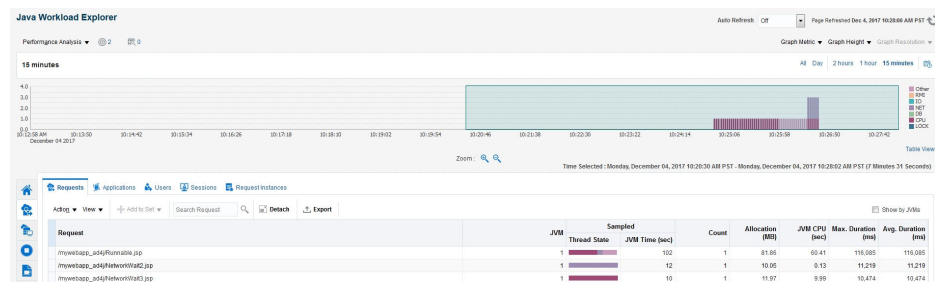


Statistics include: OSR, Context Switch (per sec), Host Memory (%), Swap Space (%), Open File Descriptors (%), Max Heap Size (MB), Min Heap Size (MB), Container Type, Container Name, and Version

Actions available are: Add to Search, Log Viewer, and Diagnostic Findings.

- Requests, ECIDs, Sessions, and Users.
 - Requests

Figure 20-9 Request Tab



Statistics include: Duration (ms), Max. Duration (ms), JVM CPU (sec), Allocation (MB), Count, JVM Time (sec), and Thread State.

Sample Request metrics (e.g. count, Allocation, Duration) are calculated based on data collected from specific instances that are **caught** while taking a sample of the stuck. The non-sample metrics are calculated based on all the instances that were executed since the previous sampling.

For example:

A Request average execution time varies from 50 to 1500ms, while the average execution time is 100ms. The request is executed 1000 times per second.

Sampling catches mainly the slow executions. The sampled count will be much smaller than 1000/s and the average execution time and other metrics will reflect the behavior of the **caught** slow executions. With the new feature introduced in 13.3, all the Request executions are counted and measured. In the example above, the count will show 100/s and the average will be 100ms.

 **Note:**

- * The **Thread state** and **JVM Time** metrics are available only as sampled data.
- * Request that are not sampled at all (not even one instance is seen the thread stuck when it is sampled by JVMD) will not show at all for that 2 second time period.

Actions available are: Add to Search, View Call Tree, Add to Set, and View Thread Transition.

– Applications

Statistics include: JVM Time (sec), Thread State, and Application Name.

Actions available are: Add to Search, Add to Set.

– Users

Statistics include JVM Time (sec), Thread State, and User.

Actions available are: Add to Search, Add to Set, and Log Viewer,

– Sessions

Statistics include: Minor GC Time (ms), Minor GC Count, Major GC Time (ms), Major GC Count, JVM Time (sec), Thread State, Number of Requests, User, and Session ID.

Actions available are: Add to Search, View Call Tree, Add to Set, and View Thread Transition.

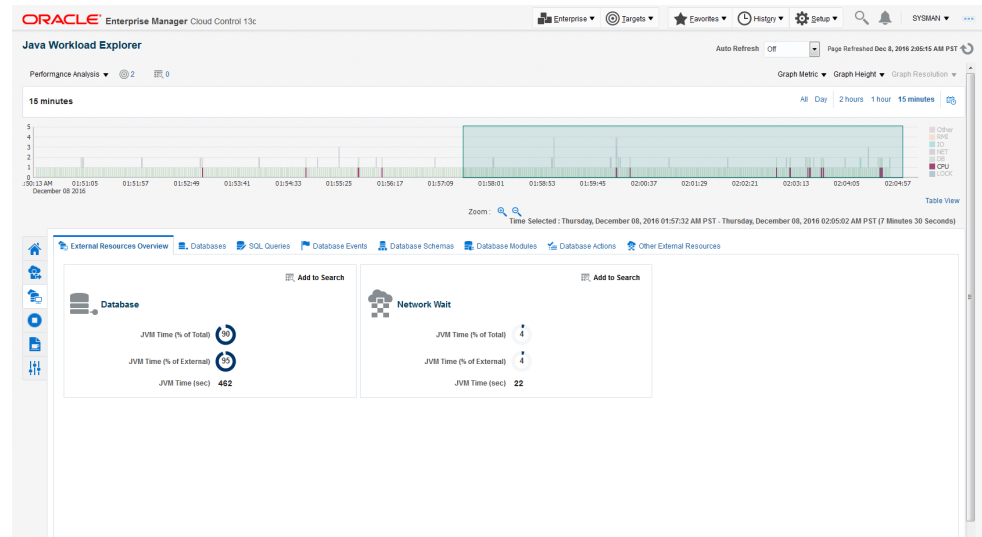
– Request Instances

Statistics include: Minor GC Time (ms), Minor GC Count, Major GC Time (ms), Major GC Count, GC Overhead (ms), Allocation (MB), JVM CPU (sec), Duration (ms), JVM Time (sec), Thread State

Actions available are: Add to Search, View Cal Tree, Add to Set, Log Viewer, and View Thread Transition, Session Diagnostics.

• External Resources

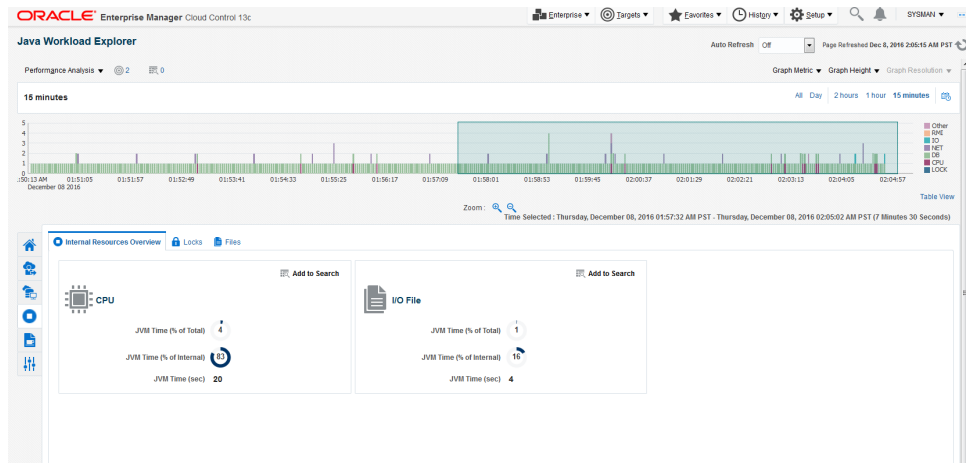
Figure 20-10 External Resources Tab



- External Resources Overview
 - Network Wait and Database (JVM Time (% of Total), JVM Time (% of Internal), and JVM Time (sec))
- Databases
 - Statistics include: Max Duration (ms), Avg. Duration (ms), Count, JVM Time (sec), and More Information, Database
 - Actions available are: Add to Search, View Call Tree, Add to Set, and Database Drill Down.
- SQL Queries
 - Statistics include: Max Duration (ms), Avg. Duration (ms), JVM Time (sec), Database, SQL ID, and SQL Statement
 - Actions available are: Add to Search, View Call Tree, Add to Set, and SQL Details
- Database Events
 - Statistics include: Max Duration (ms), Avg. Duration (ms), Count, JVM Time (sec), Database, and Database Event
 - Actions available are: Add to Search, View Call Tree, Add to Set
- Database Schemas
 - Statistics include: Max Duration (ms), Avg. Duration (ms), Count, JVM Time (sec), Database Schemas
 - Actions available are: Add to Search, View Call Tree, Add to Set
- Database Modules
 - Statistics include: Max Duration (ms), Avg. Duration (ms), Count, JVM Time (sec), Database Module
 - Actions available are: Add to Search, View Call Tree, Add to Set

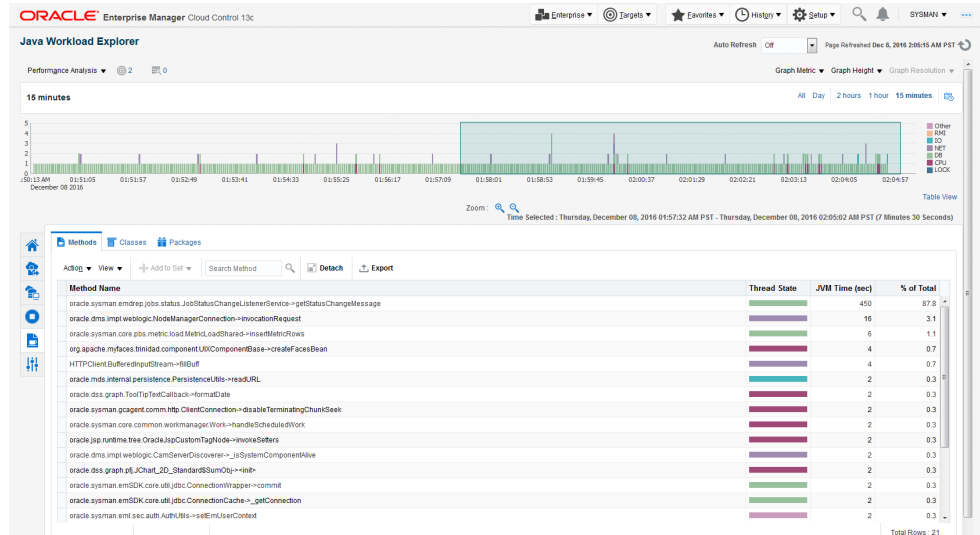
- Database Actions
 - Statistics include: Max Duration (ms), Avg. Duration (ms), Count, JVM Time (sec), Database Action
 - Actions available are: Add to Search, View Call Tree, Add to Set
- Other External Resources
 - Statistics include: JVM Time (sec), Protocol, Request
 - Actions available are: Add to Search, Add to Set
- Internal Resources

Figure 20-11 Internal Resources Tab



- Internal Resources Overview
 - * CPU - JVM Time (% of Total), JVM Time (% of Internal), and JVM Time (sec)
 - * Lock - JVM Time (% of Total, JVM Time (% of Internal), JVM Time (sec)
 - * I/O File - JVM Time (% of Total), JVM Time (% of Internal), and JVM Time (sec)
- Locks
 - Statistics include: Held Locks (JVM Time (sec), Avg. Duration (ms), Max Duration (ms) and Waiting Locks (Thread Trend, JVM Time (sec), Avg. Duration (ms), and Max Duration (ms)
 - Actions available are: Add to Search, View Details, Add to Set
- Files
 - Statistics include: I/O file and JVM Time (sec)
 - Actions available are: Add to Search, Add to Set
- Code

Figure 20-12 Code Tab



Statistics include: % of Total, JVM Time (sec), and Package

- Methods

Actions available are: Add to Search, View Call Tree, Add to Set

- Classes


Actions available are: Add to Search, Add to Set

- Packages

Actions available are: Add to Search, Add to Set

- Threads

Figure 20-13 Threads Tab



Statistics include: Write Characters, Read Characters, Wait Count, Waited Time (sec), Blocked Count, Blocked Time (ms), Hogger (%), Stuck (%), Allocation (MB per sec), JVM CPU (per minute)

Actions available are: Add to Search, View Call Tree, Add to Set, Log Viewer, View Thread Transition, Sample Analysis

20.11 Managing and Troubleshooting JVMD (Globally)

If you find that a particular JVM Diagnostics Engine is exhibiting problems, the Manage and Troubleshoot JVMD functionality provides the statistics and diagnostic aids to help resolve the issue.

To access the Manage and Troubleshoot JVMD page:

1. From the **Setup** menu select **Middleware Management**, then select **Engines and Agents**.
2. In the RUEI/JVMD Engines section, highlight the JVM Diagnostic Engine of interest and click **Troubleshoot**.

Statistics include:

- Repository Statistics

The Tablespace Growth Rate chart provides the Total Space and Used Space used by the repository over specific time intervals. The related repository tables are listed. To view the statistics of a particular repository table, highlight the table name and click **Details**.

Click the **Trend** button to view the used and allocated data for each date. This data is based on the statistics collected by the DBMS_SPACE.OBJECT_GROWTH_TREND function and enables you to see trends in the usage of the space.

Click **Export** to retrieve a table listing all the tables and their associated statistics, for example, Table Allocated Space (MB), Index Allocated Space (MB), Number of Rows, and Last Analyzed Time. **Note:** Before clicking Export, show all the columns (on the **View** menu, select **Columns**, then select **Show All**). This provides a better view of the columns in the table.

The data provided in the JVMD Operations Statistics region enables you, as the JVMD Administrator, to monitor your own applications.

- JVM Target Summary

This page provides data about the JVM Agent.

Summary section lists agent statistics such as the number of targets that are down, and the number of unassociated targets. You can manage JVMD Agents located on WebLogic Server Domains, start and stop monitoring of a JVM target, and export data.

- Engine Summary

This page provides statistics regarding the JVM engine. When you highlight the engine, the associated attributes display in the Engine Attributes table. The engine summary includes the following types of attributes: Performance, Diagnostics, and Configuration.

Also, if there are any load balancers configured, the JVMD Load Balancer table provides additional information.

Troubleshooting diagnostic aids include:

- View JVMD Health Jobs

This link directly navigates to the Job system page and by default shows all the JVMD health jobs.

The JVMDHealthReportJob job is automatically invoked every three hours. This job collects statistics for that three hour time period. Select the job for the time period of interest and click **View Results**. This is an historical view of the health of the JVMD. Name of the report is JVMD_HEATH_REPORT_AUTO.

- SR Assistance

In the event that there are issues with the JVMD, click the SR Assistance button for an explanation regarding common JVMD issues. This page also lists the statistics you need to have available before filing a Service Request.

- Generate Report

Provides the same information as is available in the Manage and Troubleshoot JVMD tabs, that is, it shows the trends of the various JVMD components. Click **Save to File** to save the information to an .html file that you can easily access at a later time.

20.12 Managing and Troubleshooting JVMD (Specific Agent)

Should you find that a particular JVM or JVM Pool is sluggish or is posing problems, the Manage and Troubleshoot JVMD functionality provides the statistics and diagnostic aids to help resolve the issue.

To access the Manage and Troubleshoot JVMD page:

1. From the **Targets** menu, select **Middleware**, then select a Java Virtual Machine or Java Virtual Machine Pool.
2. On the resulting page, select **Manage and Troubleshoot JVMD** from the Java Virtual Machine or Java Virtual Machine Pool menu.

JVM Target Summary Tab

Statistics include:

- Status and Connection

Data includes the engine host and availability, as well as JVMD Agent status, Monitoring status, and Bytecode Instrumentation (BCI) status.

- Target Attributes related to target. Attributes include: Performance, Diagnostics, and Configuration attributes.

Click the MBean Browser button to view JVMD Agent MBean data and it's live call to the JVMD Agent.

This data can be exported which is very helpful in diagnosing the JVMD Agent related issues.

- Performance and Diagnostics (Poll Interval (ms), Response Time (ms), Average Stack Depth (count), Number of Active Threads)

Troubleshooting diagnostic aids include:

- **Java Virtual Machine menu**
Provides links to performance diagnostics, thread snapshots, and configuration options.
- **Java Workload Explorer**
Provides a detailed view of all performance statistics associated with the JVM or JVM Pool.
- **Live Thread Analysis**
Shows the real-time data for all the JVMs in the selected pool or the real time data for the selected JVM.
- **SR Assistance**
In the event that there are issues with the JVMD, click the SR Assistance button for explanation regarding common JVMD issues. This page also lists the statistics you need to have available before filing a Service Request.
- **Generate Report**
Provides the same information as is available in the Manage and Troubleshoot JVMD tabs. Click **Save to File** to save the information to an .html file that you can easily access at a later time. Purpose of job is that it shows the trends of the various JVMD components.

Manage Association Tab

Target Association lists the Enterprise Manager targets with which this JVM is associated. You can associate and disassociate targets, and export the information to a spreadsheet.

20.13 Enable or Disable Monitoring of JVM Targets using EMCLI

You can also enable or disable the monitoring of JVM target using EMCLI.

Run the following command to deploy JVMD targets:

```
emcli deploy_jvmd -domain_name="/Farm03_base_domain/base_domain" -  
enableMonit="false"
```


21

Troubleshooting JVM Diagnostics

This section describes the errors you may encounter while deploying and using JVM Diagnostics and how to resolve the issues. It contains the following:

- [Cross Tier Functionality Errors](#)
- [Trace Errors](#)
- [Deployment 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](#)

21.1 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 to the following:

- Mismatched database connection information
- Insufficient user privileges

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 21-1 Live Thread Analysis (Cross Tier)

The screenshot displays the Oracle Enterprise Manager Cloud Control 13c interface. The main window shows 'Live Thread Analysis' for the JVM process 'EMGC_OMS1_jvm'. A 'Database Details' popup window is open, showing the following information:

- Database Host: 10.64.147.533
- Database Port: 1521
- Database SID: Not Available
- Database Service: N121
- Database User: SYS
- Database URL: jdbc:oracle:thin:@1066mdu302.us.oracle.com:1521:14121

The popup also contains instructions: 'To view the Database activity please ensure the following: Database target to registered with JVM Diagnostics. To view Register Database page click here. To associate a Registered database to the above Database URL click here. JVM Diagnostics Database agent is running on the database machine.'

The background interface shows a table of JVM threads. The table has columns: Request, JVM CPU, Allocation, Current Call, State, Waiting On, and Wait Request. The table contains two rows of data:

Request	JVM CPU	Allocation	Current Call	State	Waiting On	Wait Request
emfAssoc2Data	3.33	2.27	oracle.system.c2Data	Other		
ThreadName_jdb	3.7	2.49	jdb_sendData	DB Wait		SQLMsg

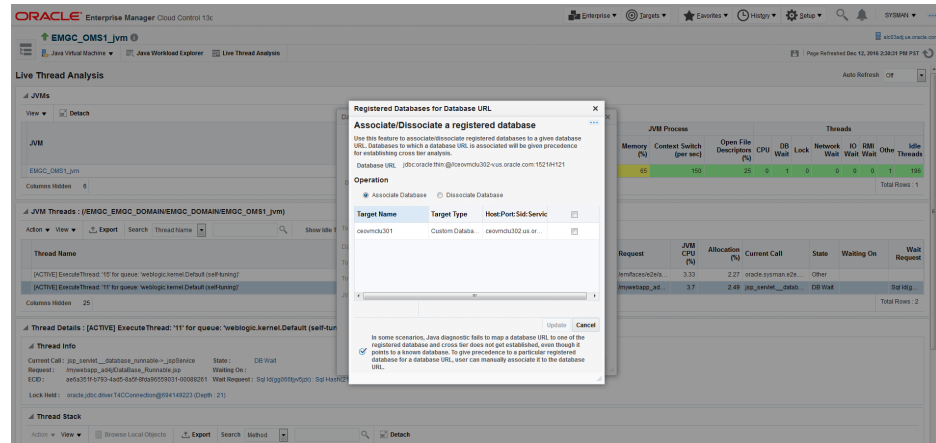
 **Note:**

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 the 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$process`, and `v$session_wait` tables.
- 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 21-2 Live Thread Analysis: Associate / Disassociate a Registered Database



- **Java Workload Explorer:** Provides a detailed view of all performance statistics associated with the JVM or JVM Pool.
- **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 21-3 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 **Engines And Agents**. 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.

21.2 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.

21.3 Deployment 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*.
 - Check if there are any pending changes. If any changes are pending, activate or undo these changes as appropriate 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(ExceptionTranslat
```

```

or.java:42)
at
weblogic.jndi.WLInitialContextFactoryDelegate.toNamingException(WLInitialContextF
actoryDelegate.java:788)
at
weblogic.jndi.WLInitialContextFactoryDelegate.pushSubject(WLInitialContextFact
oryDelegate.java:682)
atweblogic.jndi.WLInitialContextFactoryDelegate.newContext(WLInitialContextFactor
yDelegate.java:469)
at
weblogic.jndi.WLInitialContextFactoryDelegate.getInitialContext(WLInitialConte
xtFactoryDelegate.java:376)
at weblogic.jndi.Environment.getContext(Environment.java:315)
at weblogic.jndi.Environment.getContext(Environment.java:285)
at
weblogic.jndi.WLInitialContextFactory.getInitialContext(WLInitialContextFactor
y.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:** This exception may occur, either if the path to the `weblogic.jar` is invalid, or the user does not have read permissions on the `weblogic.jar` file.

```

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)

```

Solution: Ensure that the correct path is provided or the user credentials allow read access to the jar file.

- **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 the 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 that is monitoring the WebLogic Server target.

- **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:** 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://emHost.example.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/mw923/weblogic92/server/lib/weblogic.jar&&ls
/scratch/aime/mw923/weblogic92/server/lib/wljmxclient.jar&&ls
/scratch/aime/mw923/weblogic92/server/lib/wlcipher.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://emHost.example.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

```

21.4 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.

21.5 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.

21.6 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 emHost.example.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, follow these steps:
 1. From the **Setup** menu, select **Application Performance Management**.
The list of Application Performance Management Engines is displayed.
 2. Select the JVM Diagnostics Engine row, click **Configure** then click the **Register Databases** tab.
 3. 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.

21.7 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](#)

21.7.1 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
`<path to gc_inst>/em/EMGC_OMS1/sysman/log/jvmdlogs/jvmdengine.log.0`
- 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`

21.7.2 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 **Engines And Agents**.
- 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.

21.7.3 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.`

21.7.4 Monitoring Status

To verify if the JVM Diagnostics Engine is monitoring the data:

1. From the **Setup** menu, select **Middleware Diagnostics**, then click **Engines And Agents** 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.

21.7.5 JVMD SLB Configuration

The JVM Diagnostic engine may go down due to the following reasons:

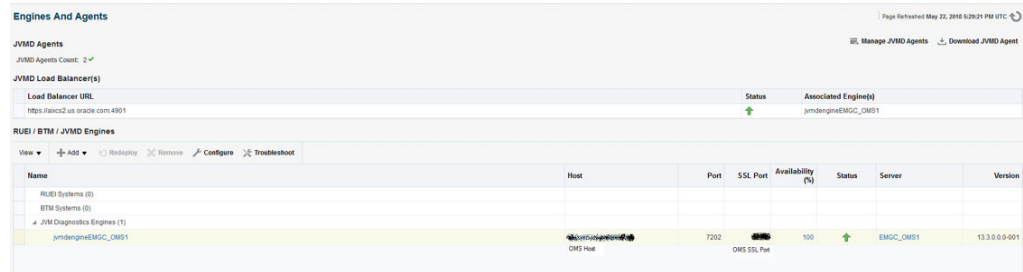
- SLB is not configured properly.

- OMS port was blocked by firewall.

To make JVMD engine accessible, OMS port must be unblocked and accessible by SLB.

The below figure shows the JVMD SLB configuration on Enterprise Manager console.

Figure 21-4 JVMD SLB Configuration



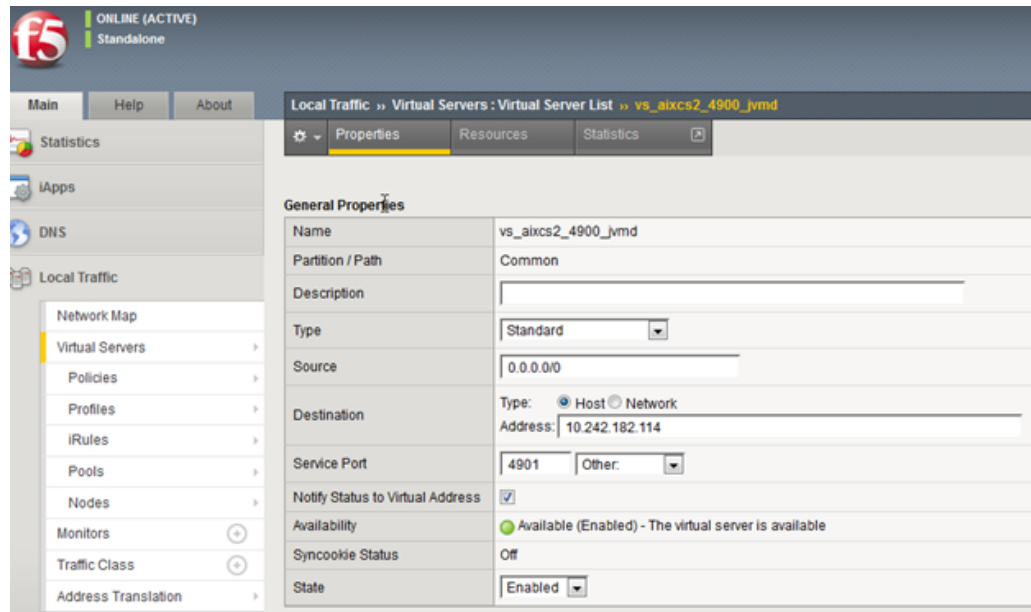
Virtual Server on SLB

SLB Virtual server port: 4901 aixcs2.us.oracle.com

IP Address configured in virtual server of SLB: 10.242.182.114

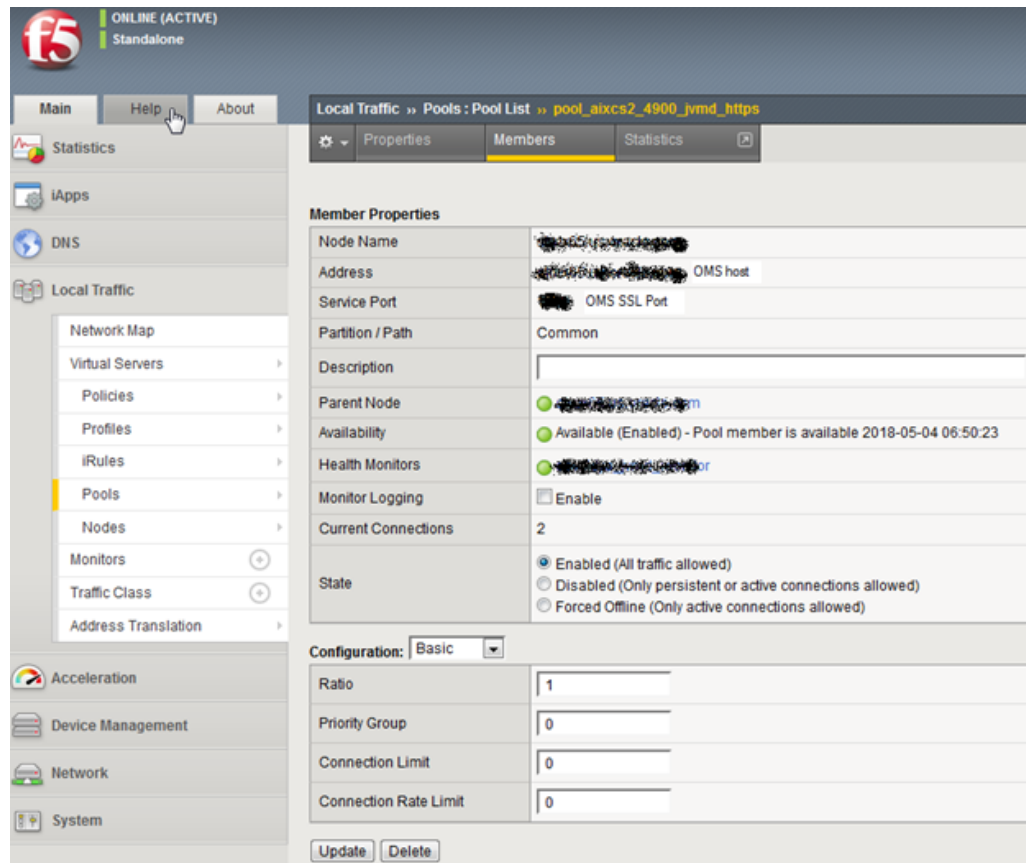
The figure below shows the virtual server configuration details on SLB:

Figure 21-5 Virtual Server on SLB



Go to **Resources** tab and note down the Default pool. Now, open the **Pools** menu in the Local Traffic panel, and go to **Members** tab. Each active member indicates an OMS configured in the EM.

For example, in the figure below, we have only one OMS. Hence, we have only one active member. You must make sure that the member host has correct properties as OMS host and OMS SSL port.



For more information about configuring the F5 SLB for EM and JVMD, see [Configuring OMS High Availability with F5 BIG-IP Local Traffic Manager](#).

21.7.6 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.

21.7.7 Usage of the Try Changing Threads Parameter

This parameter should be used only when the JVM is highly active.

21.7.8 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.

21.7.9 Custom Provisioning Agent Deployment

You can customize the JVM D 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 Management**, then click the **Engines And Agents** and on top right click **Download Jvmd Agent** to 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 JVM D Agent to all the specified servers.

21.7.10 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

21.7.11 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

This section describes Middleware Diagnostics Advisor (MDA). MDA as a part of Enterprise Manager Cloud Control analyzes the entire stack and provides diagnostic findings by identifying the root cause of a problem.

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 a JDBC connection pool that 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 for the server(s).

This section covers the following:

- [Diagnosing Performance Issues with Oracle WebLogic Server](#)
- [Diagnosing Performance Issues Using Middleware Diagnostics Advisor](#)
- [Functioning of Middleware Diagnostics Advisor](#)
- [Prerequisites for Configuring Middleware Diagnostics Advisor](#)
- [Configuring Middleware Diagnostics Advisor](#)
- [Setting Up Middleware Diagnostics Advisor \(MDA\)](#)
- [Enabling Middleware Diagnostics Advisor for a Target](#)
- [Limiting the Scope of Middleware Diagnostics Advisor](#)
- [Using Middleware Diagnostics Advisor to View and Diagnose Performance Issues](#)
- [Running an Unscheduled Middleware Diagnostics Advisor Analysis on a Target](#)
- [Troubleshooting Issues Related to Middleware Diagnostics Advisor](#)

22.1 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 the performance, and scalability of WLS, it is best to detect violations 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, see [Diagnosing Performance Issues Using Middleware Diagnostics Advisor](#).

22.2 Diagnosing Performance Issues Using Middleware Diagnostics Advisor

Middleware Diagnostics Advisor or MDA is a diagnostic module integrated within Enterprise Manager (EM) Cloud Control for diagnosing performance issues with middleware targets monitored by Enterprise Manager Cloud Control. Currently, MDA is supported for Oracle WebLogic Server 10g Release 3 (10.3) and higher. MDA monitors JDBC DataSources, 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 performance of aspects like JMS message consumption time, etc. in a runtime environment. When the performance of the aspect degrades beyond a certain limit, MDA diagnoses the issue to find the underlying cause, and the problems detected by MDA are projected as Diagnostic Findings. 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 and each of these areas is listed as a Finding Type in the Middleware Diagnostics Advisor Configuration page (see [Figure 22-1](#)):

- **JDBC Datasource Wait (JVMD mandatory)**
Investigates excessive wait for JDBC Datasource connection.
- **Delivery of Message is Delayed (JVMD mandatory)**
Investigates if the time taken to pick up a message for processing is higher than the specified time. If this is not checked it may lead to messages spending more time in the queue than expected.
- **Message Processing is slow (JVMD mandatory)**
Investigates if the queue processing is slow because of which messages are being processed much slower than they are being received. If this is not checked it may cause the queue to grow and eventually lead to out-of-memory errors.

22.3 Functioning of Middleware Diagnostics Advisor

Middleware Diagnostics Advisor functions in the following way:

1. The MDA engine starts when the Oracle Management Service (OMS) is started.
The MDA engine is responsible for executing MDA analysis tasks.
2. Run the following command in the Enterprise Manager command line interface to enable MDA.

```
emcli update_mda_properties -props="MDA_AUTO_ENABLE:1"
```


 **Note:**

MDA is not enabled in Enterprise Manager Cloud Control by default. This step is necessary to enable MDA.

If any targets have been added to MDA before performing this step, they can be explicitly enabled by executing the command:

```
emcli enable_mda_finding_types_for_targets -
finding_types="finding_type_name"-targets="target_name:target_type"
```

3. The Enterprise Manager collected metric repository data and the JVMMD agent collected repository data is accessed by MDA.
4. The MDA engine runs the analysis for each finding type as scheduled on all the applicable targets. For example, once an hour for JDBC Datasource Wait finding.

 **Note:**

When a middleware target is in Black-out or Notification Black-out state, the data will not be processed for that target. Therefore, all scheduled analysis runs for such targets are skipped.

5. During the analysis, rules are applied to see if there are issues and the results are stored in the repository.
6. The findings are displayed on the Middleware Diagnostics Advisor page.

22.4 Prerequisites for Configuring Middleware Diagnostics Advisor

Before you begin using Middleware Diagnostics Advisor (MDA) for diagnosing performance issues, ensure the following prerequisites are met:

- The WebLogic Server is discovered as a target in Enterprise Manager.
- JVMMD Manager is configured, and the JVMMD Agent is deployed on the target server.
- It is recommended to force the configuration collection for the target.

To force the configuration collection for the target, perform the following steps:

1. From the target menu, select **Configuration** and then select **Last Collected**.
2. Click **Refresh**.

22.5 Configuring Middleware Diagnostics Advisor

Middleware Diagnostic Advisor (MDA) has an auto-enable job that runs every hour to check for new targets. However, these jobs will not be enabled for MDA analysis by default unless you have run the command `emcli update_mda_properties -`

`props="MDA_AUTO_ENABLE:1"`. After the command has been executed successfully, and if the prerequisites (see [Prerequisites for Configuring Middleware Diagnostics Advisor](#))

are met, the target is automatically enabled for MDA analysis in the next run of the job. However, to immediately enable a newly discovered middleware target for MDA analysis follow the steps below:

1. On the Enterprise Manager home page, from the **Targets** menu, select **Middleware**.
2. Click the newly discovered target domain link from the Middleware targets list.
3. On the target domain home page, from the target menu select **Diagnostics**, and then select **Middleware Diagnostics Advisor Configuration**.
4. Click **Health Check** in the **Registered Finding Types** section.

A health check is performed, wherein in addition to performing a health check of MDA, any new middleware targets added in Enterprise Manager will also be discovered by MDA and enabled for all the applicable finding types.

22.6 Enabling Middleware Diagnostics Advisor for a Target

After the `MDA_AUTO_ENABLE` property has been set to "1" (enable) using the EM CLI verb `emcli update_mda_properties -props="MDA_AUTO_ENABLE:1"`, any new WLS targets discovered in EM are enabled by default. However, if you do not want to enable every WLS server managed by the EM instance, you can set `MDA_AUTO_ENABLE` to "0" (disable), and individually enable specific targets either from UI (by following the steps below) or from EM CLI with the command `emcli enable_mda_finding_types_for_targets -finding_types="finding_type_name"-targets="target_name:target_type"`.

To manually enable or disable MDA, follow these steps:

1. On the Enterprise Manager home page, from the **Targets** menu, select **Middleware**.
2. From the **Middleware Features** menu select **Middleware Diagnostics Advisor Configuration**.
3. On the Middleware Diagnostics Advisor Configuration page, from the Registered Finding Types section select any one of the Finding Type.
4. On the Targets section, select the target that you want to enable MDA for, and click **Enable**. See [Figure 22-1](#).

Note:

Unless a target is discovered by MDA it will not be displayed in the targets list. To discover a target in MDA and enable it, see [Configuring Middleware Diagnostics Advisor](#).

Figure 22-1 Middleware Diagnostics Advisor Configuration Page

The screenshot shows the 'Middleware Diagnostics Advisor Configuration' page. It features a search bar at the top left and a table of 'Registered Finding Types' on the left side. The main area displays 'Targets applicable for JDBC Datasource Wait' with a table showing target names, target types, and statuses. Below this, there is a table of 'Analysis Runs for selected targets (Latest 10 runs)' with columns for Name, Finding, Target Name, Target Type, Status, Analysis Time, and Duration (seconds).

Target Name	Target Type	Status
EMGC_DOMAIN	Oracle WebLogic Domain	2/2 Enabled

Name	Finding	Target Name	Target Type	Status	Analysis Time	Duration (seconds)
JDBC Datasource Wait: 1167.645 <=	No	EMGC_ADMINSERVER	Oracle WebLogic Server	Analysis Run C...	23-Nov-2015 10:10:55 IST	0.278
JDBC Datasource Wait: 1167.648 <=	No	EMGC_OMS1	Oracle WebLogic Server	Analysis Run C...	23-Nov-2015 10:10:55 IST	0.278
JDBC Datasource Wait: 1165.643 <=	No	EMGC_ADMINSERVER	Oracle WebLogic Server	Analysis Run C...	23-Nov-2015 09:10:52 IST	0.265
JDBC Datasource Wait: 1165.644 <=	No	EMGC_OMS1	Oracle WebLogic Server	Analysis Run C...	23-Nov-2015 09:10:52 IST	0.265
JDBC Datasource Wait: 1161.641 <=	No	EMGC_ADMINSERVER	Oracle WebLogic Server	Analysis Run C...	23-Nov-2015 08:10:49 IST	0.317
JDBC Datasource Wait: 1161.642 <=	No	EMGC_OMS1	Oracle WebLogic Server	Analysis Run C...	23-Nov-2015 08:10:49 IST	0.317
JDBC Datasource Wait: 1142.627 <=	No	EMGC_ADMINSERVER	Oracle WebLogic Server	Analysis Run C...	23-Nov-2015 07:10:47 IST	0.181

22.7 Setting Up Middleware Diagnostics Advisor (MDA)

MDA can be set according to your preference. To set up MDA, 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, you can do the following:
 - In the **Analysis Job Configuration** section, select **Skip Analysis Runs for all MDA-Enabled Servers**, if you want to skip all MDA analysis jobs.
 - In the Purge Policy section you may choose to enable purging by selecting the **Purge Data Older Than** check box.

To set your preferred frequency for purging data, enter the preferred number of days for which the data should be retained in the text box. There is an MDA job that runs every 24 hours which purges data from the repository.

Note:

This is a global setting and will be applied to all targets, and all users.

- In the **Finding Threshold Configuration** section, set the threshold or limit (in percentage) beyond which violations should result in a finding, by adjusting the **Violations Percentage**. The default value is 10%.

Note:

This setting is applicable only to JMS and JDBC findings.

- To set the wait time period (in minutes) beyond which any messages picked up will be considered as violations, adjust the **JMS Wait Time**. The default value is 1 minute.

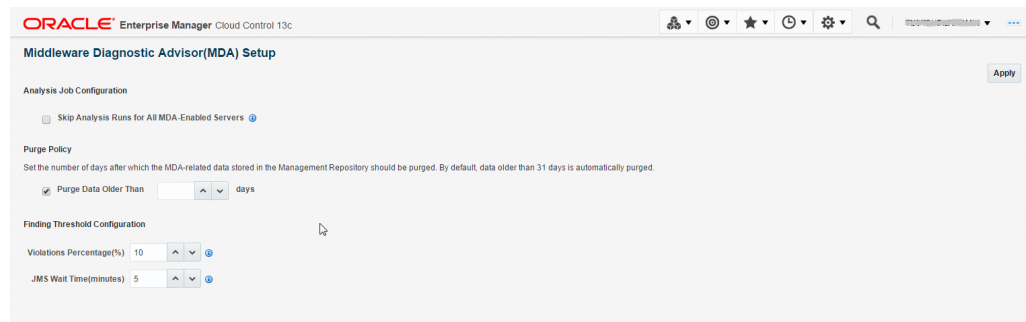
 **Note:**

This setting is applicable only to JMS wait time findings.

3. Click **Apply**.

Figure 22-2 shows the Middleware Diagnostics Advisor Setup page.

Figure 22-2 Middleware Diagnostics Advisor Setup Page



 **Note:**

Additionally, you can also use EM CLI commands to configure MDA. For details, see *Enterprise Manager Command Line Interface Guide*.

22.8 Limiting the Scope of Middleware Diagnostics Advisor

You can limit the scope of Middleware Diagnostic Advisor (MDA) by disabling MDA on some or all of the targets. If disabled, no analysis runs will be scheduled for the disabled targets. The targets can also be enabled when required.

To disable a target for MDA analysis follow the steps below:

1. On the Enterprise Manager home page, from the **Targets** menu, select **Middleware**.
2. From the **Middleware Features** menu select **Middleware Diagnostics Advisor Configuration**.
3. On the Middleware Diagnostics Advisor Configuration page, from the Registered Finding Types section select any one of the Finding Type.
4. In the Targets section, select the target that you want to disable MDA for, and click **Disable**.

To enable a target again for MDA analysis, see [Enabling Middleware Diagnostics Advisor for a Target](#).

For more information on limiting the scope of MDA, see the [Oracle Enterprise Manager Command Line Interface](#) guide.

22.9 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.



Note:

The target can be an Oracle WebLogic Domain, Oracle WebLogic Cluster or an Oracle WebLogic Server. Go to an Oracle WebLogic Server home page to see the findings for only the server. Go to a domain or a cluster home page to see the findings for all the targets under them.

2. From the target specific menu, select **Diagnostics**, then select **Middleware Diagnostics Advisor**.
3. On the Middleware Diagnostics Advisor page, view the Timeline section for findings marked against the time at which the finding was recorded.

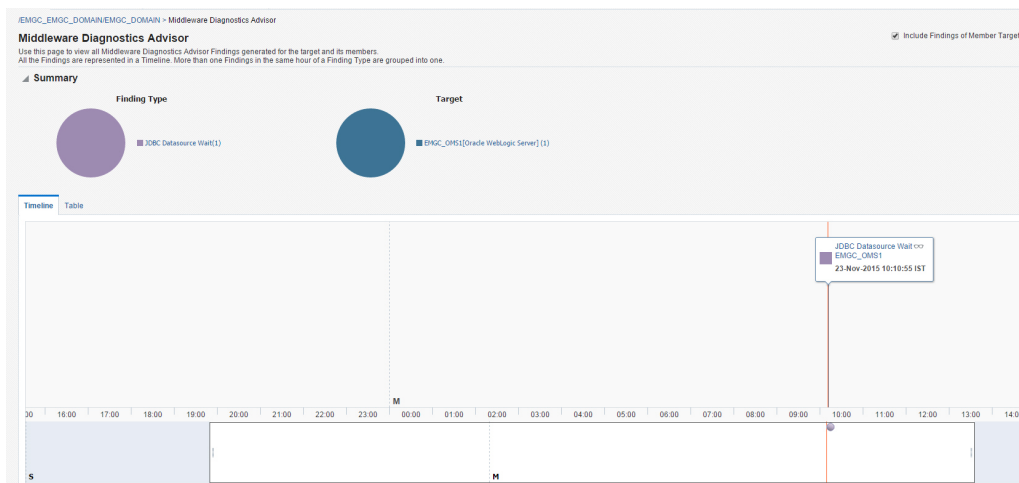
See [Figure 22-3](#) for a sample of the Middleware Diagnostics Advisor page.



Note:

To view historic data adjust the settings in the **View Data** drop-down box.

Figure 22-3 Middleware Diagnostics Advisor Page



4. Click on the **Quick View** icon next to the finding link to have a quick look at the details of the finding.

5. To know more about the findings, if any, click on the finding link in the Timeline section.

The Middleware Diagnostics Advisor (MDA) Finding Details page provides more information on the finding, and the recommended workaround. The following details are provided:

- **Finding**

The diagnostic finding for the Middleware domain.

For example: High number of messages reprocessed due to Transaction timeout.

- **Description - Target Type, Target, and Analysis Time**

The description and details related to the finding.

- **Analysis**

- **Rule**

The rule/s that are applied on the collected values to determine if there is a finding.

- **Execution Data**

The values collected during analysis.

- **Recommendations**

- **Action**

A solution or tip for the problem found.

- **Rationale**

The rationale applied for the suggested action.

- **Charts**

The Charts section contains graphs pertaining to the finding.

22.10 Running an Unscheduled Middleware Diagnostics Advisor Analysis on a Target

The MDA engine runs analysis for each finding type as scheduled on all the applicable targets. You can also run an unscheduled on-demand MDA analysis on a target. Follow the steps below to do so:

1. On the Enterprise Manager home page, from the **Targets** menu, select **Middleware**.
2. Click the target domain link from the Middleware targets list.
3. On the target domain home page, from the target menu select **Diagnostics**, and then select **Middleware Diagnostics Advisor Configuration**.
4. On the Middleware Diagnostics Advisor Configuration page select the finding type from the Registered Finding Types section to list the applicable targets in the Targets section.
5. Select any one of the targets to view the last 10 analyses run on that target in the Analysis Runs section.



Note:

Ensure that the **Show Analysis Runs with Findings Only** check box is disabled to see all the ten analysis runs.

6. Click **Run MDA Analysis** to run an MDA analysis now.



Note:

MDA runs the analysis even if the target is in blackout or notification blackout state.

22.11 Troubleshooting Issues Related to Middleware Diagnostics Advisor

To troubleshoot issues related to MDA, follow the steps mentioned in the following table.

Table 22-1 Troubleshooting Tips for Middleware Diagnostics Advisor

Tip	How to
Ensure that MDA is enabled for the specified finding type on the target.	<ol style="list-style-type: none"> 1. On the target domain home page, click the target menu. 2. Select Diagnostics menu, and then select Middleware Diagnostics Advisor Configuration. 3. Select the finding type. 4. Select the target from the Targets section. Note: If the target is not visible in the Targets section, click Health Check in the Registered Finding Types section. 5. Verify whether the status column for the target displays "Enabled" for the target. <p>If the status is disabled, enable MDA for the specified finding type for the target. Refer the <i>Enterprise Manager Command Line Interface Guide</i> for more options.</p>
Ensure that MDA auto-enable job is running properly	<ol style="list-style-type: none"> 1. From the Enterprise menu, select Job and then Activity. 2. In the Available Criteria section under Name, enter MDA_ADMIN_AUTO_ENABLE% in the text box and click the search icon. 3. In the Available Criteria section under Target Type, select Targetless. 4. In the Available Criteria section under Status, select Problems. <p>View all the jobs displayed in the table. If there are no skipped or failed jobs, it indicates that the auto-enable job is running as expected.</p>

Table 22-1 (Cont.) Troubleshooting Tips for Middleware Diagnostics Advisor

Tip	How to
<p>Ensure that MDA analysis jobs are running properly.</p>	<ol style="list-style-type: none"> 1. From the Enterprise menu, select Job and then Activity. 2. In the Available Criteria section under Name, enter MDA_ANALYSIS_RUN% in the text box and click the search icon. Note: There is one MDA_ANALYSIS_RUN job per Finding Type. By selecting MDA_ANALYSIS_RUN%, the jobs for all the Finding Types are displayed. 3. In the Available Criteria section under Target Type, select Targetless. 4. In the Available Criteria section under Status, select Problems. <p>View all the jobs displayed in the table. If there are no skipped or failed jobs, it indicates that the analysis job is running as expected.</p>
<p>Ensure that the MDA engine is running normally.</p>	<p>Refer the <i>Enterprise Manager Command Line Interface Guide</i> for the command to check the status of MDA engine.</p>

Part VIII

Managing Oracle Coherence

The chapters in this part contain information on discovering and monitoring a Coherence cluster.

The chapters are:

- [Getting Started with Management Pack for Oracle Coherence](#)
- [Monitoring a Coherence Cluster](#)
- [Administering a Coherence Cluster](#)
- [Troubleshooting and Best Practices](#)
- [Coherence Integration with JVM Diagnostics](#)

23

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 13c. The following sections are covered in this chapter:

- [About Coherence Management](#)
- [New Features for Oracle Coherence](#)
- [Configuring a Coherence Cluster](#)
- [Discovering Coherence Targets](#)
- [Enabling the Management Pack](#)

23.1 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. 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 between the caches, nodes, hosts, and Oracle WebLogic targets.

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.

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. The Topology Viewer provides a high level topology of the entire cluster and shows the relation between caches, nodes and hosts.

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.

23.2 New Features for Oracle Coherence

This section lists the new features in Oracle Enterprise Manager Cloud Control 13c. The new features are:

- **Target Navigation:** The target navigation menu is hidden by default. Click the Target Navigation icon on the left hand side of the page to display the Target Navigation tree. The navigation tree makes it easier to navigate to any node or cache from any page. Nodes are grouped based on the hosts on which they are running, and caches are grouped based on services. See [Navigation Tree](#).
- **Heat Map:** The Cluster Home page now has a Heat Map tab which provides a graphical representation of all the targets in the cluster.
- **Log Viewer:** You can now scan the log files and view the log file data.
- **Remove Down Members:** You can now delete any members of the cluster that have a **Down** status.
- **Managed Coherence Clusters:** The discovery and monitoring of managed Coherence clusters has now been integrated with Oracle Fusion Middleware/ Weblogic Domain discovery and monitoring.
- **Custom Topology Viewer:** The custom topology viewer provides a customized view of all the targets in a Coherence cluster.
- **Federated Caching:** You can synchronize cache instances across clusters by replicating updates from a cache in a source cluster to a cache in a remote cluster.
- **Discovery of Dynamic Management Node:** The coherence 12.2.1.x and above can be configured to start in dynamic management mode. Dynamic management mode automatically selects the senior cluster member as the management node (JMX cluster member). If the central management node is not operational, then the next most senior coherence node is automatically selected as the management node.
- **Federated Operations:** A new Federated Caching Operations region is added to Service Home Page. This allows the user to submit operation request. Based on operation selection, federation operation will be executed with reference to that remote participant.
- **Target Discovery:** You can now disable cache target discovery during new discovery and refresh.

23.3 Configuring a Coherence Cluster

Note:

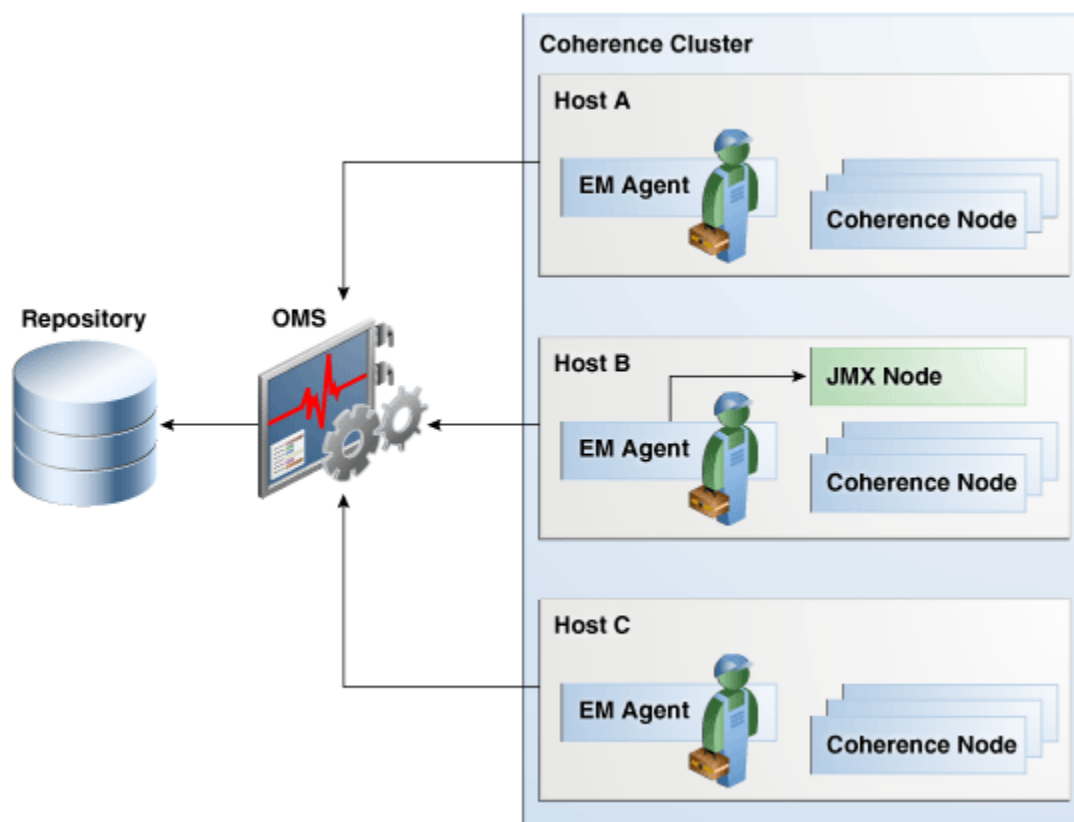
This section covers the configuration procedure for a standalone Coherence cluster.

For details on configuring a managed Coherence cluster, refer to the WebLogic documentation.

Oracle Coherence standalone deployments can be monitored using Enterprise Manager by configuring the Coherence nodes with a set of Coherence and JMX system properties (start arguments). In addition, one of the nodes will have to be configured as a central JMX management node. This JMX management node must expose all Coherence MBeans and attributes. See [Creating and Starting a JMX Management Node](#) for details. 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.

Figure 23-1 shows the configuration for monitoring standalone Coherence clusters using Enterprise Manager.

Figure 23-1 Coherence Cluster Configuration (Standalone Coherence Cluster)



As shown in the figure, Coherence Management (JMX) node's MBean server will expose MBeans for entire Coherence cluster. Enterprise Manager will connect to this management node to discover and monitor Coherence cluster.

23.3.1 Creating and 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 be 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. For more information on using JMX to manage Oracle Coherence, see [Using JMX to Manage Coherence](#) in the *Oracle Coherence Management guide*. To configure the JMX management node, you must:

- Specify Additional System Properties
- Include Additional Class Path
- Use the Enterprise Manager Custom Start Class

23.3.1.1 Specifying Additional System Properties



Note:

Oracle recommends that the management node is configured as a storage disabled node to ensure minimal performance impact on any Coherence caches.

The following start arguments must be added to one of the Coherence nodes to configure it as the JMX central management node.

- `-Dtangosol.coherence.management.extendedmbeanname=true` (allows any restarted node to be automatically detected by Enterprise Manager. This parameter is available in Coherence 3.7.1.9 and later versions)
 - If set to true, the status of the node is automatically refreshed when a node is restarted.
 - If this property is not set, you must use the Refresh Cluster option to update the status of a node when it is restarted.
 - If you start a node after setting this property to true, all nodes in the cluster must be started after the `extendedmbeanname` property is set to true.
- `-Dtangosol.coherence.management=all` (enables monitoring for all nodes)
- `-Dcom.sun.management.jmxremote.port=<port number>` (required for remote connection for coherence 12.2.1.x or older versions)
- `-Dtangosol.coherence.distributed.localstorage=false` (disables caching and ensures that the node is a dedicated monitoring node)
- `-Doracle.coherence.home=<coherence home>`
- `-Dtangosol.coherence.member=<member name>` (required for target name)

- `-Doracle.coherence.machine=<fully qualified hostname>` (must match the name of the host discovered in Enterprise Manager)

 **Note:**

If you are using JMX credentials, you must set the following additional start arguments.

- `-Dcom.sun.management.jmxremote.ssl=true`
- `-Dcom.sun.management.jmxremote.authenticate=true`

If no JMX credentials are used, you must set these arguments to **false**.

23.3.1.2 Including the Additional Class Path

You must include the path to both Enterprise Manager custom jar files, `coherenceEMIntg.jar` and `bulkoperationsmbean.jar`.

For coherence cluster versions older than 12.2.1, the jar files are available in the

`<OEM_Agent_Home>/<PLUGIN_HOME>/<MIDDLEWARE_MONITORING_PLUGIN_DIR>/archives/coherence` directory.

Coherence cluster with version 12.2.1 and above, must use the `coherenceEMIntg.jar` file available in the

`<OEM_Agent_Home>/<PLUGIN_HOME>/<MIDDLEWARE_MONITORING_PLUGIN_DIR>/archives/coherence\12.2.1`

directory.

 **Note:**

The location of the .jar files may change based on the plugin version.

23.3.1.3 Using the Custom Start Class

In addition to configuring the system properties and the class path when starting Coherence management node, it is also required that you use the Enterprise Manager `EMIntegrationServer` class as the start class. This class allows you to register the custom MBeans required for the Cache Data Management feature of Management Pack for Oracle Coherence.

23.3.1.4 Example Start Script for the Coherence Management Node

An example start script for the management node is given below:

```
#
#!/bin/sh

CP=$CP:<EM_CC_Agent_Home>/plugins/oracle.sysman.emas.agent.plugin_12.1.0.6.0/
```

```

archives/coherence/coherenceEMIntg.jar:
<EM_CC_Agent_Home>/plugins/oracle.sysman.emas.agent.plugin_12.1.0.6.0/archives/
coherence/bulkoperationsmbean.jar
COH_OPTS="$COH_OPTS -cp $CP"
$JAVA_HOME/bin/java $COH_OPTS
-Dtangosol.coherence.management.extendedmbeanname=true
-Dcom.sun.management.jmxremote.authenticate=false
-Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.ssl=false
-Dtangosol.coherence.management=all
-Dtangosol.coherence.member=<unique member name>
-Doracle.coherence.machine=<hostname_as_discovered_in_EM>
-Dcom.sun.management.jmxremote.port=<OpenTCP_Port>
-Doracle.coherence.home=$COHERENCE_HOME
-Dtangosol.coherence.distributed.localstorage=false
-Dtangosol.coherence.management.refresh.expiry=1m
-server
-Xms2048m -Xmx2048m
oracle.sysman.integration.coherence.EMIntegrationServer

```

23.3.2 Configuring All Other Nodes

In addition to configuring the Coherence JMX management node, you must configure all other Coherence cluster nodes with additional Coherence specific system properties (start arguments) used by Enterprise Manager.

23.3.2.1 Additional System Properties for All Other Coherence Nodes

The following system properties must be added to all other Coherence nodes.

```

-Dtangosol.coherence.management.extendedmbeanname=true
-Dtangosol.coherence.management.remote=true -Dtangosol.coherence.member=<unique
member name> -Doracle.coherence.home=<coherence home>
-Doracle.coherence.machine=<machine name> should be the same as the name of the host
discovered in Enterprise Manager.

```



Note:

If you are using JMX credentials, you must set the following additional start arguments.

- -Dcom.sun.management.jmxremote.ssl=true
- -Dcom.sun.management.jmxremote.authenticate=true

If no JMX credentials are used, you must set these arguments to **false**.

23.3.2.2 Example Start Script for All Other Coherence Nodes

An example start script for all other Coherence nodes is given below:

```

#!/bin/sh

COH_OPTS="$COH_OPTS -cp $CP"
$JAVA_HOME/bin/java $COH_OPTS
-Dtangosol.coherence.management.extendedmbeanname=true

```

```

-Dtangosol.coherence.management.remote=true
-Dcom.sun.management.jmxremote.authenticate=false
-Dcom.sun.management.jmxremote.ssl=false
-Doracle.coherence.home=<coherence home>
-Dtangosol.coherence.member=<unique member name>
-Doracle.coherence.machine=<hostname>
-Dcom.tangosol.net.DefaultCacheServer

```

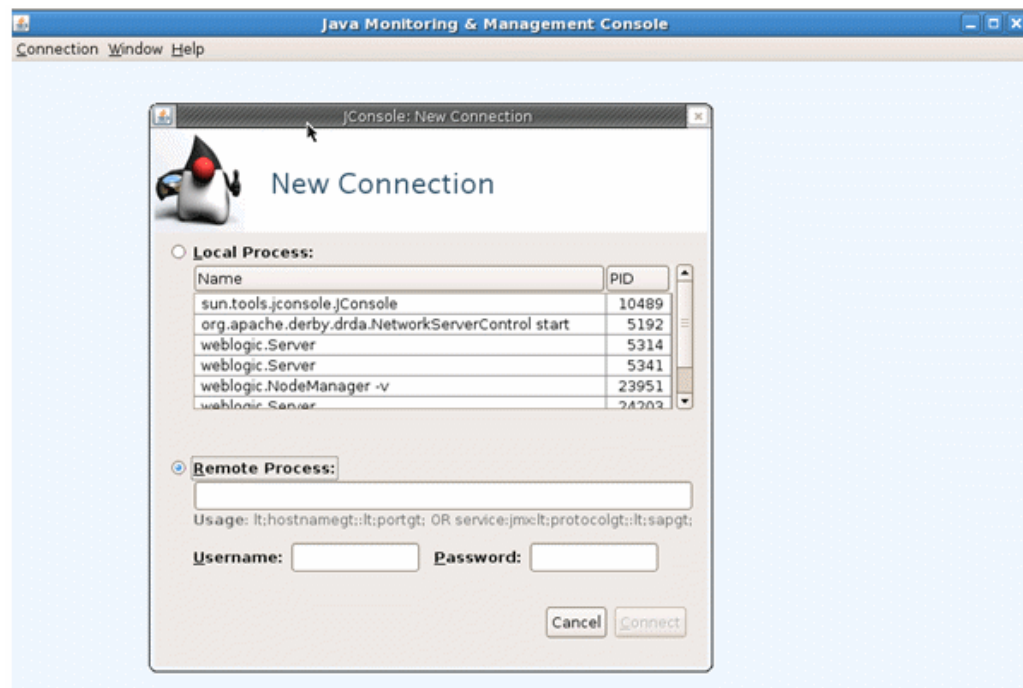
23.3.3 Testing the Configuration

To test the Coherence cluster configuration for use in Enterprise Manager, you must verify that the central management (JMX) node has information regarding the managed objects of all other Coherence cluster nodes, caches, services, and so on. Additionally, you must verify that the central management node is accessible remotely, either through `<hostname>:<port>` OR the JMX Service URL. If JMX credentials are used, they should also be specified.

23.3.3.1 Verifying Remote Access for the MBean Objects Using JConsole

JConsole is a Java tool available through JDK. You can use this to verify remote access to the MBean objects of entire Coherence cluster nodes, caches, services, and so on.

Figure 23-2 JConsole



To verify remote access, open JConsole and select "New Connection". In New Connection page, select **Remote Process** and provide connection details where `<hostname>` is the name of the machine where central management node is running, `<port>` is what you have specified in the `-Dcom.sun.management.jmxremote.port` parameter while starting the management node. If successful, you will see the MBean object tree.

Figure 23-3 MBean Object Tree

Name	Value
BufferPublishSize	45
BufferReceiveSize	1444
CpuCount	4
FlowControlEnabled	true
Id	3
LoggingDestination	stderr
LoggingFormat	{date}/{uptime} {product} {version} <{level}> {thread...}
LoggingLevel	2
LoggingLimit	2147483647
MachineId	24280
MachineName	emcc
MemberName	SNode1
MemoryAvailableMB	41
MemoryMaxMB	96
MulticastAddress	n/a
MulticastEnabled	false
MulticastPort	-1
MulticastTTL	-1
MulticastThreshold	25
NackEnabled	true
NackSent	0
PacketDeliveryEfficiency	1.0
PacketsBundled	1544
PacketsReceived	39633
PacketsRepeated	0
PacketsResent	1
PacketsResentEarly	0
PacketsResentExcess	0
PacketsSent	39562
Priority	0
ProcessName	22734
ProductEdition	Grid Edition
PublisherPacketUtilization	0.043829843
PublisherSuccessRate	0.9999747
QuorumStatus	{NullActionPolicy allowed-actions=*}

If you see MBeans for all Coherence nodes in the System MBean Browser or JConsole, you can now discover and monitor the Coherence cluster and its associated elements in Enterprise Manager.

23.4 Discovering Coherence Targets

This section covers the following:

- [Discovering a Standalone Coherence Cluster](#)
- [Discovering a Managed Coherence Cluster](#)

23.4.1 Discovering a Standalone Coherence Cluster

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.

 **Note:**

To provision new Coherence nodes, start, and stop nodes, the Management Agent must be installed on all the hosts on which the nodes are running. For more details on provisioning Coherence nodes, see the [Enterprise Manager Lifecycle Management Guide](#).

Prerequisites

Before you discover a Coherence cluster, you must have completed the following tasks:

- Created a Coherence cluster with one JMX management node and one or more other nodes.
- Started the JMX management node with the necessary parameters as defined in [Creating and Starting a JMX Management Node](#).
- Started the other nodes with the necessary parameters as defined in [Configuring All Other Nodes](#).

To discover an already running Coherence cluster, follow these steps:

1. Log in to Enterprise Manager as a user 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 **Standalone Oracle Coherence Cluster** in the **Add** drop-down box and click **Go**. The Oracle Coherence Cluster: Discover Cluster, Node, and Cache Targets page is displayed.
4. On this page, select option **Standalone Coherence cluster configured with dedicated management node** or **Standalone Coherence cluster configured with dynamic management mode** (for coherence versions 12.2.1.x and above) to specify the connection details of the Coherence MBean Server. You can select either of these two options based on coherence versions in use. This is required to discover the Coherence cluster, node and cache targets.

If you have selected option **Standalone Coherence cluster configured with dedicated management node**, you can select either of the following options to provide MBean Server details:

Figure 23-4 Add Coherence Target

Oracle Coherence Cluster Discovery

Before you discover a Coherence Cluster in Oracle Enterprise Manager, you must ensure that the following prerequisites are met:

- All nodes, including the Coherence Management Node (MBeanServer Node), must be configured as follows:
 - Each Node must be started with `tangosol.coherence.member` (MemberName) property. MemberName must be unique across the entire Coherence Cluster.
 - Each Node must be started with `oracle.coherence.machine` (Machine) property. Machine must be the same as the host name used to discover the corresponding host in Oracle Enterprise Manager.

Prerequisites for Oracle Coherence configured with dedicated management node

- The Coherence Management Node (MBean Server) must be configured as follows:
 - The Coherence Management Node must include the `coherenceEMIntgr.jar` and `bulkOperationsMBean.jar` in its classpath. These jars are available in the `<PLUGIN_HOME>\MIDDLEWARE_MONITORING_PLUGIN_DIRECTORY\archives\coherence` directory. Oracle recommends that this node be storage disabled.
 - Coherence Management Node must be started with the `oracle.sysman.integration.coherence.EMIntegrationServer` class.

Prerequisites for Oracle Coherence configured in dynamic management mode

- The Coherence Node must include the `coherenceEMIntgr.jar` and `bulkOperationsMBean.jar` in its classpath. These jars are available in the `<PLUGIN_HOME>\MIDDLEWARE_MONITORING_PLUGIN_DIRECTORY\archives\coherence` directory.
- Register following MBeans using custom MBean xml file:
 - `mbean-class=oracle.sysman.framework.bulkOperations.BulkOperationsMBeanIntgrl`, `mbean-name=CoherenceType=BulkOperations`
 - `mbean-class=oracle.sysman.integration.coherence.CacheDataManager`, `mbean-name=CoherenceType=CustomName=CacheDataManager`
 - `mbean-class=oracle.sysman.integration.coherence.EMCoherenceIntgrl`, `mbean-name=CoherenceType=CustomName=CoherenceIntgrl`

MBean Server Connection

Standalone Coherence cluster configured with dedicated management node

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

Standalone Coherence cluster configured with dynamic management mode (12.2.1.x and above)

MBean Server Credentials

If JMX authentication is enabled, you must specify both the username and password for accessing the Coherence MBean Server.

Username: Password:

Oracle Coherence Cluster Discovery and Monitoring Agent


Select a Management Agent that is running on the same host on which the Coherence MBean Server is running. This agent will be used to monitor the Coherence Cluster.

Agent:

Discovery Options

Do not discover Coherence Caches

If selected, new Coherence cache targets will not be discovered. Recommended for clusters with very large number (over 1000) of caches.

- **Host, Port, and Service:** Enter the following details:
 - **Management Node Host:** Select the host on which the Management Node is running.
 - **JMX Remote Port:** 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.
- **Service Name:** The service name used for the connection. The default is `jmxrmi`.
 - **JMX Service URL:** Service URL that will be used for the connection. If you enter the URL, the values specified in the Machine Name, Port, Communication Protocol, and ServiceName fields will be ignored. For example, `service:jmx:rmi://localhost:3000/jndi/rmi://localhost:9000/server`. For more details on the URL format, refer to <http://java.sun.com/j2se/1.5.0/docs/api/javax/management/remote/JMXServiceURL.html>
- 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.

If you have selected option **Standalone Coherence cluster configured with dynamic management mode (12.2.1.x and above)**, you must enter the following details:

Figure 23-5 Add Coherence Target

ORACLE® Enterprise Manager Cloud Control 13c Enterprise Targets

Oracle Coherence Cluster Discovery Page Refreshed Feb 1, 2018 2:20:34 AM PST

- All nodes, including the Coherence Management Node (MBeanServer Node), must be configured as follows
 - Each Node must be started with `tangosol.coherence.member` (MemberName) property. MemberName must be unique across the entire Coherence Cluster.
 - Each Node must be started with `oracle.coherence.machine` (Machine) property. Machine must be the same as the host name used to discover the corresponding host in Oracle Enterprise Manager.

Prerequisites for Oracle Coherence configured with dedicated management node

- The Coherence Management Node (MBean Server) must be configured as follows
 - The Coherence Management Node must include the `coherenceEMIntg.jar` and `bulkoperationsmbean.jar` in its classpath. These jars are available in the `<PLUGIN_HOME>/MIDDLEWARE_MONITORING_PLUGIN_DIRECTORY/archives/coherence` directory. Oracle recommends that this node be **storage disabled**.
 - Coherence Management Node must be started with the `oracle.sysman.integration.coherence.EMIntegrationServer` class.

Prerequisites for Oracle Coherence configured in dynamic management mode

- The Coherence Node must include the `coherenceEMIntg.jar` and `bulkoperationsmbean.jar` in its classpath. These jars are available in the `<PLUGIN_HOME>/MIDDLEWARE_MONITORING_PLUGIN_DIRECTORY/archives/coherence` directory.
 - Register following MBeans using custom MBean xml file.
 - `mbean-class: oracle.as.jmx.framework.bulkoperations.BulkOperationsMBeanImpl, mbean-name: Coherence.type=BulkOperations`
 - `mbean-class: oracle.sysman.integration.coherence.CacheDataManager, mbean-name: Coherence.type=Custom.name=CacheDataManager`
 - `mbean-class: oracle.sysman.integration.coherence.EMCoherenceUI, mbean-name: Coherence.type=Custom.name=CoherenceUI`

MBean Server Connection

Standalone Coherence cluster configured with dedicated management node

Standalone Coherence cluster configured with dynamic management mode (12.2.1.x and above)

Specify cluster port and hosts where nodes are started in dynamic management mode

* Cluster Port

+ Add - Delete

Host Name
<input type="text"/>

MBean Server Credentials

If JMX authentication is enabled, you must specify both the username and password for accessing the Coherence MBean Server.

Username Password

Oracle Coherence Cluster Discovery and Monitoring Agent

Select a Management Agent that is running on the same host on which the Coherence MBean Server is running. This agent will be used to monitor the Coherence Cluster.

* Agent

Discovery Options

Do not discover Coherence Caches

Cluster Port: In dynamic management mode any coherence node configured with dynamic managed mode can be a Management Node. Hence, you must enter coherence Cluster Port to discover Coherence target instead of JMX port.

Host Name: You must enter the host name where nodes with dynamic management mode are running. You can see a list of host names on which coherence nodes with dynamic management mode are running. If the current management node is down then the next management node is discovered using these host names.

For more information, see [Specifying a Cluster's Multicast Address and Port](#) in the *Oracle® Fusion Middleware Developing Applications with Oracle Coherence guide*.

- MBean Server Credentials:** If JMX authentication is used, specify the username and password required to access the MBean Server.
- Select the Management Agent that will be used to monitor the Coherence target.
- Select the **Do not discover Coherence Caches** checkbox to skip Coherence cache targets discovery and click **Continue**.

If this checkbox is selected, new Coherence cache targets will not be discovered. It is recommended to skip cache discovery for clusters with large number (over 1000) of caches.

- The details of the discovered targets are displayed. Click **Add Targets** to add these targets to Enterprise Manager.

 **Note:**

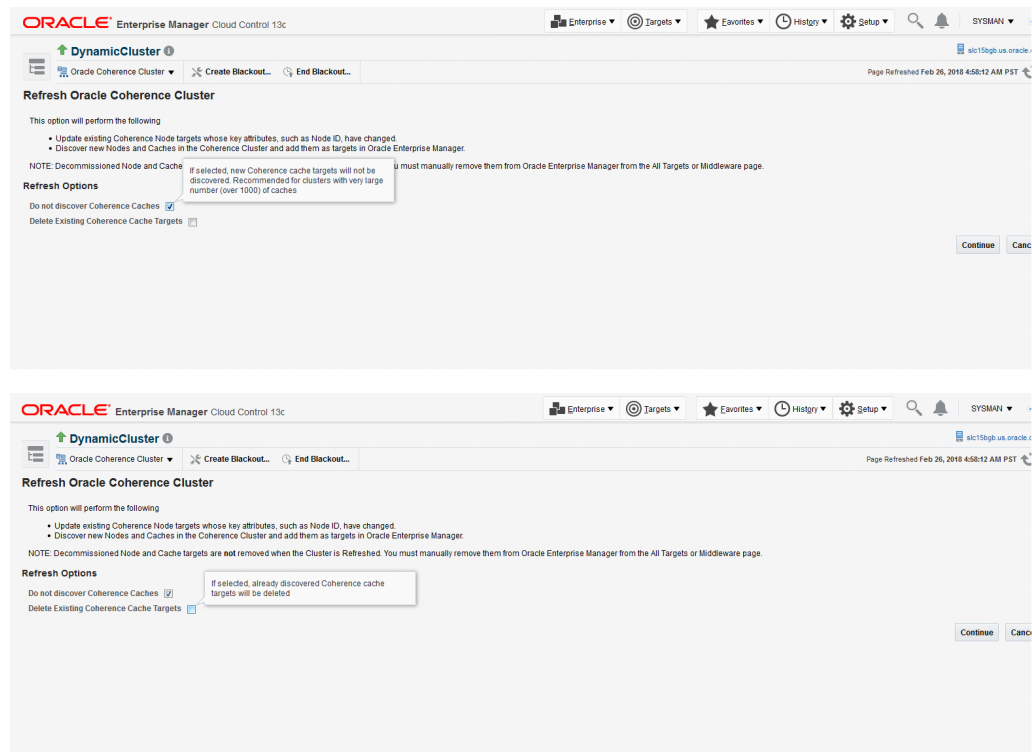
To automatically discover a new node or target in Enterprise Manager, you must refresh the cluster as described in [Refreshing a Cluster](#).

23.4.1.1 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.

Figure 23-6 Refresh Cluster



Click **Close**. The list of nodes and caches that can be added are displayed.

If you want to remove already discovered caches, select **Do not discover Coherence Caches** checkbox, and then select **Delete Existing Coherence Cache Targets**.

Click **Add Targets** to add the targets to the cluster.

 **Note:**

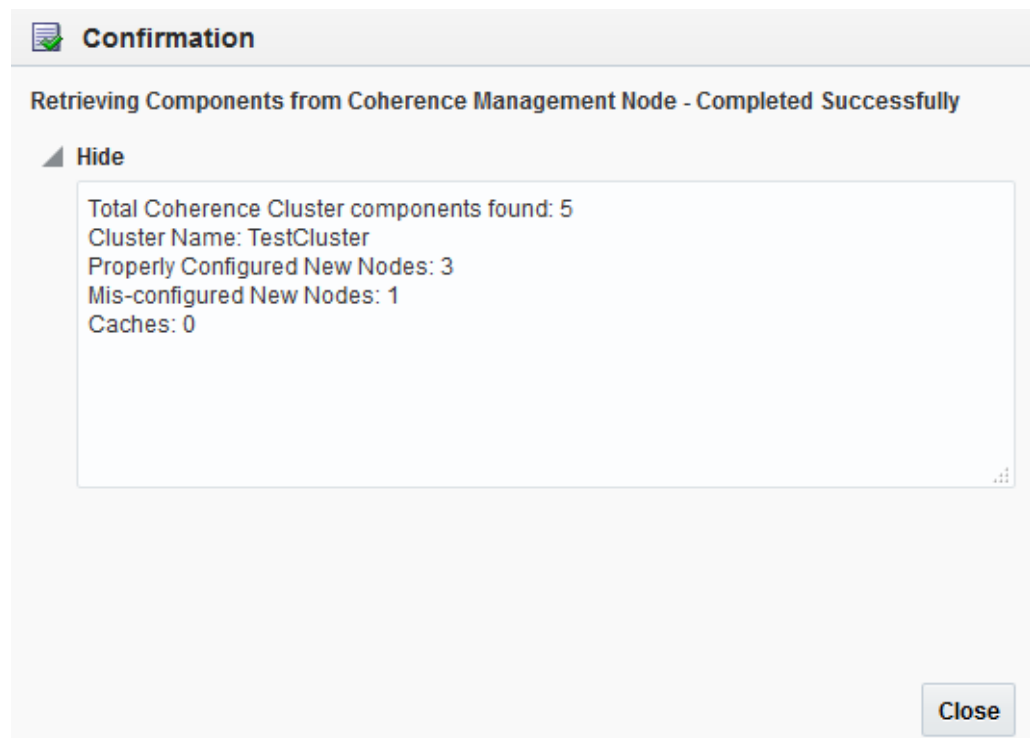
Decommissioned nodes and caches will not be removed during the **Refresh** process. You must remove them manually.

23.4.1.2 Managing Mis-configured Nodes

While discovering a Coherence cluster, all nodes must be started with the proper guidelines as described in [Configuring a Coherence Cluster](#).

If a node is improperly configured or has been started without the necessary guidelines, it will be categorized as a mis-configured node and will not be a part of the newly discovered cluster. During discovery, if any improperly configured nodes are present in the cluster, you will see the following screen:

Figure 23-7 Mis-Configured Nodes



This indicates that there are some improperly configured nodes in the cluster. Click **Close**. The following page is displayed.

Figure 23-8 Mis-configured Nodes II

Warning
Mis-configured Coherence Nodes were found in Oracle Coherence Cluster. You can either proceed with the discovery of properly configured Nodes or cancel the discovery process and fix the mis-configured Nodes. If you proceed with discovery, you can invoke Refresh Cluster after fixing mis-configured Nodes at a later time.

Oracle Coherence Cluster
 Oracle Coherence Cluster Target Name TestCluster Oracle Coherence Cluster Version 3.7.1.15
 Monitoring Agent [REDACTED] Extended MBean true
 Oracle Coherence Cluster Name TestCluster

Oracle Coherence Nodes

Mis-configured Nodes
 These Nodes are mis-configured and cannot be added as targets in Oracle Enterprise Manager.
 Total Mis-configured Nodes 1

PID	Member Name	Machine	ID	Reason for Discovery Failure
JV10	C10	[REDACTED]	4	Node is not started with Extended MBean parameter

Node Targets
 The following Coherence Nodes will be added as targets in Oracle Enterprise Manager.
 Total Node Targets 3

Target Name	Member Name	Machine	ID	PID
TestClusterC9	C9	[REDACTED]	3	JV9
TestClusterC8	C8	[REDACTED]	2	JV8

The list of improperly configured nodes along with the reasons for their failure is listed in this page. You can either choose to cancel the discovery process and fix these nodes or continue with the discovery with the properly configured nodes.

If you wish to continue with the discovery process, follow the steps listed in [Discovering Coherence Targets](#).

If you click **Cancel**, the discovery process is aborted and the cluster is not refreshed. If mis-configured nodes are found during the Refresh process, they must be fixed before you can run the Refresh operation again. For more information, see [Refreshing a Cluster](#) and then discover the cluster.

23.4.2 Discovering a Managed Coherence Cluster

You can discover a managed Coherence cluster while discovering an Oracle Fusion Middleware / WebLogic Domain by following these steps:

1. Login to Enterprise Manager as a user with the **Add Target** privilege.
2. From the **Targets** menu, select **Middleware**. You will see a list of Middleware targets.
3. Select Oracle Fusion Middleware / WebLogic Domain from the Add drop down menu and click **Go**.

Figure 23-9 Add Oracle Fusion Middleware / WebLogic Domain: Find Targets

4. Enter the following details
 - Administration Server Host: Enter the host name on which the Administration Server is installed.
 - Port: Enter the WebLogic Administration Server port.
 - Username and Password: Enter the user name and password for the WebLogic Administration Server.
 - Agent: Enter The host name for a Management Agent that will be used to discover the Fusion Middleware targets.
5. Click **Continue**. You will see a window indicating that the targets are being discovered. Click **Close**. Any Coherence clusters that are present in the WebLogic Domain will listed.

Figure 23-10 Targets and Agents Assignments

Target Name	Target Type	Host	Configured Agent	Status
AdminServer	Oracle Coherence Node		[Inherited From Parent]	New Target
opes-rest	Domain Application Deployment		8336	New Target
uifwksamples-coexist	Domain Application Deployment		8336	New Target
uifwksamples-encust	Domain Application Deployment		8336	New Target
mds-owsm	Metadata Repository	.com	8336	New Target
defaultCoherenceCluster	Oracle Coherence Cluster	.com	8336	New Target
*oracle.mds.localcache.scopeoracle	Oracle Coherence Cache		8336	New Target
*oracle.mds.localcache.scopeoracle	Oracle Coherence Cache		8336	New Target
*oracle.mds.localcache.scopeoracle	Oracle Coherence Cache		8336	New Target
*oracle.mds.localcache.scopeoracle	Oracle Coherence Cache		8336	New Target
*oracle.mds.localcache.scopeoracle	Oracle Coherence Cache		8336	New Target
*oracle.mds.localcache.scopeoracle	Oracle Coherence Cache		8336	New Target
*oracle.mds.localcache.scopeoracle	Oracle Coherence Cache		8336	New Target
*oracle.mds.localcache.scopeoracle	Oracle Coherence Cache		8336	New Target
*oracle.wls.internal.wsmc.localCache	Oracle Coherence Cache		8336	New Target
AdminServer	Oracle Coherence Node		[Inherited From Parent]	New Target

6. Click **Add Targets** to add these targets to Enterprise Manager and click **OK** to return to the Middleware page.

For more details on Oracle Fusion Middleware / WebLogic Domain discovery, see the *Oracle Fusion Middleware* documentation.

 **Note:**

If the management node of a 12.2.1 managed cluster is restarted, you must manually refresh the WebLogic Domain before you can continue monitoring the cluster.

23.5 Enabling the Management Pack

 **Note:**

For managed Coherence clusters, you must enable the Oracle Cloud Management Pack for Oracle Fusion Middleware.

You must enable the Management Pack for Oracle Coherence if you want to use any custom features. If the management pack is not enabled, you can access only the Home pages and base platform features. 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.

 **Note:**

Apart from enabling the Management Pack, you must grant `VIEW` privileges to all users on the Management Agent that is monitoring the Coherence targets. This ensures that all targets being monitored by the Management Agent are visible to the user.

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, proxies, and connections.

This chapter contains the following sections:

- [Understanding the Page Layout](#)
- [Viewing the Home Pages](#)
- [Viewing the Summary Pages](#)
- [Log Viewer](#)
- [Viewing the Performance Pages](#)
- [Removing Down Members](#)
- [Topology Viewer](#)
- [Viewing Incidents](#)

Before you start monitoring a cluster in Enterprise Manager, you must perform the following tasks:

- Install the 13.3.0.0.0 Management Agent on all hosts where Coherence nodes are running.
- Deploy the 13.3.0.0.0 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 the [Testing the Configuration](#).

**Note:**

If the Management Agent is upgraded to 13.3.0.0.0, you must ensure that the Fusion Middleware Plug-in is also upgraded to 13.3.0.0.0.

24.1 Understanding the Page Layout

This section describes the layout of the Coherence pages in Enterprise Manager and how the pages can be customized. It contains the following sections:

- Navigation Tree
- Personalization

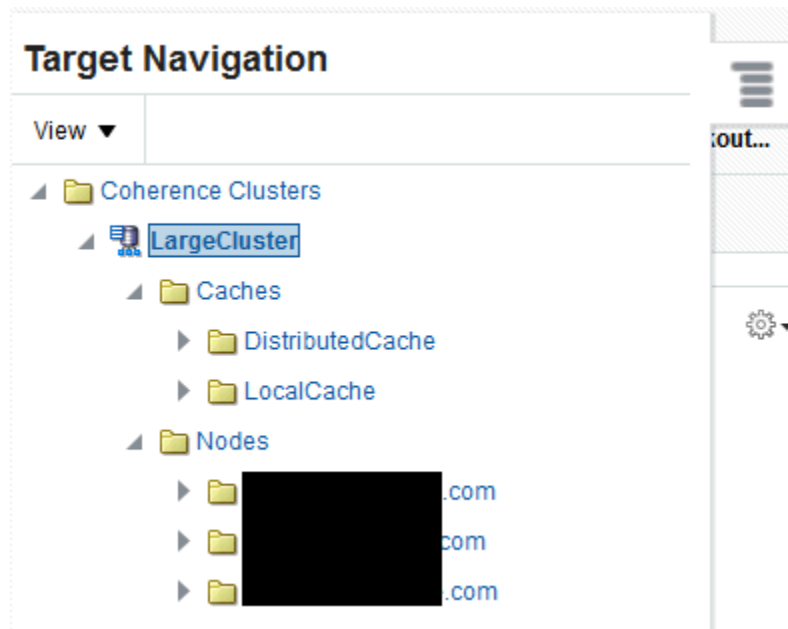
24.1.1 Navigation Tree

All Coherence pages in Enterprise Manager contain a navigation tree in the left panel of the page. The navigation tree is hidden by default but can be displayed by clicking the Target Navigation icon. The navigation tree displays all the entities in a selected cluster with the Cluster at the top level, followed by caches and nodes as the children entities. The entities are grouped as follows:

- All caches that belong to a particular cluster are listed under the Caches folder in the navigation tree.
- Cache targets of a service type are grouped together.
- The Nodes folder contains host names on which the nodes are running as children entities.
- Nodes that are running a particular host are grouped together.

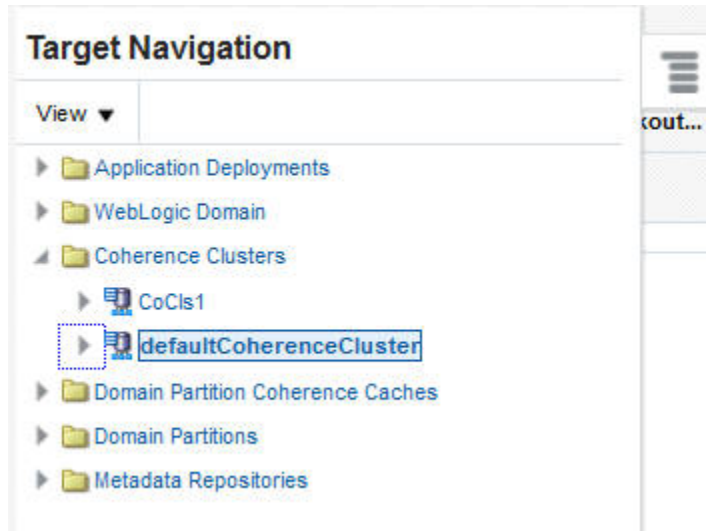
You can expand or collapse any entity in the navigation tree by clicking on the Expand/Collapse icon. Click on an entity such as a node, cache, or service in the tree to view the associated home page on the right hand side. A snap shot of the navigation is shown below.

Figure 24-1 Navigation Tree



For a managed Coherence cluster, all the Coherence clusters in a domain are included in the Coherence Clusters folder. This folder appears at the same level as the WebLogic Domain folder. If multi-tenancy is supported for a target, you will also see the Domain Partition Coherence Caches folder.

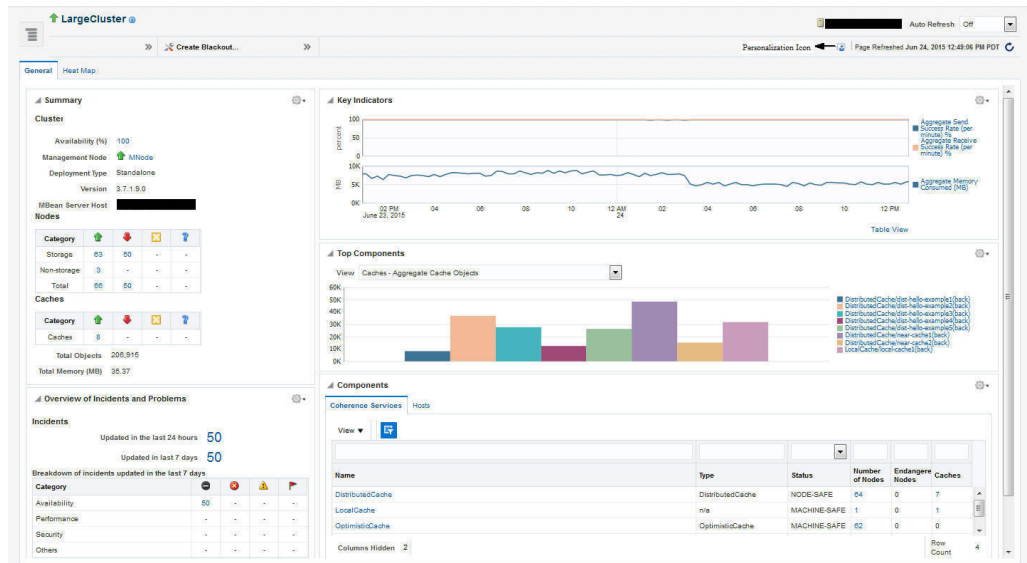
Figure 24-2 Navigation Tree (Managed Clusters)



24.1.2 Personalization

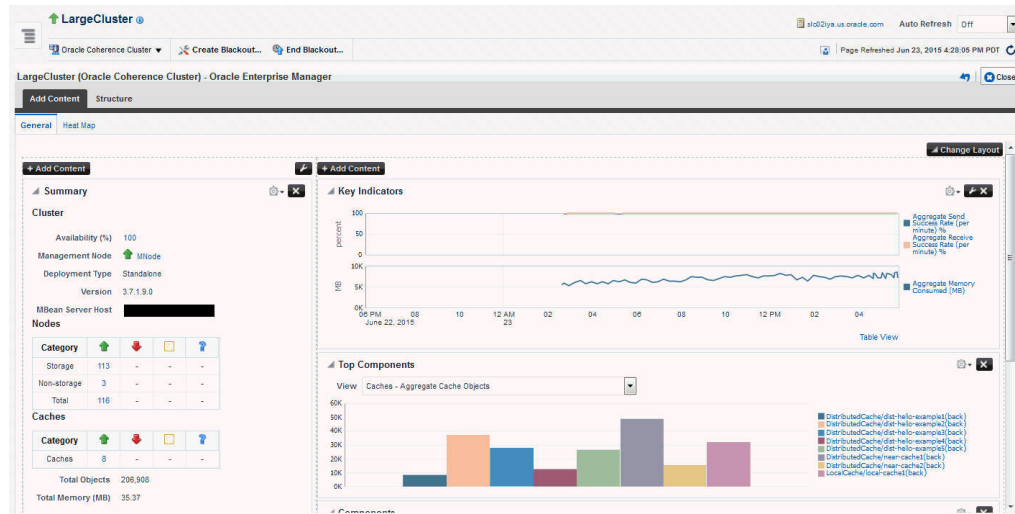
You can personalize any of the Coherence pages and select the regions to be displayed, the order in which they are displayed, the metrics to be included in the charts and so on. Click the **Personalization** icon on a page to view the page in Edit mode.

Figure 24-3 Cluster Home Page (Personalization Icon)



You will see the page in Edit mode as shown below.

Figure 24-4 Cluster Home Page (Edit Mode)



In the Edit mode, you can do the following:

- **Change Layout:** Click **Change Layout** and select a different layout for the page.
- **Add Content:** Click **Add Content**. The regions that can be displayed on the page are displayed. Select a region, click **Add**, then click **Close** to return to the previous page.
- **Edit Regions:** Click the Edit icon for a region to add or delete any parameters or metrics being displayed in the region.
- **Move Up / Move Down:** You can change the location of a region on a page by using the Move Up / Down icon.

After you have made all the changes, click **Close** to apply the changes or click **Reset Page** to return to the default mode.

24.2 Viewing the Home Pages

When you discover a Coherence cluster, a Coherence cluster target, caches, and properly configured nodes are created. Each of these entities collect a rich set of metrics. From the Home pages, you can view the overall cluster summary and key indicators from components such as nodes, caches, and services.

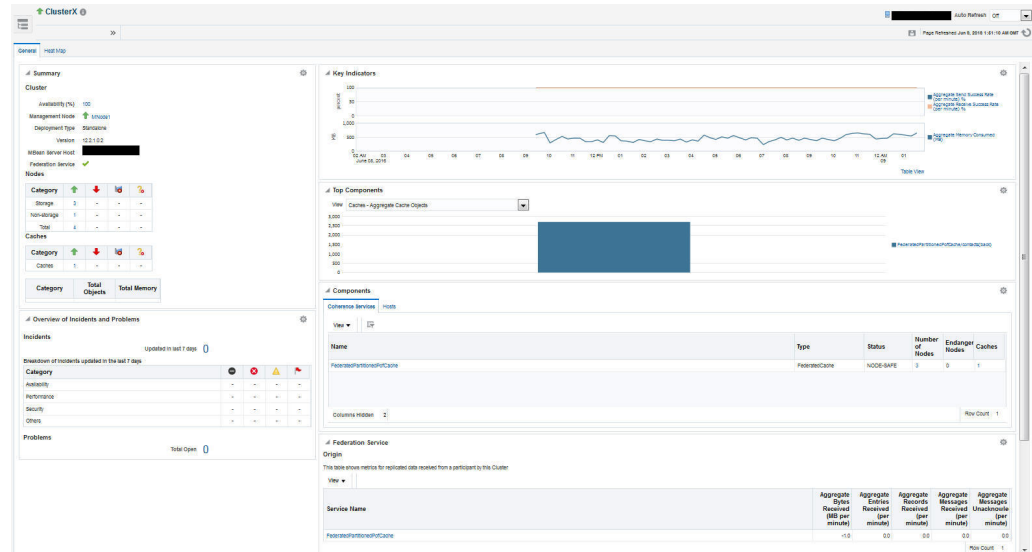
24.2.1 Coherence Cluster Home Page

Note:

The data shown on this page is not real time data but is based on the latest data available from the OMS repository. After the Coherence cluster has been discovered, the most recent data is displayed only after the performance and configuration collection has been completed for the cluster and its members.

To see a global view of the cluster, from the **Targets** menu, select **Middleware**, then click on a **Coherence Cluster** target. The Coherence Cluster Home page appears:

Figure 24-5 Coherence Cluster Home Page



To view details about the cluster, click the **Target Information** icon next to the cluster name at the top left hand corner of the page. The following details are displayed in the Target Information popup window:

- **Up Since:** The date and time from which the cluster is up and running.
- **Availability%:** The percentage of time that the management agent was able to communicate with the cluster. Click the link to view the availability details over the past 24 hours.
- **Version:** The version of the Coherence software obtained from the Cluster MBean.
- **Oracle Home:** The location of the Oracle Home.
- **Agent:** The Management Agent that Oracle Enterprise Manager is using to communicate with the MBean Server. Click on the link to drill down to the Agent Home page.
- **Host:** The host on which the cluster is running. Click on the link to drill down to the Host Home page.
- **Time Zone:** Displays the time zone for the target.
- **Name:** This is the actual name of the cluster that is discovered and may be different from the name of the cluster target in Enterprise Manager.
- **Auto Detected Restarted Nodes:** If the cluster has been started with an extendedMBean property, the Auto Detect Restarted Nodes property (`tangosol.coherence.management.extendedmbeanname`) is enabled and this field is set to **True**.

The Cluster Home page contains the **General** and **Heatmap** tabs.

24.2.1.1 General Tab

- **Summary:** The following details are displayed:
 - **Cluster**
 - * Availability (%): The availability of the cluster over the last 24 hours.
 - * Management Node: Shows the name of the management node and its status. Click on the link to drill down to the Node Home page.
 - * Deployment Type: Indicates if this is a standalone Coherence cluster or a managed Coherence cluster.
 - * Version: The Coherence software version.
 - * MBean Server Host: Shows the host on which the Coherence management node with Mbean Server is running.

If the node on the MBean Server Host is not accessible, the monitoring capability of the node will be affected. To avoid this, we recommend that at least two management nodes are running in 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.
 - * Federation Service: Indicates if the cluster is participating in data federation.
 - **Nodes**
 - * Storage Nodes: The number of storage enabled nodes in the cluster. Click on the link to drill down to the Storage Nodes page.

Note: The Number of Nodes and Storage Nodes listed here may be different from the number of node targets that have been discovered. As a result, when you click on the link, the number of nodes displayed may be lesser than the nodes shown in this table.
 - * Non Storage Nodes: The nodes that are not storage enabled such as proxy, client nodes, and so on.
 - * Total Nodes: The total number of nodes in the cluster. Click on the link to drill down to the All Nodes page.
 - **Caches**
 - * Caches: The total number of caches in the cluster. Click on the link to drill down to the All Caches page.
 - * Total Objects: The total number of objects stored across all the back caches in the cluster.
 - * Total Memory: The total memory in MB used by all the objects in the back caches. A numeric value is displayed only if a Binary calculator is used in cache configuration. If a Binary calculator is not used, a **N/A** will be displayed in this field.
- **Overview of Incidents and Problems:** This region lists any incidents that have occurred over the last 7 days and any problems in the cluster and its associated targets (nodes, caches, and hosts). Click on the link to drill down to the Incident Manager page.

- **Key Indicators:** This region displays graphs with key metrics that indicate the health and performance of the cluster. You can use the Personalization feature to specify the key metrics that are to be included in the charts.
- **Top Components:** This region contains a graphical representation of the top 10 performing targets for a selected metric based on the latest available data from the OMS repository. The top components are listed in ascending or descending order depending on the metric selected and indicates how the top component data has been collected. Select a metric from the View drop down list to see a graphical representation of the top 10 targets for the selected metric. For example, if you select the Cache - Cache Objects metric, the graph displays the top 10 cache targets. Click on the graph or legend to drill down to the detail pages.
- **Components:** This is a tabbed region with Coherence Services tab showing the Coherence Cluster Services and the Hosts table showing the list of hosts on which the cluster nodes are running. A detailed description of each tab is given below:
 - **Coherence Service:** This tab shows all the services in the Coherence cluster. For a multi-tenant managed Coherence Cluster, a Domain Partition column is also displayed.
 - **Hosts:** This tab shows the hosts on which the nodes are running. It contains the following details:
 - * **Host:** The host on which the node is present. The Host Name link is displayed if: only if the Machine Name property has been defined for the node.
 - * The host on which the nodes are running is monitored by Enterprise Manager.
 - * The name of the discovered host target must be the same as the name specified in the `oracle.coherence.machine` system property.
 - * **Number of Nodes:** The number of nodes present on each host. Click the link to drill-down to the Node Performance page.
 - * **CPU Used%:** The percentage of CPU used on the host.
 - * **Memory Used%:** The percentage of memory used on the host.
- **Federation Service:** If the cluster is participating in data federation, the table Origin and Destination are displayed. Also, if the Federation Service is running on Domain Partition Caches, then the Domain Partition column will be displayed in these tables:
 - **Origin:** This table shows details of the data being received. The Aggregate Entries Received, Aggregate Bytes Received, Aggregate Records Received, Aggregate Messages Received, and Aggregate Messages Unacknowledged are displayed.
 - **Destination:** This table shows details of the data being sent. The Aggregate Entries Sent, Aggregate Bytes Sent, Aggregate Messages Unacknowledged, Aggregate Messages Sent, and the Aggregate Records Sent are displayed.

24.2.1.1.1 Cluster Management Operations

You can perform cluster management operations if you meet the following prerequisites:

- 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 the `-Doracle.coherence.machine` Java option 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 to start this node must 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.

The operations you can perform are:

- **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. You can start multiple nodes on multiple remote hosts in one operation. Select the hosts on which the new node is to be started and click **Start New Nodes**. You will see the Start New Nodes page where you can add one or more nodes.
- **Stop Nodes:** You can stop all the nodes on a specific host. Select a host and click **Stop Nodes**. You will see the Stop Nodes page where the details of the nodes being stopped are displayed.

 **Note:**

- The **Start New Nodes** and **Stop Nodes** options will be available only if the hosts on which the nodes are running are monitored by Enterprise Manager. An asterisk indicates hosts that are not monitored by Enterprise Manager.
- Information about a newly started node is uploaded into the repository only after one regular agent metric collection i.e. by default value of 5 minutes.

24.2.1.2 Heatmap

The **Heatmap** tab provides a graphical representation of all targets in the cluster.

The data shown in the heatmap is based on the following criteria:

- **View By:** You can choose to **View By** Nodes, Caches, Services, Hosts, or Partition Caches (if available) target type.
- **Block Size:** This parameter allows you to draw the size of the cell to be displayed in the heatmap.
- **Block Color:** This parameter allows you to select the metric by which the Heat Map is rendered. The metric you can select is based on the target selected in the **View By** field. Depending the metric value and whether it is within the threshold, the block color can be green, orange, yellow, or red. You can use the color palette to change the thresholds to monitor the targets that are at critical or warning threshold levels.

You can hover over a cell in the heat map to view the target details and selected filters. Click on a cell in the heat map to view detailed information for the target along with a link that allows you to drill down to the Home page for the target.

24.2.1.3 Cluster Menu Navigation

The following key menu options are available from the Coherence Cluster Home menu:


- **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. See [Performance Summary 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.
- **Logs:** You can use the log viewer to view log messages. For more details, see [Log Viewer](#).
- **Members:** You can navigate to the following pages from this menu:
 - **Coherence Topology:** The Coherence Topology Viewer provides a visual layout of the Coherence deployment and shows the Coherence cluster and its associated nodes and caches. See [Topology Viewer](#) for details.
 - **Nodes:** This page lists all the nodes in the cluster. See [Nodes Page](#) for details.
 - **Caches:** This page lists all the caches in the cluster. See [Caches Page](#) for details.
 - **Services:** This page lists all services in the cluster. See [Services Page](#) for details.
 - **Applications:** This page lists all applications in the cluster. See [Applications Page](#) for details.
 - **Proxies:** This page lists all connection managers in the cluster. See [Proxies Page](#)
- **Cluster Administration:** From the Oracle Coherence Cluster menu, select Administration. See [Cluster Administration Page](#) for details.
- **Refresh 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.
- **Remove Down Members:** You can delete any member in the cluster that is not available. For details, see [Removing Down Members](#).
- **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. For more information, see Deployment Procedure in the *Enterprise Manager Lifecycle Management Administrator's Guide*.
- **Latest Configuration:** From the **Oracle Coherence Cluster** menu, select **Configuration**, then select **Latest** to view the latest configuration data for the Coherence cluster.
- **JVM Diagnostics:** From the **Oracle Coherence Cluster** menu, select **JVM Diagnostics** to view the Coherence Cluster JVM Diagnostics Pool Drill Down page. This option is available only if the cluster has been configured for JVM

Diagnostics and the WLS Management Pack EE Management Pack has been included. See [Coherence Integration with JVM Diagnostics](#) for details.

24.2.2 Node Home Page

This page provides details of a selected node in the cluster. From the **Coherence Cluster** menu, select **Nodes**, and click on a specific node to drill down to the Node Home page.

Figure 24-6 Coherence Node Home Page



To view details about the node, click the **Target Information** icon next to the Node name at the top left hand corner of the page. The following details are displayed in the Target Information popup window:

- **Up Since:** The date and time from which the node is up and running.
- **Availability%:** The percentage of time that the management agent was able to communicate with the node. Click the link to view the availability details over the past 24 hours.
- **Version:** The version of the Coherence software obtained from the Cluster MBean.
- **Oracle Home:** The location of the Oracle Home.
- **Agent:** The Management Agent that Oracle Enterprise Manager is using to communicate with the MBean Server. Click on the link to drill down to the Agent Home page.
- **Host:** The host on which the cluster is running. Click on the link to drill down to the Host Home page.
- **Time Zone:** Displays the time zone for the target.
- **Member Of:** The cluster to which this node belongs.
- **Cluster Name:** This is the actual name of the cluster that is discovered and may be different from the name of the cluster target in Enterprise Manager.

- Auto Detected Restarted Nodes: If the node has been started with an extendedMBean property, the Auto Detect Restarted Nodes property is enabled and this field is set to **True**.

The Node Home page contains the following regions:

- **Summary**

- **General**

- * Availability: The availability of the node over the last 24 hours.
 - * Coherence Cluster: The cluster with which this node is associated.
 - * Auto Detect Restarted Nodes: If the node has been started with an extendedMBean flag, this flag is enabled and a check mark is displayed.
 - * Federation Service: Indicates if this node is participating in data federation.

- **CPU**

- * CPU (%): The CPU percentage used.

- **Cache Size**

- * Objects: The aggregate number of objects in the cache.
 - * Units: The aggregate number of units in the cache.
 - * Memory: The aggregate memory used by the cache.

- **Cache Usage**

- * Caches: The total number of caches in the cluster.
 - * Total Gets: The total number of get() operations over the last 24 hours.
 - * Total Misses: The total number of cache misses in the last 24 hours.
 - * Total Puts: The total number of put() operations over the last 24 hours.

 **Note:**

If you are monitoring a multi-tenant managed Coherence cluster, you will see an additional row with the total number of partition caches in the cluster and the cache usage data for these caches.

- **Services**

- * Services: The total number of services running on the cache.
 - * Task Backlog: The size of the backlog queue that holds tasks scheduled to be executed by one of the service pool threads.

 **Note:**

If you are monitoring a multi-tenant managed Coherence cluster, you will see an additional row for partition caches with the total number of services running on the partition caches and the task backlog.

- **Storage Manager**

- * **Total Evictions:** The total number of evictions from the backing map managed by this Storage Manager.
- * **Total Events Dispatched:** The total number of events dispatched by the Storage Manager per minute.

 **Note:**

If you are monitoring a multi-tenant managed Coherence cluster, you will see an additional row for partition caches.

- **Overview of Incidents and Problems**

This region lists any incidents that have occurred over the last 7 days and any problems in the node and its associated host target. Click on the link to drill down to the Incident Manager page.

- **Key Indicators**

This region displays graphs with key metrics that indicate the health and performance of the node over the last 24 hours. You can customize the metrics specify the key metrics that are to be included in the charts by selecting them from the metric palette.

- **Top Components**

This region contains a graphical representation of the top 10 performing targets for a selected metric from the last metric collection. The graph does not display real time data. The top components are listed in ascending or descending order depending on the metric selected and indicates how the top component data has been collected. Select a metric from the View drop down list to see a graphical representation of the top 10 targets for the selected metric. For example, if you select the Cache - Cache Objects metric, the graph displays the top 10 cache targets.

- **Components**

This region lists the components associated with the node such as caches, services, connections, connection managers, and applications. the cluster. The table displays the name and type of the component.

- **Federation Service:**

If this node is participating in data federation, the table Origin and Destination are displayed. Also, if the Federation Service is running on Domain Partition Caches, then the Domain Partition column will be displayed in these tables:

- **Origin:** This table shows metrics for replicated data received from a participant by this cluster. The following metrics are displayed: Bytes Received Per Second, Total Bytes Received, Messages Received Per Second, Record Backlog Delay Time, Total Entries Received, Message Apply Time, Total Records Received, and Remote Participant.
- **Destination:** This table shows metrics of the replicated data sent to a participant by this cluster. The following metrics are displayed: Bytes Sent Per Second, Current Bandwidth, Total Entries Sent, Total Bytes Sent, Total Messages Unacknowledged, State, Status, and Remote Participant.

24.2.2.1 Node Menu Navigation

The following key menu options are available from the Oracle Coherence Node Home page:

- **Performance Summary:** From the **Oracle Coherence Node** menu, select **Monitoring**, then select **Performance Summary**. You can view the performance of the cluster on this page. See [Performance Summary Page](#).
- **Metric and Collection Settings:** From the **Oracle Coherence Node** menu, select **Monitoring**, then select **Metric and Collection Settings**. You can set up corrective actions to add nodes and caches as Enterprise Manager targets.
- **Logs:** You can use the log viewer to view log messages for the selected node. For more details, see [Log Viewer](#).


Note: This option is not available for managed Coherence clusters.

- **Components:** From this menu, you can navigate to the following pages:
 - **Coherence Topology:** The Topology Viewer provides a visual layout of the Coherence deployment and shows the Coherence cluster and its associated nodes and caches. See [Topology Viewer](#) for details.
 - **Caches:** This page lists all the caches associated with the selected node. See [Caches Page](#) for details.
 - **Services:** This page lists all services associated with the selected node. See [Services Page](#) for details.
 - **Administration:** From the **Oracle Coherence Node** menu, select **Administration**. See [Node Administration Page](#) for details.
- Note:** This option is not available for managed Coherence clusters.
- **Latest Configuration:** From the **Oracle Coherence Node** menu, select **Configuration**, then select **Latest** to view the latest configuration data for the Coherence cluster.
 - **JVM Diagnostics:** From the **Oracle Coherence Node** menu, select **JVM Diagnostics** to view the Coherence Node JVM Pool Drill Down page. This option is available only if the node has been configured for JVM Diagnostics and the WLS Management Pack EE Management Pack has been included. See [Coherence Integration with JVM Diagnostics](#) for details.

24.2.3 Cache Home Page

This page provides detailed information of a selected cache. From the **Coherence Cluster** menu, select **Caches**, and click on a specific cache to drill down to the Cache Home page.

Figure 24-7 Cache Home Page



It contains the following regions:

- **Summary**

- **General**

- * Availability: The availability of the cache over the last 24 hours.
- * Coherence Cluster: The cluster with which this cache is associated.
- * Federation Service: Indicates if this cache is participating in data federation. A federated cache permits the capture and later replay of operations, performed against cache entries, across a federation.

- **Nodes**

- * Total Nodes: The total number of nodes in the cluster. Click on the link to drill down to the All Nodes page.
- * Storage Nodes: The number of storage enabled nodes in the cluster. Click on the link to drill down to the Storage Nodes page.

 **Note:**

New storage enabled nodes are not automatically added to the cluster. You must refresh the cluster to add node targets for physical nodes added to cluster.

- * Non Storage Nodes: The nodes that are not storage enabled such as proxy, client nodes, and so on. These are relevant for front caches only. See [Near Cache](#) for details.

- **Cache Size:**

- * Objects: The aggregate number of objects in the cache.
- * Units: The aggregate number of units in the cache.

- * Memory: The aggregate memory used by the cache.
- * Total High / Low Units: This represents the high and low units configured for the cache. If this parameter has not been configured, an **n/a** will be displayed.
- **Cache Usage**
 - * Total Gets: The aggregate number of get operations across all nodes supporting this cache in the last 24 hours.
 - * Total Misses: The aggregate number of cache misses across all nodes supporting this cache in the last 24 hours.
 - * Total Puts: The aggregate number of put operations across all nodes supporting this cache in the last 24 hours.
- **Queries**
 - * Non Optimized Queries: The total execution time, in milliseconds for queries that could not be resolved per minute.
 - * Optimized Queries: The total number of parallel queries that were fully resolved using indexes per minute.
- **Service**
 - * Service: The service supporting this cache.
 - * Task Backlog: The size of the backlog queue that holds tasks scheduled to be executed across all services.
- **Storage Manager** (These metrics are applicable only for Back caches.)
 - * **Total Evictions:** The aggregate number of evictions from the backing map managed by this Storage Manager.
 - * **Total Events Dispatched:** The total number of events dispatched by the Storage Manager per minute.
- **Overview of Incidents and Problems**

This region lists any incidents that have occurred over the last 7 days and any problems in the node and its associated host target. Click on the link to drill down to the Incident Manager page.
- **Key Indicators**

This region displays graphs with key metrics that indicate the health and performance of the node over the last 24 hours. You can customize the metrics that are charted by selecting them from metric palette.
- **Top Components**

This region contains a graphical representation of the top 10 performing targets for a selected metric from the last configuration collection. The top components are listed in ascending or descending order depending on the metric selected and indicates how the top component data has been collected. Select a metric from the View drop down list to see a graphical representation of the top 10 targets for the selected metric. For example, if you select the Cache - Cache Objects metric, the graph displays the top 10 cache targets.
- **Components**

This region lists the nodes with which the cache is associated. Click on the Name link to drill down to the Node Home page.

24.2.3.1 Near Cache

A near cache is a hybrid cache; it typically fronts a distributed cache or a remote cache with a local cache. A **near cache** invalidates front cache entries, using a configured invalidation strategy, and provides excellent performance and synchronization. Near cache backed by a partitioned cache offers zero-millisecond local access for repeat data access, while enabling concurrency and ensuring coherency and fail over, effectively combining the best attributes of replicated and partitioned caches.

The objective of a **near cache** is to provide the best of both worlds between the extreme performance of the Replicated Cache and the extreme scalability of the Distributed Cache by providing fast read access to Most Recently Used (MRU) and Most Frequently Used (MFU) data. Therefore, the **near cache** is an implementation that wraps two caches: a "front cache" and a "back cache" that automatically and transparently communicate with each other by using a read-through/write-through approach. The "front cache" provides local cache access. It is assumed to be inexpensive, in that it is fast, and is limited in terms of size. The "back cache" can be a centralized or multitiered cache that can load-on-demand in case of local cache misses. The "back cache" is assumed to be complete and correct in that it has much higher capacity, but more expensive in terms of access speed.

If a **near cache** is present in the cluster, you will see a tabbed Cache Home, one for each of back and front caches respectively.

Figure 24-8 Near Cache (Back Cache)

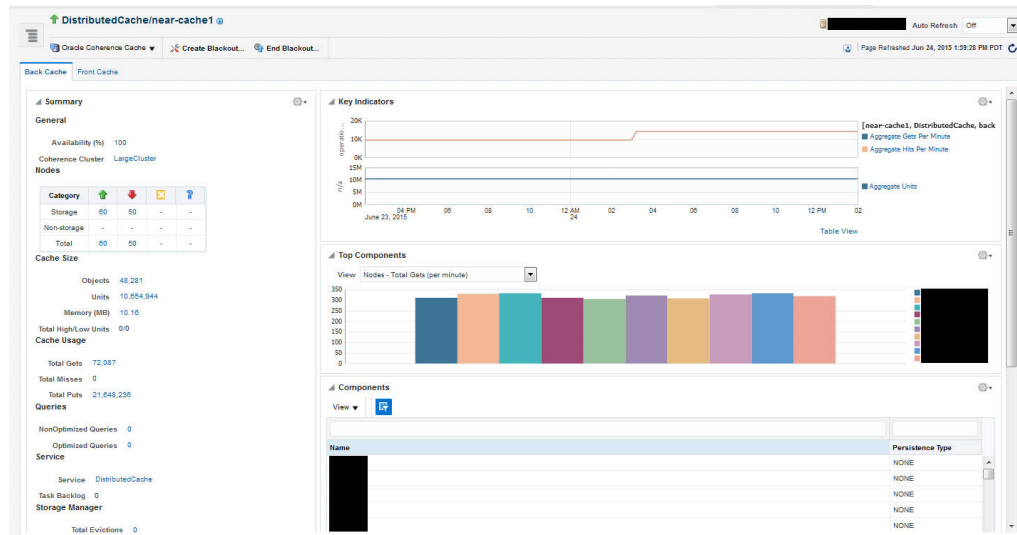



Figure 24-9 Near Cache (Front Cache)



24.2.3.2 Cache Menu Navigation

The following key menu options are available from the Coherence Cache Home page:

- **Performance Summary:** From the **Oracle Coherence Cache** menu, select **Monitoring**, then select **Performance Summary**. You can view the performance of the cluster on this page. See [Performance Summary Page](#).
- **Metric and Collection Settings:** From the **Oracle Coherence Cache** menu, select **Monitoring**, then select **Metric and Collection Settings**. You can set up corrective actions to add nodes and caches as Enterprise Manager targets.
- **Components:** You can navigate to the following pages from this menu:
 - **Coherence Topology:** The Topology Viewer provides a visual layout of the Coherence deployment and shows the Coherence cluster and its associated nodes and caches. See [Topology Viewer](#) for details.
 - **Nodes:** This page lists all the nodes associated with the selected cache. See [Nodes Page](#) for details.
 - **Services:** This page lists all services associated with the selected cache. See [Services Page](#) for details.
- **Administration:** From the **Oracle Coherence Cache** menu, select **Administration**. See [Cache Administration Page](#) for details.
- **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. See [Cache Data Management](#) for details.
- **Latest Configuration:** From the **Oracle Coherence Cluster** menu, select **Configuration**, then select **Latest** to view the latest configuration data for the Coherence cluster.
- **JVM Diagnostics:** From the **Oracle Coherence Cache** menu, select **JVM Diagnostics** to view the Coherence Cache JVM Diagnostics Pool Drill Down page. This option is available only if the cluster has been configured for JVM Diagnostics and the WLS Management Pack EE Management Pack has been included. See [Coherence Integration with JVM Diagnostics](#) for details.

24.2.4 Partition Cache Home Page

This page provides detailed information for a selected partition cache. From the **Coherence Cluster** menu, select **Caches**, and click on a specific partition cache to drill down to the Cache Home page. This page contains the following regions:

- **Summary**
 - General
 - * Coherence Cluster: The cluster with which this cache is associated.
 - * Domain Partition: The name of the domain partition with which the cache is associated is displayed here.
 - * Federation Service: Indicates if this cache is participating in data federation. A federated cache permits the capture and later replay of operations, performed against cache entries, across a federation.
 - Cache Size:
 - * Objects: The total number of objects in the cache.
 - * Units: The amount of memory used by the cache in units.
 - * Memory: The aggregate memory used by the cache.
 - * Total High / Low Units: This represents the high and low units configured for the cache. If this parameter has not been configured, an n/a will be displayed.
 - Cache Usage
 - * Total Gets: The aggregate number of get operations across all nodes supporting this cache in the last 24 hours.
 - * Total Misses: The aggregate number of cache misses across all nodes supporting this cache in the last 24 hours.
 - * Total Puts: The aggregate number of put operations across all nodes supporting this cache in the last 24 hours.
 - Queries
 - * Non Optimized Queries: The total execution time, in milliseconds for queries that could not be resolved per minute.
 - * Optimized Queries: The total number of parallel queries that were fully resolved using indexes per minute.
 - Service
 - * Task Backlog: The size of the backlog queue that holds tasks scheduled to be executed by one of the service pool threads.
- **Key Indicators**

This region displays graphs with key metrics that indicate the health and performance of the node over the last 24 hours. You can specify the key metrics that are to be included in the charts.
- **Top Components**

This region contains a graphical representation of the top 10 performing targets for a selected metric from the last configuration collection. Select a metric from the

View drop down list to see a graphical representation of the top 10 targets for the selected metric. For example, if you select the Cache-Cache Objects metric, the graph displays the top 10 cache targets.

24.2.4.1 Cache Menu Navigation

The following key menu options are available from the Coherence Cache Home page:

- **Performance Summary:** From the **Oracle Coherence Cache** menu, select **Monitoring**, then select **Performance Summary**. You can view the performance of the cluster on this page. See [Performance Summary Page](#) for details.
- **Metric and Collection Settings:** From the **Oracle Coherence Cache** menu, select **Monitoring**, then select **Metric and Collection Settings**. You can set up corrective actions to add nodes and caches as Enterprise Manager targets.
- **Latest Configuration:** From the **Oracle Coherence Cluster** menu, select **Configuration**, then select **Latest** to view the latest configuration data for the Coherence cluster.

24.2.5 Application Home Page

This page allows you to view and monitor the application data stored in various types of caches. To view this page, select the **Applications** option from the **Oracle Coherence Cluster** menu.

If an application contains multiple web modules, the application data for each module is displayed. Click **Reset Statistics** to reset the session management statistics.

The following graphs are displayed:

- Local Attribute Count: Shows the local attribute count.
- Local Session Count: Shows the local session count.
- Overflow Updates: Shows the number of overflow updates per minute.
- Session Updates: Shows the number of session updates per minute
- Reap Duration: Shows the average reap duration in milliseconds.
- Reap Session: Shows the average number of reaped sessions in a reap cycle.

Overflow Cache

This table contains the following details:

- Module: The name of the Coherence cluster with the application.
- Node ID: This is the node target name. Click on the link to drill down to the Node Home page.
- Cache: This is the name of the cache target. Click on the link to drill down to the Cache Home page.
- Average Size: The average size (in bytes) of a session object placed in the session storage clustered cache since the last time statistics were reset.
- Max Size: The maximum size (in bytes) of a session object placed in the session storage clustered cache since the last time statistics were reset.

- **Threshold:** The minimum length (in bytes) that the serialized form of an attribute value must be in order for that attribute value to be stored in the separate "overflow" cache that is reserved for large attributes.
- **Overflow Updates:** The number of updates to session attributes stored in the "overflow" clustered cache since the last time statistics were reset.

Clustered Session Cache

- **Module:** The name of the Coherence cluster with the application.
- **Node ID:** This is the node target name. Click on the link to drill down to the Node Home page.
- **Cache:** This is the name of the cache target. Click on the link to drill down to the Cache Home page.
- **Average Size:** The average size (in bytes) of a session object placed in the session storage clustered cache since the last time statistics were reset.
- **Min Size:** The minimum size (in bytes) of a session object placed in the session storage clustered cache since the last time statistics were reset.
- **Max Size:** The maximum size (in bytes) of a session object placed in the session storage clustered cache since the last time statistics were reset.
- **Session ID Length:** The length of the generated session IDs.
- **Timeout:** The session expiration time (in seconds) or -1 if sessions never expire.
- **Session Updates:** The number of updates of session objects stored in the session storage clustered cache per minute.
- **Pinned Objects:** The number of session objects that are pinned to this instance of the web application or -1 if sticky session optimizations are disabled.

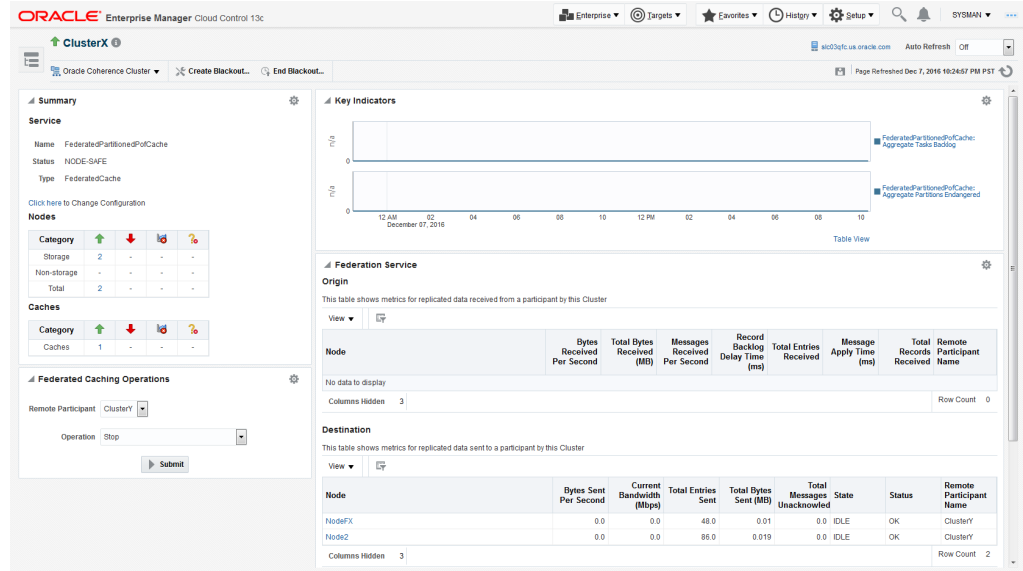
Reaped Sessions

- **Module:** The name of the Coherence cluster with the application.
- **Node ID:** This is the name of the node target. Click on the link to drill down to the Node Home page.
- **Average Reap Duration:** The average reap duration in minutes.
- **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.

24.2.6 Service Home Page

This page shows all the details of a service in a coherence cluster. From the **Coherence Cluster** menu, select **Members**, and click **Services**. In **All Services** page, select a **FederatedCache** type service.

Figure 24-10 Service Home Page



It contains the following regions:

- **Name:** The name assigned to the service.
- **Nodes:** The number of nodes in the service.
- **Type:** Some of the service types available are:
 - **Cluster Service:** This service is started when a cluster node needs to join the cluster. It keeps track of the membership and services in the cluster.
 - **Distributed Cache Service:** Allows cluster nodes to distribute (partition) data across the cluster so that each piece of data in the cache is managed (held) by only one cluster node.
 - **Invocation Service:** This service provides clustered invocation and supports grid-computing architecture.
 - **Replicated Cache Service:** This is the synchronized replicated cache service, which fully replicates all of its data to all cluster nodes that are running the service.
 - **Federated Cache Service:** This service is a version of the distributed cache service that replicates and synchronizes cached data across geographically dispersed clusters that are participants in a federation.
- **Key Indicators:** This region displays graphs with key metrics that indicate the health and performance of the service over the last 24 hours. You can specify the key metrics that are to be included in the charts.
- **Top Components:** This region contains a graphical representation of the top 10 performing targets for a selected metric from the last configuration collection. Select a metric from the View drop down list to see a graphical representation of the top 10 targets for the selected metric.
- **Federation Service:** If the cluster is participating in data federation, then the table Origin and Destination are displayed. Also, if the Federation Service is running on

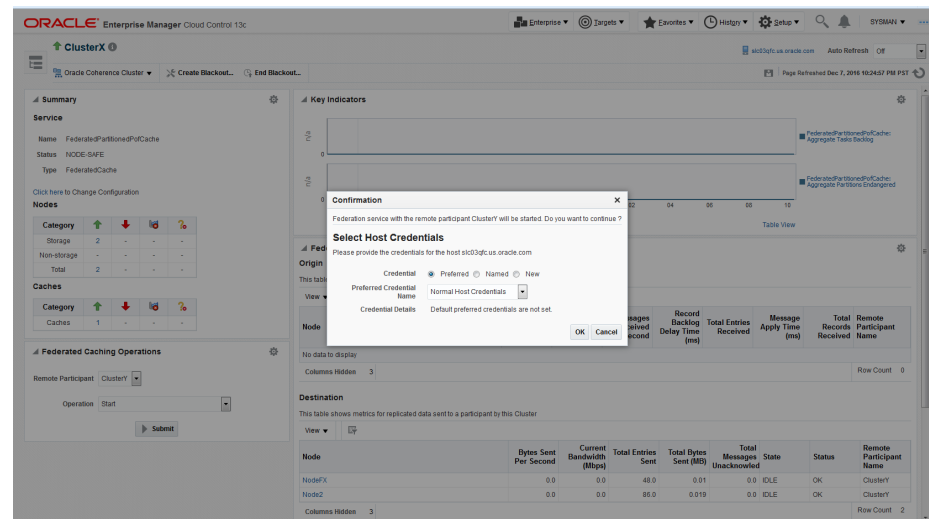
Domain Partition Caches, then the Domain Partition column will be displayed in these tables:

- **Origin:** This table shows metrics for replicated data received from a participant by the cluster. The following metrics are displayed: Bytes Received Per Second, Total Bytes Received, Messages Received Per Second, Record Backlog Delay Time, Total Entries Received, Message Apply Time, and Remote Participant.
- **Destination:** This table shows metrics for replicated data sent to a participant by this cluster. The following metrics are displayed: Bytes Sent Per Second, Current Bandwidth, Total Entries Sent, Total Messages Unacknowledged, State, Status, and Remote Participant.
- **Federated Caching Operations:** If the Service type is FederatedCache, then the following federation specific operations are supported:
 - **Remote Participant:** Displays the remote participant to submit for operation request.
 - **Operation:** Displays the coherence federation coordinator operations. You can select any of the following corresponding operation to perform:
 - * Stop
 - * Start
 - * Pause
 - * Replicate All
 - * Retrieve State
 - * Retrieve Pending Incoming Messages
 - * Retrieve Pending Outgoing Messages

Depending on selected federated operation, it will be executed with reference to that remote participant.

A confirmation pop-up will be shown where you will be prompted to provide credentials.


Figure 24-11 Confirmation Pop Up



24.2.7 Connection Manager Home Page

Use this page to view the Connection Manager details in the Coherence cluster.

Figure 24-12 Connection Manager Home Page



This page contains the following sections:

- **General**
 - Service Name: The unique name assigned to the service.
 - Node ID: This is the node target name.
 - Connection Count: The number of connections associated with the connection manager instance.
 - Incidents: Any incidents that have occurred.
 - Host IP: The IP address of the host machine.
- **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.
 - Node ID: This is the node target name.
 - 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.

24.3 Viewing the Summary Pages

These pages describe the target pages such as nodes, caches, services, and so on associated with the cluster.


24.3.1 Nodes Page

This page lists all the discovered node targets that belong to the cluster, support a cache, or a service. The list of nodes displayed will vary depending on how you have navigated to this page.

This is a master detail page where you can select a node in the master table to view the key performance metrics in the Details region. The list of nodes displayed here can vary based on how you have navigated to this page. To view this page, perform the following steps:

- From the **Targets** menu, select **Middleware**, then click on a Coherence Cluster. In the Oracle Coherence Cluster Home page, select **Nodes** from the **Oracle Coherence Cluster** menu. You can also navigate to this page from the Cache Home page.
- Click **Storage**, **Non Storage Nodes**, or the **Number of Nodes** link in the Oracle Coherence Cluster Home page.

Figure 24-13 All Nodes Page



The following details are displayed by default. To display the hidden fields, from the **View** menu, select **Columns**, then select **Manage Columns**. In the Manage Columns table, select one or more columns from the **Hidden Column** list, move them to the **Visible Columns** list and click **OK**. The selected fields will be displayed in the table.

 **Note:**

You can filter the list of nodes displayed in the table by specifying values in the Query by Example fields at the top of the table. If you want to see a list of nodes that are running on **xyz** host for instance, you can enter '**xyz**' in the Host query field.

- **Name:** This is the name of the node target. Click on the link to drill down to the Node Home page.
- **Status:** Shows whether the node is Up, Down, in an Error, or Unknown status.
- **Host:** The host on which node is running. If the host is a monitored target in Enterprise Manager, you can click on the link to drill down to the Host Home page.
- **Caches:** The total number of cache targets that this node supports.
- **Receiver Success (%):** The percentage of received packets out of the total packets sent.
- **Publisher Success (%):** This is the rate at which the publisher transmits packets on the network.
- **Memory Available (MB):** The memory available on this node.
- **Total Puts:** The aggregate number of put operations.
- **Total Gets:** The aggregate number of get operations.

The following federation metrics are displayed only for nodes participating in data federation:

- **Total Messages Received:** The total number of messages received.
- **Total Bytes Received (MB):** The total number of bytes received.
- **Total Messages Sent:** The total number of messages sent.
- **Total Bytes Sent (MB):** The total number of bytes sent.

Select a node in the table to view a detailed graphical representation of the node. The following graphs are displayed.

- **Node Memory Available:** This graph shows the nodes that have lowest available memory over the last 24 hours.
- **Aggregate Gets Per Minute:** This graph displays the aggregate get operations across all the caches supported by the selected node.
- **Aggregate Puts Per Minute:** This graph displays the aggregate put operations across all the caches supported by the selected node.
- **Publisher Success Rate:** These graphs show the rate at which the publisher transmits packets on the network.
- **Receiver Success Rate:** The percentage of received packets out of the total packets sent.
- **Node Memory Used (MB):** The total memory used by the node.
- **CPU Usage (%):** The CPU percentage used.



Note:

You can use the Personalization feature to customize these charts.

You can perform the following actions:


- **Start:** You can start any node that has a **Down** status. This option is available only if the node is running on an Enterprise Manager monitored host.
- **Stop:** You can stop any node that has a **Up** status. This option is available only if the node is running on an Enterprise Manager monitored host.
- **Start New Nodes:** You can start new nodes on the same host on which a selected node is running. The host must be monitored by Enterprise Manager.
- **Reset Statistics:** Select a node and click **Reset Statistics**. You are prompted for the password for the host on which the node is running. Enter the password and click OK to reset the statistics. This option is available only for nodes with an Up status.
- **Query by Example:** Click the **Query by Example** icon. In the Query row that appears, enter a query string in any of the columns to search for. All nodes that meet the specified criteria are displayed.

24.3.2 Caches Page

This page lists all the discovered cache targets that belong to the cluster. This is a master detail page where you can select a cache in the master table to view the key performance metrics in the Details region. The list of nodes displayed here can vary based on how you have navigated to this page. To view this page, you can:

- From the **Targets** menu, select **Middleware**, then click on a Coherence Cluster. In the Oracle Coherence Cluster Home page, select **Nodes** from the **Oracle Coherence Cluster** menu.
- Click on the **Caches** link in the Oracle Coherence Cluster Home page.

Figure 24-14 All Caches Page



For each cache, the following details are displayed:

- **Name:** This is the name of the cache target. Click on the link to drill down to the Cache Home page.
- **Domain Partition:** If you are monitoring a multi-tenant managed Coherence cluster, the domain partition with which the cache is associated is displayed here.
- **Service:** The name of the caching service used by the cache.
- **Tier:** The back tier is displayed for most caches. For a Near Cache, the cache can have front and back tiers. In this case, multiple rows for the same cache with unique tier values will be displayed.
- **Objects:** The number of objects in the cache.
- **Gets:** The aggregate number of get() operations in the cache.
- **Hits:** The aggregate number of successful fetches of cached objects.
- **Misses:** The aggregate number of failed fetches of cached objects.
- **Reads:** The aggregate number of reads to a data store.
- **Writes:** The aggregate number of writes to a data store.

Select a cache in the table to view a detailed graphical representation of the aggregated values across all the nodes supporting a cache. For example, Aggregate Puts Per Minute is the per minute value computed for put operations aggregated across all nodes supporting a cache.

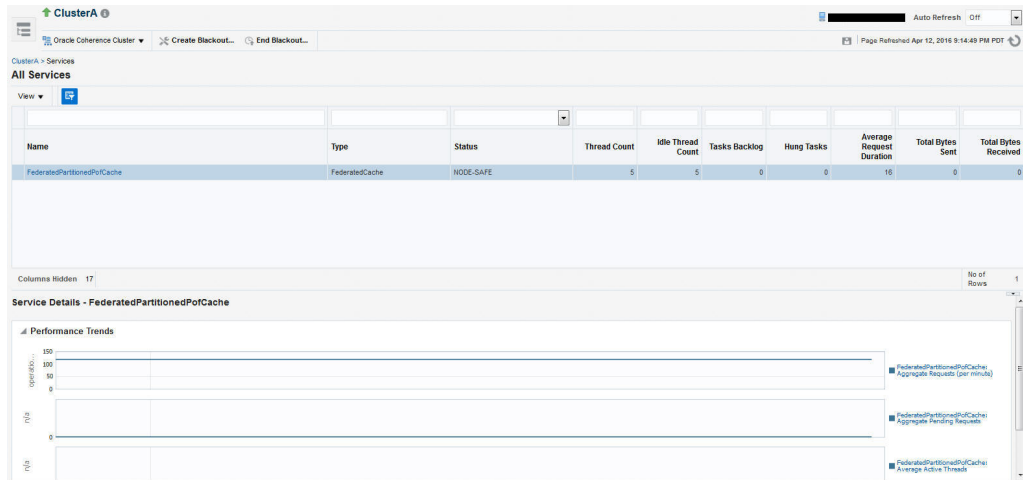
By default, the following graphs are displayed but this can be customized. Click the Personalize button and select the graphs to be displayed and the metrics to be included in each graph.

- **Aggregated Puts Per Minute:** The aggregate number of put operations per minute across all the nodes supporting this cache.
- **Aggregated Hits Per Minute:** The aggregate number of get operations per minute across all the nodes supporting this cache.
- **Aggregated Misses Per Minute:** The aggregate number of failed fetches of the cached objects per minute across all the nodes supporting this cache.
- **Aggregated Evictions Per Minute:** The aggregate number of eviction operations per minute across all the nodes supporting this cache.
- **Aggregate Inserts Per Minute:** The aggregate number of insert operations per minute across all the nodes supporting this cache.
- **Aggregate Removes Per Minute:** The aggregate number of delete operations per minute across all the nodes supporting this cache.

24.3.3 Services Page

This page lists all the discovered service targets that belong to the cluster. This is a master detail page where you can select a service in the master table to view the key performance metrics in the Details region. The list of nodes displayed here can vary based on how you have navigated to this page. To view this page, select the **Services** option from the **Oracle Coherence Cluster** menu.

Figure 24-15 Services Page



For each service, the following details are displayed:

- Name: The name assigned to the service. Click on the link to drill down to the Service Home page.
- Type: Some of the service types available are:
 - Cluster Service: This service is started when a cluster node needs to join the cluster. It keeps track of the membership and services in the cluster.
 - Distributed Cache Service: Allows cluster nodes to distribute (partition) data across the cluster so that each piece of data in the cache is managed (held) by only one cluster node.
 - Invocation Service: This service provides clustered invocation and supports grid-computing architecture.
 - Replicated Cache Service: This is the synchronized replicated cache service, which fully replicates all of its data to all cluster nodes that are running the service.
 - Federated Cache Service: This service is a version of the distributed cache service that replicates and synchronizes cached data across geographically dispersed clusters that are participants in a federation.
- Domain Partition: If you are monitoring a multi-tenant managed Coherence cluster, the domain partition with which the cache service is associated is displayed.
- Status: The High Availability status for this service. This can be:
 - MACHINE-SAFE: This means that all the cluster nodes running on any given machine could be stopped at once without data loss.
 - NODE-SAFE: This means that any cluster node could be stopped without data loss.
 - ENDANGERED: This indicates that termination of any cluster node that runs this service may cause data loss.
 - RACK-SAFE: This status indicates that a rack can be stopped without any data loss.

- **SITE-SAFE:** This status indicates that a site can be stopped without any data loss.
- **Thread Count:** The number of threads in the service thread pool.
- **Idle Thread Count:** The number of currently idle threads in the service thread pool.
- **Tasks Backlog:** The size of the backlog queue that holds tasks scheduled to be executed by one of the service pool threads.
- **Hung Tasks:** The id of the of the longest currently executing hung task.
- **Average Request Duration:** The average duration (in milliseconds) of an individual synchronous request issued by the service.

If the service is participating in data federation, the following metrics are displayed:

- **Total Messages Received:** The total number of messages received.
- **Total Bytes Received (MB):** The total number of bytes received.
- **Total Messages Sent:** The total number of messages sent.
- **Total Bytes Sent (MB):** The total number of bytes sent.

Select a service in the table to view a detailed graphical representation of the aggregated values across all the nodes supporting the service. The following graphs are displayed.

- **Aggregated Requests Per Minute:** The total number of the synchronous requests issued by the service.
- **Aggregated Pending Requests:** This graph displays the aggregate number of pending requests issued by the service.
- **Average Active Threads:** This graph displays the average number of active threads in the service thread pool.

24.3.4 Applications Page

This page lists all the applications associated with the cluster. For each application, the following details are displayed:

- Local Attribute Cache
- Local Session Cache
- Overflow Cache
- Clustered Session Cache

Click on the Application Name link to drill down to the Application Home page.

24.3.5 Proxies Page

This page shows the performance of all connection managers and connections in the cluster. To view this page, select **Proxies** from the **Coherence Cluster** menu. The following Connection Manager graphs are displayed:

- **Top Connection Managers with Most Bytes Sent** since the connection manager's statistics were last reset.
- **Top Connection Managers with Most Bytes Received** since the connection manager's statistics were last reset.

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:** This is the node target name.
- **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's statistics were last reset.
- **Top Connections with Highest / Most Bytes Received** since the connection's statistics were last reset.

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:** 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:** This is the node target name.
- **Bytes Sent:** The number of bytes sent per minute.
- **Bytes Received:** The number of bytes received per minute.
- **Connection Time:** The connection time in milliseconds.
- **Outgoing Message Backlog:** The number of outgoing messages in the backlog.
- **Outgoing Byte Backlog:** The number of outgoing bytes in the backlog.

24.4 Log Viewer

The Log Viewer scans the log files produced by the nodes and this log data is shown in the log viewer. This section describes the following:

- [Configuring the Log Location Settings](#)
- [Viewing the Log Messages](#)

24.4.1 Configuring the Log Location Settings

Before you can view the log data, you must configure the log file location. To configure the log file location, follow these steps:

1. Navigate to the Coherence Cluster Home page. From the **Oracle Coherence Cluster** menu, select **Logs**, then select **Configure Log Location Settings**.

Figure 24-16 Configure Log Location Settings

2. Select a target and click **Assign Host Credentials** and provide the host name and login credentials. These credentials are used to access the host and retrieve the log locations.
3. Select the **Apply Above Host Credentials to Child Targets** checkbox to apply these credentials to all related child targets.
4. Click **Assign Log Location** and select the directory in which the log files are to be stored. Click **OK** to return to the Configure Log Location Settings page.

24.4.2 Viewing the Log Messages

After the log file has been configured, you can view the log messages. From the **Oracle Coherence Node** menu, select **Logs**, then select **View Log Messages**. Click the **Search** icon and specify the date range, the message type, and any other additional search criteria. Click **Search**. The list of messages that meet the search criteria are displayed.

Figure 24-17 Log Messages

Select a message to view more details of the message. Click **Export Messages to File** and specify the format which can .txt, .xml, or .csv. The messages will be exported to a file in the selected format. Click the **Log File** link to drill down to the log details page. Click **Download** to download the log file data to an external file.

24.5 Viewing the Performance Pages

This section describes the Performance Summary page, and the connection manager performance pages.

24.5.1 Performance Summary Page

The Performance Summary page can be used to monitor the performance of the selected component or application. To view this page, select **Monitoring**, then **Performance Summary** from the menu for any Coherence target such as cluster, node, cache, or domain partition cache. The performance page typically contains:

- A set of default performance charts that shows the values of specific performance metrics over time. You can customize these charts to help you isolate potential performance issues.
- A series of regions that is specific to the component or application. For example, the Oracle Cache Performance Summary page displays metrics such as Aggregate Cache Objects, Aggregate Evictions, Maximum Query Duration, and so on. These sections will vary from component to component.

24.5.1.1 Customizing the Performance Page Charts

The Performance page is configured to provide a default set of metric charts, but you can customize the charts in different ways. You can identify potential performance issues by correlating and comparing specific metric data. To customize the charts, some of the actions you can perform are:

- Click **Show Metric Palette** to display a hierarchical tree, containing all the metrics for the selected component or application. The tree organizes the performance metrics into various categories of performance data.
- Select a metric in the palette to display a performance chart that shows the changes in the metric value over time. The chart will continue to refresh automatically to show updated data.
- Click the "x" icon on the chart to close a chart. Click and drag the right side of the chart to move the chart to a new position on the page.
- Drag and drop a metric from the metric palette and drop it on top of an existing chart. The existing chart will show the data for both metrics.

See the Enterprise Manager Online Help for more details on customizing the Performance Page.

24.5.2 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.
- Bytes Received: The number of bytes received per minute.
- Bytes Sent: The number of bytes sent per minute.

24.6 Removing Down Members

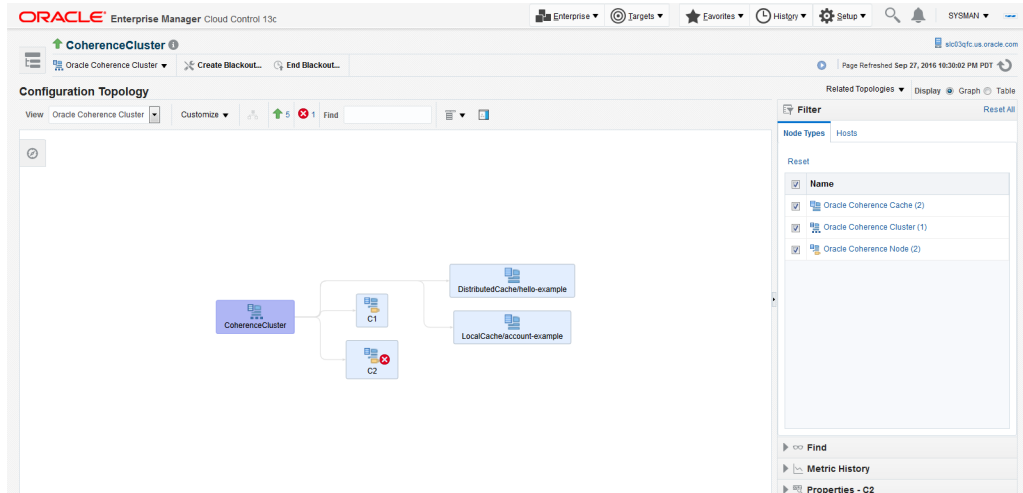
You can delete any members in the cluster that have a **Down** status. From the **Oracle Coherence Cluster** menu, select **Remove Down Members**. You will see a list of targets that are down. Select the target from the list and click **Remove**. A confirmation message is displayed. Click **OK** to delete the member from the cluster.

24.7 Topology Viewer

The topology viewer provides a customized view of all the targets in a Coherence cluster. You can view the topology for a cluster, node, or cache. To view the topology, select the Topology Viewer option from the Oracle Coherence Cluster, Oracle Coherence Node, or Oracle Coherence Cache menu. The topology graph is rendered based on the context of the selected target. If you launch the topology viewer from:

- Cluster Home Page: The topology for the entire cluster is displayed.
- Cache: The topology of the hosts and nodes on which the selected cache is running is displayed.
- Node: The topology of the node and the caches running on the node is displayed.

Figure 24-18 Coherence Cluster Topology



The topology is organized into 3 tiers, cluster, hosts, and services. All nodes running on the host are grouped under host group. All caches of a service are grouped under the service folder. When you select a node, links to all the caches it supports are displayed. If you select a cache, links to all nodes on which the cache is configured are displayed.

If you hover over a target, you can view the detailed information for the target. For example, if you hover over a cluster target, you can see the name of the cluster, status, the host, nodes, and incidents if any. Click on the name link to drill down to the Home page for the target. The relationship between the targets is depicted by arrows. For example, if you click on a service target, you can see arrows showing the nodes on which cache is running.

24.8 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.

25

Administering a Coherence Cluster

This chapter describes the administration options available for the various Coherence targets. It contains the following sections:

- [Cluster Administration Page](#)
- [Node Administration Page](#)
- [Cache Administration Page](#)
- [Service Administration Page](#)
- [Cache Data Management](#)


25.1 Cluster Administration Page

This page allows you to change the configuration of nodes, caches, and services.

Note:

Any changes made to the configuration are applied to the active cluster but will not be saved.

Figure 25-1 Cluster Administration Page



The screenshot shows the 'Administration' page for 'Coherence: Cluster1'. It is divided into three sections: 'Node', 'Service', and 'Cache'. Each section has a 'Change Configuration' header and a 'Select a [entity] and its [entity] to change configuration' instruction. The 'Node' section has a 'Node' dropdown menu with 'CoherenceTestApp' selected and a 'Go' button. The 'Service' section has a 'Service' dropdown menu with 'DistributedCacheWithPublishingCacheStore' selected, a 'Node' dropdown menu with 'CoherenceTestApp' selected, and a 'Go' button. The 'Cache' section has a 'Service' dropdown menu with 'DistributedCache' selected, a 'Cache' dropdown menu with 'dist-cache1' selected, a 'Node' dropdown menu with a blacked-out selection, a 'Tier' dropdown menu with 'back' selected, a 'Loader' dropdown menu with 'None' selected, and a 'Go' button.

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.

25.1.1 Changing the Node Configuration

To change the node configuration, select a node from the Node drop down list and click **Go**. The Change Configuration on Node page appears.

Figure 25-2 Change Node Configuration

The screenshot shows the 'Change Configuration on Node' page for a node in the Oracle Coherence Cluster. The page is titled 'Change Configuration on Node: skl03q/c65/skl03q/c' and includes a breadcrumb trail 'Coherence: LargeCluster > Nodes: [redacted]'. The page is divided into several sections:

- Connection:** Multicast Threshold (%) is set to 25.
- Network:** Resend Delay (ms) is 200, Send Ack Delay (ms) is 16, Traffic Jam Count is 8192, and Traffic Jam Delay (ms) is 10.
- Logging:** Logging Limit is 2147483647 and Logging Level is 9.
- Credentials:** Select credential from one of the following options: Credential (selected), Preferred, Named, New. Preferred Credential Name is Normal Host Credentials. Credential Details: Default preferred credentials are not set.

You can change the following values:

Connection

- **Multicast Threshold:** The percentage (0 to 100) of the servers in the cluster that a packet will be sent to, above which the packet will be multicasted and below which it will be unicasted.

Network

- **Resend Delay:** The minimum number of milliseconds that a packet will remain queued in the Publisher's re-send queue before it is resent to the recipient(s) if the packet has not been acknowledged.
- **Traffic Jam Count:** The maximum total number of packets in the send and resend queues that forces the publisher to pause client threads. Zero means no limit.
- **Send Ack. Delay:** The minimum number of milliseconds between the queuing of an Ack packet and the sending of the same. This value should be not more than a half of the Resend Delay value.
- **Traffic Jam Delay:** The number of milliseconds to pause client threads when a traffic jam condition has been reached. Anything less than one (e.g. zero) is treated as one millisecond.

Logging

- **Logging Level:** Specifies which logged messages will be output to the log destination.
- **Logging Limit:** The maximum number of characters that the logger daemon will process from the message queue before discarding all remaining messages in the queue.

Credentials

Specify the credentials of the host on which the Management Agent is running.

Usage Tips

- Click **Update** to save the changes of the selected node and click **Return** to return to the Node Administration page.
- Click **Update All** to update all the nodes in the cluster and click Return to return to the Node Administration page.

25.1.2 Changing the Cache Configuration

Use this page to modify the configuration of the selected node of the cache. You can change the following values:

- **High Units:** The limit of the cache size measured in units. The cache will prune itself automatically once it reaches its maximum unit level.
- **Low Units:** The number of units to which the cache will shrink when it prunes.
- **Expiry Delay:** The time-to-live for cache entries in milliseconds. Value of zero indicates that the automatic expiry is disabled.

Usage Tips

- Click **Update** to save the changes of the selected node and return to the Cache Administration page.
- Click **Update All** to update all the nodes that support the selected cache.
- Click **Return** to return to the previous page.

25.1.3 Changing the Service Configuration

Use this page to modify the configuration of the selected node of the service. You can change the following values:

- **Request Timeout:** The request execution timeout value.

After you have modified the value of the parameters, you must specify the credentials. You can do either of the following:

- Click **Update** to save the changes of the selected node and click **Return** to return to the Service Administration page.
- Click **Update All Nodes** to update all the nodes that support the selected service and click **Return** to return to the **Service Administration** page.

25.2 Node Administration Page

On this page, you can perform the following administration tasks:

- **Change the Node Configuration:** Click **Change Configuration** to modify the configuration of the node. The Change Configuration page is displayed. Enter the new values and click **Update** to save the values and return to the Coherence Node Administration page. See [Changing the Node Configuration](#) for details.
- **Setup Log Alerts:** You can set up each Coherence node to log all its messages into a log file on the host on which this node is running. Click the **Log Alert Setup** link to drill down to the Metric and Policy Settings page. Configure the Log File Pattern Matched Line Count metric to specify a specific string pattern in the log file name. This metric should be set up on the host on which the Coherence node is running. The log file related alerts will be displayed in the Node Details Home page.

 **Note:**

You can set up Log Alerts only for nodes that are running on hosts monitored by Enterprise Manager.

25.3 Cache Administration Page

This page allows you to perform the following cache related administration tasks.

- **Cache Data Management:** Click **Go** to perform cache data management operations. The Cache Data Management page is displayed where you can perform operations like view, export, import, insert, update, purge, add, and remove indexes from Cache Data Management page. See [Cache Data Management](#) for details.
- **Changing the Cache Configuration:** Select a node from the list, Tier, Loader, 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 Cache Administration page. See [Changing the Cache Configuration](#) for details.

 **Note:**

To perform the Cache Data Management and Change Cache Configuration tasks, you need to login as a user with Administrator privileges.

25.4 Service Administration Page

This page allows you to change the configuration of a service. Select a node from the list 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 Service Administration page. See [Changing the Service Configuration](#) for details.

25.5 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.

Prerequisites for Cache Data Management

Before performing any Cache Data Management operation in Oracle Enterprise Manager, the following prerequisites must be met:

- **Prerequisites for Oracle Coherence configured with dedicated management node:**

The Coherence Management Node must include the `coherenceEMIntg.jar` from the location `<PLUGIN_HOME>/<MIDDLEWARE_MONITORING_PLUG-IN_DIRECTORY>/archives/coherence/` in its classpath.

- **Prerequisites for Oracle Coherence configured in dynamic management mode:**

The Coherence Node must include the `coherenceEMIntg.jar` from the location `<PLUGIN_HOME>/<MIDDLEWARE_MONITORING_PLUG-IN_DIRECTORY>/archives/coherence/12.2.1/` in its classpath.

- **Prerequisites for Oracle Managed Coherence Cluster with dedicated management node:**

The Coherence Management Node must include the `coherenceEMIntg.jar` from the location `<PLUGIN_HOME>/<MIDDLEWARE_MONITORING_PLUG-IN_DIRECTORY>/archives/coherence/` in its classpath.

Include `custom-mbeans.xml` file with below content in its classpath.

```
<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>
</mbeans>
```

- **Prerequisites for Oracle Managed Coherence Cluster with dynamic management node:**

The Coherence Node must include the `coherenceEMIntg.jar` from the location `<PLUGIN_HOME>/<MIDDLEWARE_MONITORING_PLUG-IN_DIRECTORY>/archives/coherence/12.2.1/` in its classpath.

Include `custom-mbeans.xml` file with below content in its classpath.

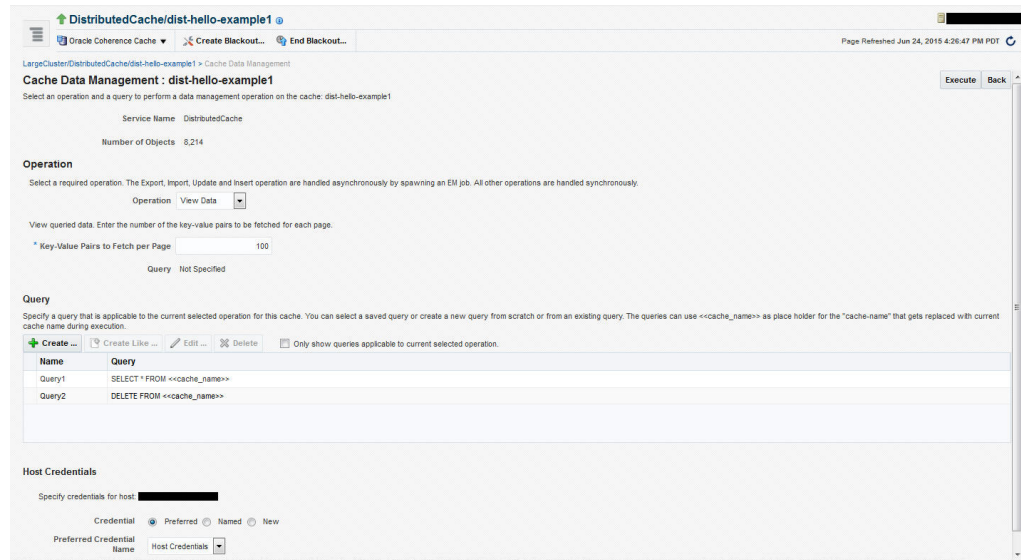
```
<mbeans>
<mbean id="100">
<mbean-class>oracle.sysman.integration.coherence.CacheDataManager</mbean-
class>
<mbean-name>type=Custom,name=CacheDataManager</mbean-name>
<enabled>>true</enabled>
</mbean>
</mbeans>
```

 **Note:**

Users with Administration privileges can use Cache Data Management only if the Cache Data Management MBean has been registered in the Coherence JMX management node.

To perform cache data management operations, select the **Cache Data Management** menu option from the **Oracle Coherence Cache** menu. You can also arrive at this page by selecting the **Administration** menu option from the **Oracle Coherence Cache** menu. Then on the Cache Administration page, click **Go** in the Cache Data Management section.

Figure 25-3 Cache Data Management



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 Data**: Select **View Data** 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.
- **Explain Plan**: Select **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.
- **Trace**: Select **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.

25.5.1 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 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 Explain Plan option can be used with Coherence version 3.7.1 or later.

25.5.2 Trace

The 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.

26

Troubleshooting and Best Practices

This chapter lists a few tips for troubleshooting Coherence and some Coherence best practices. It contains the following sections:

- [Troubleshooting Coherence](#)
- [Best Practices](#)

26.1 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.
- **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.

26.2 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:

- [Monitoring Templates](#)

26.2.1 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.

Coherence Integration with JVM Diagnostics

This chapter describes the JVM Diagnostics integration with Coherence. It contains the following sections:

- [Overview](#)
- [Configuring Coherence Nodes for JVM Diagnostics Integration](#)
- [Accessing JVM Diagnostics from Coherence Targets](#)
- [Including the JVM Diagnostics Regions in the Coherence Target Home Pages](#)

27.1 Overview

JVM Diagnostics provides deep visibility into the runtime of the JVM. It 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.

You can also use JVM Diagnostics to diagnose performance issues in Oracle Coherence cluster nodes. 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. To diagnose performance issue in a Coherence node, you must configure the node so that it can be monitored by JVM Diagnostics.

 **Note:**

JVM Diagnostics is a part of the WLS Management Pack EE Management Pack.

27.2 Configuring Coherence Nodes for JVM Diagnostics Integration

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 class path.

Additionally, you must add the `Doracle.coherence.jamjvmid` system property. The value of this property must match the value specified for `jamjvmid`. For more details on setting up the `jamjvmid` property, see the *Enterprise Manager Cloud Control Administrator's Guide*.

27.2.1 Example Start Script for Coherence Management Node

An example start script is given below.

```
#!/bin/sh

CP=$CP:<Path to jamagent.war>:<EM CC_Agent_Home>/plugins/
oracle.sysman.emas.agent.plugin_
12.1.0.6.0/archives/coherence/coherenceEMIntg.jar:
<EM CC_Agent_Home>/plugins/oracle.sysman.emas.agent.plugin_
12.1.0.6.0/archives/coherence/bulkoperationsmbean.jar
COH_OPTS="$COH_OPTS -cp $CP"

JVM_ID=<coherence_cluster_name/node_member_name>

JAM_TARGET="jamagent.jamrun"

JAM_ARGS=""
JAM_ARGS="$JAM_ARGS jamconshost=<oms_host>"
JAM_ARGS="$JAM_ARGS jamconspport=<oms_port>"
JAM_ARGS="$JAM_ARGS jamjvmid=$JVM_ID"
JAM_ARGS="$JAM_ARGS jampool=<coherence_cluster_name>"

$JAVA_HOME/bin/java $COH_OPTS
-Dtangosol.coherence.management.extendedmbeanname=true
-Dcom.sun.management.jmxremote.authenticate=false
-Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.ssl=false
-Dtangosol.coherence.management=all
-Dtangosol.coherence.member=<unique member name>
-Doracle.coherence.machine=<hostname_as_discovered_in_EM>
-Dcom.sun.management.jmxremote.port=<OpenTCP_Port>
-Doracle.coherence.home=$COHERENCE_HOME
-Dtangosol.coherence.distributed.localstorage=false
-Dtangosol.coherence.management.refresh.expiry=1m
-Doracle.coherence.jamjvmid=$JVM_ID
$JAM_TARGET $JAM_ARGS
-server
-Xms2048m -Xmx2048m
oracle.sysman.integration.coherence.EMIntegrationServer
```

27.2.2 Example Start Script for All Other Nodes

An example start script for all other nodes is given below.

```
#!/bin/sh

JVM_ID=<coherence_cluster_name/node_member_name>

JAM_TARGET="jamagent.jamrun"

JAM_ARGS=""
JAM_ARGS="$JAM_ARGS jamconshost=<oms_host>"
JAM_ARGS="$JAM_ARGS jamconspport=<oms_port>"
```

```
JAM_ARGS="$JAM_ARGS jamjvmid=$JVM_ID"
JAM_ARGS="$JAM_ARGS jampool=<coherence_cluster_name>"

COH_OPTS="$COH_OPTS -cp $CP"
$JAVA_HOME/bin/java $COH_OPTS
-Dtangosol.coherence.management.extendedmbeanname=true
-Dtangosol.coherence.management.remote=true
-Dcom.sun.management.jmxremote.authenticate=false
-Dcom.sun.management.jmxremote.ssl=false
-Doracle.coherence.home=<coherence home>
-Dtangosol.coherence.member=<unique member name>
-Doracle.coherence.machine=<hostname_as_discovered_in_EM>
-Doracle.coherence.jamjvmid=$JVM_ID
$JAM_TARGET $JAM_ARGS
com.tangosol.net.DefaultCacheServer
```

27.3 Accessing JVM Diagnostics from Coherence Targets

If the Coherence nodes have been correctly configured for JVM Diagnostics, menu items for JVM Diagnostics will be available from each of the Oracle Coherence Node, Oracle Coherence Cache and Oracle Coherence Cluster targets.

27.3.1 Accessing JVM Diagnostics from Oracle Coherence Node Menu

From the Oracle Coherence Node Home page, select **JVM Diagnostics** from the **Oracle Coherence Node** menu. The drill down page for the JVM corresponding to the Coherence node is displayed.

27.3.2 Accessing JVM Diagnostics from Oracle Coherence Cache Menu

From the Oracle Coherence Cache Home page, select **JVM Diagnostics** from the **Oracle Coherence Cache** menu. The Java Workload Explorer page that shows a summary of JVMs for the nodes that supports the cache appears.

27.3.3 Accessing JVM Diagnostics from Oracle Coherence Cluster Menu

From the Oracle Coherence Cluster Home page, select **JVM Diagnostics** from the **Oracle Coherence Cluster** menu. The Java Workload Explorer page which will show a summary of JVMs for the nodes that supports the cache is displayed.

27.4 Including the JVM Diagnostics Regions in the Coherence Target Home Pages

If the Coherence cluster nodes have been configured with JVM Diagnostics, the JVM Diagnostics regions can be included in the Coherence cluster and node Home pages. For more information about adding these regions, see [Personalization](#).

Part IX

Using Identity Management

The chapters in this part provide a brief introduction to the Management Pack Plus for Identity Management. The chapters guide 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:

- [Getting Started with Oracle Identity Management](#)
- [Prerequisites for Discovering Oracle Identity Management Targets](#)
- [Discovering and Configuring Oracle Identity Management Targets](#)

Getting Started with Oracle Identity Management

This section explains the benefits and features of using Oracle Enterprise Manager to monitor Oracle Identity Management systems.

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.

To view a video about managing Oracle Identity Management, click [here](#).

28.1 Benefits of Using the Identity Management Pack

The benefits of using the Identity Management Pack include:

- Using a centralized systems management solution to efficiently manage multiple Oracle Identity Management deployments including testing, staging, and production environments from a single console
- Gaining 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
- Automating configuration management to accelerate problem resolution
- Recording synthetic Web transactions (or service tests) to monitor Identity Management Service availability and analyze end user response times
- Defining 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

28.2 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 11g and Identity Management 12c 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.

28.2.1 New Features for this Release

New features for Identify Management Pack include:

- Problem Analysis

Problem analysis is now available for IDM targets.
- Performance Page

This page shows the performance of the database corresponding to the Oracle Access Manager (OAM) Enterprise Manager target. Using this data, the OAM administrator can identify problems that causes performance bottlenecks.
- Configuration Compare Templates

Using a template, you can remove properties that typically signal "false positives" in comparisons by setting flags to ignore differences. When comparing hosts, for example, you know that host names will be different, so you can indicate to ignore differences on the name property value.
- 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.

28.3 Monitoring Oracle Identity Management Components in Enterprise Manager

You can use Enterprise Manager to monitor the following Identity Management 11g components ([Table 28-1](#)).

Table 28-1 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>

Table 28-1 (Cont.) Licensed Targets for Identity Management 11g Targets

Enterprise Manager Target Type	Purpose
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 monitored targets in the Identity Management pack associated with release 11g are summarized in [Table 28-2](#).

Table 28-2 Targets Associated with Identity Management 11g Targets

Enterprise Manager Target Type	Purpose
Generic Service	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.
Host	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.
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 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.

The monitored targets in the Identity Management pack associated with release 12c are summarized in this table.

Table 28-3 Targets Associated with Identity Management 11g and 12c Targets

Enterprise Manager Target Type	Purpose
Oracle Identity Federation Oracle Identity Manager Oracle Unified Directory Oracle Directory Integration Platform Oracle Internet 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 Identity Federations, Oracle Unified Directories, Oracle Directory Integration Platforms, 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

29

Prerequisites for Discovering Oracle Identity Management Targets

This section lists the system requirements and prerequisites needed to discover identity management targets.

29.1 System Requirements

Table 29-1 lists the supported Oracle Identity Management products in the Management Pack Plus for Identity Management in Enterprise Manager Cloud Control 13c.

Note: For the most up-to-date list of supported platforms, check My Oracle Support Certification Matrix on My Oracle Support (<http://support.oracle.com>).

Table 29-1 Supported Identity Management Products and Platforms in Enterprise Manager Cloud Control

Product	Application Server	Directory Server/Database
Oracle Access Manager	Oracle WebLogic Server	Oracle Internet Directory; Microsoft Active Directory
Oracle Access Manager	Oracle WebLogic Server	Oracle Database
Oracle Adaptive Access Manager	Oracle WebLogic Server	Oracle Database
Oracle Directory Integration Platform	Oracle WebLogic Server	Oracle Database
Oracle Directory Server Enterprise Edition	Not Applicable	Not Applicable
Oracle Identity Federation	Oracle WebLogic Server	Oracle Internet Directory
Oracle Identity Manager	Oracle WebLogic Server; Oracle SOA Suite	Oracle Database
Oracle Internet Directory	Oracle WebLogic Server	Oracle Database
Oracle Unified Directory	Not Applicable	Not Applicable
Oracle Virtual Directory	Oracle WebLogic Server	Not Applicable

Table 29-2 Supported Identity Management Products and Platforms in Enterprise Manager Cloud Control Release 3

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 Identity Federation	10.1.4.2; 10.1.4.3.0	Oracle Application Server	Oracle Internet Directory 10.1.4.x

Table 29-2 (Cont.) Supported Identity Management Products and Platforms in Enterprise Manager Cloud Control Release 3

Product	Version	Application Server	Directory Server/ Database
Oracle Identity Federation	11g PS1 (11.1.2.0); 11g PS2 (11.1.1.3.0); 11g PS2-11.1.1.2.0; 11g PS3-11.1.1.2.0; 11g PS4-11.1.1.2.0; 11g PS5-11.1.1.2.0	Oracle WebLogic Server 10.3	Oracle Internet Directory 11g PS 1 (11.1.1.2.0); 11g PS2 (11.1.1.3.0)
Oracle Identity Manager	9.1.0.1	Oracle WebLogic Server 10.3; JBoss Application Server	Oracle 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 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
Oracle Internet Directory	11g PS1-11.1.1.2.0; 11g PS2-11.1.1.3.0; 11g PS2-11.1.1.4.0; 11g PS4-11.1.1.5.0; 11g PS5-11.1.1.6.0	Oracle WebLogic Server 10.3	Oracle Database
Directory Integration Platform	11g PS1 (11.1.2.0); 11g PS2 (11.1.1.3.0); 11g PS1-11.1.1.2.0; 11g PS2-11.1.1.2.0; 11g PS3-11.1.1.2.0; 11g PS4-11.1.1.2.0; 11g PS5-11.1.1.2.0	Oracle WebLogic Server 10.3	Oracle Database
Oracle Virtual Directory	11g PS1 (11.1.2.0); 11g PS2 (11.1.1.3.0); 11g PS2-11.1.1.4.0; 11g PS4-11.1.1.5.0; 11g PS5-11.1.1.6.0	Oracle WebLogic Server 10.3	Not Applicable
Oracle Access Manager	11g (11.1.1.3.0); 11gR1-11.1.1.3.0; 11gPS1-11.1.1.5.0	Oracle WebLogic Server 10.3	Oracle Database
Oracle Adaptive Access Manager	11g (11.1.1.3.0); 11gR1-11.1.1.3.0; 11gPS1-11.1.1.5.0	Oracle WebLogic Server 10.3	Oracle Database

Table 29-2 (Cont.) Supported Identity Management Products and Platforms in Enterprise Manager Cloud Control Release 3

Product	Version	Application Server	Directory Server/ Database
Oracle Identity Manager	11g (11.1.1.3.0); 11gR1-11.1.1.3.0; 11gPS1-11.1.1.5.0	Oracle WebLogic Server 10.3; Oracle SOA Suite 11.1.1.3.0	Oracle Database
Oracle Directory Server Enterprise Edition	6.x; 7.x; 11g (11.1.1.3.0); 11.1.1.3.0; 11.1.1.5.0	Not Applicable	Not Applicable
Oracle Access Manager	12.1.2.3.0	Oracle WebLogic Server 10.3	Oracle Database
Oracle Identity Manager	12.1.2.3.0	Oracle WebLogic Server 10.3	Oracle Database
Oracle Internet Directory	12.1.2.3.0	Oracle WebLogic Server 10.3	Oracle Database
Oracle Directory Integration Platform	12.1.2.3.0	Oracle WebLogic Server 10.3	Oracle Database

29.2 Installing Oracle Enterprise Manager Cloud Control 13c

Before you begin configuring Cloud Control 13c to manage your Identity Management components, you must install and configure Cloud Control 13c 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 emHost1.example.com, then install and configure the Oracle Management Service and Oracle Management Repository on emHost2.example.com. Install the Cloud Control 13c Management Agent on every host that includes the components you want to manage with Cloud Control.

For more information, see *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.

29.3 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 13c Agent on each of the hosts that run Oracle Identity Management components.

If you would like to monitor additional targets, such as 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 13c 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 29-1](#).

Figure 29-1 Plug-Ins Deploy On Options

ORACLE Enterprise Manager Cloud Control 12c

Enterprise ▾ Targets ▾ Favorites ▾ History ▾

Plug-ins

This page displays the list of plug-ins available, downloaded and deployed in the Enterprise Manager environment. Plug-in lifecycle

Actions ▾ View ▾ Deploy On Undeploy From Check Updates Deployment Activities

Name	Management Servers...		Version	
	Management Agent...	able	Latest Downloaded	On Management Server
<ul style="list-style-type: none"> Applications <ul style="list-style-type: none"> Oracle Fusion Applications 12.1.0.3.0 Oracle Siebel 12.1.0.2.0 Oracle Database 12.1.0.2.0 [u120625] Oracle Fusion Middleware 12.1.0.3.0 Servers, Storage and Network <ul style="list-style-type: none"> Oracle Beacon 12.1.0.2.0 Oracle Chargeback and Capacity P 12.1.0.3.0 Oracle Exadata 12.1.0.3.0 Oracle MOS (My Oracle Support) 12.1.0.2.0 				

Oracle Fusion Middleware

30

Discovering and Configuring Oracle Identity Management Targets

This section provides the information needed to discover and configure Oracle Identity Management targets.

30.1 Discovering Identity Management Targets

This section describes how to discover Identity Management targets.

30.1.1 Discovering Identity Management 11g and 12c

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) and 12c targets. The Discovery wizard collects details about Oracle Identity Management 11g and 12c targets including information about the host, WebLogic User Name/Password, and other details.

 **Note:**

Before discovering the targets associated with Oracle Access Manager 11g, and 12c download and install patch 10094106.

To discover Oracle Identity Management 11g and 12c (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 and targets 12c.

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.
Password	WebLogic domain password.

Field	Description
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	A signal to discover the servers that are down.
JMX Service URI	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.

4. A list of all the Identity Management targets is displayed. 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.
5. 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**.

30.1.2 Discovering Oracle Directory Server Enterprise Edition 11g and 12c

To discover Oracle Directory Server Enterprise Edition 11g and 12c targets, perform the following steps:

1. Log in to Enterprise Manager. Select **Targets**, then select **Middleware**.
2. From the **Add** menu, select **Oracle Directory Server Enterprise Edition**.
3. Enter the information requested.
 - a. Oracle Directory Server Enterprise Edition Registry Host: Host of the Directory Server Control Center Registry.
 - b. Oracle Directory Server Enterprise Edition Registry Port: Port of the Directory Server Control Center Registry.
 - c. Directory Server User Name - for example CN=Directory Manager.
 - d. 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.

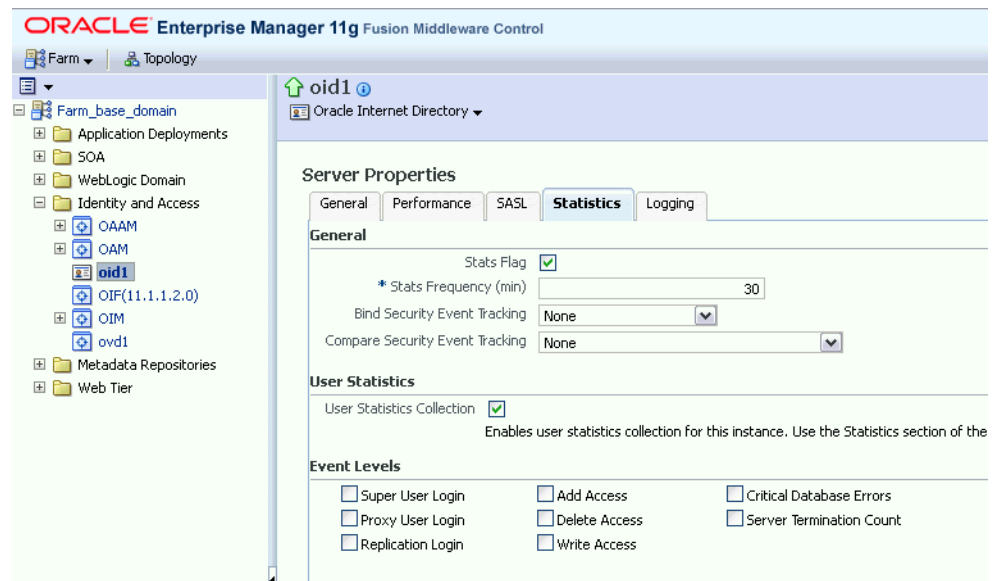
30.2 Collecting User Statistics for Oracle Internet Directory

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:

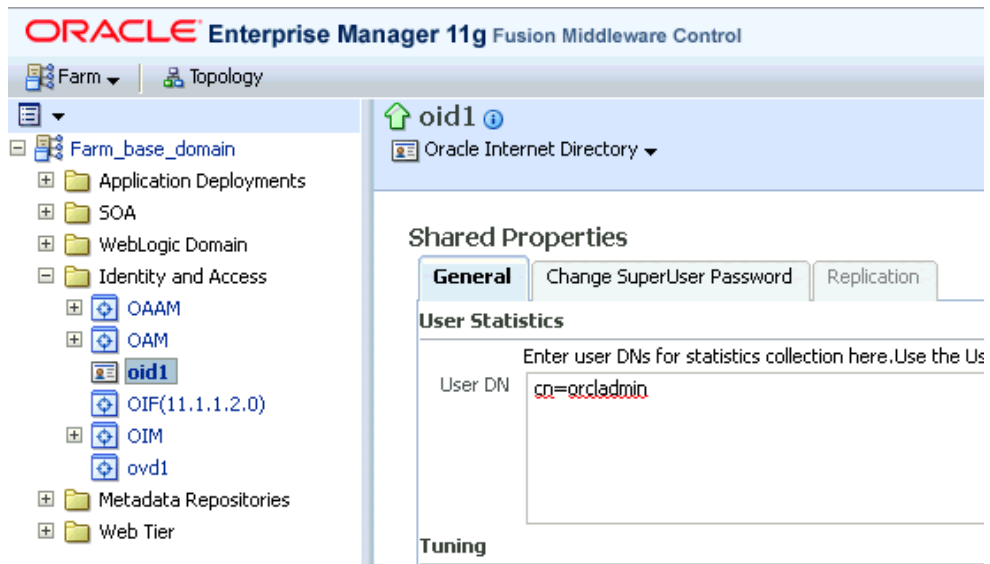
1. From the **Targets** menu, select **Middleware**. From the **Middleware Features** menu, select **Identity and Access**.
2. Select the discovered Oracle Internet Directory target.
3. From the **Oracle Internet Directory** menu, select **Fusion Middleware Control**.
4. From the **Targets** menu in 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 30-1](#).

Figure 30-1 Server Properties - Statistics Tab



5. 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 30-2](#).

Figure 30-2 Shared Properties - General Tab



30.3 Creating Identity Management Elements

This section describes how to create Identity Management elements.

30.3.1 Creating Identity and Access System Target

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.

30.3.2 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.

30.3.3 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 after 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 X

Discovering and Monitoring Non-Oracle Middleware

The chapters in this part describe how to discover and monitor non-Oracle middleware components.

The chapters are:

- [Discovering and Monitoring IBM WebSphere MQ](#)
- [Discovering and Monitoring IBM WebSphere Application Servers, Clusters, and Cells](#)
- [Discovering and Monitoring JBoss Application Server](#)
- [Discovering and Monitoring Apache HTTP Server](#)

31

Discovering and Monitoring IBM WebSphere MQ

This section describes how you can discover and monitor IBM WebSphere MQ targets in Enterprise Manager Cloud Control.

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 are supported by a command line tool, a user interface, and programmable command format messages.

IBM WebSphere MQ enables administrators to derive instant value, while giving them the flexibility to fine-tune thresholds according to their specific operational requirements.

Note:

The support for discovering and monitoring IBM WebSphere MQ targets is offered via the System Monitoring Plug-ins for Non-Oracle Middleware.

IBM WebSphere MQ V6 is only supported with Agent 12.1.0.5. For more information, refer to the MOS certifications page to know about the relevant middleware software version.

The following topics are discussed in this document:

- [Introduction](#)
- [Prerequisites](#)
- [Understanding Discovery](#)
- [Monitoring](#)

31.1 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](#)
- [Enhance Service Modeling and Perform Comprehensive Root Cause Analysis](#)

31.1.1 Out-of-Box Availability and Performance Monitoring

You can see immediate value through out-of-box availability and performance monitoring. Some of the key areas of more than 60 performance indicators monitored

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 monitoring of IBM WebSphere MQ targets includes various out-of-box reports, summarizing key information about availability and performance. Some of the 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. For more information, see *Using Notifications in the Enterprise Manager Cloud Administrator guide*.
- **Groups/Systems**
Simplify management of large numbers of components, allowing administrators to *manage many-as-one*.

31.1.2 Centralized Monitoring of all Information in a Single Console

IBM WebSphere MQ provides a consolidated view of the entire enterprise, enabling administrators to monitor and manage all of their components from a central place.

Having such an integrated console 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.

31.1.3 Enhance Service Modeling and Perform Comprehensive Root Cause Analysis

Enterprise Manager Cloud Control'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 monitoring of IBM WebSphere MQ targets in Enterprise Manager Cloud Control enables IT organizations running applications on top of Oracle and IBM to derive greater value from the Service Level Management features 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.

31.2 Prerequisites

This section covers the following:

- [Basic Prerequisites](#)
- [JAR File Requirements \(for Local Monitoring and Remote Monitoring\)](#)

31.2.1 Basic Prerequisites

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.SVRCONN channel must be available

Oracle Management Agent (Management Agent) must be up and running, either locally or remotely, and must be able to upload data successfully to Oracle Management Repository. 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 Management Agent user or a user part of the same group (primary gid).

31.2.2 JAR File Requirements (for Local Monitoring and Remote Monitoring)

For local monitoring, the Management Agent OS user should have read privileges over the following JAR files, which are needed for the discovery of the target IBM WebSphere MQ.

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 which is accessible to the Management Agent, and the OS user starting this Management Agent must have *read* privileges on these files. For example, create a directory / new_dir/sysman/mq_jar_files and copy the JAR files into this directory.

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

- `$MQ_HOME/java/lib/com.ibm.mq.jar`
- `$MQ_HOME/java/lib/com.ibm.mq.jmqi.jar`
- `$MQ_HOME/java/lib/com.ibm.mq.commonservices.jar`
- `$MQ_HOME/java/lib/com.ibm.mq.headers.jar`
- `$MQ_HOME/java/lib/com.ibm.mq.pcf.jar`
- `$MQ_HOME/java/lib/connector.jar`

31.3 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 include 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.

The topics covered under this section are:

- [Discovery Prerequisites for Local Agent](#)
- [Discovery Prerequisites for Remote Agent](#)
- [Queue Manager Cluster Discovery](#)
- [Standalone Queue Manager Discovery](#)

31.3.1 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.SVRCONN` channel should also be available.

31.3.2 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.SVRCONN` channel should also be available.

31.3.3 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.

The 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 Enterprise Manager Cloud Control, you cannot explicitly remove any member queue managers from that cluster.

To manually add targets, complete the following:

1. From the **Setup** menu, select **Add Target**, then select **Add Target Manually**.
2. Select **Add Targets Using Guided Process (Also Adds Related Targets)**.
3. From the **Target Types** list, select **IBM WebSphere MQ Queue Manager**, 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. Server connection channel to be used for monitoring.
 - d. Jar path location (location must be accessible to the OEM agent).

Mention the full path of all the required individual JARs in the **Jar Path** field as shown below:

```
$MQ_HOME/java/lib/com.ibm.mq.jar:$MQ_HOME/java/lib/  
com.ibm.mq.jmqi.jar:$MQ_HOME/java/lib/  
com.ibm.mq.commonservices.jar:$MQ_HOME/java/lib/  
com.ibm.mq.headers.jar:$MQ_HOME/java/lib/  
com.ibm.mq.pcf.jar:$MQ_HOME/java/lib/connector.jar
```

Note:

For local monitoring replace \$MQ_HOME in the above classpath with the actual IBM WebSphere MQ home directory path, and for remote monitoring, replace \$MQ_HOME with the path where all the JAR files are stored. In case of a Windows agent, replace the ':' (colons) in the above classpath with ';' (semi-colons).

- e. Provide the agent host details used to monitor IBM WebSphere MQ.
- f. MQ administration credentials.

- 5. Select the cluster you want to discover, and click **Next**.

The following image shows a Queue Manager cluster:


Target Name	Type	MQ Listener IP Address/Host	Port	Target Exists
Cluster 1				No
QMGR4	Queue Manager	141.111.111.111	1417	No
Qmgr3	Queue Manager	141.111.111.111	1416	No

- 6. Click **OK** to finish discovery.

31.3.4 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. For more information, see [Queue Manager Cluster Discovery](#). The

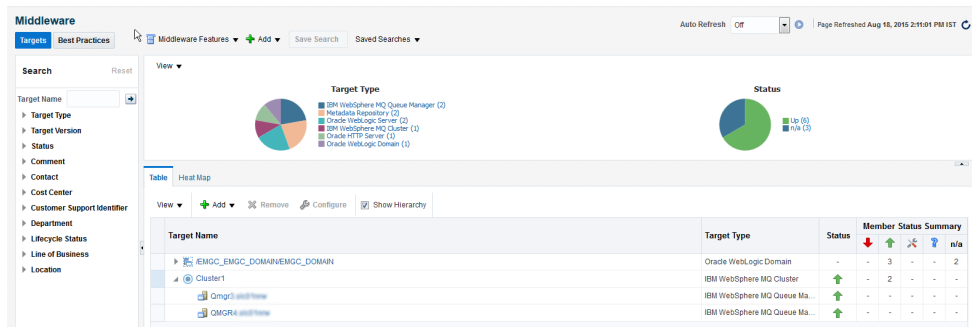
only difference is that you select a standalone queue manager from the list of possible targets on the Add IBM WebSphere MQ Targets: Available Targets page.



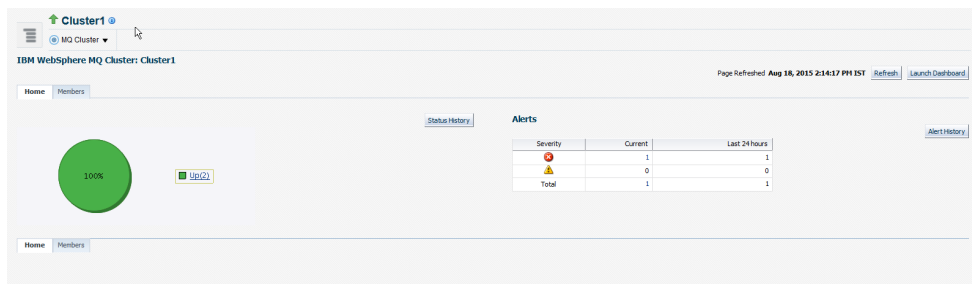
31.4 Monitoring

The following illustrates some of the different methods used to monitor the performance of the IBM WebSphere MQ targets:

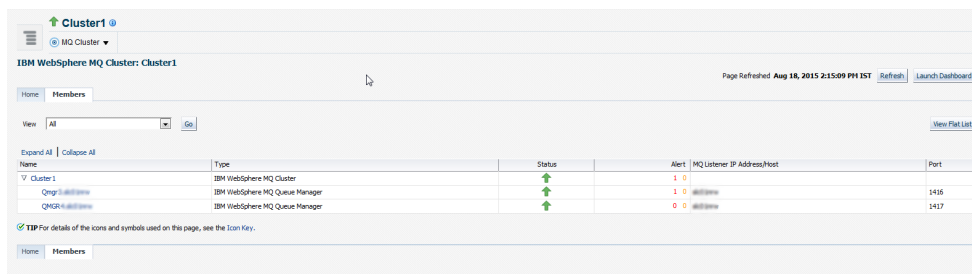
- Middleware page displaying the discovered IBM MQ clusters and queue-managers.



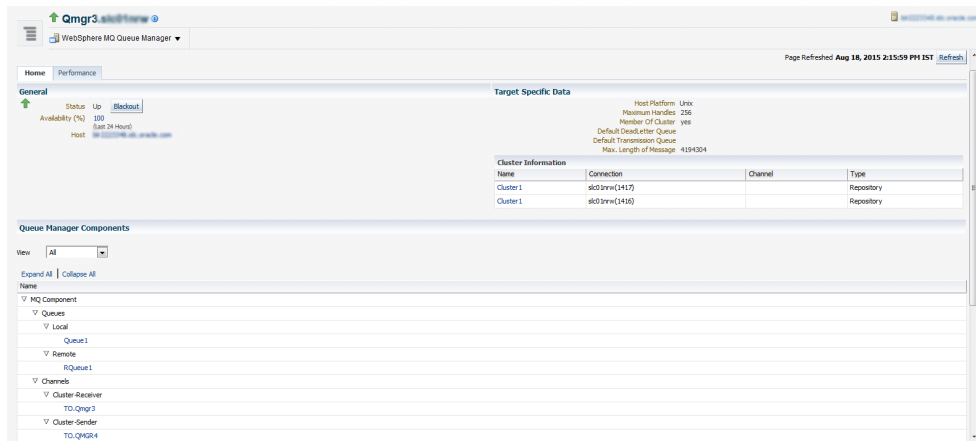
- Cluster page displaying information about a cluster.



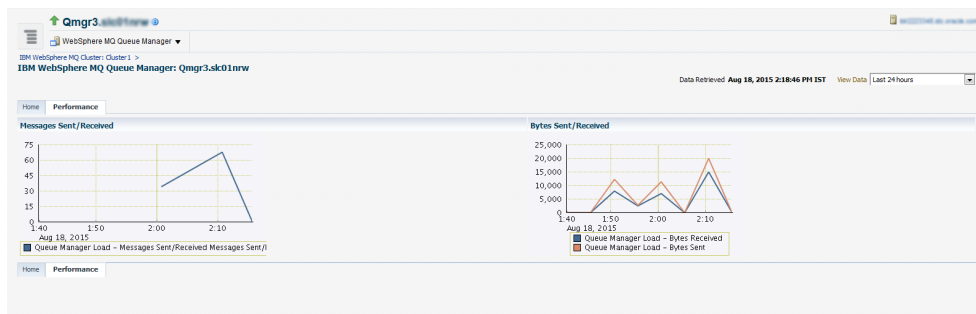
- Cluster Member's page displaying information about members of the cluster.



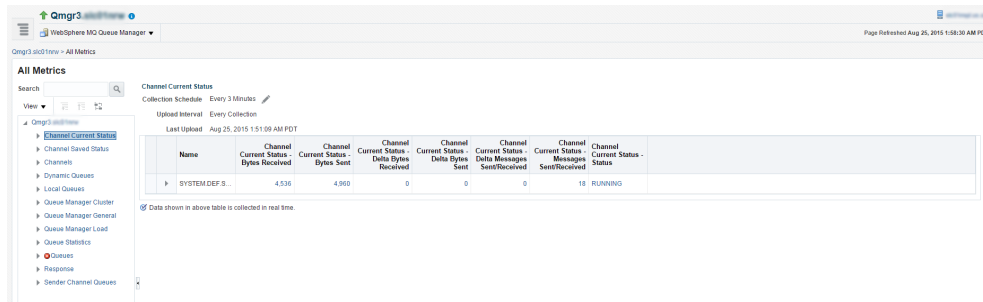
- Queue Manager page displaying information about queues and channels.




- 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.




- Performance Summary page displays the overall performance of IBM WebSphere MQ Metrics.

You can use the IBM WebSphere MQ Metric Performance Summary page to do the following:

- Monitor the Metric performance for preferred metrics

Using the Performance Summary page, you can monitor the overall performance of the IBM WAS MQ metrics.



- Select a preferred chart set

From the **Chart Set** list, select the preferred chart set.

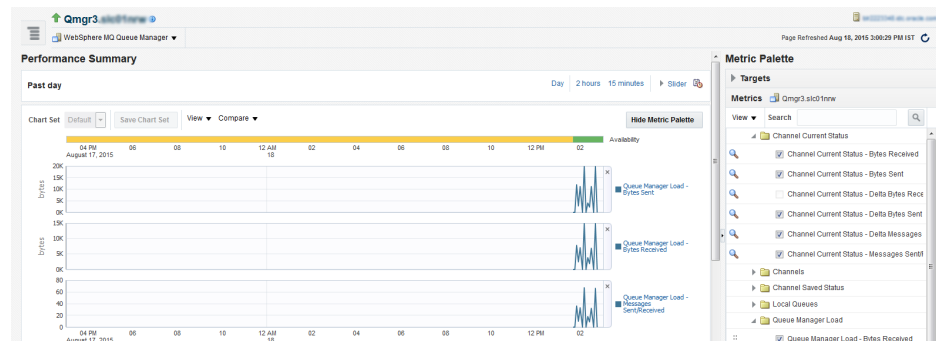
- Update an existing chart set

You can customize a chart set according to your requirements, and save changes by clicking **Save Charts**. You can create a new chart set by using the Save Charts option.

- Change time frames

You can use the slider to set the time frame manually, or you can select from the default values provided.

- Show/Hide the Metrics Palette



- Delete metric performance charts
You can delete metric performance charts either by clicking the close button on the chart itself, or by deselecting the metric name on the Metric Palette.
- Create new metric performance charts
You can create new metric performance charts by selecting the preferred metrics. The charts are automatically created once the metrics are selected from the Metric Palette.
- Drag and drop metrics
You can drag and drop the metrics from a particular metric group to the same chart.
- View individual metrics on a chart
When you hover the cursor over a particular metric on the chart, the other metrics are greyed out.

Discovering and Monitoring IBM WebSphere Application Servers, Clusters, and Cells

This section describes how you can discover and monitor IBM WebSphere application servers, clusters, and cells in Enterprise Manager Cloud Control.

IBM WebSphere Application Server is an application server developed and maintained by IBM Corporation. It offers options for a faster, more flexible Java application server runtime environment with enhanced reliability and resiliency. It supports single server environments and medium-sized configurations, as well as dynamic web applications requiring web tier clustering over multiple application server instances.

IBM WebSphere Application Server Cell is a high-level, logical grouping of IBM WebSphere Application Server Clusters within your enterprise configuration. Each IBM WebSphere Application Server Cluster is a composite target, or in other words, a logical target, comprising one or more individual IBM WebSphere Application Servers.

Enterprise Manager Cloud Control enables you to discover IBM WebSphere Application Servers, IBM WebSphere Application Server Clusters, and IBM WebSphere Application Server Cells in your environment, and add them for central monitoring and management.

This chapter describes how you can discover and monitor these IBM WebSphere Application Server targets in Enterprise Manager Cloud Control. In particular, this chapter covers the following:

- [About Managing IBM WebSphere Application Servers, Clusters, and Cells](#)
- [Supported Versions for Discovery and Monitoring](#)
- [Prerequisites for Discovering IBM WebSphere Application Servers, Clusters, and Cells](#)
- [Discovering IBM WebSphere Application Servers, Clusters, and Cells](#)
- [Monitoring IBM WebSphere Application Servers](#)
- [Monitoring IBM WebSphere Application Server Clusters](#)
- [Monitoring IBM WebSphere Application Server Cells](#)
- [Troubleshooting IBM WebSphere Application Server Discovery and Monitoring Issues](#)

32.1 About Managing IBM WebSphere Application Servers, Clusters, and Cells

Using Enterprise Manager Cloud Control, you can do the following with IBM WebSphere Application Server targets:

- Discover the following for central monitoring and management:

- IBM WebSphere Application Servers
- IBM WebSphere Application Server Clusters

When you discover an IBM WebSphere Application Server that is part of an IBM WebSphere Application Server Cluster, the IBM WebSphere Application Server Cluster and all other IBM WebSphere Application Servers that are part of that cluster get automatically discovered and added to Enterprise Manager Cloud Control.

- IBM WebSphere Application Server Cells

When you discover an IBM WebSphere Application Server that is part of an IBM WebSphere Application Server Cell, the IBM WebSphere Application Server Cell and all other IBM WebSphere Application Server Clusters and IBM WebSphere Application Servers that are part of that cell get automatically discovered and added to Enterprise Manager Cloud Control.

- Monitor the status, the availability percentage, the CPU usage, the heap usage, and so on.
- View a summary of incidents and problems occurred for a given interval.
- Monitor the status and the overall health of the member application servers that are part of IBM WebSphere Application Server Clusters and IBM WebSphere Application Server Cells.
- Diagnose, notify, and correct performance and availability problems with the help of GUI-rich, intuitive graphs and illustrations.
- Monitor the status of deployed applications.
- Monitor the status of most requested Servlets, EJBs, and JSPs in the last 24 hours.
- Create or end blackouts to suspend or resume the collection of metric data, respectively.
- Monitor and manage configuration details that were last collected and also the ones that were saved at a given point of time.
- Compare the configuration between:
 - A last collected configuration of a server instance with a saved configuration of the same server instance or a different server instance.
 - A last collected configuration of a server instance with a last collected configuration of a different server instance.
 - A saved configuration of a server instance with another saved configuration of the same server instance or a different server instance.
 - A saved configuration of a server instance with a last collected configuration of the same server instance or a different server instance.
- View compliance-related information, such as the compliance standards and frameworks associated with the server, the real-time observations, the evaluation results, and so on.
- View a list of metrics, their collection interval, and the last upload for each metric.

32.2 Supported Versions for Discovery and Monitoring

To search for the IBM WebSphere Application Server versions that are supported for discovery and monitoring in Enterprise Manager Cloud Control, follow these steps:

1. Log into <https://support.oracle.com/>.
2. On the My Oracle Support home page, select **Certifications** tab.
3. On the Certifications page, enter the following search criteria in the Certification Search section.
 - Enter the product name **Enterprise Manager Base Platform - OMS** in the Product field.
 - Select the applicable Enterprise Manager Cloud Control release number from the Release list.
4. Click **Search**.
5. In the Certification Results section, expand the **Application Server** menu to view the certified IBM WebSphere Application Server versions.

Certified With	Number of Releases / Versions
> Operating Systems (8 Items)	
> Agents (1 Item)	
▼ Application Servers (10 Items)	
Apache Tomcat (Managed Target)	2 Releases (17, 4.0.5, 5.5.7, 5.5.9, 5.5.10)
IBM WebSphere Application Server (Managed Target)	4 Releases (8.5, 7.5, 8.5, 8.5.5)
JBoss Application Server (Managed Target)	4 Releases (5.7, 5.0, 5.1, 4.2.7, 4.8.7)
Oracle Application Server (Managed Target)	2 Releases (10.1.2.0.0, 10.1.2.0.1, 10.1.2.0.2)
Oracle Application Server Single Signon (Managed Target)	1 Release (10.1.4.1.0)
Oracle GlassFish Server (Managed Target)	2 Releases (10.1.2.0.0, 10.1.2.0.1)
Oracle WebLogic Portal (Managed Target)	4 Releases (10.1.2.0.0, 10.1.2.0.1, 10.1.2.0.2, 10.1.2.0.3)
Oracle WebLogic Server (Infrastructure)	2 Releases (10.1.2.0.0, 10.1.2.0.1)
Oracle WebLogic Server (Managed Target)	7 Releases (10.1.2.0.0, 10.1.2.0.1, 10.1.2.0.2, 10.1.2.0.3, 10.1.2.0.4, 10.1.2.0.5, 10.1.2.0.6)
WebLogic Server (Managed Target)	7 Releases (10.1.2.0.0, 10.1.2.0.1, 10.1.2.0.2, 10.1.2.0.3, 10.1.2.0.4, 10.1.2.0.5, 10.1.2.0.6)
> Databases (8 Items)	
> Desktop Applications, Browsers and Clients (5 Items)	
> Directory/LDAP Services (6 Items)	
> Enterprise Applications (150 Items)	
> Management and Development Tools (49 Items)	
> Middleware (28 Items)	
> Other (3 Items)	
> Server (2 Items)	
> Virtualization Software (1 Item)	

32.3 Prerequisites for Discovering IBM WebSphere Application Servers, Clusters, and Cells

Meet the following prerequisites for discovering IBM WebSphere Application Server, IBM WebSphere Application Server Clusters, and IBM WebSphere Application Server Cells.

- Ensure that the Deployment Manager is running.
- For standalone IBM WebSphere Application Servers, ensure that the particular Server is running.
- Ensure that the SOAP connector port of the IBM WebSphere Application Server or the Deployment Manager is open to the Management Agent host.

To find the SOAP connector ports, search for the keyword `SOAP_CONNECTOR_ADDRESS` in the following location:

```
$<WEBSPPHERE_HOME>/AppServer/profiles/<PROFILE>/config/cells/<cellname>/nodes/
<nodename>/serverindex.xml
```

- Ensure that the PMI Service is enabled. To do so, follow these steps:
 - For IBM WebSphere Application Server 7.x, and 8.x:
 1. Log in to the Integrated Solutions Console.
 2. From the **Monitoring and Tuning** menu, select **Performance Monitoring Infrastructure (PMI)**.
 3. Select the application server instance.
 4. From the **Configuration** tab, under **General Properties**, enable PMI by select the check box for **Enable Performance Monitoring Infrastructure (PMI)**.
 5. From **Currently Monitored Static Set**, select **Custom**. Click the **Custom** link, and specify the list of metrics that are to be enabled. Click **OK**.
 6. Click **Save**, and restart the server.

 **Note:**

For a clustered configuration, enable PMI for each server individually.

- Ensure that when Administrative Security is enabled with the absolute path, a Java Trust Keystore is provided during the discovery.
- For local monitoring, you must have *read* privileges over the following IBM WebSphere directories and JAR files:
 - For IBM WebSphere 7.0


```
$<WEBSPPHERE_HOME>/runtimes/
com.ibm.ws.admin.client_7.0.0.jar
```
 - For IBM WebSphere 8.0


```
$<WEBSPPHERE_HOME>/runtimes/
com.ibm.ws.admin.client_8.0.x.jar
```
 - For IBM WebSphere 8.5


```
$<WEBSPPHERE_HOME>/runtimes/
com.ibm.ws.admin.client_8.5.x.jar
```
- For remote monitoring, you must copy the required WebSphere JARs and the Trusted Keystore file to a folder on the remote Management Agent.
 - For IBM WebSphere 7.0.x:
 1. Create a dummy WebSphere home directory on the remote Agent host; for example `/scratch/WebSphere7Jars/AppServer` and under it, create the following directory structure:


```
WAS_HOME
  /trustedKeyStore
  /runtimes
```

```
/plugins
```

```
/java/jre/lib/ext
```

2. Copy the jar files listed below from the WebSphere host to the remote Agent host (in the similar locations of the actual WAS_HOME):

```
WAS_HOME/runtimes
```

```
com.ibm.ws.admin.client_7.0.0.jar
```

```
WAS_HOME/java/jre/lib/ext
```

```
ibmkeycert.jar
```

```
WAS_HOME/java/jre/lib
```

```
ibmjgssprovider.jar
```

- For IBM WebSphere 8.0.x:

1. Create a dummy WebSphere home directory on the remote Agent host; for example `/scratch/WebSphere8Jars/AppServer` and under it, create the following directory structure:

```
WAS_HOME
```

```
/trustedKeyStore
```

```
/runtimes
```

```
/plugins
```

```
/java/jre/lib/ext
```

2. Copy the jar files listed below from the WebSphere host to the remote Agent host (in the similar locations of the actual WAS_HOME):

```
WAS_HOME
```

```
/runtimes
```

```
com.ibm.ws.admin.client_8.0.0.jar
```

```
WAS_HOME/java/jre/lib/ext
```

```
ibmkeycert.jar
```

```
WAS_HOME/java/jre/lib
```

```
ibmjgssprovider.jar
```

```
ibmorb.jar
```

- For IBM WebSphere 8.5.x:

1. Create a dummy WebSphere home directory on the remote Agent host; for example `/scratch/WebSphere8Jars/AppServer` and under it, create the following directory structure:

```
WAS_HOME
```

```
/trustedKeyStore
```

```
/runtimes
```

```
/plugins
```

```
/java/jre/lib/ext
```

2. Copy the jar files listed below from the WebSphere host to the remote Agent host (in the similar locations of the actual WAS_HOME):

```

WAS_HOME
  /runtimes
    com.ibm.ws.admin.client_8.5.0.jar
  WAS_HOME/java/jre/lib/ext
    ibmkeycert.jar
  WAS_HOME/java/jre/lib
    ibmjgssprovider.jar
    ibmorb.jar

```

- Create MBean resources for applications.

The management data for the OEM metrics that relate to various application components such as EJBs, Web Modules, Servlets, and so on. is available only if the MBean resources are created for each of the application deployed on the Server.

During the deployment of the application on the server, the WebSphere Application Deployment UI allows user to deselect this feature. By default, these mbeans are created but user can change the behavior by deselecting this option. Hence, it should be made sure that this is enabled for each of the application deployed on each server before using OEM to monitor the corresponding server and application.

The procedure to enable this feature is as follows:

1. For each server in the cell, open the Installed Applications page.
2. Click on the application.
3. Click on the **Configuration** tab and select **Start up behavior** link.
4. Enable **Create MBean Resources** check box and click **Apply**.
5. Save the settings and restart the server.

This procedure should be followed for each of the application deployed on every server to be monitored.

- If the websphere cell or server is SSL enabled, perform the following steps before discovering the target.

1. Export the trusted signer certificate of the server and import it in to the AgentTrust.jks file.

To export the trusted signer certificate, perform the following steps on the host where the target is installed:

- a. Navigate to location of trust store file of the WebSphere deployment manager/server.
- b. Run the following command:

```

keytool -export -v -alias <alias_name_cert> -file <output_file> -
keystore <Trust Store File Name> -storepass <Password for the trust>

```

For example,


```
keytool -export -v -alias default_signer -file rootca04.cert -keystore
<WAS_HOME>/profiles/Dmgr04/etc/DummyClientTrustFile.jks -storepass
WebAS
```

To import the trusted signer certificate in to the `AgentTrust.jks` file, run the following command on the host where the agent is installed:

```
keytool -import -v -trustcacerts -alias <unique_alias_name> -file
<certificate_location> -keystore <AgentTrust.jks File> -storepass
<password for the agent trust>
```

For example,

```
keytool -import -v -trustcacerts -alias signer_02 -file <WAS_HOME>/
AppServer/profiles/Dmgr02/etc/rootca02.cert -keystore AgentTrust.jks -
storepass welcome
```

2. Confirm that all certificates have been added in `AgentTrust.jks` file, by running the command:

```
keytool -v -list -keystore <AgentTrust.jks Filename >
```

Troubleshooting Tip: If you encounter the following error when entering the password, it indicates that more than one keytools exists in the box:

```
keytool error: gnu.javax.crypto.keyring.MalformedKeyringException:
incorrect magic
```

If there are more than one keytools existing in the `AgentTrust.jks` file, perform the following steps:

- a. Run the command `locate keytool` to find the relevant keytool to be used.
- b. Use the full path of the keytool to run the keytool command. For example:

```
/scratch/agents/slc02ktx/agent_13.1.0.0/oracle_common/jdk/bin/keytool
-list -keystore <AgentTrust.jks Filename>
```

3. Restart the agent.

During target discovery, on the discovery host page, provide the location of the `AgentTrust.jks` file in the **Agent Trust Filename** field. The input parameter is not needed if cell/server is not SSL enabled.

32.4 Discovering IBM WebSphere Application Servers, Clusters, and Cells

Enterprise Manager Cloud Control enables you to discover and add IBM WebSphere Application Servers (and their associated clusters and cells) for central monitoring and management.

To add an IBM WebSphere Application Server (and their associated clusters and cells), follow these steps:

1. Meet the prerequisites. For details, see [Prerequisites for Discovering IBM WebSphere Application Servers, Clusters, and Cells](#).
2. From the **Targets** menu, select **Middleware**.
3. On the Middleware page, from the **Add** menu, select **IBM WebSphere Application Server**.

4. Click **Go**.

Enterprise Manager Cloud Control displays the IBM WebSphere Application Server Discovery Wizard.

5. On the Host page, enter the details of the host on which the IBM WebSphere Application Server and Oracle Management Agent are running. In case of IBM WebSphere Application Server Cell-based installation, enter the details of the Deployment Manager so that all IBM WebSphere Application Servers present under the cell are automatically discovered.

Figure 32-1 IBM WebSphere Application Server Host Page

The screenshot shows the 'Host' page of the 'IBM WebSphere Application Server Discovery' wizard. The page title is 'IBM WebSphere Application Server Discovery: Host'. Below the title, there is a progress bar with 'Host' selected. The main content area contains a form with the following fields:

- WebSphere Application Server Host:** ip@hostname.oracle.com
- SOAP Connector Port:** 8880
- Version:** 8.5.x
- User Name:** was
- Password:** masked with three dots
- Agent Trust Filename:** <Agent_HOME>/stage/sysman/config/montrust/AgentTrust.jks
- Server Home Directory:** ip@hostname.oracle.com/oracle/oh4/IBM/WebSphere
- Agent:** ip@hostname.oracle.com:1183

Link	Description
WebSphere Application Server Host	Enter the name of the host on which the IBM WebSphere Application Server or the IBM WebSphere Deployment Manager is installed.
SOAP Connector Port	Enter the SOAP connector port on which the IBM WebSphere Application Server or the IBM WebSphere Deployment Manager is listening.
Version	Select the version of the IBM WebSphere Application Server.
User Name	Enter the user name to access the IBM WebSphere Application Server or the IBM WebSphere Deployment Manager.
Password	Enter the password to access the IBM WebSphere Application Server or the IBM WebSphere Deployment Manager.
Agent Trust Filename	If the port is SSL enabled, then enter the absolute path to the agent trust file. Agent trust file is a protected database that holds keys and certificates for an enterprise. Ensure that the path leads up to the file name. For example, <Agent_home>/stage/sysman/config/montrust/AgentTrust.jks
Server Home Directory	Enter the absolute path to the Oracle home directory where the IBM WebSphere Application Server or the IBM WebSphere Deployment Manager is installed. For example, /net/host1/software/IBM/WebSphere/AppServer/

Link	Description
Agent	Click the search icon and select the Oracle Management Agent (Management Agent) that is monitoring the IBM WebSphere Application Server or the IBM WebSphere Deployment Manager. The Management Agent can be local or remote to the IBM WebSphere Application Server or the IBM WebSphere Deployment Manager.

6. On the Select Servers page, select the IBM WebSphere Application Servers and/or the IBM WebSphere Application Server Clusters that you want to monitor in Enterprise Manager Cloud Control.

On selection of an IBM WebSphere Application Server Cluster, all the IBM WebSphere Application Servers that are part of the cluster are automatically selected and added to Enterprise Manager Cloud Control for monitoring.

7. On the Review page, review the information you have provided in the previous screens for discovering IBM WebSphere Application Servers and IBM WebSphere Application Server Cells.

If you want to modify any information, click **Back** repeatedly to reach the page where you want to make some changes. If you are satisfied with the information, click **Submit**.

32.5 Monitoring IBM WebSphere Application Servers

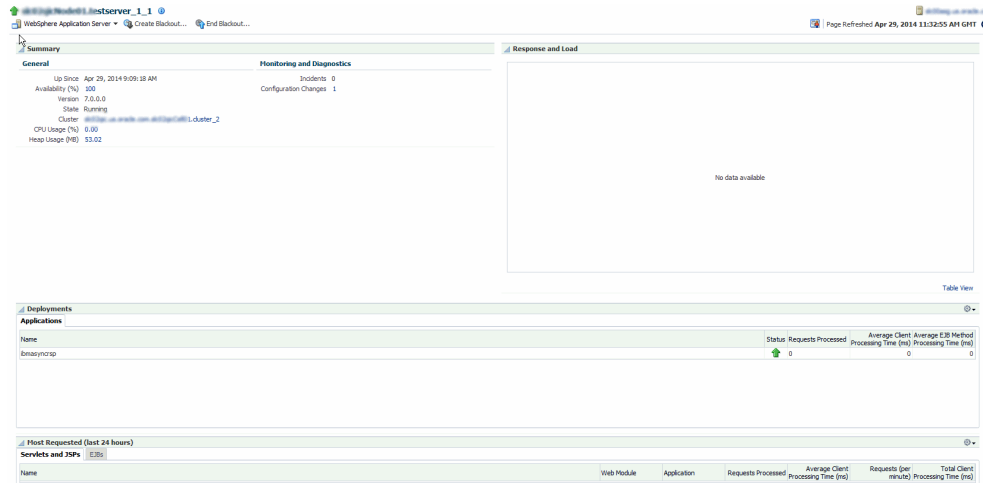
This section covers the following:

- [Monitoring IBM WebSphere Application Servers](#)
- [Administering IBM WebSphere Application Servers](#)
- [Monitoring the Performance of IBM WebSphere Application Servers](#)
- [Monitoring the Applications Deployed to IBM WebSphere Application Servers](#)
- [Viewing the Top EJBs of IBM WebSphere Application Servers](#)
- [Viewing the Top Servlets and JSPs of IBM WebSphere Application Servers](#)
- [Viewing IBM WebSphere Application Server Metrics](#)

32.5.1 Monitoring IBM WebSphere Application Servers

To monitor IBM WebSphere Application Servers, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired **IBM WebSphere Application Server**.
3. On the IBM WebSphere Application Server Home page, you can monitor the availability, usage, and performance of the selected server at a high level.



The IBM WebSphere Application Server Home page has the following sections:

32.5.1.1 General Section

Provides general information about the health of the server.

Element	Description
Up/Down/Pending Since	Date and time when the status was last determined.
Availability (%)	Percentage of time that the server was up during the last 24 hours. Click the percentage link to view availability details for the past 24 hours.
Version	Version of the server that is being monitored.
State	Current state of the server, whether it is running or shut down.
CPU Usage (%)	CPU consumption as a percentage of CPU time at any given moment in time. Click the percentage link to view availability details for the past 24 hours.
Heap Usage (MB)	Current JVM Memory heap Usage in MB as per the last Metric collection.

32.5.1.2 Monitoring and Diagnostics Section

Provides a summary of incidents and configuration changes made to the server. Use this information to diagnose and troubleshoot performance issues with the server.

Element	Description
Incidents	Number of unresolved issues that require your attention and corrective action. Click the value to drill down and view more detailed information.
Configuration Changes	Number of incidents related to the applications. The displayed integer is also a link to the Incident Manager page. Click the value to drill down and view more detailed information.

32.5.1.3 Response and Load Section

Provides a graphical representation of the server's performance, measuring request-processing time for a given interval. To switch to a tabular format, click **Table View**. To drill down and view more detailed metric-related information and to diagnose issues by looking at other related infrastructure metrics, click the server names in the legend and select an appropriate option in the Additional Information message.

32.5.1.4 Applications Tab

Provides critical information about the applications deployed to the server. For more details, see [Monitoring the Applications Deployed to IBM WebSphere Application Servers](#).

32.5.1.5 Servlets and JSPs Tab

Provides details of the most requested Servlets and JSPs in the last 24 hours.

32.5.1.6 EJBs Tab

Provides details of the most requested EJBs in the last 24 hours.

32.5.2 Administering IBM WebSphere Application Servers

To administer IBM WebSphere Application Servers, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired IBM WebSphere Application Server.
3. On the IBM WebSphere Application Server Home page, you can view high-level information pertaining to the selected server.

To perform administrative tasks on the IBM WebSphere Application Server, from the **WebSphere Application Server** menu, select any of the following according to your needs:

- **Monitoring**, to monitor the performance of the target, view metric details, view status information, view incidents and alerts raised so far for the target, and view blackouts created for the target.
- **Diagnostics**, to analyze and diagnose performance issues.
- **Control**, to create or end blackouts.
- **Job Activity**, to view details of the jobs created for the target.
- **Information Publisher Reports**, to view reports.
- **Administer**, to directly administer the IBM WebSphere Application Server using the IBM WebSphere Application Server Console.
- **Configuration**, to search, view, and compare configuration details.
- **Compliance**, to view and create compliance standards.
- **Target Setup**, to view monitoring configuration details and target properties, to remove the target or add it to a group, to view the properties of the target.

- **Target Sitemap**, to view the overall topology of the target.
- **Target Information**, to view general information about the target.

32.5.3 Monitoring the Performance of IBM WebSphere Application Servers

Enterprise Manager Cloud Control provides several key performance charts that can help you quickly assess the health of your IBM WebSphere Application Server.

To check the performance of an IBM WebSphere Application Server, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired **IBM WebSphere Application Server**.
3. On the IBM WebSphere Application Server Home page, from the **WebSphere Application Server** menu, select **Monitoring**, then select **Performance Summary**.
4. On the Performance Summary page, you can do the following:
 - View a set of performance charts, monitor the performance over a given interval, and diagnose and correct problems.
 - Customize the set of performance charts that appear on the page. To do so, click **Show Metric Palette**, and select the charts you want to add to the page.
 - Show or hide the Metrics Palette. To do so, click **Show Metric Palette** or **Hide Metric Palette**, respectively.
 - Reorder the performance charts. To do so, from the **View** menu, select **Reorder Charts**.
 - Customize the performance charts to show or hide availability and threshold details, and grid lines. To do so, from the **View** menu, select **Availability**, **Thresholds**, or **Grid Lines**, respectively.
 - Draw a comparison with another IBM WebSphere Application Server's performance, or with the previous day's performance. To do so, from the **Compare** menu, select **With Another IBM WebSphere Application Server** or **Today with Yesterday**, respectively.
 - Remove comparison. To do so, from the **Compare** menu, select **Remove Comparison**.
 - Create or delete baselines. To do so, from the **Compare** menu, select **Create Baseline** or **Delete Baseline**, respectively.
 - Delete metric performance charts either by clicking the close button on the chart itself, or by deselecting the metric name in the Metric Palette.
 - Change time frames using the slider, or set a default value.
 - Create new metric performance charts by selecting the preferred metrics from the Metric Palette. The charts are automatically created once the metrics are selected.
 - Drag and drop the metrics from a particular metric group to the same chart.

32.5.4 Monitoring the Applications Deployed to IBM WebSphere Application Servers

To monitor the applications running on a IBM WebSphere Application Server, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired **IBM WebSphere Application Server**.
3. On the IBM WebSphere Application Server Home page, in the **Deployments** section, in the **Applications** region, view the following details about applications deployed to the server.

Column	Description
Name	Name of the application deployed to the server.
Status	Status of the application, either Up or Down.
Number of Requests Processed	Total number of requests processed by the application in the last 24 hours.
Average Request Processing Time	Average time taken to service the requests.
Average Request Processing Time by EJB Method	Average time taken by the EJB methods to service the requests.

32.5.5 Viewing the Top EJBs of IBM WebSphere Application Servers

To view the top or the most requested EJBs of an IBM WebSphere Application Server, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired **IBM WebSphere Application Server**.
3. On the IBM WebSphere Application Server Home page, in the **EJBs** section, view a list of EJBs that were most requested in the last 24 hours.

32.5.6 Viewing the Top Servlets and JSPs of IBM WebSphere Application Servers

To view the top or the most requested Servlets and JSPs of an IBM WebSphere Application Server, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired **IBM WebSphere Application Server**.
3. On the IBM WebSphere Application Server Home page, in the **Servlets and JSPs** section, view a list of Servlets and JSPs that were most requested in the last 24 hours.

32.5.7 Viewing IBM WebSphere Application Server Metrics

To view all IBM WebSphere Application Server metrics, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired **IBM WebSphere Application Server**.
3. On the IBM WebSphere Application Server Home page, from the **WebSphere Application Server** menu, select **Monitoring**, then select **All Metrics**.

32.6 Monitoring IBM WebSphere Application Server Clusters

This section covers the following:

- [Monitoring IBM WebSphere Application Server Clusters](#)
- [Administering IBM WebSphere Application Server Clusters](#)
- [Viewing IBM WebSphere Application Server Cluster Members](#)
- [Viewing IBM WebSphere Application Server Cluster Metrics](#)

32.6.1 Monitoring IBM WebSphere Application Server Clusters

To monitor IBM WebSphere Application Server Clusters, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired **IBM WebSphere Application Server Cluster**.
3. On the IBM WebSphere Application Server Cluster Home page, you can monitor the availability, usage, and performance of the selected cluster at a high level.

The status of an IBM WebSphere Application Server Cluster depends on the status of all its members, that is the individual IBM WebSphere Application Servers within the cluster. IBM WebSphere Application Server is an application server developed and maintained by IBM Corporation.

The IBM WebSphere Application Server Cluster Home page has the following sections:

32.6.1.1 Summary Section

Provides a quick, high-level, graphical summary of the availability of the cluster in the last 24 hours.

To view the status (either up or down), hover your mouse over the timeline bar. To zoom in and review the hours of a particular time period, place the cursor at one particular hour of the timeline bar, and with the mouse key pressed, drag the cursor to another hour of interest. You will see that the timeline bar zooms in and displays the hours, minutes, and seconds within that particular time period. This helps when you want to identify the exact time when the cluster went down.

32.6.1.2 Monitoring and Diagnostics Section

Provides a summary of incidents, descendant target incidents, and configuration changes made to the server. Use this information to diagnose and troubleshoot performance issues with the server.

Element	Description
Incidents	Number of unresolved issues that require your attention and corrective action. Click the value to drill down and view more detailed information.
Descendant Target Incidents	Number of changes made to the server configuration in the last 7 days. Click the value to drill down and view more detailed information.
Configuration Changes	Number of incidents related to the applications. The displayed integer is also a link to the Incident Manager page. Click the value to drill down and view more detailed information.

32.6.1.3 Servers Section

Provides information about the members of the cluster, mainly the IBM WebSphere Application Servers that are part of the cluster.

Column	Description
Name	Name of the IBM WebSphere Application Server that is part of the cluster. To navigate to the home page of the server, click the member or server name.
Status	Current status of the IBM WebSphere Application Server. To drill down to the Status History page, click the status icon. You will be taken to the Status History (Availability) page, which shows the availability of the server along with the availability history of the constituents that are used to compute its availability.
Active Sessions	Total number of active or live HTTP sessions in the server over 24 hours. The value appears as a link. To view the number of active sessions for each hour of the 24-hour scale, click the value (link). A graph appears depicting the active sessions for each hour. To drill down further and view more detailed metric-related information, click Metric Details .
Request Processing Time (ms)	Average time taken (in milliseconds) to service a request in the last 24 hours.

32.6.1.4 Resource Usage Section

Provides a graphical representation of the CPU utilization rate and the memory used by JVM for a given interval. To switch to a tabular format, click **Table View**. To drill down and view more detailed metric-related information and to diagnose issues by looking at other related infrastructure metrics, click the server names in the legend and select an appropriate option in the Additional Information message.

32.6.2 Administering IBM WebSphere Application Server Clusters

To administer IBM WebSphere Application Server Clusters, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired IBM WebSphere Application Server Cluster.
3. On the IBM WebSphere Application Server Cluster Home page, you can view high-level information pertaining to the selected server.

To perform administrative tasks on the IBM WebSphere Application Server Cluster, from the **WebSphere Cluster** menu, select any of the following according to your needs:

- **Monitoring**, to monitor the performance of the target, view metric details, view status information, view incidents and alerts raised so far for the target, and view blackouts created for the target.
- **Diagnostics**, to analyze and diagnose performance issues.
- **Control**, to create or end blackouts.
- **Job Activity**, to view details of the jobs created for the target.
- **Information Publisher Reports**, to view reports.
- **Members**, to view and monitor the health of the members of the IBM WebSphere Application Server Cluster. The members are typically the IBM WebSphere Application Servers that are part of the cluster.
- **Configuration**, to search, view, and compare configuration details.
- **Compliance**, to view and create compliance standards.
- **Target Setup**, to view monitoring configuration details and target properties, to remove the target or add it to a group, to view the properties of the target.
- **Target Sitemap**, to view the overall topology of the target.
- **Target Information**, to view general information about the target.

32.6.3 Viewing IBM WebSphere Application Server Cluster Members

Enterprise Manager Cloud Control helps you view the members of an IBM WebSphere Application Server Cluster. You can see what type of members form the cluster, monitor their status, and perform various administrative operations.

To view the members of an IBM WebSphere Application Server Cluster, follow these steps:

1. From the **Targets** menu, click **Middleware**.
2. On the Middleware page, click the desired IBM WebSphere Application Server Cluster.
3. On the IBM WebSphere Application Server Cluster Home page, from the **WebSphere Cluster** menu, select **Members**, then select **Show All** to view the following details of the members.

Column	Description
Name	Name of the IBM WebSphere Application Server that is part of the IBM WebSphere Application Server Cluster. Click the name to access the home page of that IBM WebSphere Application Server.
Type	Type of the member. Typically, IBM WebSphere Application Server.
Status	Current status of the member. Click the status icon to see a consolidated availability summary. You can see the current and past availability status within the last 24 hours, 7 days, or month (31 days).
Incidents	Number of fatal, critical, and warning incidents that occurred for the member server. To drill down to the Incident Manager page and view more detailed information about the incident, click the count.

To search for a particular member, use the **Search** menu.

By default, all members of the IBM WebSphere Application Server Cluster are listed in the table. To refresh the table and view only a particular type of members, select either **Direct Members** or **Indirect Members** from the **View** section.

To save the information about members in a file and to download that file to your local disk, click **Export**.

32.6.4 Viewing IBM WebSphere Application Server Cluster Metrics

To view all IBM WebSphere Application Server metrics, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired **IBM WebSphere Application Server Cluster**.
3. On the IBM WebSphere Application Server Cluster Home page, from the **WebSphere Cluster** menu, select **Monitoring**, then select **All Metrics**.

32.7 Monitoring IBM WebSphere Application Server Cells

This section covers the following:

- [Monitoring IBM WebSphere Application Server Cells](#)
- [Administering IBM WebSphere Application Server Cells](#)
- [Viewing IBM WebSphere Application Server Cell Members](#)

32.7.1 Monitoring IBM WebSphere Application Server Cells

To monitor IBM WebSphere Application Server Cells, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired **IBM WebSphere Application Server Cell**.
3. On the IBM WebSphere Application Server Cell Home page, you can monitor the availability, usage, and performance of the selected cell at a high level.

The status of an IBM WebSphere Application Server Cell depends on the status of all its members, that is the individual IBM WebSphere Application Server Clusters within the cell. The status of each IBM WebSphere Application Server Cluster depends on the status of each IBM WebSphere Application Server within that cluster.

The IBM WebSphere Application Server Cell Home page has the following sections:

32.7.1.1 General Section

Column	Description
WebSphere Cell	Name of the IBM WebSphere Application Server Cell.
Deployment Manager	Name of the Deployment Manager that manages the operations of the IBM WebSphere Application Server Cell being monitored.
Deployment Manager Host	Name of the host where the Deployment Manager is running.
Deployment Manager SOAP Connector Port	SOAP connector port number used to connect to the Deployment Manager.
WebSphere Cell Refreshed	Date and time when the membership of the IBM WebSphere Application Server Cell was last refreshed.

32.7.1.2 Incidents Summary Section

Provides a summary of the fatal, critical, warning, and escalated incidents and problems that occurred on the IBM WebSphere Application Server Cell.

To filter and view a particular category of incidents and problems, from the **Category** list, select a particular category. The table automatically refreshes and lists the incidents and problems pertaining to the selected category.

To hide, unhide, and reorder columns, and to filter and view either all incidents, all incidents without symptoms, or only causes, from the **View** menu, select an appropriate option.

Column	Description
Summary	Intuitive message indicating what the incident is about.
Target	Target type on which the incident or problem occurred.
Severity	Severity of the incident or problem. The severity is either Fatal, Critical, or Warning.
Status	Status of the incident or problem. The status can be either New, Work in Progress, Closed, or Resolved.
Escalation Level	Escalation level signifying the level of attention required on the incident. The escalation level can be either None, which means it is not escalated, or Level 1 through Level 5.
Type	Type of incident or problem being reported.
Time Since Last Update	Date and time the incident was last updated or when the incident was closed.

32.7.1.3 Clusters Section

Provides availability information about the IBM WebSphere Application Server Cluster member targets that are part of the IBM WebSphere Application Server Cell.

Column	Description
Name	Name of the IBM WebSphere Application Server Cluster member target that is part of the IBM WebSphere Application Server Cell. To navigate to the home page of the cluster, click the cluster name.
Status	Current status of the IBM WebSphere Application Server Cluster member target.
Number of Servers	Number of IBM WebSphere Application Servers that are part of the IBM WebSphere Application Server Cluster.

32.7.1.4 Servers Section

Provides availability information about the IBM WebSphere Application Server member targets that are part of the IBM WebSphere Application Server Cell.

Column	Description
Name	Name of the IBM WebSphere Application Server member target that is part of the IBM WebSphere Application Server Cell. To navigate to the home page of a member target, click the member name.
Status	Current status of the IBM WebSphere Application Server member target.
Cluster	Name of the IBM WebSphere Application Server Cluster to which the IBM WebSphere Application Server member target belongs. This column applies only to server targets that are part of a cluster target.

32.7.2 Administering IBM WebSphere Application Server Cells

To administer IBM WebSphere Application Server Cells, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired IBM WebSphere Application Server Cell.
3. On the IBM WebSphere Application Server Cell Home page, you can view high-level information pertaining to the selected cell.

To perform administrative tasks on the IBM WebSphere Application Server Cell, from the **WebSphere Cell** menu, select any of the following according to your needs:

- **Monitoring**, to monitor the performance of the target, view metric details, view status information, view incidents and alerts raised so far for the target, and view blackouts created for the target.
- **Diagnostics**, to analyze and diagnose performance issues.

- **Control**, to create or end blackouts.
- **Job Activity**, to view details of the jobs created for the target.
- **Information Publisher Reports**, to view reports.
- **Members**, to view and monitor the health of the members of the IBM WebSphere Application Server Cell. The members are typically IBM WebSphere Application Cells, IBM WebSphere Application Server Clusters, and IBM WebSphere Application Servers that are part of the cell.
- **Configuration**, to search, view, and compare configuration details.
- **Compliance**, to view and create compliance standards.
- **Target Setup**, to view monitoring configuration details and target properties, to remove the target or add it to a group, to view the properties of the target.
- **Target Sitemap**, to view the overall topology of the target.
- **Target Information**, to view general information about the target.

32.7.3 Viewing IBM WebSphere Application Server Cell Members

Enterprise Manager Cloud Control helps you view the members of an IBM WebSphere Application Server Cell. You can see what type of members form the cell, monitor their status, and perform various administrative operations.

To view the members of an IBM WebSphere Application Server Cell, follow these steps:

1. From the **Targets** menu, click **Middleware**.
2. On the Middleware page, click the desired IBM WebSphere Application Server Cell.
3. On the IBM WebSphere Application Server Cell Home page, from the **WebSphere Cell** menu, select **Members**, then select **Show All** to view the following details of the members.

Column	Description
Name	Name of the IBM WebSphere Application Server Cluster or the IBM WebSphere Application Server that is part of the IBM WebSphere Application Server Cell. To navigate to the home page of a member, click the member name.
Type	Type of the member. Typically, IBM WebSphere Application Server Cluster or IBM WebSphere Application Server.
Status	Current status of the member. Click the status icon to see a consolidated availability summary. You can see the current and past availability status within the last 24 hours, 7 days, or month (31 days).
Incidents	Number of fatal, critical, and warning incidents that occurred for the member server. To drill down to the Incident Manager page and view more detailed information about the incident, click the count.

To search for a particular member, use the **Search** menu.

By default, all members of the IBM WebSphere Application Server Cell are listed in the table. To refresh the table and view only a particular type of members, select either **Direct Members** or **Indirect Members** from the **View** section.

To save the information about members in a file and to download that file to your local disk, click **Export**.

32.8 Troubleshooting IBM WebSphere Application Server Discovery and Monitoring Issues

This section provides troubleshooting tips for the issues encountered while discovering or monitoring IBM WebSphere Application Servers.

- [Troubleshooting Discovery Issues](#)
- [Troubleshooting Monitoring Issues](#)

32.8.1 Troubleshooting Discovery Issues

1. Problem Description

The discovery of a target IBM WebSphere fails at the Host Credentials phase. The discovery of IBM WebSphere fails when you click next after having entered valid target properties for discovery with the following error message.

```
Could not find the required library, specify the home directory.
```

This message is expected at this step as the Agent does not know the `WAS_HOME` directory. However, when you enter the `WAS_HOME` directory, you still get the same error.

Root Cause

This issue is a known issue.

Action

Apply the following workaround.

- a. Create a directory without any space in it and copy the jar files required for discovery in this directory as mentioned in [Prerequisites for Discovering IBM WebSphere Application Servers, Clusters, and Cells](#). Remember to create these directories logged as the OS user you defined in the Agent Host Preferred Credentials.
- b. Select "Agent is running on a host other than the Deployment Manager" as if it was remote monitoring and provide the correct path to the jar files.

2. Problem Description

The IBM WebSphere Application Server still reports Metric Collection errors even after the Agent has been stopped and re-started.

Root Cause

The PMI Service for IBM WebSphere Application Server has not been enabled or the same agent is used to monitor other application servers such as WebLogic or tomcat.

Action

Enable the PMI Service for the WebSphere server that is being Monitored.

Make sure the same agent is not already in use to monitor other application servers such as WebLogic or tomcat. Use a different agent or install a new agent to monitor WebSphere server.

3. Problem Description

In the server home page, select the Applications tab from the IBM WebSphere Application Server Home Page then Applications, you do not see any application listed.

Also, in the all metrics page when you click on a particular metric, you see no data instead of some values.

Root Cause

If you don't see any data in the applications tab or in any particular metric, it just means that there is no load on the Deployed Applications. But, if the load is there and still the data is not seen, the required resources are not created on the server.

Action

None except if there is load on Deployed Applications.

Else enable the option "Create MBeans for Resources" for the application in question from the IBM WebSphere Console.

4. Problem Description

The discovery of IBM WebSphere Application server (as well as other Third Party Application Servers) passes successfully all discovery phases.

It fails only when the **Finish** is clicked and the following error is displayed:

Discovery failed unknown error.

You may be redirected automatically to the first step of the Discovery Wizard.

Root Cause

You were not logged in Cloud Control as a Super User. As stated in the Pre-Requisites, you must be logged with a Super User account (like SYSMAN) in order to successfully discover a target IBM WebSphere Application Server (Cell or Standalone).

Action

Logout of Cloud Control and log in with a Super User account.

5. Problem Description

The discovery of IBM WebSphere Application Server or Application Server Cell fails with the following message displayed:

Error:

No application servers were found on the host <host>. If the port is SSL enabled, specify the port number and the Trusted Keystore file name.

The OMS trace file \$ORACLE_HOME/sysman/log/emoms.trc includes:

Caused by:

```
com.ibm.websphere.management.exception.ConnectorNotAvailableException:
[SOAPException: faultCode=SOAP-ENV:Client; msg=Error opening socket:
javax.net.ssl.SSLHandshakeException: com.ibm.jsse2.util.h: No trusted
certificate found; targetException=java.lang.IllegalArgumentException: Error
```



```
opening socket: javax.net.ssl.SSLHandshakeException: com.ibm.jsse2.util.h: No
trusted certificate found]

at
com.ibm.ws.management.connector.soap.SOAPConnectorClient.reconnect(SOAPConnect
orClient.java:344)

at
com.ibm.ws.management.connector.soap.SOAPConnectorClient.<init>(SOAPConnectorC
lient.java:177)

... 6 more

Caused by: [SOAPException: faultCode=SOAP-ENV:Client; msg=Error opening
socket: javax.net.ssl.SSLHandshakeException: com.ibm.jsse2.util.h: No trusted
certificate found; targetException=java.lang.IllegalArgumentException: Error
opening socket: javax.net.ssl.SSLHandshakeException: com.ibm.jsse2.util.h: No
trusted certificate found]
```

Potential Cause

The SOAP port provided for the discovery process could be incorrect.

Action

Find the correct SOAP port for the node or cell that needs to be discovered.

```
<WAS_HOME>/profiles/<PROFILE>/config/cells/<CELL_NAME>/nodes/<NODE_NAME>/
serverindex.xml
```

The SOAP port is defined within the following XML tags:

```
<specialEndpoints xmi:id="NamedEndPoint_4"
endPointName="SOAP_CONNECTOR_ADDRESS">
<endPoint xmi:id="EndPoint_4" host="celtpvm4.us.example.com" port="8879"/>
</specialEndpoints>
```

In this example the node or cell SOAP port is 8879.

This is the value that should be used for 'SOAP connector port' in the discovery form.

6. Problem Description

After having discovered the WebSphere instance, the following metric collection error is returned:

```
oracle.sysman.emSDK.emd.fetchlet.FetchletException:
java.lang.NoClassDefFoundError:

Could not initialize class com.ibm.websphere.management.AdminClientFactory
```

Root Cause

Caused by an incorrect class path used during discovery.

Action

From the **Targets** tab, select **Middleware**, and then select **IBM WebSphere Server Target**.

From the **Target Setup** menu, select **Monitoring Configuration**, and enter the correct WebSphere Home path, and click **OK**.

7. Problem Description

The discovery of IBM WebSphere Application Server or Application Server Cell fails with unknown error.

After following the above trouble shooting sections if the discovery issue is still not resolved, we recommend to run the discovery from UI with the following property set on the agent.

Potential Cause

There could be various causes for the discovery failure; looking at the log file - emagent_perl.trc with the following property set, will help to identify the root cause for the discovery failure.

Action

- a. Add the following property in the file `$AGENT_INSTALL_HOME/sysman/config/emd.properties`:

```
EMAGENT_PERL_TRACE_LEVEL=DEBUG
```

- b. Perform the discovery from the UI.
- c. Look at the log file `$AGENT_INSTALL_HOME/sysman/log/emagent_perl.trc`.

for the xml output beginning with the tag

```
<Targets>
```

```
.....
```

```
</Targets>
```

The error message encoded in the xml output will help identify the root cause of the discovery.

8. Problem Description

An error is displayed after initiating a refresh of an upgraded WebSphere Cell target from EM.

After a WebSphere AS instance is upgraded to a newer version and when a refresh action of the existing WebSphere Cell targets on EM is performed, the following error is displayed:

```
<DiscoveryWarning DISCOVERY_SCRIPT="666">Version incorrect</DiscoveryWarning>
```

Note:

The error mentioned above can be seen on the Management Agent from `emagent_perl.trc` log file when debug mode is enabled.

Root Cause

The version of WebSphere AS captured as a part of **Monitoring Configuration** of all the WebSphere AS targets within the affected WebSphere Cell needs to be updated to reflect the upgraded version before initiating a refresh of the WebSphere Cell.

Action

The **Version** field captured as part of the Monitoring Configuration of any WebSphere AS target can be updated in one of the two ways mentioned below:

Method 1

- a. In the EM Cloud Control console, click **WebSphere Application Server** drop-down button and select **Target Setup**.
- b. Select **Monitoring Configuration**.
- c. Edit the **Version** field to display the new version.

Note:

The **Version** field is editable only in EM Cloud Control 12c FMW Plug-in 12.1.0.5.x and below. For versions 12.1.0.6.0 and above, this field is read-only on the EM CC console. Hence for versions 12.1.0.6.0 and above use the 'emcli modify_target' command to update the this field.

Method 2

Use `emcli modify_target` command. Enter the following command:

```
emcli modify_target -name="was_target_name" -type="websphere_j2eeserver" -  
properties="version:x.x.x.x"
```

The above command modifies the `version` field of **Monitoring configuration** for the specified WebSphere AS Target.


To view the list of WebSphere AS targets currently monitored by EM enter the following command:

```
emcli get_targets -targets="websphere_j2eeserver"
```

32.8.2 Troubleshooting Monitoring Issues

The names of the IBM WebSphere Application Servers discovered in Enterprise Manager Cloud Control 12c Release 2 (12.1.0.2) or lower, appear only with the server name. For example, `testserver_1_1`. However, the names of the IBM WebSphere Application Servers discovered in Enterprise Manager Cloud Control 12c Release 3 (12.1.0.3) or higher, appear with the node name and server name. For example, `exampleNode01.testserver_1_2`. See [Figure 32-2](#).

Figure 32-2 Issue with the Display Name of IBM WebSphere Application Server



Discovering and Monitoring JBoss Application Server

This section describes how you can discover and monitor JBoss application servers in Enterprise Manager Cloud Control.

JBoss Application Server is the market-leading, open source Java Platform, Enterprise Edition (Java EE) application server, delivering a high-performance and enterprise-class platform for e-business applications. JBoss provides enterprise-class security, transaction support, resource management, load balancing, and clustering.

Enterprise Manager Cloud Control enables you to discover JBoss Application Servers in your environment and add them to Cloud Control for central monitoring and management.

This chapter describes how you can discover and monitor these JBoss Application Server targets in Enterprise Manager Cloud Control. In particular, this chapter covers the following:

- [About Managing JBoss Application Servers, JBoss Domains, and JBoss Partitions](#)
- [Finding Out the Supported Versions for Discovery and Monitoring](#)
- [Prerequisites for Discovering JBoss Application Servers, Domains, and Partitions](#)
- [Discovering JBoss Application Servers 7.x and JBoss Domains](#)
- [Discovering JBoss Application Servers 6.x and JBoss Partitions](#)
- [Monitoring JBoss Application Servers](#)
- [Monitoring JBoss Domains](#)
- [Monitoring JBoss Partitions](#)
- [Deploying JVMD on JBoss Application Server 7.x and 6.x to Diagnose Issues](#)
- [Troubleshooting JBoss Application Server Discovery and Monitoring Issues](#)

33.1 About Managing JBoss Application Servers, JBoss Domains, and JBoss Partitions

Using Enterprise Manager Cloud Control, you can do the following with JBoss Application Server targets:

- Discover the following for central monitoring and management:
 - JBoss Application Servers
 - JBoss Server Groups and Domains (for version 7.x)

JBoss Domain is a logical grouping of JBoss Application Servers within your enterprise configuration for JBoss Application Servers version 7.x. JBoss Application Servers are grouped into Server Groups and one or more Server

Groups form a JBoss Domain. The Domain Controller runs and manages the collection of JBoss Application Servers within a domain.

When you discover a JBoss Domain, the JBoss Domain, the JBoss Server Groups and all other up and running JBoss Application Servers that are part of that JBoss Domain get automatically discovered. You can select the JBoss Application Servers that you want to add to Enterprise Manager Cloud Control for monitoring, and the associated JBoss Server Groups are automatically added.

– JBoss Partitions (for version 6.x)

JBoss Partition is a logical grouping of JBoss Application Servers within your enterprise configuration for JBoss Application Servers version 6.x.

When you discover a JBoss Application Server that is part of a JBoss Partition, the JBoss Partition and all other JBoss Application Servers that are part of that JBoss Partition gets automatically discovered and added to Enterprise Manager Cloud Control.

- Monitor the status, the availability percentage, the CPU usage, the heap usage, the Java vendor and version used, and so on.
- Monitor the status and the overall health of the member application servers that are part of JBoss Domains and Partitions.
- Monitor the performance by measuring the load and the request processing time for a given interval.
- Diagnose, notify, and correct performance and availability problems with the help of GUI-rich, intuitive graphs and illustrations.
- Monitor the status of the deployed applications.
- Monitor the Servlets and JSPs running on the application servers, including the most requested Servlets in the last 24 hours - applicable only for JBoss Application Servers version 6.x.
- View details about the associated JVM threads and data sources.
- Create or end blackouts as well as notification blackouts to suspend or resume the collection of metric data, respectively.
- Monitor and manage configuration details in JBoss version 6 that were last collected and also the ones that were saved at a given point of time.
- Compare the configuration between in JBoss version 6:
 - A last collected configuration of a server instance with a saved configuration of the same server instance or a different server instance.
 - A last collected configuration of a server instance with a last collected configuration of a different server instance.
 - A saved configuration of a server instance with another saved configuration of the same server instance or a different server instance.
 - A saved configuration of a server instance with a last collected configuration of the same server instance or a different server instance.
- View compliance-related information, such as the compliance standards and frameworks associated with the server, the real-time observations, the evaluation results, and so on.
- View a list of metrics, their collection interval, and the last upload for each metric.

33.2 Finding Out the Supported Versions for Discovery and Monitoring

To search for the JBoss Application Server versions that are supported for discovery and monitoring in Enterprise Manager Cloud Control, follow these steps:

1. Log in to <https://support.oracle.com/>
2. On the My Oracle Support home page, select **Certifications** tab.
3. On the Certifications page, enter the following search criteria in the Certification Search section.
 - Enter the product name **Enterprise Manager Base Platform - OMS** in the Product field.
 - Select the release number **13.3.1.0.0** from the Release list.
4. Click **Search**.
5. In the Certification Results section, expand the **Application Server** menu to view the certified JBoss Application Server versions.

Certified With	Number of Releases / Versions
> Operating Systems (8 Items)	
> Agents (1 Item)	
∨ Application Servers (10 Items)	
Apache Tomcat (Managed Target)	7 Releases (11, 6.0.1, 6.0.2, 6.0.3, 6.0.4)
IBM WebSphere Application Server (Managed Target)	4 Releases (8.5, 8.5.5, 8.5.6, 8.5.7)
JBoss Application Server (Managed Target)	4 Releases (5.1, 5.1.1, 5.1.2, 5.1.3)
Oracle Application Server (Managed Target)	2 Releases (10.1.2.0.0, 10.1.2.0.1)
Oracle Application Server Single Signon (Managed Target)	1 Release (10.1.4.2.0)
Oracle GlassFish Server (Managed Target)	1 Release (10.1.2.0.0)
Oracle WebLogic Portal (Managed Target)	4 Releases (10.1.2.0.0, 10.1.2.0.1, 10.1.2.0.2, 10.1.2.0.3)
Oracle WebLogic Server (Infrastructure)	2 Releases (10.1.2.0.0, 10.1.2.0.1)
Oracle WebLogic Server (Managed Target)	7 Releases (10.1.2.0.0, 10.1.2.0.1, 10.1.2.0.2, 10.1.2.0.3, 10.1.2.0.4, 10.1.2.0.5, 10.1.2.0.6)
WebLogic Server (Managed Target)	7 Releases (10.1.2.0.0, 10.1.2.0.1, 10.1.2.0.2, 10.1.2.0.3, 10.1.2.0.4, 10.1.2.0.5, 10.1.2.0.6)
> Databases (8 Items)	
> Desktop Applications, Browsers and Clients (5 Items)	
> Directory/LDAP Services (6 Items)	
> Enterprise Applications (150 Items)	
> Management and Development Tools (49 Items)	
> Middleware (28 Items)	
> Other (3 Items)	
> Server (2 Items)	
> Virtualization Software (1 Item)	

33.3 Prerequisites for Discovering JBoss Application Servers, Domains, and Partitions

Meet the following prerequisites for discovering JBoss Application Servers and JBoss Partitions.

- Ensure that you download and extract the JBoss Application Server installable ZIP file available on JBoss site, and set the `JBOSS_HOME` and `PATH` environment variables as follows:

```
setenv JBOSS_HOME <jboss_install_location>
setenv PATH "${PATH}:/${JBOSS_HOME}/bin"
```

- Ensure that you start the JBoss Application Server or the JBoss Partition by running the following command from the `bin` directory:

```
./run.sh -c <deployment_profile> -b <binding_address> [-Djboss.partition.name=<partition_name>]
```

Here, `<deployment_profile>` indicates whether you are starting a standalone JBoss Application Server or a JBoss Partition. The `<binding_address>` is the host name or the IP address running the JBoss Application Server. The `<partition_name>` is the partition name from where the JBoss Application Servers must start. By default, they start as part of `DefaultPartition`.

For example,

```
./run.sh -c node1 -Djboss.service.binding.set=ports-01 -b example.oracle.com
```

Note:

- To start a standalone JBoss Application Server, set the `<deployment_profile>` to `default`.

For example,

```
./run.sh -c default -b <binding address>
```

- To start a JBoss Partition, enable the JBoss clustering service, set the `<deployment_profile>` to `all`.

For example,

```
./run.sh -c all -b <binding address>
```

- To start multiple server instances on the same host, complete the following:

(a) Create multiple deployment profiles as per your requirements.

(b) Use a different port-set to start the individual servers. Note that `ports-01`, `ports-02`, `ports-03`, and `ports-04` are predefined port-sets.

For example,

```
-Djboss.service.binding.set=ports-01
```

33.4 Discovering JBoss Application Servers 7.x and JBoss Domains

To discover JBoss Application Servers and JBoss Domains, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, from the **Add** menu, select **JBoss Application Server**. Enterprise Manager Cloud Control displays the JBoss Discovery Wizard.
3. On the Host page, enter details about the host on which the JBoss Application Server is running.

Figure 33-1 JBoss Application Server Host Page

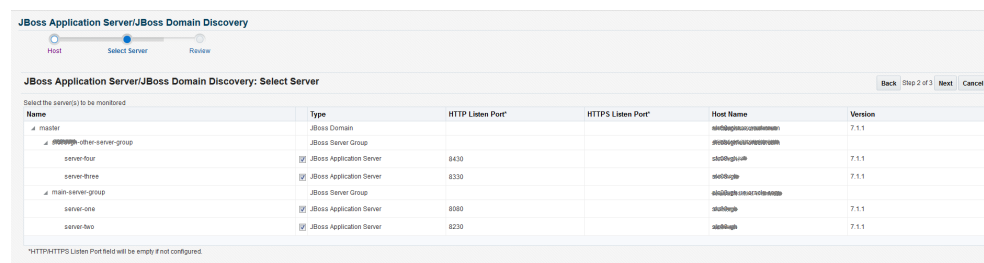


Element	Description
Version	Select the version of the JBoss Application Server you want to discover.
Discovery Mode	Select either Standalone or Domain.
JBoss Application Server Host /JBoss Domain Controller Host	If you have selected standalone server enter the host address where the JBoss Server is running, and if you have selected domain enter the host address where the JBoss Domain Controller is running.
HTTP Management Port	Depending on whether you have selected a standalone server or domain, specify the HTTP Management port of the JBoss Application Server or the Domain Controller.
Agent	Select the Management Agent that is installed on the host where the JBoss Application Server is running.
Authentication Type	Select the authentication type. The available options are None, Basic and HTTP Digest.
Username	Enter the user name for authentication.
Password	Enter the password for authentication.

Click **Next**.

- On the Select Server page, view a list of JBoss Domains, JBoss Server Groups and standalone JBoss Application Servers discovered on the host you specified, and to select the ones you want to add and monitor in Enterprise Manager Cloud Control.

Figure 33-2 JBoss Application Server Select Server Page



Column	Description
Name	Name of the JBoss Domains and Server Groups and standalone JBoss Application Servers discovered on the specified host.
Select	Select the JBoss Application Server(s) that you want to add and monitor in Enterprise Manager Cloud Control. Note: When you select a JBoss Application Server, the associated JBoss Domain and Server Group also gets selected for monitoring.
Type	Type of JBoss target discovered on the specified host. They can be either JBoss Domain, JBoss Server Group or JBoss Application Server.
HTTP Listen Port	The HTTP listener port that is configured.
HTTPS Listen Port	The HTTPS listener port that is configured.
Host Name	The name of the JBoss Application Server host.
Version	Version of the JBoss Application Server discovered on the specified host.

Click **Next**.

5. On the Review page, review the details you have provided for discovering and adding JBoss targets to Enterprise Manager Cloud Control. Click **Submit** to discover the JBoss targets.



Note:

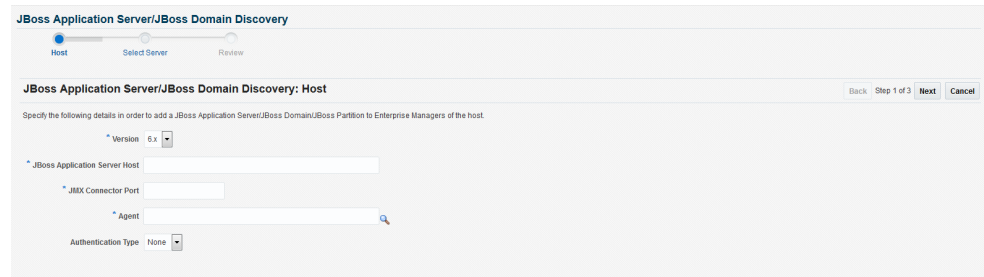
When you discover a JBoss Domain all the JBoss Application Servers part of that domain get automatically discovered and added to Enterprise Manager Cloud Control. At any point after discovering a JBoss Domain, if new JBoss Application Servers are added to the domain, then you can refresh the JBoss Domain as described in [Refreshing JBoss Partitions](#).

33.5 Discovering JBoss Application Servers 6.x and JBoss Partitions

To discover JBoss Application Servers and JBoss Partitions, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, from the **Add** menu, select **JBoss Application Server**. Enterprise Manager Cloud Control displays the JBoss Discovery Wizard.
3. On the Host page, enter details about the host on which the JBoss Application Server is running.

Figure 33-3 JBoss Application Server Host Page

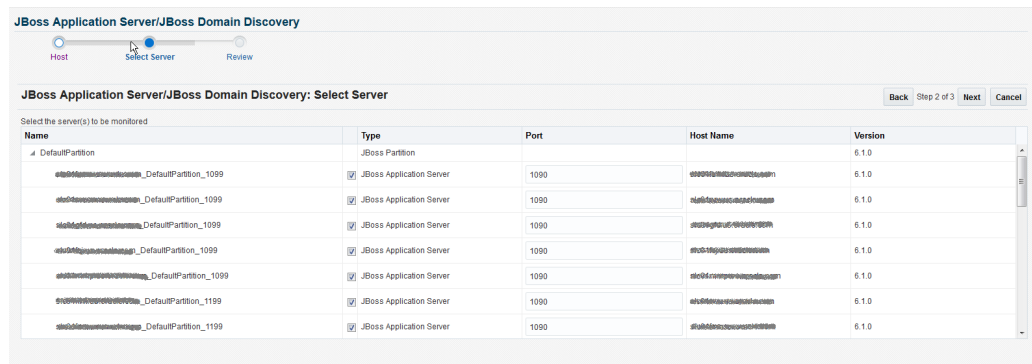


Element	Description
Version	Select the version of the JBoss Application Server you want to discover.
JBoss Application Server Host	Enter the host address where the JBoss Server is running.
JMX Connector Port	Enter the JMX connector port number.
Agent	Select the Management Agent that is installed on the host where the JBoss Application Server is running.
Authentication Type	Select the authentication type. The available options are None, and Basic.
Username	Enter the JMX user name for authentication.
Password	Enter the JMX password for authentication.

Click **Next** after you are done.

- On the Select Server page, view a list of JBoss Partitions and standalone JBoss Application Servers discovered on the host you specified, and to select the ones you want to add and monitor in Enterprise Manager Cloud Control.

Figure 33-4 JBoss Application Server Select Server Page



Column	Description
Name	Name of the JBoss Partitions and standalone JBoss Application Servers discovered on the specified host.

Column	Description
Select	Select the JBoss Partition or standalone JBoss Application Server that you want to add and monitor in Enterprise Manager Cloud Control. Note: When you select a JBoss Application Server, the associated JBoss Partition also gets selected for monitoring.
Type	Type of JBoss target discovered on the specified host. They can be either JBoss Partition or JBoss Application Server.
Port	Enter the JMX connector port of the JBoss Application Server. By default, the port number appears for a JBoss Application Server target if you entered it in the previous page of the wizard. Otherwise, the field is blank.
Host Name	The name of the JBoss Application Server host.
Version	Version of the JBoss Application Server discovered on the specified host.

 **Note:**

Enterprise Manager Cloud Control does not validate the values you provide for User Name and Password. So if you provide incorrect values, Enterprise Manager will add the JBoss target without displaying any errors now, but will eventually show the status as *Down*.

Click **Next** after you are done.

5. On the Review page, review the details you have provided for discovering and adding JBoss targets to Enterprise Manager Cloud Control. Click **Submit** to discover the JBoss targets.

 **Note:**

When you discover a JBoss Application Server that is part of a JBoss Partition, the JBoss Partition and all other JBoss Application Servers part of that partition get automatically discovered and added to Enterprise Manager Cloud Control. At any point after discovering a JBoss Partition, if new JBoss Application Servers are added to the partition, then you can refresh the JBoss Partition as described in [Refreshing JBoss Partitions](#).

33.6 Monitoring JBoss Application Servers

This section covers the following:

- [Monitoring JBoss Application Servers 7.x](#)
- [Monitoring JBoss Application Servers 6.x](#)
- [Administering JBoss Application Servers 7.x and 6.x](#)
- [Monitoring Applications Deployed to JBoss Application Servers 7.x and 6.x](#)
- [Monitoring the Performance of JBoss Application Servers 7.x and 6.x](#)

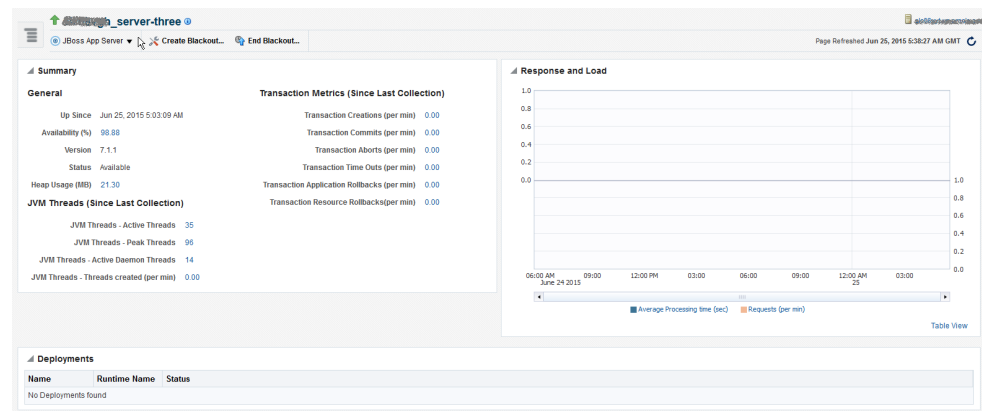
- [Monitoring Servlets and JSPs Running on JBoss Application Servers 6.x](#)
- [Viewing JBoss Application Server Metrics](#)
- [Analyzing Problems Using Metric Correlation](#)

33.6.1 Monitoring JBoss Application Servers 7.x

To monitor JBoss Application Servers version 7.x, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired **JBoss Application Servers**.
3. On the JBoss Application Server Home page, you can view a summary of the most critical information pertaining to JBoss Application Server. You can view general information about the server, information about the JVM threads running on the server, and performance summary in terms of load and response time.

Figure 33-5 JBoss Application Server 7.x Home Page



The JBoss Partition Home page has the following sections:

- **Summary Section:**
 - [General](#)
 - [JVM Threads](#)
 - [Transaction Metrics](#)
- [Response and Load Section](#)
- [Deployments Section](#)

33.6.1.1 General

Element	Description
Up/Down/Pending Since	Date and time when the status was last determined.

Element	Description
Availability (%)	Indicates whether the JBoss Application Server is available and the availability percentage over the last 24 hours. To drill down to the Status History (Availability) page, click the link. The Status History page displays the availability of the JBoss Application Server along with the availability history of the constituents that are used to compute its availability.
Version	Version of the JBoss Application Server.
Status	Current status of the JBoss Application Server. The status can be down even if incorrect credentials were provided while discovering the JBoss target.
Heap Usage (MB)	Amount of heap space (in MB) used by the JBoss Application Server since the last collection.

33.6.1.2 JVM Threads

Element	Description
JVM Threads - Active Threads	Number of active JVM threads, including both daemon and non-daemon threads.
JVM Threads - Peak Threads	Number of peak active JVM threads since the Java Virtual Machine started or peak was reset.
JVM Threads - Active Daemon Threads	Number of active daemon JVM threads.
JVM Threads - Threads Created (per min)	Number of JVM threads created per minute.

33.6.1.3 Transaction Metrics

Element	Description
Transaction Creations (per min)	Number of new transactions created per minute.
Transaction Commits (per min)	Number of transactions committed per minute.
Transaction Aborts (per min)	Number of transactions aborted per minute.
Transaction Time Outs (per min)	Number of transactions timed-out, per minute.
Transaction Application Rollbacks (per min)	Number of transactions rolled back by the application, per minute.
Transaction Resource Rollbacks (per min)	Number of transactions rolled back by the resource, per minute.

33.6.1.4 Response and Load Section

Provides a graphical representation of the server's performance, measuring request-processing time for a given interval. To switch to a tabular format, click **Table View**. To drill down and view more detailed metric-related information and to diagnose issues by looking at other related infrastructure metrics, click the metric names in the legend and select an appropriate option in the Additional Information message.

33.6.1.5 Deployments Section


Provides the deployments on the JBoss Application Server and their statuses.

33.6.2 Monitoring JBoss Application Servers 6.x

To monitor JBoss Application Servers version 6.x, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired **JBoss Application Servers**.
3. On the JBoss Application Server Home page, you can view a summary of the most critical information pertaining to JBoss Application Server. You can view general information about the server, information about the Servlets and JVM threads running on the server, and performance summary in terms of load and response time.

Figure 33-6 JBoss Application Server 6.x Home Page



The JBoss Application Server Home page has the following sections:

- Summary Section:
 - General
 - Servlets/JSPs
 - JVM Threads
 - Datasource
- Response and Load Section
- Most Requested Servlets/JSPs Section

33.6.2.1 General

Element	Description
Up/Down/Pending Since	Date and time when the status was last determined.

Element	Description
Availability (%)	Indicates whether the JBoss Application Server is available and the availability percentage over the last 24 hours. To drill down to the Status History (Availability) page, click the link. The Status History page displays the availability of the JBoss Application Server along with the availability history of the constituents that are used to compute its availability.
Version	Version of the JBoss Application Server.
Status	Current status of the JBoss Application Server. The status can be down even if incorrect JMX credentials were provided while discovering the JBoss target.
Server CPU Usage (%)	Percentage of CPU time used by the JBoss Application Server.
Heap Usage (MB)	Amount of heap space (in MB) used by the JBoss Application Server over a given interval.
Java Vendor	Vendor of the Java Virtual Machine that this JBoss Application Server runs.
Java Version	Version of the Java Virtual Machine that this JBoss Application Server runs.

33.6.2.2 Servlets/JSPs

Element	Description
Active Sessions	Number of active servlet and JSP sessions.
Request Processing Time (ms)	Average time taken (in milliseconds) to service a request in the last 24 hours.
Requests (per min)	Number of requests serviced per minute in the last 24 hours.

33.6.2.3 JVM Threads

Element	Description
JVM Threads - Active Threads	Number of active JVM threads, including both daemon and non-daemon threads.
JVM Threads - Peak Threads	Number of peak active JVM threads since the Java Virtual Machine started or peak was reset.
JVM Threads - Active Daemon Threads	Number of active daemon JVM threads.
JVM Threads - Threads Created (per min)	Number of JVM threads created per minute.

33.6.2.4 Datasource

Element	Description
Total connections in Use	Number of active database connections in this instance of the data source since the data source was instantiated.

Element	Description
Total connection Pool Size	Total size of the connection pool.
Total connections Created (per min)	Number of connections created per minute for this data source.
Total connections Destroyed (per min)	Number of connections closed per minute for this data source.

33.6.2.5 Response and Load Section

Provides a graphical representation of the server's performance, measuring request-processing time for a given interval. To switch to a tabular format, click **Table View**. To drill down and view more detailed metric-related information and to diagnose issues by looking at other related infrastructure metrics, click the metric names in the legend and select an appropriate option in the Additional Information message.

33.6.2.6 Most Requested Servlets/JSPs Section

Provides details of the most requested servlets and JSPs in the last 24 hours.

33.6.3 Administering JBoss Application Servers 7.x and 6.x

To administer JBoss Application Servers, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired JBoss Application Server.
3. On the JBoss Application Server Home page, you can view high-level information pertaining to the selected JBoss Application Server.

To perform administrative tasks on the JBoss Application Server, from the JBoss Server menu, select any of the following according to your needs:

- **Monitoring**, to monitor the performance of the target, view metric details, view status information, view incidents and alerts raised so far for the target, and view blackouts and notification blackouts created for the target.
- **Diagnostics**, to analyze and diagnose performance issues.
- **Control**, to create or end blackouts and notification blackouts.
- **Job Activity**, to view details of the jobs created for the target.
- **Information Publisher Reports**, to view reports.
- **Configuration**, to search, view, and compare configuration details in JBoss version 6.
- **Compliance**, to view and create compliance standards.
- **Target Setup**, to view monitoring configuration details and target properties, to remove the target or add it to a group, to migrate and use JMX.
- **Target Sitemap**, to view the overall topology of the target.
- **Target Information**, to view general information about the target.

33.6.4 Monitoring Applications Deployed to JBoss Application Servers 7.x and 6.x

To monitor the applications deployed to JBoss Application Servers 7.x and 6.x, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the JBoss Application Server where the Servlets and JSPs are deployed.
3. On the JBoss Application Server Home page, from the **JBoss Server** menu, select **Monitoring**, then select **Performance Summary**.
4. On the Performance Summary page, scroll down to the **Applications** section.

33.6.5 Monitoring the Performance of JBoss Application Servers 7.x and 6.x

Enterprise Manager Cloud Control helps you monitor the overall performance of JBoss Application Servers. You can view the graphs that depict their memory usage and heap usage. This helps you gauge the performance and perform a root cause analysis to drill down to the problem areas and fix them before they affect the end users.

To monitor the performance of a JBoss Application Server:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the JBoss Application Server whose performance you want to monitor.
3. On the JBoss Application Server Home page, from the **JBoss Server** menu, select **Monitoring**, then select **Performance Summary**.
4. On the Performance Summary page, you can do the following:
 - View a set of performance charts, monitor the performance over a given interval, and diagnose and correct problems.
 - Customize the set of performance charts that appear on the page. To do so, click **Show Metric Palette**, and select the charts you want to add to the page.
 - Show or hide the Metrics Palette. To do so, click **Show Metric Palette** or **Hide Metric Palette**, respectively.
 - Reorder the performance charts. To do so, from the **View** menu, select **Reorder Charts**.
 - Customize the performance charts to show or hide availability and threshold details, and grid lines. To do so, from the **View** menu, select **Availability**, **Thresholds**, or **Grid Lines**, respectively.
 - Draw a comparison with another IBM WebSphere Application Server's performance, or with the previous day's performance. To do so, from the **Compare** menu, select **With Another JBoss Application Server** or **Today with Yesterday**, respectively.
 - Remove comparison. To do so, from the **Compare** menu, select **Remove Comparison**.

- Create or delete baselines. To do so, from the **Compare** menu, select **Create Baseline** or **Delete Baseline**, respectively.
- Delete metric performance charts either by clicking the close button on the chart itself, or by deselecting the metric name in the Metric Palette.
- Change time frames using the slider, or set a default value.
- Create new metric performance charts by selecting the preferred metrics from the Metric Palette. The charts are automatically created once the metrics are selected.
- Drag and drop the metrics from a particular metric group to the same chart.

33.6.6 Monitoring Servlets and JSPs Running on JBoss Application Servers 6.x

Enterprise Manager Cloud Control helps you monitor the Servlets and JSPs that are running on JBoss Application Servers. You can not only view high-level information about them but also a performance summary that reflects their response time and load. You can also drill down and diagnose issues by viewing related infrastructure metrics, alert history, and so on.

To monitor the Servlets and JSPs running on JBoss Application Servers, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the **JBoss Application Server** where the Servlets and JSPs are deployed.
3. On the JBoss Application Server Home page, do these:
 - a. To view high-level information about the Servlets, see the **Servlets** region. To understand the metric details displayed in this region, click **Help**.
 - b. To monitor the performance of Servlets and JSPs, see the response and load graphic.
 - c. To drill down and diagnose issues, click a metric name in the legend. From the pop-up message, click **Problem Analysis**.
 - d. To view metric statistics, thresholds, and metric value history, click a metric name in the legend. From the pop-up message, click **Metric Details**.

33.6.7 Viewing JBoss Application Server Metrics

To view all JBoss Application Server metrics, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired **JBoss Application Server**.
3. On the JBoss Application Server Home page, from the **JBoss Server** menu, select **Monitoring**, then select **All Metrics**.

33.6.8 Analyzing Problems Using Metric Correlation


For information on spikes in the performance of metrics, you can use the Problem Analysis page to compare results between the source metric and related metrics. Currently, problem analysis is only available for the following metrics.

- Server CPU Usage
- Servlet and JSP - Request Processing Time

To access the Problem Analysis page, follow these steps.

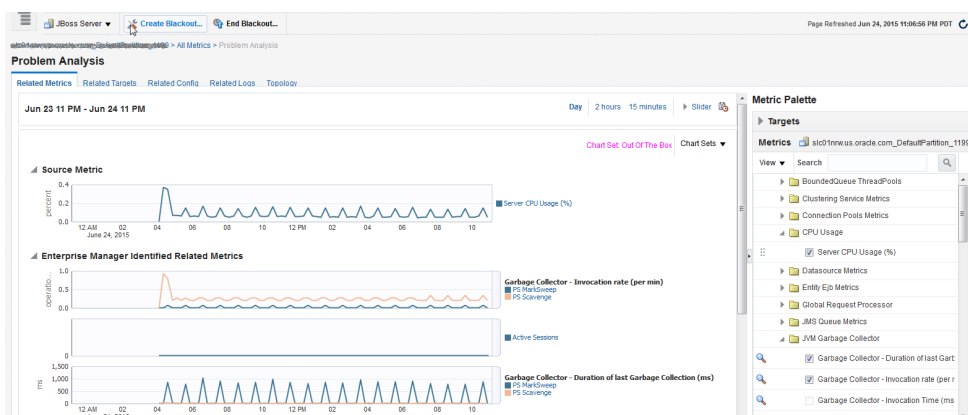
1. In the JBoss Application Server home page from the JBoss Server menu, select **Monitoring** and then select **Performance Summary**.
2. From the **Performance Summary** page, click the name of the metric next to the performance chart.
3. From the window that pops up, select **Problem Analysis**.

Figure 33-7 JBoss Application Server Performance Page



4. On the Problem Analysis page, you can compare the results of the Source Metric and the Related Metrics.

Figure 33-8 JBoss Application Server Problem Analysis Page



33.7 Monitoring JBoss Domains

This section covers the following:

- [Monitoring JBoss Domains](#)
- [Monitoring JBoss Server Groups](#)
- [Administering JBoss Domains](#)
- [Viewing JBoss Domain Members](#)
- [Refreshing JBoss Domains](#)

33.7.1 Monitoring JBoss Domains

To monitor JBoss Domains, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired **JBoss Domain**.
3. On the JBoss Domain Home page, you can do the following:
 - View a summary of the most critical information pertaining to the JBoss Domain.
 - Monitor the status and availability of all the members within the JBoss Domain.
 - View the incidents reported on the domain.

Using the JBoss Domain Home page, you can monitor the individual status of each of the members and also refresh the domain to update the membership and reflect the current deployment state. You can also view the member application servers' resource usage, availability, performance, configuration information, and reports with or without historical data.

Figure 33-9 JBoss Domain Home Page

The screenshot displays the JBoss Domain Home Page for 'master_9990'. The page is divided into several sections:

- Summary:** Shows 'JBoss Domain Refreshed Jun 24, 2015 4:36:04 AM PDT', 'Version 7.1.1', and 'Agent'.
- General:** Displays 'JBoss Domain Refreshed Jun 24, 2015 4:36:04 AM PDT', 'Version 7.1.1', and 'Agent'.
- Incidents:** A table with columns: Summary, Target, Severity, Status, Escalation Level, Type, and Time Since Last Update. It shows two incidents with 'New' status and 'Incident' type.
- Servers:** A section showing a green circle indicating 100% availability and a table listing server groups and individual servers with their status and version.

Name	Type	Status	Version
/jbossdomain/master_9990	JBoss Domain		
/jbossdomain/master_9990/slc0@high-other-server-group	JBoss Server G...		
/jbossdomain/master_9990/slc0@high/server-three	JBoss Applicat...	↑	7.1.1
/jbossdomain/master_9990/slc0@high/server-four	JBoss Applicat...	↑	7.1.1
/jbossdomain/master_9990/main-server-group	JBoss Server G...		
/jbossdomain/master_9990/slc0@high/server-one	JBoss Applicat...	↑	7.1.1
/jbossdomain/master_9990/slc0@high/server-two	JBoss Applicat...	↑	7.1.1

The JBoss Domain Home page has the following sections:

- [Summary Section](#)
- [Servers Section](#)

- [Incidents Section](#)

33.7.1.1 Summary Section

The General sub-section under the Summary section provides general information about JBoss Domain.

Element	Description
JBoss Domain Refreshed	Indicates the time and date when the JBoss Domain was last refreshed.
Version	Version of the JBoss Domain.
Agent	Management Agent used for discovering the JBoss Application Servers that are part of the JBoss Domain. To drill down to the Management Agent home page, click the link.

33.7.1.2 Servers Section

The Servers section provides a real-time view of the status and availability of all the members within the JBoss Domain. For example, if there are four JBoss Application Servers within a JBoss Domain, and if three of these are up, then the pie chart shows 75% up and 25% down status. Accordingly, the legend shows 3 against the Up status to indicate the JBoss Application Servers that are up, and 1 against the Down status to indicate the server that is down.

The table provides the complete hierarchy of the JBoss Domain. You can drill down and view information about a JBoss Server Group or a JBoss Application Server, by clicking on it. To view more information about the status, click the status icon.

33.7.1.3 Incidents Section

The Incidents section provides a summary of all the incidents reported on the domain. Use the **View** menu to sort, filter and organize the Incidents table. Click the incident for further information related to each incident.

Switch between **Current Target** and **Hardware Targets** tabs to view the related incidents.

33.7.2 Monitoring JBoss Server Groups

To monitor JBoss Domains, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the pointer before the JBoss Domain to see the Server Groups under it.
3. Click the desired Server Group to open the Server Group Home page.
4. On the JBoss Server Group Home page, you can do the following:
 - View a summary of the most critical information pertaining to the JBoss Server Group.
 - Monitor the status and availability of all the members within the JBoss Server Group.

33.7.2.3 Incidents Section

The Incidents section provides a summary of all the incidents reported on the Server Group. Use the **View** menu to sort, filter and organize the Incidents table. Click the incident for further information related to each incident.

Switch between **Current Target** and **Hardware Targets** tabs to view the related incidents.

33.7.3 Administering JBoss Domains

To administer JBoss Domains, follow these steps:

1. From the **Targets** menu, click **Middleware**.
2. On the Middleware page, click the desired JBoss Domain target.
3. On the JBoss Domain home page, you can view high-level information pertaining to the selected JBoss Domain.

To perform administrative tasks on the JBoss Domain, from the **JBoss Domain** menu, select any of the following according to your needs:

- **Monitoring**, to monitor the performance of the target, view metric details, view status information, view incidents and alerts raised so far for the target, and view blackouts and notification blackouts created for the target.
- **Control**, to create or end blackouts and notification blackouts.
- **Members**, to view details of the JBoss Application Servers that are part of the JBoss Domain.
- **Refresh JBoss Domain**, to refresh the JBoss Domain and to add any new JBoss Application Servers added to the domain.
- **Configuration**, to search, view, and compare configuration details.
- **Compliance**, to view and create compliance standards.
- **Target Setup**, to view monitoring configuration details and target properties, to remove the target or add it to a group.
- **Target Sitemap**, to view the overall topology of the target.
- **Target Information**, to view general information about the target.

33.7.4 Viewing JBoss Domain Members

Enterprise Manager Cloud Control helps you view the members of a JBoss Domain. You can see what type of members form the domain, monitor their status, and perform various administrative operations

To view a list of members, follow these steps:

1. From the **Targets** menu, click **Middleware**.
2. On the Middleware page, click the desired JBoss Domain target.
3. On the JBoss Domain Home page, from the **JBoss Domain** menu, select **Members**, then select **Show All** to see the following details.

Column	Description
Name	Name of the JBoss Application Server that is part of the JBoss Domain. Click the name to access the home page of that JBoss Application Server.
Type	Type of the member.
Status	Current status of the member. Click the status icon to see a consolidated availability summary. You can see the current and past availability status within the last 24 hours, 7 days, or month (31 days).
Incidents	Number of critical, warning, and error alerts generated for the past 24 hours. Click the alert links to drill down and see more detailed information.

To search for a particular member, use the **Search** menu.

By default, all members of the JBoss Partition are listed in the table. To refresh the table and view only a particular type of members, select either **Direct Members** or **Indirect Members** from the **View** section.

To capture the membership configuration details in a spreadsheet, click **Export**.

33.7.5 Refreshing JBoss Domains

Enterprise Manager Cloud Control allows you to refresh the membership of a JBoss Domain so that it can reflect the current deployment state. This helps you add additional JBoss Application Servers to the existing JBoss Domain.

To refresh a JBoss Domain, follow these steps:

1. From the **Targets** menu, click **Middleware**.
2. On the Middleware page, click the desired JBoss Domain target.
3. On the JBoss Domain Home page, click the JBoss Domain menu and select **Refresh JBoss Domain**.

Enterprise Manager Cloud Control takes you to the Refresh Domain home page.

4. Select the JBoss Application Servers you want to add and click **Submit**.

33.8 Monitoring JBoss Partitions

This section covers the following:

- [Monitoring JBoss Partitions](#)
- [Administering JBoss Partitions](#)
- [Viewing JBoss Partition Members](#)
- [Refreshing JBoss Partitions](#)

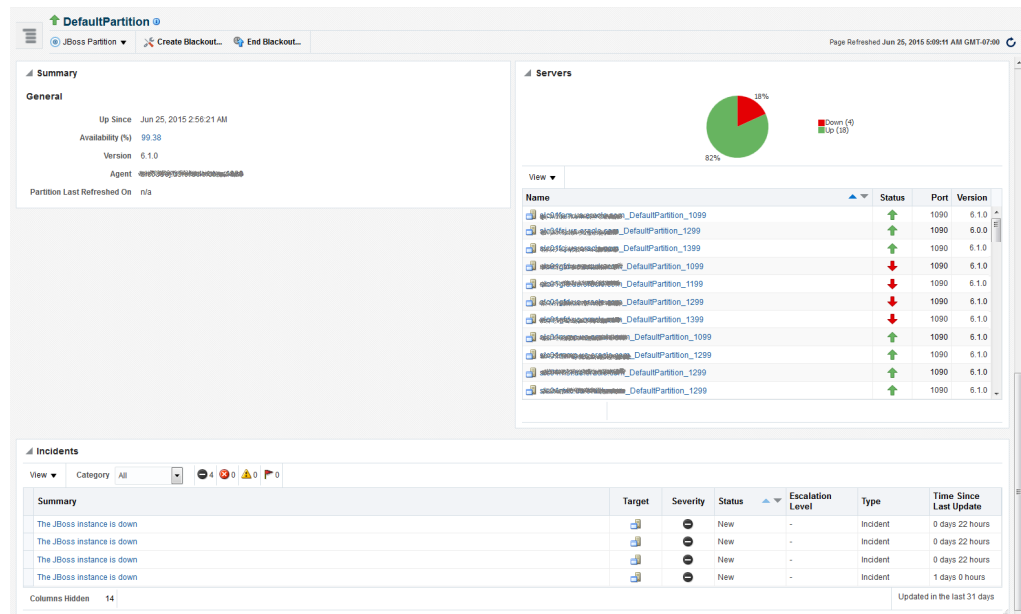
33.8.1 Monitoring JBoss Partitions

To monitor JBoss Partitions, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, click the desired **JBoss Partition**.
3. On the JBoss Partition Home page, you can do the following:
 - View a summary of the most critical information pertaining to the JBoss Partition.
 - Monitor the status and availability of all the members within the JBoss Partition.
 - View the incidents reported on the partition.

Using the JBoss Partition Home page, you can not only monitor the collective status of the partition but also the individual status of each of the members. You can also refresh the partition to update the membership and reflect the current deployment state. You can also view the member application servers' resource usage, availability, performance, configuration information, and reports with or without historical data.

Figure 33-11 JBoss Partition Home Page



The JBoss Partition Home page has the following sections:

- [Summary Section](#)
- [Servers Section](#)
- [Incidents Section](#)

33.8.1.1 Summary Section

The General sub-section in the Summary section provides general information about JBoss Partition.

Element	Description
Up/Down/Pending Since	Date and time when the status was last determined.
Availability (%)	Availability rate for the last 24 hours, considering the status of all the members of the JBoss Partition. For example, if there are four JBoss Application Servers within a partition, and if only three of them are up, then the pie chart shows 75% up and 25% down status.
Version	Version of the JBoss Application Servers that are part of the JBoss Partition.
Agent	Management Agent used for discovering the JBoss Application Servers that are part of the JBoss Partition. To drill down to the Management Agent home page, click the link.
Partition Last Refreshed On	Date and time when the partition was last refreshed.

33.8.1.2 Servers Section

The Servers section provides a real-time view of the status and availability of all the members within the JBoss Partition. For example, if there are four JBoss Application Servers within a JBoss Partition, and if three of these are up, then the pie chart shows 75% up and 25% down status. Accordingly, the legend shows 3 against the Up status to indicate the JBoss Application Servers that are up, and 1 against the Down status to indicate the server that is down.

The table provides high-level details of the JBoss Application Servers that are part of the JBoss Partition. To drill down and view information about a JBoss Application Server, click the JBoss Application Server name. To view more information about the status, click the status icon.

33.8.1.3 Incidents Section

The Incidents section provides a summary of all the incidents reported on the partition. Use the **View** menu to sort, filter and organize the Incidents table. Click the incident for further information related to each incident.

33.8.2 Administering JBoss Partitions

To administer JBoss Partitions, follow these steps:

1. From the **Targets** menu, click **Middleware**.
2. On the Middleware page, click the desired JBoss Partition target.
3. On the JBoss Partition home page, you can view high-level information pertaining to the selected JBoss Partition.

To perform administrative tasks on the JBoss Partition, from the **JBoss Partition** menu, select any of the following according to your needs:

- **Monitoring**, to monitor the performance of the target, view metric details, view status information, view incidents and alerts raised so far for the target, and view blackouts and notification blackouts created for the target.
- **Diagnostics**, to analyze and diagnose performance issues.

- **Control**, to create or end blackouts and notification blackouts.
- **Information Publisher Reports**, to view reports.
- **Members**, to view details of the JBoss Application Servers that are part of the JBoss Partition.
- **Refresh JBoss Partition**, to refresh the JBoss Partition and to add any new JBoss Application Servers added to the partition.
- **Configuration**, to search, view, and compare configuration details in JBoss version 6.
- **Compliance**, to view and create compliance standards.
- **Target Setup**, to view monitoring configuration details and target properties, to remove the target or add it to a group, to migrate and use JMX.
- **Target Sitemap**, to view the overall topology of the target.
- **Target Information**, to view general information about the target.

33.8.3 Viewing JBoss Partition Members

Enterprise Manager Cloud Control helps you view the members of a JBoss Partition. You can see what type of members form the partition, monitor their status, and perform various administrative operations

To view a list of members, follow these steps:

1. From the **Targets** menu, click **Middleware**.
2. On the Middleware page, click the desired JBoss Partition target.
3. On the JBoss Partition Home page, from the **JBoss Partition** menu, select **Members**, then select **Show All** to view the following details of the members.

Column	Description
Name	Name of the JBoss Application Server that is part of the JBoss Partition. Click the name to access the home page of that JBoss Application Server.
Type	Type of the member.
Status	Current status of the member. Click the status icon to see a consolidated availability summary. You can see the current and past availability status within the last 24 hours, 7 days, or month (31 days).
Incidents	Number of critical, warning, and error alerts generated for the past 24 hours. Click the alert links to drill down and see more detailed information.

To search for a particular member, use the **Search** menu.

By default, all members of the JBoss Partition are listed in the table. To refresh the table and view only a particular type of members, select either **Direct Members** or **Indirect Members** from the **View** section.

To capture the membership configuration details in a spreadsheet, click **Export**.

33.8.4 Refreshing JBoss Partitions

Enterprise Manager Cloud Control allows you to refresh the membership of a JBoss Partition so that it can reflect the current deployment state. This helps you add additional JBoss Application Servers to the existing JBoss Partition.

To refresh a JBoss Partition, follow these steps:

1. From the **Targets** menu, click **Middleware**.
2. On the Middleware page, click the desired JBoss Partition target.
3. On the JBoss Partition home page, click the JBoss Partition menu and select **Refresh JBoss Partition**.

Enterprise Manager Cloud Control takes you to the Refresh Partition page.

4. On the Refresh Partition page, select the JBoss Application Servers that should be added to the JBoss Partition, and click **Submit**.

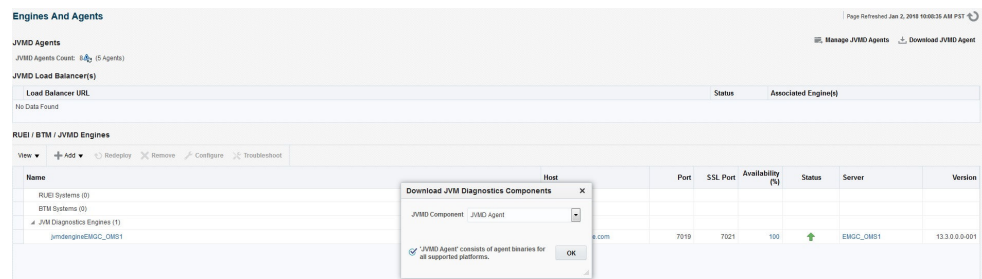
33.9 Deploying JVMD on JBoss Application Server 7.x and 6.x to Diagnose Issues

Oracle Enterprise Manager Cloud Control 13c's JVM Diagnostics enables administrators to diagnose performance problems in a 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. The correlation between the JBoss target and the JVMD/JVM target enables administrators to navigate to the JVM in context of a JBoss Application Server.

To deploy JVMD agents to JBoss, complete the following:

1. User can download `jamagent.war` from **Setup--> Middleware Management --> Engines And Agents--> Download Jvmd Agent**.

Figure 33-12 Download JVM Diagnostics Components



2. From the **Download JVM Diagnostics Components** window, select **JVMD Agent**, then click **OK**.
3. From the JVM Diagnostics Agent web.xml Parameters window, select the appropriate JVMD manager from the Available Managers list.

Figure 33-13 JVM Diagnostics Agent web.xml Parameters

JVM Diagnostics Agent web.xml Parameters

Configure Hybrid Cloud Agent

Available Engines

Tuning/Timeouts Parameters

Connection Retry Time (secs) Monitoring Enabled

Long Request Timeout (secs)

GC Wait Timeout (secs)

Idle Agent Timeout (secs)

Target Association Parameters

Group ID Property

Pool Name

JVM ID

Library Location

JAM Lib Directory

JAM Lib Win Directory

Logging Parameters

Agent Log Level

JVMMD Agent deployment to a server running on IBM JDK requires additional manual steps. Copy the libjamcapability.so library on the server host, and restart the server with the following java option added -agentpath:<absolute_path_to_libjamcapability.so>. To include -agentpath in the Java option, check deploy container/server documentation.

4. 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 where the JBoss is running).

5. [Optional] Copy the following script to the same location on the target machine where the jamagent.war. file was copied:

```
#!/bin/sh

# Edit this if jamagent.war is not in the current dir & point it to its new
location

WARFILE="jamagent.war"

#UNPACK="jar xf"

UNPACK="unzip -qo"

#PACK="jar uf"

PACK="zip -q"

#if no arguments are passed, print usage statement & exit

[ -n "$1" ] || ( echo "Usage: ${0} [node1] [node2] ... [node n]"; exit 0 ; )

WEBXMLFILE="WEB-INF/web.xml"

SINGLEINDENT="\ \ \ \ \ \ \ \ "

DOUBLEINDENT="\ \ \ \ \ \ \ \ \ \ \ \ "

INSERTPREFIX=

"${SINGLEINDENT}<init-param>\n${DOUBLEINDENT}<param-name>jamjvmid</param-name>\n${
DOUBLEINDENT}<param-value>"

INSERTSUFFIX="</param-name>\n${DOUBLEINDENT}<description>The ID of this
JVM</description>\n${SINGLEINDENT}</init-param>"

INSERTBEFORE="<load-on-startup>"

# Iterate over each node name (i.e. command line argument)

for NODE_NAME in "$@"

do

    echo "Processing ${NODE_NAME}"

    # Make a copy of the war file that is node specific

    cp ${WARFILE} ${NODE_NAME}-${WARFILE}

    # Extract the web.xml file

    ${UNPACK} ${NODE_NAME}-${WARFILE} ${WEBXMLFILE}

    # Construct the XML block to be inserted using the node name
```

```

INSERTBLOCK=${INSERTPREFIX}${NODE_NAME}${INSERTSUFFIX}

# Insert the xml block into web.xml

sed -i "${INSERTBEFORE}/i ${INSERTBLOCK}" ${WEBXMLFILE}

# Add the web.xml back into the war file

${PACK} ${NODE_NAME}-${WARFILE} ${WEBXMLFILE}

#tail -15 ${WEBXMLFILE}

# Cleanup all the temp stuff

rm ${WEBXMLFILE}

rmdir WEB-INF/

# TIP: You could try to deploy the war file directly from here using something
like

#cp ${NODE_NAME}-${WARFILE} ${DEPLOYMENT-DIR-OF-THIS-NODE}

done

```

6. 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.

7. 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.

If you are using JBoss AS 7 modular class loading model, you must perform following configuration changes to make JVMD agent work:

1. Enable `sun.instrument` package loading as follows:

Add `<path name="sun/instrument"/>` in file `$(JBOSS_HOME)/modules/system/layers/base/sun/jdk/main/module.xml`

2. Load BCI and JVMD classes:

Add the `jvmd.jar` and `wldf.jar` to the class path as follows:

```
JAVA_OPTS="$JAVA_OPTS -cp :$(JBOSS_HOME)/wldf.jar,$JBOSS_HOME/jvmd.jar"
```

The `wldf.jar` and `jvmd.jar` are bundled with `jamagent.war`. You can download the `jamagent.war` from Enterprise Manager Cloud Control console (**Setup--> Middleware Management --> Engines And Agents--> Download Jvmd Agent**).

After you download the `jamagent.war` file, you can extract the `jamagent.war` and go to `jamagent.war/WEB-INF/libbci/` and copy both `wldf.jar` and `jvmd.jar` files at required location.

For example, in case of standalone server modify `$(JBOSS_HOME)/bin/standalone.conf` as follows:

```

if [ "x$JBASS_MODULES_SYSTEM_PKGS" = "x" ]; then
JBASS_MODULES_SYSTEM_PKGS="org.jboss.byteman,com.oracle.jvmd.repackaged,oracle
.jvmd.agent"
fi
JAVA_OPTS="$JAVA_OPTS -cp :$JBASS_HOME/wldf.jar,$JBASS_HOME/jvmd.jar"

```

33.10 Troubleshooting JBoss Application Server Discovery and Monitoring Issues

This section provides troubleshooting tips for the issues encountered while discovering or monitoring JBoss Application Servers.

- [Troubleshooting Monitoring Issues](#)
- [Troubleshooting Discovery Issues](#)
- [Additional Useful Resources](#)

33.10.1 Troubleshooting Monitoring Issues

- If the target status is DOWN after discovery, check the Monitoring properties page for the JBoss Application Server and verify the following:
 - Local discovery: only JBoss home is present
 - Remote discovery: only Library path is present
- 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 Management Agent (such as Weblogic).

The following are the various log locations:

- JBoss server logs: `$JBASS_HOME/server/<config_mode>/log`
- OMS logs: `emoms.trc` (under `$OMS_HOME`)
- Agent logs: `$AGENT_STATE_DIR/sysman/log`

33.10.2 Troubleshooting Discovery Issues

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=<AGENT_PLUGIN_LOCATION> \
-cp \
<AGENT_PLUGIN_LOCATION>/lib/xmlparserv2.jar:\
<AGENT_PLUGIN_LOCATION>/jlib/emConfigInstall.jar:\
<AGENT_PLUGIN_LOCATION>/sysman/jlib/log4j-core.jar:\
<AGENT_PLUGIN_LOCATION>/modules/oracle.http_client.11.1.1.jar:\
<DISCOVERY_PLUGIN_LOCATION>/archives/em-as-thirdparty-discovery.jar \
oracle.sysman.emas.thirdparty.discovery.jboss.JBossDiscovery \
<JMX_PORT> <SERVER_HOST> " "

```


33.10.3 Additional Useful 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](#)
- [JConsole](#)
- [JBoss server logs](#)

34

Discovering and Monitoring Apache HTTP Server

Enterprise Manager Cloud Control enables you to discover Apache HTTP Servers in your environment, and add them for central monitoring and management. This chapter describes how to discover and monitor these Apache HTTP Server targets.

In particular, this chapter covers the following topics:

- [Introduction to HTTP Servers](#)
- [Supported Versions of Apache HTTP Server for Discovery and Monitoring](#)
- [Prerequisites for Discovering and Monitoring Apache HTTP Server](#)
- [Discovering Apache HTTP Servers](#)
- [Monitoring Apache HTTP Servers](#)
- [Configuration Management for Apache HTTP Servers](#)
- [Troubleshooting Apache HTTP Server Issues](#)

34.1 Introduction to HTTP Servers

Using Enterprise Manager Cloud Control, you can do the following with Apache HTTP Server targets:

- Discover the Apache HTTP Server targets for real-time and historical availability monitoring.
- Create or end blackouts and notification blackouts to suspend or resume the collection of metric data, respectively.
- View a list of metrics, their collection interval, and the last upload for each metric.
- Create monitoring templates that can be used as a source for all the future installations, so that they follow a standard, consistent configuration.
- Generate availability and event reports.

34.2 Supported Versions of Apache HTTP Server for Discovery and Monitoring

To search for the Apache HTTP Server versions that are supported for discovery and monitoring in Enterprise Manager Cloud Control, follow these steps:

1. Log in to <https://support.oracle.com/>.
2. On the My Oracle Support home page, select the **Certifications** tab.
3. On the Certifications page, enter the following search criteria in the Certification Search section.

- Enter the product name **Enterprise Manager Base Platform - OMS** in the Product field.
 - Select the appropriate release number from the Release list.
4. Click **Search**.
 5. In the Certification Results section, expand the **Middleware** menu to view the certified Apache HTTP Server versions.

Certified With	Number of Releases / Versions
> Operating Systems (9 Items)	
> Agents (1 Item)	
> Application Servers (10 Items)	
> Databases (9 Items)	
> Desktop Applications, Browsers and Clients (5 Items)	
> Directory/LDAP Services (5 Items)	
> Enterprise Applications (11 Items)	
> Management and Development Tools (60 Items)	
▼ Middleware (28 Items)	
Apache HTTP Server (Managed Target)	1 Release (2017.007.000)
Exalogic Elastic Cloud Software (Managed Target)	1 Release (2016.003.000, 2016.004.000, 2016.005.000, 2016.006.000)
IBM WebSphere MQ (Managed Target)	1 Release (2017.007.000)
IBM WebSphere Portal Server (Managed Target)	1 Release (2016.003.000, 2016.004.000)
Microsoft .NET Framework (Managed Target)	1 Release (2017.007.000)
Oracle Access Manager (Managed Target)	7 Releases (2016.003.000, 2016.004.000, 2016.005.000, 2016.006.000, 2016.007.000, 2016.008.000, 2016.009.000)
Oracle Adaptive Access Manager (Managed Target)	1 Release (2016.003.000, 2016.004.000, 2016.005.000)
Oracle BPPEL Process Manager (Managed Target)	1 Release (2016.003.000, 2016.004.000, 2016.005.000)
Oracle Business Intelligence Enterprise Edition (Managed Target)	1 Release (2016.003.000, 2016.004.000, 2016.005.000, 2016.006.000)
Oracle Business Intelligence Publisher (Infrastructure)	1 Release (2016.003.000)
Oracle Business Process Management (Managed Target)	1 Release (2016.003.000)
Oracle Coherence (Managed Target)	1 Release (2016.003.000, 2016.004.000, 2016.005.000, 2016.006.000, 2016.007.000, 2016.008.000)
Oracle Data Integrator Agent (Managed Target)	1 Release (2016.003.000, 2016.004.000, 2016.005.000)
Oracle Enterprise Content Management (Managed Target)	1 Release (2016.003.000, 2016.004.000, 2016.005.000, 2016.006.000, 2016.007.000, 2016.008.000)
Oracle Forms (Managed Target)	1 Release (2016.003.000, 2016.004.000, 2016.005.000, 2016.006.000, 2016.007.000, 2016.008.000)
Oracle Fusion Middleware (Infrastructure)	1 Release (2016.003.000)
Oracle Fusion Middleware 12c: Infrastructure (Managed Target)	1 Release (2016.003.000)
Oracle HTTP Server (Managed Target)	1 Release (2016.003.000, 2016.004.000, 2016.005.000, 2016.006.000, 2016.007.000, 2016.008.000, 2016.009.000)
Oracle Identity Management (Managed Target)	1 Release (2016.003.000, 2016.004.000, 2016.005.000, 2016.006.000, 2016.007.000, 2016.008.000)

34.3 Prerequisites for Discovering and Monitoring Apache HTTP Server

Meet the following prerequisites for discovering Apache HTTP Servers:

- The Management Agent must be installed and running on the same host where the Apache HTTP Server is being configured. Remote agent is not supported.
- Ensure that the same user/role is used to install the Management Agent and the Apache HTTP Server.

34.4 Discovering Apache HTTP Servers

To discover Apache HTTP Server Servers, follow these steps:

1. In Cloud Control, from **Setup** menu, select **Add Target**, then select **Add Targets Manually**.
2. On the Add Targets Manually page, click **Add Targets Declaratively**.
3. On the Add Targets Declaratively page enter the **Host** and select **Apache HTTP Server** from the Target Type table, and then click **Add**.
4. On the Add: Apache HTTP Server page, provide the target name, the directory location where the `httpd.conf` file has been installed, and the directory location where the Apache binaries (like the bin folder) are stored. Click **OK**.

Add: Apache HTTP Server

Add a target to be monitored by Enterprise Manager by specifying target monitoring properties

Target

* Target Name

Target Type Apache HTTP Server

Host

Agent

Properties

* Absolute path of httpd.conf

* Apache Binaries Home

▶ Global Properties

34.5 Monitoring Apache HTTP Servers

After adding the Apache HTTP Server target, it becomes automatically available for monitoring. For this target, only the response metrics and configuration metrics are collected or monitored.

After discovery, to access the Apache HTTP Server targets, from **Targets** menu, select **All targets**. From the Refine Search section on the left hand pane, expand **Middleware**. From the list, select **Apache HTTP Server**. Click on the target name to view the status of the target.

↑
Apache_24
i

Apache HTTP Server
▼

Home

General

↑

Status Up Blackout

Availability (%) 98
(Last 24 Hours)

Apache Home /usr/lib/httpd/bin

Version 2.4.20 (Unix)

Host 192.168.1.100

Incidents

Severity	Message	Created At	Last Updated At
No incidents found.			

Host Incidents

Metric Collection Errors 7

Severity	Message	Created At	Last Updated At
No incidents found.			

Home

On the Apache HTTP Server home page, you can view general information about the server, information about the status of the server, the availability, the absolute path to the Apache server binaries, and so on.

34.6 Configuration Management for Apache HTTP Servers

The configuration data for the Apache HTTP server is collected on a daily basis.

To view the configuration data, on the Apache HTTP Server home page, from **Apache HTTP Server** menu, select **Configuration**, and then click **Last Collected**.

The screenshot shows the 'Latest Configuration' view for 'Apache_24'. The left sidebar lists various configuration categories like General, Routing Information, Performance Related, etc. The main content area displays 'Configuration Properties' for 'Apache_24', including a table of property names and values, and a 'Last collected at' timestamp of Jun 1, 2016 1:31:12 PM.

Property Name	Property Value
Target Version	2.4.20 (Unix)
Operating System	Linux
Platform	x86_64
Absolute path of httpd.conf	/usr/local/apache2/conf/httpd.conf
Apache Binaries Home	/usr/local/apache2/bin

The following configuration details are collected for Apache HTTP server:

- Generic information like server name, listen port, and so on.
- General Routing information for WebLogic/WebSphere requests.
- Apache Server listen host ports and protocol.
- Virtual host information which is used for routing the requests that come to Apache Server to particular host port.

34.7 Troubleshooting Apache HTTP Server Issues

Issue: Response and Configuration Metrics collection for Apache HTTP Server fails.

Problem: If the process owner (Apache installation owner) is different from Management Agent user, then Apache HTTP Server target will be discovered, but the response and configuration metrics will not be collected.

Workaround: Ensure that the same user/role is used to install the Management Agent and the Apache HTTP Server.

 **Note:**

The file which the Management Agent accesses to draw information is `httpd.conf`.

Part XI

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:

- [Configuring and Monitoring Oracle Data Integrator](#)

Configuring and Monitoring Oracle Data Integrator

This section describes Oracle Data Integrator (ODI). ODI as a part of Enterprise Manager Cloud Control 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.

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 following ODI components that can be managed using Enterprise Manager Cloud Control.

- One Master and one or more Work repositories attached to it.
- 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 (ODI 11g), Colocated Standalone Agents (ODI 12c), or Java EE Agents (ODI 11g or 12c).
- One or several Oracle Data Integrator Console applications. An Oracle Data Integrator Console application is used to browse Master and Work repositories.

This chapter describes how you can set up and manage ODI targets using Enterprise Manager Cloud Control:

- [Prerequisites for Monitoring Oracle Data Integrator](#)
- [Monitoring Oracle Data Integrator](#)
- [Administering Oracle Data Integrator](#)
- [Creating Alerts and Notifications](#)
- [Monitoring Run-Time Agents](#)
- [Configuring Oracle Data Integrator Console](#)
- [Configuring an Oracle Data Integrator Domain](#)

35.1 Prerequisites for Monitoring Oracle Data Integrator

Before you start managing ODI with Enterprise Manager, you must do the following:

- Deploy the Oracle Management Agent

Oracle Management Agents must be installed on the servers which have the databases hosting the ODI repositories. Optionally, an Oracle Management Agent can also be installed on a machine hosting an ODI Agent.

For more information, see Installing the Oracle Management Agent in the *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.

- Discover ODI Targets

ODI targets are discovered along with the WebLogic domain linked to them. Use the Fusion Middleware discovery to discover your WebLogic domain. This in turn discovers three types of ODI targets, ODI Standalone Agent (ODI 11g), ODI Colocated Standalone Agent (ODI 12c), and ODI Java EE Agent (ODI 11g or 12c).

For more information, see Fusion Middleware discovery in *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

- Discover the Databases Hosting ODI Repositories

All database instances need to be discovered since more than one database could be hosting the ODI repositories.

See Oracle® Enterprise Manager Cloud Control Advanced Installation and Configuration Guide.

- User login credentials should be setup in the Enterprise Manager console.

All the operations are available out-of-box in Enterprise Manager.



Note:

ODI supports the discovery of the following data servers: Oracle, Microsoft SQL Server, and IBM DB2 UDB.

35.2 Monitoring Oracle Data Integrator

This section describes the following:

- [Monitoring Oracle Data Integrator](#)
- [Monitoring ODI Agents](#)
- [Monitoring Repositories](#)
- [Monitoring Load Plan Executions and Sessions](#)

35.2.1 Monitoring Oracle Data Integrator

To monitor ODI, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, from the **Middleware Features** menu, select **ODI Home**.
3. On the ODI Home page, click the **Dashboard** tab.

The Dashboard page displays a seven-day outage view for all the objects on the page when the main Up/Down nodes are expanded. There are seven squares. If the square is green, there were no alerts that day. If the square is red, there was an alert.

The Dashboard tab has the following regions:

35.2.1.1 Master Repositories Health

This region reports the following:

- Number of master repositories that are either up or down. Click the number for a list of the repositories.
- Number of master repositories with incidents. Click the number to find out which repositories have incidents.

 **Note:**

Starting with Oracle Fusion Middleware Plug-in (12.1.0.6), you can monitor the repositories that are configured even with Microsoft SQL Server and IBM DB2. However, as a prerequisite, make sure you first deploy the Microsoft SQL Server Plug-in and IBM DB2 Plug-in, respectively, and then discover those database instances as targets in Enterprise Manager Cloud Control.

The database information that is stored in the ODI does not use local host or IP address to identify the database. It only uses the host name of the database. The host name is derived from the URL of the application. Ensure that the host name in the ODI is consistent with the host name stored in EMCC. Also, check the JDBC data sources defined in WLS for the Master and Work repositories. They should match the information stored in the ODI.

The supported JDBC patterns are:

- `jdbc:oracle:thin:@//adc2120612.us.example.com:19016/db8482.us.example.com`
- `jdbc:oracle:thin:@adc2120612.us.example.com:19016:db8482`
- `jdbc:weblogic:sqlserver://adc6140804.us.example.com:50457;databaseName=ODI_REPOSITORY`
- `jdbc:weblogic:db2://slc02pfl.us.example.com:5031/orcl993`

To resolve issues reported in this section:

- If the ODI repositories are down, then act based on the statuses by either bringing up the databases, which are hosting the repositories, or troubleshooting why they are down and resolving the issues.
- If there are any repositories that are undiscovered, then discover the databases, which are hosting the repositories, in Enterprise Manager Cloud Control.
- If there are any repositories with alerts, then identify the root cause for those alerts and resolve the issues.

35.2.1.2 ODI Agents Health

This region reports the following:

- 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, then act based on the statuses by either bringing up the Agents, which are down, or troubleshooting why they are down and resolving the issues.
- If there are any Agents that are undiscovered, then either discover the Agents or refresh the Oracle WebLogic Domain that is linked to those Agents.
- If there are any Agents with alerts, then identify the root cause for those alerts and resolve the issues.

35.2.1.3 Work Repositories Health

This region reports the following:

- 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 in Enterprise Manager. Click the number of a list of the work repositories that have not been discovered.
- Number of 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, then act based on the statuses by either bringing up the databases, which are hosting the repositories, or troubleshooting why they are down and resolving the issues.
- If there are any repositories that are undiscovered, then discover the repositories in Enterprise Manager Cloud Control.
- If there are any repositories with alerts, then identify the root cause for those alerts and resolve the issues.

35.2.1.4 Data Servers Health

This region reports the following:

- 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 in Enterprise Manager. Click the number of a list of the data servers that have not been discovered.
- Number of 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, then act based on the statuses by either bringing up the databases used by the data servers, or troubleshooting why they are down and resolving the issues.

- If there are any data servers that are undiscovered, then discover the databases, which are used by the data servers, in Enterprise Manager Cloud Control.
- If there are any data servers with alerts, then identify the root cause for those alerts and resolve the issues.

35.2.1.5 Sessions/Load Plan Executions

This region reports the following:

- Number of sessions in error across all discovered ODI environments.
- Number of sessions with error records across all discovered ODI environments.
- Number of load plan executions in error across all discovered ODI environments.
- Number of load plan executions with error records across all discovered ODI environments.

35.2.2 Monitoring ODI Agents

To monitor the ODI Agents, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, from the **Middleware Features** menu, select **ODI Home**.
3. On the ODI Home page, click the **ODI Agents** tab.

The ODI Agents tab has the following regions:

35.2.2.1 Search Agents

Use this region to search for agents for all Java EE, Colocated Standalone, 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.

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.

35.2.2.2 ODI Agents

Use this region to view information about the ODI 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.
Status	Displays the current status of the Agent: Up, Down.
Discovery Status	A blue tick indicates that the Agent is discovered as a target in Enterprise Manager. A clock indicates that the Agent is not discovered as a 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 • Number of queued sessions • Number of running sessions • Number of waiting sessions • Number of successful sessions • Number of failed sessions
Active Sessions	Displays the number of active sessions.
Master Repository	A check mark indicates that the Master Repository is discovered. A clock indicates that the Master Repository is not discovered.
Version	Displays the version and date of the Agent.
Response Time (ms)	Displays the repository database response time (in milliseconds).
Create Alert	Redirects users to the Metric and Collection Settings page of a particular agent where they can set up their alerts.
User Defined Alerts	Displays the number of Critical and Warning alerts. Click the number to view the alerts in the Incident Manager page.

35.2.3 Monitoring Repositories

To monitor the ODI repositories, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, from the **Middleware Features** menu, select **ODI Home**.
3. On the ODI Home page, click the **Repositories** tab.

 **Note:**

- The ODI database credentials have to be selected for this region to display. There are different credentials for different repositories. Choose the credentials based on your need.
- Starting with Oracle Fusion Middleware Plug-in 13.1.1.0, you can also monitor RAC databases along with other databases. However, for other repositories that are configured with Microsoft SQL Server and IBM DB2, as a prerequisite, make sure you first deploy the Microsoft SQL Server Plug-in and IBM DB2 Plug-in, respectively, and then discover those database instances as targets in Enterprise Manager Cloud Control.

The Repositories tab has the following regions:

35.2.3.1 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.

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.

35.2.3.2 Repositories

Use this region to view details of the work repositories.

Element	Description
Name	<p>Displays the name of the Master and Work Repository. A star icon against the name of the repository indicates that it is a non-Oracle Database repository.</p> <ul style="list-style-type: none"> • To view the Work Repositories under a particular Master Repository, expand the Master Repository name. • To drill down and access the respective database home page for more details, click the repository name. • For more details on a particular repository, select the row of that repository to see the Database Details table appear. For non-Oracle Database repositories, Enterprise Manager Cloud Control might not be able to display data for all the metrics.

Element	Description
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.
Technology	Displays the technology used.
Host	Displays the name of the host on which the repository resides.
Port	Displays the port of the host on which the repository resides.
SID/Database Instance	Displays the system identifier of the repository or the database instance name.
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/File Group	Displays the total rows and segment size (in GB).
Purge	Click the icon to purge the ODI logs. <ul style="list-style-type: none"> For 12.1.3 ODI Agents monitored with Oracle Fusion Middleware Plug-in (12.1.0.6), a separate dialog appears where you can provide the required information, and click Purge. The ODI logs will be deleted from within the Enterprise Manager Cloud Control Console. For 12.1.3 ODI Agents monitored with Oracle Fusion Middleware Plug-in (12.1.0.5) or lower, and for all 12.1.2 or lower ODI Agents, a separate browser window with the ODI Console appears. Log in to the console, and delete the unwanted ODI logs.

35.2.3.3 Cluster Databases

The Cluster Databases region is displayed only if the Repository happens to have a cluster database. This section provided details of the databases within the selected cluster.

Element	Description
Name	Displays the name of the database.
Status	Displays the status of the database. <ul style="list-style-type: none"> Up (green arrow): on Down (red arrow): off
Host	Displays the name of the host on which the database resides.
SID	Database SID.

35.2.3.4 Database Details

By looking at the database details, you have a clear picture of how your database is performing. For example, if the database tablespace 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: For this region to appear, you must select the credentials and the repository. The credentials must be of a DBA user and must be of the type *Global*. The credentials are required to depict the tablespace and schema-related charts.

 **Note:**

For non-Oracle Database repositories, Enterprise Manager Cloud Control might not be able to display data for all the metrics

35.2.3.5 Tablespace/File Group 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.

 **Note:**

For non-Oracle Database repositories, Enterprise Manager Cloud Control might not be able to display data for all the metrics.

35.2.4 Monitoring Load Plan Executions and Sessions

To monitor the load plan executions and sessions, follow these steps:

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

2. On the Middleware page, from the **Middleware Features** menu, select **ODI Home**.
3. On the ODI Home page, click the **Load Plan Executions/Sessions** tab.

The Load Plan Executions/Sessions tab enables you to search and view information about the load plan executions and sessions executed by the Agent. This tab has the following regions:

Expand a session and review the Steps and Tasks information. For example if an ODI Mapping was executed, you can review each task that this mapping executed, view the generated code, and drill down to the database execution details.



Note:

Oracle Database Diagnostics and Tuning Packs are required to be able to use the Database Execution Details link and drill down into the Oracle Database monitoring pages.

The Load Plan Executions/Sessions tab has the following regions.

35.2.4.1 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.

Element	Description
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.

35.2.4.2 Load Plan Executions/Sessions

Use this region to view execution details of the Load Plan Executions and Sessions executed by the Agent.

To view more details such as hierarchy, status of each step, the start and end time of each step, and so on, for a particular Load Plan Execution or Session, select the row in the table and scroll down the page to see the Load Plan Executions/Session Detail table.

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 an icon to indicate the status of the Load Plan Execution run or Session executed. Hover your mouse over the icon to understand the status and view more details if there is an error. The status can be one of the following: <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.
Error Records	Displays the number of error records.
Execution Type	Displays the Load Plan or Sessions type, for example, Scenario.
Work Repository Name	Displays the name of the Work Repository into which this Load Plan/Session run execution information is stored.

Element	Description
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.

35.2.4.3 Load Plan Executions/Session Detail

Use this region to view more detailed information on the Load Plan Executions and Sessions executed by the Agent.

Element	Description
Load Plan Executions/ Session Hierarchy	Displays the hierarchy of the Load Plan Execution or Session. Click and expand the Load Plan Execution or Session name to view the complete hierarchy.
Status	Displays an icon to indicate the status of the Load Plan Execution or Session step. Hover your mouse over the icon to understand the status and view more details if there is an error.
Source Code	Displays the code executed on the source database. Click the icon to view details of the executed code. If the source and target databases are Oracle Databases, which have been discovered in Enterprise Manager Cloud Control, then you will see a Database Execution Details hyperlink. Click the link to drill down to the ASH Analytics page and view information about the active sessions run for a particular time period.
Target Code	Displays the code executed in the target database. Click the icon to view details of the executed code. If the source and target databases are Oracle Databases, which have been discovered in Enterprise Manager Cloud Control, then you will see a Database Execution Details hyperlink. Click the link to drill down to the ASH Analytics page and view information about the active sessions run for a particular time period.
Step Task Type	Displays the type of task performed by the step. The task type value is a hyperlink when the source and target systems are database systems. In that case, click the task type to view details of the source database and the target database that exchanged data.
Started On	Displays the date and time when the step started.
Ended On	Displays the date and time when the step ended.
Duration	Displays the time taken (in seconds) to execute the task.
Updates	Displays the number of updates or changes done to a row per task.
Inserts	Displays the number of data insertions done per task.
Error Records	Displays the number of error records reported per task.
Deletes	Displays the number of data deletions done per task.

35.3 Administering Oracle Data Integrator

You can perform the following operations while administering Oracle Data Integrator:

- [Starting Up, Shutting Down, and Restarting Oracle Data Integrator Agents](#)
- [Managing Agent Status and Activities](#)
- [Searching Sessions and Load Plan Executions](#)
- [Viewing Log Messages](#)

35.3.1 Starting Up, Shutting Down, and Restarting Oracle Data Integrator Agents

 **Note:**

- Oracle Process Manager and Notification (OPMN) is used for release 11g Standalone Agents. WebLogic Management Framework is used for release 12c Colocated Standalone Agents only.
- Only *Start* and *Stop* operations are supported for ODI Java EE Agents.
- *Start* and *Stop* operations are supported for all ODI Standalone Agents managed by WebLogic Management Framework and OPMN instances. *Restart* operation is supported only for 11g Standalone Agents managed by OPMN instances, and not for 12c Colocated Standalone Agents managed by WebLogic Management Framework instances.

To start, stop, and restart Oracle Data Integrator Agents, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, from the **Middleware Features** menu, select **ODI Home**.
3. On the Oracle Data Integrator Home page, click the **ODI Agents** tab.
4. In the ODI Agents tab, search for the ODI agents. Then, in the ODI Agents table, click the name of an Agent.
5. On the ODI Agent Home page, from the **ODI Agent** menu, select **Control**, then select either **Start Up**, **Shut Down**, or **Restart**.

 **Note:**

If you want to start or stop ODI Standalone Agents, that are not managed by OPMN or WebLogic Management Framework, you must use the Agent's startup and shutdown scripts. For more information about how to start and shut down Agents, see Managing Agents in the *Oracle Fusion Middleware Developer's Guide for Oracle Data Integrator*.

35.3.2 Managing Agent Status and Activities

To manage the agent status and monitor its activities, follow these steps:

1. Click the target link corresponding to your JEE, Standalone, or Colocated Standalone Agent either in the target navigation pane or in the ODI Home Page. The Java EE Application Page for this agent appears.
2. From the **Agent Page** menu, select **Monitoring** then select **Performance Summary**.

Enterprise Manager Cloud Control displays the Performance Summary page, which enables you to view and customize the metrics and charts.

35.3.3 Searching Sessions and Load Plan Executions

To sessions and load plan executions, follow these steps:

1. From the **Targets** menu on Enterprise Manager, select **Middleware**.
2. In the Middleware Features menu, select **ODI Home**.
3. Click the **LPE/Sessions** tab. For more information on the tab, click **Help**.

35.3.4 Viewing Log Messages

You can view log messages of Java EE agents in Enterprise Manager Cloud Control.

The steps for this process are:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, from the **Middleware Features** menu, select **ODI Home**.
3. On the Oracle Data Integrator Home page, click the **ODI Agents** tab.
4. In the ODI Agents tab, search for the ODI agents. Then, in the ODI Agents table, click the name of an Agent.
5. On the ODI Agent Home page, 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.

35.4 Creating Alerts and Notifications

For detailed information on alerts and notifications, see *Using Incident Management* and *Using Notifications* chapters in the *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

As an example, to create an alert for the Master Repository status, see the instructions below:

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

2. On the Middleware page, from the **Middleware Features** menu, select **ODI Home**.
3. On the Oracle Data Integrator Home page, click the **ODI Agents** tab.
4. In the ODI Agents tab, search for the ODI agents. Then, in the ODI Agents table, click the name of an Agent.
5. On the ODI Agent Home page, from the **ODI Agent** menu, select **Monitoring**, then select **Metric and Collection Settings**.
6. In the Metric column, expand Master Repositories to see the Status row.
7. In the Critical Threshold text field, in the Status row, enter **0**.

0 indicates that EM will generate an alert when the Master Repository is down, whereas **1** will generate an alert when the Master Repository is up.

 **Note:**

Similarly, you can create warning or critical alerts for other rows mentioned in the Metric column.

35.5 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 WebLogic Management Framework 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 11.1.1.9. WebLogic Management Framework will be used to manage ODI Standalone agents from ODI release 12c and later.

The Management Pack for ODI can monitor and manage the following ODI Agent types:

- 11g: Java EE Agents and Standalone Agents managed by OPMN.
- 12c: Java EE Agents and Collocated Standalone Agents managed by the WebLogic Management Framework.

To access the ODI Agent Home page, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. On the Middleware page, from the **Middleware Features** menu, select **ODI Home**.
3. On the ODI Home page, click the **ODI Agents** tab.

4. In the ODI Agents tab, search for ODI Agents, and in the search results table, click the name of the ODI Agent that interests you.

For further details on the agent home page, see [Agent Home Page](#).

35.5.1 Agent Home Page

The Agent Home page is arranged in the following order:

- [General Info](#)
- [Load](#)
- [Target Incidents](#)
- [LPEs/Sessions Execution Incidents](#)
- [Load Balancing Agents](#)

35.5.1.1 General Info

The General Info region displays general information about this Agent.

Element	Description
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.

35.5.1.2 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.

35.5.1.3 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.

35.5.1.4 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.
ID	Incident ID.
Summary	Summary description of the incident.
Category	Classification of an incident, for example, Error.

35.5.1.5 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.
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. This is a record which did not meet the requirements to be inserted into the target system by ODI. ODI captures these records when moving the data and stores them into an error table. <ul style="list-style-type: none"> Error Records - Records which did not meet the requirements to be inserted into the target system by ODI. ODI captures these records when moving the data and stores them into an error table. 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.

35.6 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 Enterprise Manager Cloud Control to define the linking between Enterprise Manager 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**.

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.

3. 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

35.7 Configuring an Oracle Data Integrator Domain

An Oracle Data Integrator (ODI) domain contains the Oracle Data Integrator 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.

Installing and configuring components for an Oracle Data Integrator domain is described in the *Oracle Fusion Middleware Installing and Configuring Oracle Data Integrator*.

Index

A

accessing
 JVM Diagnostics pages, [20-7](#)

active threads in JVM Diagnostics, [20-36](#)

adding
 domain certificate in Oracle GlassFish Server, [8-1](#)
 files to Support Workbench package, [2-36](#)
 JVM pool to group, [20-13](#)
 JVM to group, [20-35](#)
 more files to incident, [2-35](#)
 Oracle Traffic Director
 target configurations, [6-3](#)
 targets, [6-7](#)
 to Exalogic target, [6-2](#)

administration servers
 ending blackouts after starting up, [2-18](#)
 restart method, [2-19](#)
 restart time limit, [2-18](#)
 restarting, [2-17](#)
 shutdown errors, [2-20](#), [2-21](#)
 shutdown method, [2-19](#)
 shutdown time limit, [2-18](#)
 shutting down, [2-17](#)
 starting up, [2-17](#)
 starting up while the domain is started, [2-18](#)
 startup errors, [2-20](#), [2-21](#)
 startup method, [2-19](#)
 startup time limit, [2-18](#)
 stopping errors, [2-20](#), [2-21](#)
 stopping while the domain is stopped, [2-18](#)

agent deployment, [13-6](#)

agent status
 JVM Diagnostics, [21-10](#)

Agent truststore
 updating, [7-1](#)

Aggregated Diagnostic Summary page, [2-33](#)

alert notifications, [2-11](#)

analyzing
 heat snapshots, [20-29](#)
 JVM Diagnostics snapshot, [20-39](#)

anti-pattern report in JVM Diagnostics, [20-34](#)

Application Replay
 analyzing replay results, [3-21](#)

Application Replay (*continued*)
 creating captures, [3-6](#)
 Importing Divergences, [3-22](#)
 introduction, [3-1](#)
 monitoring capture process, [3-12](#)
 OpenScript, [3-22](#)
 prerequisites, [3-3](#)
 replaying captures, [3-13](#)
 testing against real workloads, [3-1](#)
 troubleshooting, [3-23](#)

Application Server
 extensible monitoring, [2-13](#)
 managing configurations, [2-23](#)

B

blackouts
 creating blackouts before stopping targets, [2-18](#)
 ending blackouts after starting the targets, [2-18](#)
 monitoring, [2-12](#)

boot.properties file, [2-19](#)

business application, [13-2](#)
 creating, [16-10](#)
 home page, [16-15](#)
 key components, [16-2](#)
 monitoring, [16-14](#)
 overview, [16-1](#)
 sample, [16-2](#)
 setting up, [13-9](#)
 target type, [16-2](#)

Business Transaction Management
 ECID, use of, [13-4](#)

C

CA certificate, [7-1](#)
 importing, [7-2](#)

CAs (Certificate Authorities), [7-1](#)
 importing, [7-1](#)

class histograms, viewing, [20-37](#)

clusters, Oracle GlassFish Server, [8-4](#), [8-13](#)

comparing
 class histograms, [20-27](#)

- comparing (*continued*)
 - heat snapshots, [20-33](#)
- compliance management, [2-24](#)
- composite applications, [4-1](#)
 - creating, [4-3](#)
 - dashboard, [4-1](#)
 - editing, [4-4](#)
 - editing home page, [4-5](#)
 - viewing, [4-5](#)
- configuration
 - adding Oracle Traffic Director, [6-3](#)
 - managing, [2-23](#)
 - Oracle Traffic Director, [6-2](#), [6-3](#)
- configuring
 - heap analysis hosts, [20-6](#)
 - JVM, [20-35](#)
 - JVM Diagnostics engine, [20-2](#)
 - JVM pools, [20-4](#), [20-12](#)
 - JVMs, [20-4](#)
 - Oracle Identity Management targets, [30-1](#)
 - Oracle Traffic Director, [6-2](#)
 - SOA Suite, [11-6](#)
- CPU usage, Middleware targets, [2-6](#)
- create_jvm_diagnostic_db_user.sh script, [21-12](#)
- creating
 - blackouts, [2-12](#)
 - composite applications, [4-3](#)
 - Generic Service for Oracle Identity Management, [30-5](#)
 - Identity and Access System target, [30-4](#)
 - JVM Diagnostic snapshot, [20-38](#)
 - Oracle GlassFish Server configuration comparison template, [8-15](#)
 - Oracle Identity Management elements, [30-4](#)
 - Service Dashboard report, [30-5](#)
 - service request, [2-37](#)
 - Support Workbench package, [2-35](#)
 - Web Application targets for Oracle Identity Management, [30-5](#)
- cross tier
 - analysis in JVM Diagnostics, [20-18](#)
 - correlation in JVM Diagnostics, [19-2](#)
 - functionality errors in JVM Diagnostics, [21-1](#)

D

- databases
 - registering in JVM Diagnostics, [20-5](#)
- deleting
 - class histograms, [20-27](#)
 - Oracle Traffic Director targets, [6-7](#), [6-8](#)
- Deploying JVM Diagnostics, [33-25](#)
- diagnosing
 - performance problems, [2-15](#)

- diagnostic snapshots
 - available tasks, [2-15](#)
 - definition of, [2-15](#)
 - usage of, [2-15](#)
- discovered, Oracle Traffic Director targets, [6-5](#)
- discovering
 - Oracle Directory Server, [30-2](#)
 - Oracle Essbase targets, [12-4](#)
 - Oracle Identity Management target prerequisites, [29-3](#)
 - Oracle Identity Management targets, [29-1](#), [30-1](#)
- domain
 - adding Oracle GlassFish Server, [8-4](#)
 - Oracle GlassFish Server, [8-2](#)
- domain certificate
 - adding to Oracle GlassFish Server, [8-1](#)
- Domain Home page
 - Oracle GlassFish Server, [8-2](#)

E

- ECID,
 - JVMD displays, use of, [15-1](#)
 - request instance diagnostics, [15-3](#)
 - tracking requests, [13-4](#)
- editing
 - composite application home page, [4-5](#)
 - composite applications, [4-4](#)
 - JVM pool thresholds, [20-12](#)
- Enterprise Manager
 - agent deployment, [13-6](#)
 - JVMD, accessing from console, [15-2](#)
 - managing Middleware, [1-1](#)
 - registering RUEI with, [16-5](#)
 - services, [16-2](#)
 - setting up, [13-8](#)
 - systems, [16-2](#)
 - targets, [16-2](#)
- Error Hospital
 - customize report, [11-48](#)
 - generate report, [11-47](#)
- errors
 - JVM Diagnostics
 - cross-tier functionality, [21-1](#)
 - deployment execution, [21-4](#)
 - heap dump, [21-8](#)
 - loadheap, [21-7](#)
 - tracing, [21-4](#)
 - UI, [21-8](#)
- Exalytics target
 - monitoring, [5-1](#)
- execution context, [13-4](#)
- execution context ID
 - See ECID

extensible monitoring, [2-13](#)

F

features

Oracle Fusion Middleware Management, [1-2](#)

frequently asked questions

JVM Diagnostics, [21-9](#)

Fusion Middleware, [1-3](#)

See also Oracle Fusion Middleware
Components

Fusion Middleware

managing using Fusion Middleware Control,
[1-3](#)

H

heap analysis hosts

configuring, [20-6](#)

heap snapshots

in JVM Diagnostics, [20-29](#)

taking, [20-27](#)

viewing, [20-37](#)

heap usage by objects

viewing, [20-33](#)

heat map

Middleware targets, [2-6](#)

histograms, in JVM Diagnostics, [20-26](#)

historical diagnostics using JVM Diagnostics,
[19-3](#)

I

IBM WebSphere Application Server, [32-1](#)

managing, [32-1](#)

supported versions, [32-3](#)

IBM WebSphere Application Server cells, [32-1](#),
[32-2](#)

administering, [32-19](#)

monitoring, [32-17](#)

viewing members, [32-20](#)

IBM WebSphere Application Server clusters,
[32-1](#)

administering, [32-16](#)

monitoring, [32-14](#)

viewing, [32-16](#)

viewing metrics, [32-17](#)

IBM WebSphere Application Servers

administering, [32-11](#)

discovering, [32-7](#)

prerequisites, [32-3](#)

monitoring, [32-9](#)

monitoring applications, [32-13](#)

monitoring performance, [32-12](#)

IBM WebSphere Application Servers (*continued*)

troubleshooting

discovery, [32-21](#)

monitoring, [32-25](#)

viewing metrics, [32-14](#)

viewing the top EJBs, [32-13](#)

viewing the top servlets and JSPs, [32-13](#)

IBM WebSphere MQ, [31-1](#)

discovery prerequisites

local agent, [31-4](#)

remote agent, [31-4](#)

monitoring, [31-7](#)

prerequisites, [31-3](#)

queue manager cluster discovery, [31-5](#)

standalone queue manager discovery, [31-6](#)

understanding discovery, [31-4](#)

Identity and Access System target

creating, [30-4](#)

installing

Oracle Enterprise Manager

in Oracle Identity Management, [29-3](#)

instance

of Oracle Traffic Director, [6-6](#)

IWS

Configuring, [11-11](#)

Disabling, [11-11](#)

Enabling, [11-11](#)

Generating an IWS report, [11-12](#)

J

Java EE, [33-1](#)

Java EE application responsiveness, monitoring,
[2-10](#)

Java Keystore, [7-1](#)

Java Platform, Enterprise Edition, [33-1](#)

Java Virtual Machine Diagnostics

See JVM Diagnostics

Java Virtual Machines See JVMs, [20-35](#)

JBoss Application Server, [33-1](#)

administering, [33-13](#)

analyzing problems, [33-16](#)

discovering, [34-2](#)

managing, [33-1](#)

monitoring

applications, [33-14](#)

performance, [33-14](#)

Servlets and JSPs, [33-15](#)

prerequisites for discovery, [33-3](#)

supported versions, [33-3](#), [34-1](#)

troubleshooting, [33-29](#)

viewing metrics, [33-15](#)

JBoss Application Server 6.x

discovering, [33-6](#)

monitoring, [33-11](#)

- JBoss Application Server 7.x
 - discovering, [33-4](#)
 - monitoring, [33-9](#)
- JBoss Domains
 - administering, [33-20](#)
 - discovering, [33-4](#)
 - managing, [33-1](#)
 - monitoring, [33-17](#)
 - prerequisites for discovery, [33-3](#)
 - refreshing, [33-21](#)
 - supported versions, [33-3](#)
 - viewing members, [33-20](#)
- JBoss Partitions
 - administering, [33-23](#)
 - discovering, [33-6](#), [34-2](#)
 - managing, [33-1](#)
 - monitoring, [33-21](#)
 - prerequisites for discovery, [33-3](#)
 - refreshing, [33-25](#)
 - supported versions, [33-3](#), [34-1](#)
 - viewing members, [33-24](#)
- JBoss Server Groups
 - monitoring, [33-18](#)
- JFR snapshots, managing, [20-34](#)
- JKS (Java Keystore), [7-1](#)
- Job System, monitoring, [2-27](#)
- JRockit Flight Recorder See JFR, [20-34](#)
- JVM Diagnostics,
 - accessing, [15-1](#)
 - accessing pages, [20-7](#)
 - agent deployment, [13-6](#)
 - class histograms, [20-26](#)
 - features
 - cross tier correlation, [19-2](#)
 - in-depth visibility, [19-2](#)
 - JVM pooling, [19-3](#)
 - low overhead, [19-1](#)
 - memory leak detection, [19-2](#)
 - new features, [19-3](#)
 - real-time and historical diagnostics, [19-3](#)
 - real-time transaction tracing, [19-2](#)
 - supported platforms and JVMs, [19-3](#)
 - user roles, [19-4](#)
 - heap object information
 - heap objects, [20-32](#)
 - heap snapshots, [20-29](#)
 - comparing, [20-33](#)
 - heap usage by roots, [20-30](#)
 - top 40 objects, [20-31](#)
 - live thread analysis, [15-2](#)
 - location of logs, [21-9](#)
 - managing JVM pools, [20-1](#)
 - JVM Pool Home page, [20-8](#)
 - live thread analysis, [20-9](#)
- JVM Diagnostics (*continued*)
 - managing JVMs
 - offline diagnostics, [20-38](#)
 - Oracle Real Application Cluster drill-down, [20-20](#)
 - overview, [15-1](#), [19-1](#)
 - request instance diagnostics, [15-3](#)
 - sample analyzer, [15-2](#)
 - setting up, [13-8](#), [20-1](#)
 - Snapshots page, [20-39](#)
 - thread snapshots
 - analyzing trace diagnostic images, [20-37](#)
 - Thread Stat transition chart, [15-2](#)
 - threshold violations, [20-40](#)
 - view, initial, [15-1](#)
- JVM Diagnostics troubleshooting
 - agent status, [21-10](#)
 - cross tier functionality errors, [21-1](#)
 - customizing provisioning agent deployment, [21-13](#)
 - deployment script execution errors, [21-4](#)
 - engine status, [21-10](#)
 - frequently asked questions, [21-9](#)
 - heap dump errors, [21-8](#)
 - loadheap errors, [21-7](#)
 - log manager level, [21-13](#)
 - monitoring status, [21-10](#)
 - optimization levels, [21-12](#)
 - repository space requirements, [21-13](#)
 - running create_jvm_diagnostic_db_user.sh script, [21-12](#)
 - trace errors, [21-4](#)
 - Try Changing Threads parameter, [21-12](#)
 - user interface errors, [21-8](#)
- JVM pools, [19-3](#)
 - adding to group, [20-13](#)
 - configuring, [20-4](#), [20-12](#)
 - managing, [20-8](#)
 - removing, [20-13](#)
 - thresholds, editing, [20-12](#)
- JVMs
 - configuring, [20-4](#), [20-35](#)
 - managers, viewing registered, [20-7](#)
 - managing, [20-13](#)
 - JVM Home page, [20-14](#)
 - live time thread analysis, [20-15](#)
 - offline diagnostics, [20-38](#), [20-39](#)
 - performance summary, [20-14](#)
 - performance metrics, collecting, [7-3](#)
 - removing, [20-35](#)

K

- key components, [16-2](#)

keytool utility
changing passwords, 7-2

KPIs
monitoring, 16-30
RUEI, 14-7

L

lifecycle management
managing configurations, 2-23
monitoring, 2-21, 2-22
live thread analysis, 21-1
cross tier, 21-1
logs
searching, 2-16
low overhead in JVM Diagnostics, 19-1

M

managed servers
ending blackouts after starting up, 2-18
restart method, 2-19
restart time limit, 2-18
restarting, 2-17
shutdown errors, 2-20, 2-21
shutdown method, 2-19
shutdown time limit, 2-18
shutting down, 2-17
starting up, 2-17
startup errors, 2-20, 2-21
startup method, 2-19
startup time limit, 2-18
stopping errors, 2-20, 2-21
managing
blackouts, 2-12
configurations, 2-23
JFR snapshots, 20-34
JVM pools, 20-8
JVMs, 20-13
thread snapshots, 20-35
memory leak detection using JVM Diagnostics, 19-2
memory leak report, 20-34
metric thresholds, 2-11
Middleware
analyzing, 2-27
Middleware Applications
troubleshooting, 18-1
Middleware Diagnostics Advisor, 22-2
diagnosing performance issues, 22-2, 22-7
enabling, 22-4
functions, 22-2
limiting scope of, 22-6
overview, 22-1
prerequisites, 22-3

Middleware Diagnostics Advisor (*continued*)
purging data, 22-5
troubleshooting issues, 22-9

Middleware management, 1-1, 2-1
using Enterprise Manager, 1-1

Middleware targets, 2-1
administering, 2-16
ending blackouts after starting up, 2-18
heat map, 2-6
monitoring, 2-5
restart method, 2-19
restart time limit, 2-18
restarting, 2-17
searching, 2-7
shutdown errors, 2-20, 2-21
shutdown method, 2-19
shutdown time limit, 2-18
shutting down, 2-17
starting up, 2-17
startup errors, 2-20, 2-21
startup method, 2-19
startup time limit, 2-18
status and CPU usage, 2-6
stopping errors, 2-20, 2-21

monitoring
administer Middleware targets, 2-16
Ealytics target, 5-1
extensible, Application Server, 2-13
Job System, 2-27
lifecycle management, 2-21, 2-22
compliance management, 2-24
managing configurations, 2-23
patch management, 2-24
provisioning, 2-25
managing service levels, 2-25
Middleware targets, 2-1, 2-5
non-Oracle Middleware components, 2-4
ODI agents, 35-5
ODI repositories, 35-6
Oracle Application Server components, 2-4
Oracle Identity Management components, 28-3
out-of-box monitoring
blackouts, 2-12
extending, 2-13
historical performance, 2-11
metric thresholds, 2-11
monitoring templates, 2-12
performance problems, 2-15
diagnostics snapshots, 2-15
Home page, 2-15
predefined metrics, 2-10
Routing Topology Viewer, 2-27
status in JVM Diagnostics, 21-10
Support Workbench, 2-30

monitoring Oracle Essbase targets, [12-5](#)

N

named credentials

Support Workbench, [2-31](#)

new features

Oracle Coherence, [23-2](#)

SOA Suite, [10-1](#)

O

ODI (Oracle Data Integrator), [35-1](#)

See also Oracle Data Integrator, [35-1](#)

Oracle Application Server

components, [2-4](#)

Oracle Business Analytics, [12-1](#)

Oracle Business Intelligence, [2-1](#), [12-1](#)

Oracle Business Intelligence Instance, [12-2](#)

component failovers, [12-23](#)

dashboard reports, [12-16](#)

discovering, [12-4](#)

monitoring, [12-5](#)

monitoring credentials, [12-24](#)

scheduler reports, [12-18](#)

Oracle Business Intelligence Instance

components, [12-2](#)

BI Cluster Controller, [12-2](#)

BI Java Host, [12-2](#)

BI Presentation Server, [12-2](#)

BI Scheduler, [12-2](#)

BI Server, [12-2](#)

Oracle Business Intelligence targets, [12-5](#)

alerts, [12-11](#)

availability, [12-7](#)

blackouts, [12-22](#)

compliance, [12-15](#)

configuration, [12-14](#)

health, [12-11](#)

incidents, [12-12](#)

job activity, [12-15](#)

logs, [12-12](#)

metrics, [12-10](#)

monitoring configuration, [12-22](#)

performance, [12-8](#)

resource usage, [12-8](#)

Oracle Coherence, [2-1](#), [23-1](#)

administration

cache data management, [25-4](#)

change cache configuration, [25-3](#), [25-4](#)

change node configuration, [25-1](#)

change service configuration, [25-3](#), [25-4](#)

cluster administration, [25-1](#)

node administration, [25-3](#)

setup log alerts, [25-3](#)

Oracle Coherence (*continued*)

best practices

monitoring templates, [26-1](#)

cluster management

start new nodes, [24-8](#)

stop nodes, [24-8](#)

Coherence Cluster, [23-1](#)

discover cluster, [23-9](#)

JVM Diagnostics integration, [27-1](#)

access JVM Diagnostics, [27-3](#)

configure coherence node, [27-1](#)

manage mis-configured nodes, [23-13](#)

monitor

application home page, [24-19](#)

applications page, [24-29](#)

cache data management, [24-17](#)

cache home page, [24-13](#)

caches page, [24-26](#)

cluster home page, [24-5](#)

cluster management, [24-7](#)

connection manager home page, [24-23](#)

connection manager performance page, [24-32](#)

high availability status, [24-28](#)

near cache, [24-16](#)

node home page, [24-10](#)

nodes page, [24-24](#)

performance summary page, [24-32](#)

proxies page, [24-29](#)

reset statistics, [24-24](#)

service home page, [24-20](#)

services page, [24-27](#)

start node, [24-24](#)

stop nodes, [24-24](#)

navigation tree, [24-2](#)

new features, [23-2](#)

personalization, [24-3](#)

Refresh Cluster, [23-12](#)

standalone cluster, [23-3](#)

JMX management node, [23-3](#)

management node sample start script, [23-5](#)

sample start script (other nodes), [23-6](#)

start JMX management node, [23-4](#)

troubleshooting

collecting metric data, [26-1](#)

dynamic client nodes, [26-1](#)

target proliferation of nodes, [26-1](#)

Oracle Data Integrator

administering, [35-13](#)

configuring console, [35-18](#)

configuring domain, [35-19](#)

load plan executions, [35-14](#)

managing agent activities, [35-14](#)

managing agent status, [35-14](#)

- Oracle Data Integrator (*continued*)
 - monitoring, [35-2](#)
 - load plan executions and sessions, [35-9](#)
 - monitoring agents, [35-5](#)
 - monitoring prerequisites, [35-1](#)
 - monitoring repositories, [35-6](#)
 - restarting, [35-13](#)
 - searching sessions, [35-14](#)
 - shutting down, [35-13](#)
 - viewing log messages, [35-14](#)
- Oracle Data Integrator Agents
 - ending blackouts after starting up, [2-18](#)
 - restart method, [2-19](#)
 - restart time limit, [2-18](#)
 - restarting, [2-17](#)
 - shutdown errors, [2-20](#), [2-21](#)
 - shutdown method, [2-19](#)
 - shutdown time limit, [2-18](#)
 - shutting down, [2-17](#)
 - starting up, [2-17](#), [35-13](#)
 - startup errors, [2-20](#), [2-21](#)
 - startup method, [2-19](#)
 - startup time limit, [2-18](#)
 - stopping errors, [2-20](#), [2-21](#)
- Oracle Essbase, [12-2](#)
 - applications, [12-19](#)
 - discovering, [12-4](#)
 - monitoring, [12-5](#)
- Oracle Forms Services, [2-1](#)
- Oracle Fusion Middleware Components, [1-3](#), [2-1](#)
 - Oracle Business Intelligence, [2-1](#), [12-1](#), [12-3](#)
 - Oracle Coherence, [2-1](#)
 - Oracle Forms Services, [2-1](#)
 - Oracle Identity Management, [2-3](#)
 - Oracle Portal, [2-1](#)
 - Oracle SOA Suite, [2-1](#)
 - Oracle Web Tier, [2-1](#)
 - Oracle WebCenter, [2-2](#)
 - Oracle WebCenter Content, [2-1](#)
 - Oracle WebLogic Domains, Clusters, and Managed Servers, [2-1](#)
 - See also *Middleware targets*
- Oracle Fusion Middleware Management
 - features, [1-2](#)
- Oracle GlassFish Server
 - before getting started, [8-1](#)
 - cluster home page, [8-13](#)
 - creating configuration comparison template, [8-15](#)
 - domain, [8-2](#)
 - adding, [8-4](#)
 - displaying results, [8-9](#)
 - finding and assigning targets, [8-5](#)
 - Domain Home page, [8-2](#)
 - home page, [8-11](#)
- Oracle GlassFish Server (*continued*)
 - how to access, [8-12](#)
 - how to access cluster, [8-14](#)
 - how to access domain, [8-9](#)
 - overview, [8-1](#)
 - refreshing domain, [8-10](#)
 - roles and privileges, [8-1](#)
 - start and stop procedures, [8-1](#)
 - viewing configuration data, [8-15](#)
- Oracle HTTP Server, [2-1](#)
 - ending blackouts after starting up, [2-18](#)
 - restart method, [2-19](#)
 - restart time limit, [2-18](#)
 - restarting, [2-17](#)
 - shutdown errors, [2-20](#), [2-21](#)
 - shutdown method, [2-19](#)
 - shutdown time limit, [2-18](#)
 - shutting down, [2-17](#)
 - starting up, [2-17](#)
 - startup errors, [2-20](#), [2-21](#)
 - startup method, [2-19](#)
 - startup time limit, [2-18](#)
 - stopping errors, [2-20](#), [2-21](#)
- Oracle HTTP Server session volumes, [2-10](#)
- Oracle Identity Management
 - creating elements, [30-4](#)
 - discovering and configuring targets, [30-1](#)
 - discovering targets, [29-1](#)
 - features, [28-1](#)
 - getting started with, [28-1](#)
 - installing Oracle Enterprise Manager, [29-3](#)
 - licensed targets, [28-3](#)
 - monitoring components, [28-3](#)
 - system requirements, [29-1](#)
- Oracle Internet Directory
 - collecting statistics, [30-3](#)
- Oracle Portal, [2-1](#)
- Oracle Real Application Cluster
 - JVM Diagnostics, [20-20](#)
- Oracle Service Bus
 - discovery, [10-1–10-3](#)
 - enabling Management Packs, [10-5](#)
 - supported versions, [10-1](#)
 - troubleshooting, [10-7](#)
- Oracle Traffic Director
 - adding Exalogic target, [6-2](#)
 - configuration, [6-2](#), [6-3](#)
 - configuring for SNMP monitoring, [6-2](#)
 - discovered targets, [6-5](#)
 - instance, [6-6](#)
 - overview, [6-1](#)
 - refresh flow, [6-6](#)
- Oracle Web Cache, [2-1](#)
- Oracle WebCenter Content, [2-1](#)

Oracle WebLogic Domain

- ending blackouts after starting up, [2-18](#)
- restart method, [2-19](#)
- restart time limit, [2-18](#)
- restarting, [2-17](#)
- shutdown errors, [2-20](#), [2-21](#)
- shutdown method, [2-19](#)
- shutdown time limit, [2-18](#)
- shutting down, [2-17](#)
- starting up, [2-17](#)
- startup errors, [2-20](#), [2-21](#)
- startup method, [2-19](#)
- startup time limit, [2-18](#)
- stopping errors, [2-20](#), [2-21](#)

Oracle WebLogic Servers, [22-1](#)

- ending blackouts after starting up, [2-18](#)
- restart method, [2-19](#)
- restart time limit, [2-18](#)
- restarting, [2-17](#)
- shutdown errors, [2-20](#), [2-21](#)
- shutdown method, [2-19](#)
- shutdown time limit, [2-18](#)
- shutting down, [2-17](#)
- starting up, [2-17](#)
- startup errors, [2-20](#), [2-21](#)
- startup method, [2-19](#)
- startup time limit, [2-18](#)
- stopping errors, [2-20](#), [2-21](#)

P

package

- creating in Support Workbench, [2-35](#)
- uploading to Oracle support, [2-36](#)

patch management, [2-24](#)

performance

- diagnostics
 - JVM Diagnostics, [21-1](#)
- problems, diagnosing, [2-15](#)

performance monitoring

- data collection, [13-6](#)
- end-to-end, [13-1](#)
- example of end-to-end, [17-1](#)
- overview, [13-1](#)
- processing engines, [13-6](#)
- RUEI, and JVMD, [13-1](#)
- setting up, [13-6](#)
- troubleshooting, [17-6](#)
- user roles, [13-9](#)
- views and dimensions, [13-1](#)

platform MBeans

- activating, [7-3](#)

PlatformMBeanServerUsed

- setting attribute, [7-3](#)

platforms supported in JVM Diagnostics, [19-3](#)

preferred credentials

- Support Workbench, [2-32](#)

prerequisites

- discovering Oracle Identity Management targets, [29-3](#)

privileges and roles

- Oracle GlassFish Server, [8-1](#)

Problem Analysis

- using, [2-27](#)

problems

- annotating, [2-34](#)
- closing in Support Workbench, [2-37](#)
- searching for, [2-34](#)

promoting

- JVM Diagnostics events, [20-9](#)

provisioning

- lifecycle management, [2-25](#)

R

Real User Experience Insight, [14-10](#)

- accessing from Enterprise Manager console, [14-9](#)
- accessing JVMD from, [14-10](#)
- application, [14-2](#)
- collector, [13-6](#)
- dashboards, [14-3](#)
- data analysis, [14-3](#)
- data collection, [14-1](#)
- ECID, use of, [13-4](#)
- exporting sessions, [16-24](#)
- features, [14-8](#)
- KPIs, [14-7](#)
- monitoring data, [16-17](#)
- monitoring metrics, [16-26](#)
- overview, [14-1](#)
- registering with Enterprise Manager, [16-5](#)
- reports, [14-4](#)
- requirements for using in Enterprise Manager, [16-4](#)
- service level agreements, [14-3](#), [14-7](#)
- session diagnostics, [14-5](#), [16-19](#)
- session replay, [16-23](#)
- setting up, [13-8](#)
- top users, [16-17](#)
- troubleshooting, [17-1](#)
- user flows, [14-3](#), [14-5](#)

real-time diagnostics using JVM Diagnostics, [19-3](#)

refresh flow

- Oracle Traffic Director, [6-6](#)

registering

- databases in JVM Diagnostics, [20-5](#)

removing

- JVM pools, [20-13](#)

removing (*continued*)
 JVMs, [20-35](#)
 reports
 anti-pattern in JVM Diagnostics, [20-34](#)
 memory leak in JVM Diagnostics, [20-34](#)
 SOA Suite, [11-20](#)
 reports (RUEI), [14-4](#)
 repository
 space requirements in JVM Diagnostics,
 [21-13](#)
 roles and privileges
 Oracle GlassFish Server, [8-1](#)
 Routing Topology Viewer, [2-27](#)

S

saving
 JVM Diagnostics class histogram, [20-26](#)
 scheduling
 JVM Diagnostics histogram job, [20-26](#)
 scripts
 create_jvm_diagnostic_db_user.sh, [21-12](#)
 startManagedWeblogic, [2-19](#)
 stopManagedWeblogic, [2-19](#)
 searching
 logs, [2-16](#)
 Middleware targets, [2-7](#)
 Secure Socket Layer, [7-1](#)
 service level agreements
 RUEI, [14-7](#)
 service levels
 managing, [2-25](#)
 service request, creating, [2-37](#)
 session diagnostics, [17-2](#)
 session diagnostics (RUEI), [14-5](#), [16-19](#)
 setting up
 JVM Diagnostics, [20-1](#)
 SNMP monitoring
 Oracle Traffic Director, [6-2](#)
 SOA faults
 bulk recovery, [11-32](#)
 bulk recovery from Error Hospital, [11-37](#)
 bulk recovery from Faults and Rejected
 Messages, [11-36](#)
 bulk recovery from Jobs page, [11-33](#)
 bulk recovery workflow, [11-42](#)
 overview, [11-25](#)
 recovery, [11-26](#)
 search and view, [11-27](#)
 simple recovery, [11-31](#)
 total faults, [11-29](#)
 SOA instance tracing
 across SOA infrastructures, [11-17](#)
 within SOA infrastructure, [11-17](#)

SOA Management Pack Enterprise Edition, [9-1](#)
 BPEL Process Manager, [9-1](#)
 Central Management Console, [9-1](#)
 Error Hospital, [9-2](#)
 Service Bus, [9-1](#)
 SOA Composite, [9-1](#)
 SOA Infrastructure, [9-1](#)
 SOA reports
 SOA diagnostic reports, [11-22](#)
 using BI Publisher, [11-20](#)
 using IP, [11-22](#)
 SOA Suite
 bulk recovery, [11-32](#)
 configuring, [11-6](#)
 Dehydration Store, [11-18](#)
 Error Hospital, [11-44](#)
 faults and recovery, [11-25](#)
 metrics and collection, [11-8](#)
 new features, [10-1](#)
 simple recovery, [11-31](#)
 SOA artifacts and composites, [11-24](#)
 SOA Instance Tracing, [11-13](#)
 SOA reports, [11-20](#)
 supported versions, [11-1](#)
 tracking bulk recovery, [11-38](#)
 troubleshooting, [11-49](#)
 UDDI publishing, [11-20](#)
 starting, Oracle GlassFish Server, [8-1](#)
 startManagedWeblogic script, [2-19](#)
 status
 Middleware targets, [2-6](#)
 stopManagedWeblogic script, [2-19](#)
 stopping
 Oracle GlassFish Server, [8-1](#)
 Support Workbench
 accessing and logging into, [2-31](#)
 adding files to package, [2-36](#)
 adding more diagnosability information, [2-35](#)
 Aggregated Diagnostic Summary page, [2-33](#)
 aggregated diagnostics, viewing, [2-33](#)
 annotating problems, [2-34](#)
 closing problems, [2-37](#)
 compatibility with Oracle Fusion Middleware
 components, [2-30](#)
 diagnostics, viewing, [2-32](#)
 named credentials, [2-31](#)
 new credentials, [2-32](#)
 overview and purpose of, [2-30](#)
 preferred credentials, [2-32](#)
 searching for problems, [2-34](#)
 uploading package to Oracle support, [2-36](#)
 work flows, [2-30](#)
 system requirements
 Oracle Identify Management, [29-1](#)

T

- thread snapshots
 - analyzing trace diagnostic images, [20-37](#)
 - managing, [20-35](#)
- threshold violations
 - in JVM Diagnostics, [20-40](#)
- trace errors
 - JVM Diagnostics, [21-4](#)
- tracing
 - active threads in JVM Diagnostics, [20-36](#)
- transactions
 - tracing in real time using JVM Diagnostics, [19-2](#)
- troubleshooting
 - JVM Diagnostics, [21-9](#)
 - Oracle Service Bus, [10-7](#)
 - SOA Suite, [11-49](#)

U

- user
 - experience, [14-1](#)
 - flows, [14-5](#)
 - privileges, [13-9](#)
 - roles, required to use JVM Diagnostics, [19-4](#)
- user roles, [13-9](#)

V

- viewing
 - class histograms, [20-37](#)
 - composite applications, [4-5](#)
 - composite applications dashboard, [4-1](#)
 - heap snapshots, [20-37](#)
 - heap usage by objects, [20-33](#)
 - heap usage by roots, [20-30](#)
 - JVM Diagnostics class histogram, [20-26](#)
 - JVM Diagnostics Performance Summary page, [20-14](#)
 - JVM Diagnostics snapshot, [20-40](#)
 - JVM Diagnostics threshold violations, [20-40](#)
 - JVM Home page, [20-14](#)
 - JVM Live Thread Analysis page, [20-15](#)
 - JVM pool live thread analysis, [20-9](#)
 - registered JVM managers, [20-7](#)
 - registered JVMs, [20-7](#)

W

- WebLogic Domains, [12-2](#)
 - monitoring, [7-1](#)
 - Oracle Business Intelligence Instance, [12-2](#)
 - Oracle Essbase, [12-2](#)
- WebLogic Servers
 - collecting metrics, [7-3](#)