Oracle® Enterprise Manager

Middleware Management Guide

24ai Release 1 (24.1)
F97230-01
December 2024

Oracle Enterprise Manager Middleware Management Guide, 24ai Release 1 (24.1)

F97230-01

Contents

Preface

Part I Managing Oracle Fusion Middleware

1 Introduction to Middleware Management

2 Managing Middleware Targets

Part II Monitoring Exalytics Target

Part III Monitoring Oracle WebLogic Domains

4 Monitoring WebLogic Domains

Part IV Managing Oracle SOA

5 Overview of Oracle SOA Management

6 Discovering and Monitoring Service Bus

7 Discovering and Monitoring the SOA Suite

Part V Managing Oracle Business Intelligence

8 Discovering and Monitoring Oracle Business Intelligence Instance and Oracle Essbase

9 Introduction to JVM Diagnostics

10 Using JVM Diagnostics

Part VII Managing Oracle Coherence

12 Getting Started with Management Pack for Oracle Coherence

13 Monitoring a Coherence Cluster

Part X Managing Oracle Data Integrator

20 Configuring and Monitoring Oracle Data Integrator

Preface

This document describes how to use Oracle Enterprise Manager 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 documents available in the Oracle Enterprise Manager documentation library: Enterprise Manager Documentation.

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.
italic	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:

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

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)
- 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.

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, and Apache Tomcat.
- 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

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:

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

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:

- Oracle Fusion Middleware software
- Non-Oracle Middleware software

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 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 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.

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.

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 (`https://support.oracle.com`).

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.

Middleware Summary Page

Enterprise Manager provides centralized monitoring across domains, configuration management, provisioning, real time and historical performance analysis. Fusion Middleware Plug-in administration features are exposed within the console. These features enable you to perform configuration changes directly from the 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 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 Enterprise Manager; in some cases, you still need to drill down to the administration consoles.

> **Note:**
>
> WebLogic Server Admin console features that are accessible from the Cloud Control console, are not available for WebLogic Server targets version 14.1 and higher.

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*.

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.

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
```

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

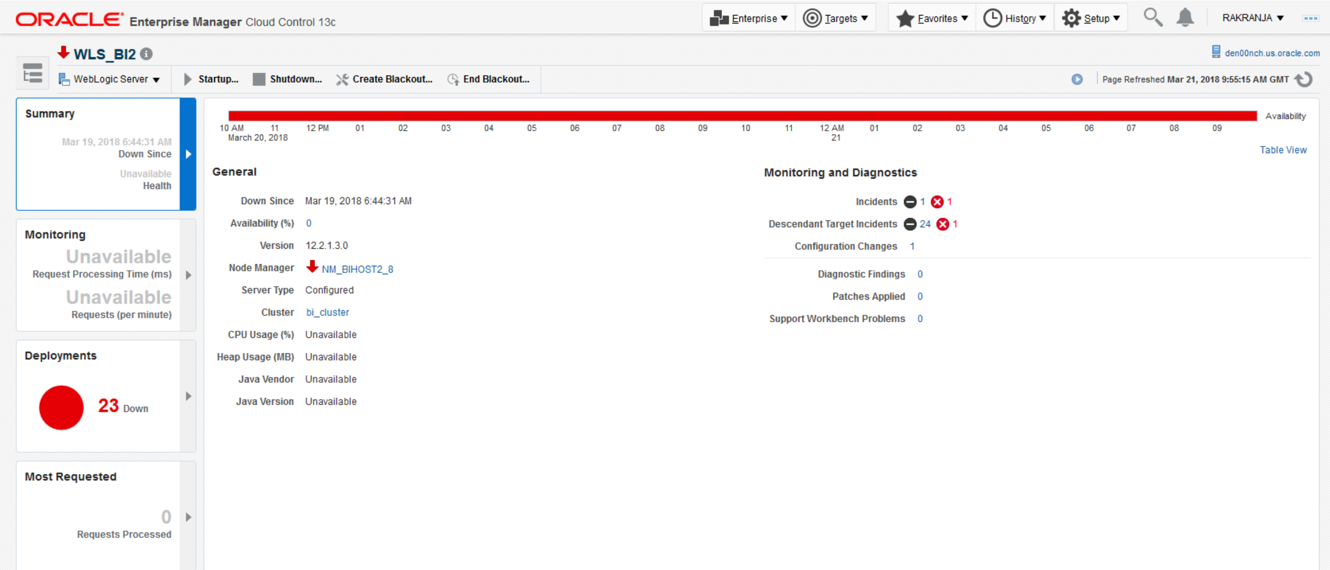

Summary

Mar 19, 2018 6:44:31 AM
Down Since

Unavailable
Health

Monitoring

Unavailable
Request Processing Time (ms)

Unavailable
Requests (per minute)

Deployments

23 Down

Most Requested

0
Requests Processed

General

Down Since Mar 19, 2018 6:44:31 AM
Availability (%) 0
Version 12.2.1.3.0
Node Manager NM_BIHOST2_8
Server Type Configured
Cluster bi_cluster
CPU Usage (%) Unavailable
Heap Usage (MB) Unavailable
Java Vendor Unavailable
Java Version Unavailable

Monitoring and Diagnostics

Incidents 1 1
Descendant Target Incidents 24 1
Configuration Changes 1
Diagnostic Findings 0
Patches Applied 0
Support Workbench Problems 0

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.)

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

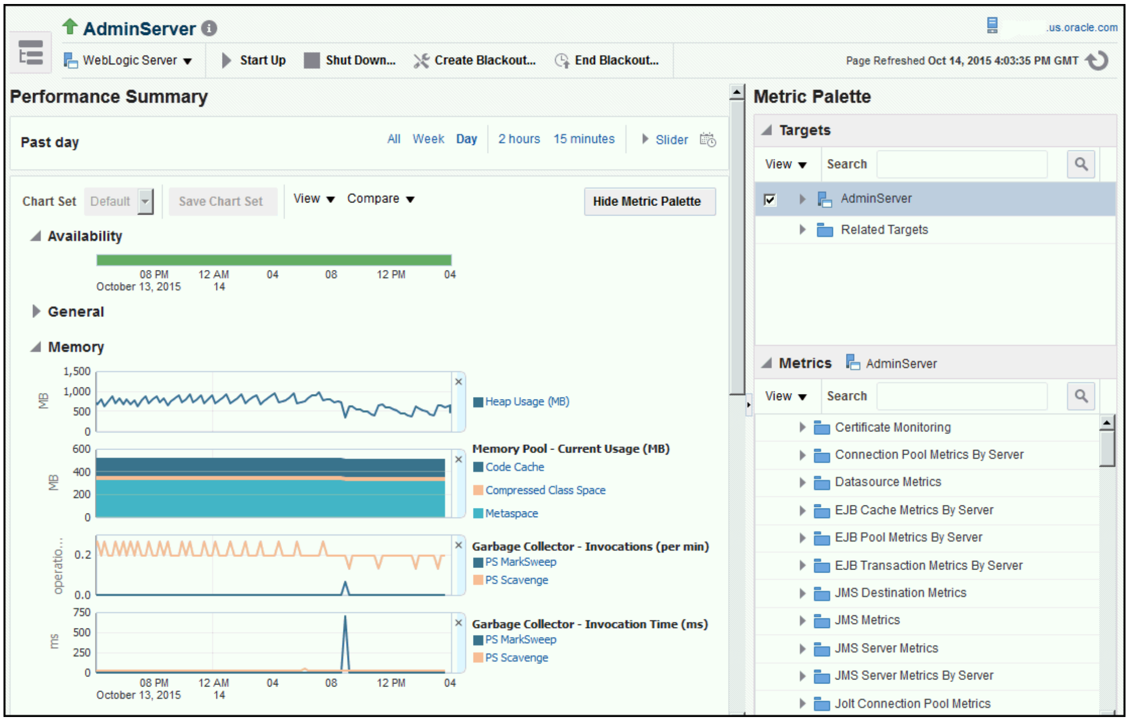

You can change which charts are displayed on the performance page and then save the changes on a per-user, per-target-type basis. You can also save multiple customized versions of a performance page, giving each version a name. This will save time by allowing quick access to previously created version 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.

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*.

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.

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.

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.

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*.

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.

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.

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.

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
```

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.

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.

> **Note:**
>
> WebLogic Server Admin console features that are accessible from the Cloud Control console, are not available for WebLogic Server targets version 14.1 and higher.

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.x), 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).

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 *(Appears only for shutdown operation)*	Selected	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)). 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. **Note:** • 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 *(Appears only for start up operation)*	Selected	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. **Note:** • 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 *(Appears only for Oracle WebLogic Domains)*	Not Selected	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. **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.
Time Out After (in minutes)	5 Minutes Per Target	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. 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.

Element	Default Value	Description
Process Control Method *(Appears only for Oracle WebLogic Domains, Oracle WebLogic Clusters, Oracle WebLogic Servers)*	Administration Server	Select one of the following ways in which the shutdown, start-up, or restart operation can be performed: **Note:** Options not applicable to a particular target type are disabled. • **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 `startManagedWeblogic` script and the `stopManagedWeblogic` script located in the `<DOMAIN_HOME>/bin` 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 `boot.properties` file for the server. • **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 `boot.properties` file for the server.
Credentials	Preferred	If default script or custom script is selected then Administration Server Credentials are not required, only agent host credentials are required. 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.
Targets	Not Selected	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.

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 &";
   }
   ```

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_DOMAIN_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.

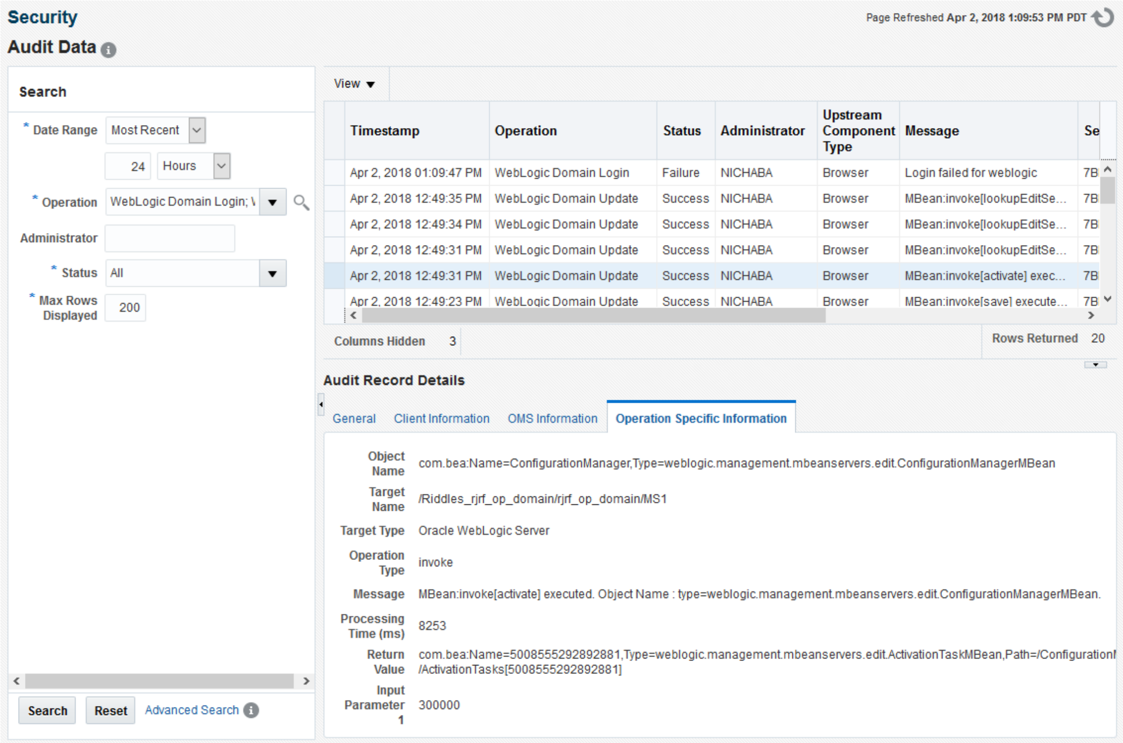

Security

Audit Data

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

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

Object Name	com.bea:Name=ConfigurationManager,Type=weblogic.management.mbeanservers.edit.ConfigurationManagerMBean
Target Name	/Riddles_rjrf_op_domain/rjrf_op_domain/MS1
Target Type	Oracle WebLogic Server
Operation Type	invoke
Message	MBean:invoke[activate] executed. Object Name : type=weblogic.management.mbeanservers.edit.ConfigurationManagerMBean.
Processing Time (ms)	8253
Return Value	com.bea:Name=5008555292892881,Type=weblogic.management.mbeanservers.edit.ActivationTaskMBean,Path=/Configuration /ActivationTasks[5008555292892881]
Input Parameter 1	300000

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.

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.

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*.

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*.

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*.

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*.

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.

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.

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*.

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*.

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.

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.

Accessing and Logging In To Support Workbench

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

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**.

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.

Using Fusion Middleware Support Workbench

You can use Support Workbench within Fusion Middleware to:

- Viewing an Aggregated Diagnostic Summary
- Create a problem, package it, and upload it to Oracle Support

The following sections provide procedures for these diagnostic tasks.

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.

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

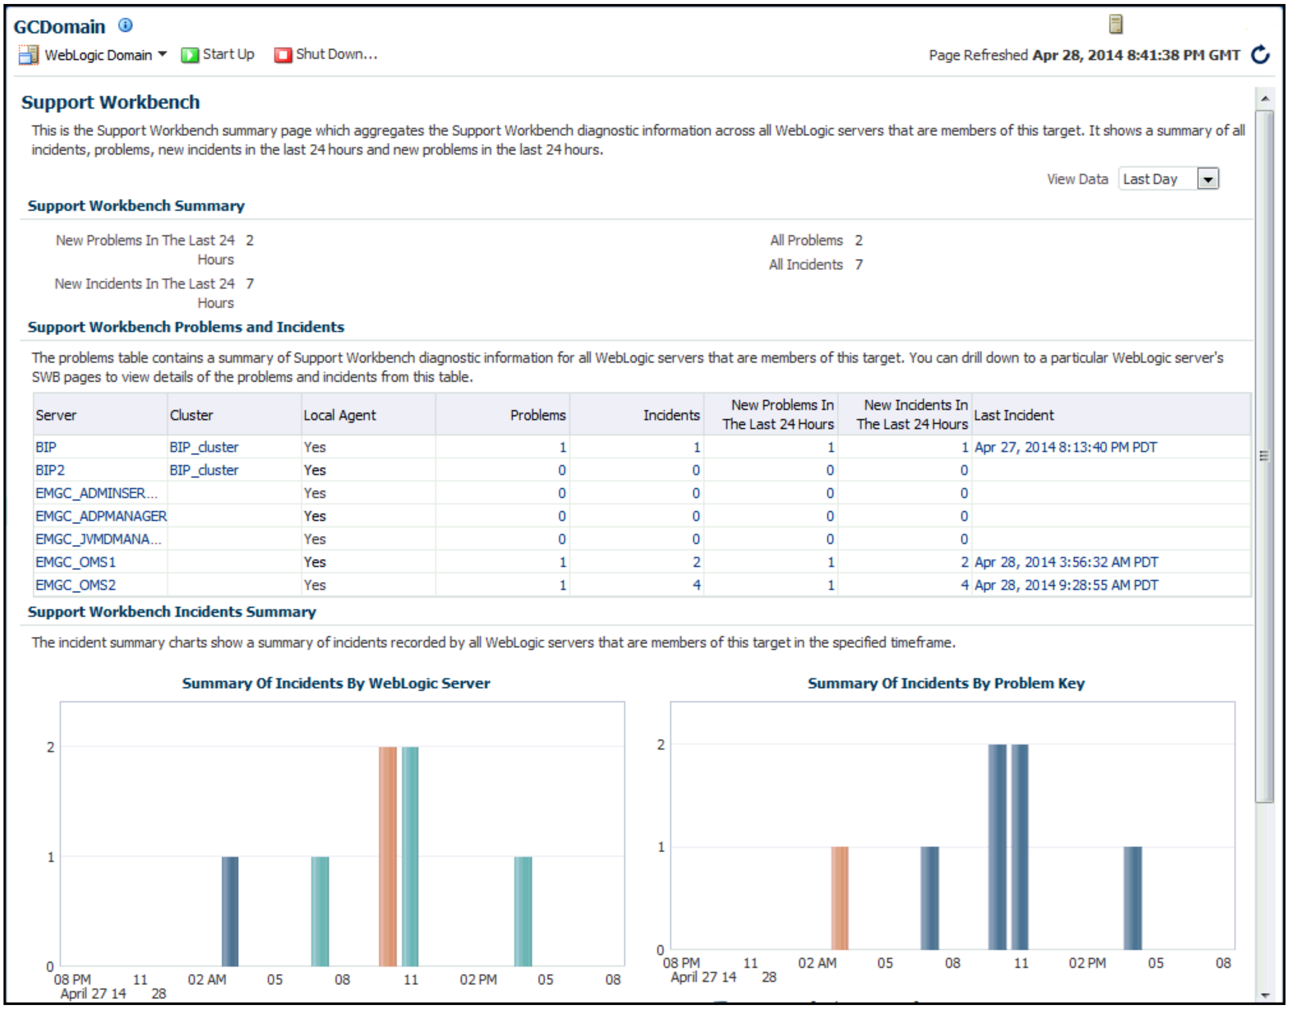

GCDomain

WebLogic Domain ▾   Start Up   Shut Down...

Page Refreshed **Apr 28, 2014 8:41:38 PM GMT**

Support Workbench

This is the Support Workbench summary page which aggregates the Support Workbench diagnostic information across all WebLogic servers that are members of this target. It shows a summary of all incidents, problems, new incidents in the last 24 hours and new problems in the last 24 hours.

View Data: Last Day

Support Workbench Summary

New Problems In The Last 24 Hours	2	All Problems	2
New Incidents In The Last 24 Hours	7	All Incidents	7

Support Workbench Problems and Incidents

The problems table contains a summary of Support Workbench diagnostic information for all WebLogic servers that are members of this target. You can drill down to a particular WebLogic server's SWB pages to view details of the problems and incidents from this table.

Server	Cluster	Local Agent	Problems	Incidents	New Problems In The Last 24 Hours	New Incidents In The Last 24 Hours	Last Incident
BIP	BIP_cluster	Yes	1	1	1	1	Apr 27, 2014 8:13:40 PM PDT
BIP2	BIP_cluster	Yes	0	0	0	0	
EMGC_ADMINSER...		Yes	0	0	0	0	
EMGC_ADPMANAGER		Yes	0	0	0	0	
EMGC_JVMDMANA...		Yes	0	0	0	0	
EMGC_OMS1		Yes	1	2	1	2	Apr 28, 2014 3:56:32 AM PDT
EMGC_OMS2		Yes	1	4	1	4	Apr 28, 2014 9:28:55 AM PDT

Support Workbench Incidents Summary

The incident summary charts show a summary of incidents recorded by all WebLogic servers that are members of this target in the specified timeframe.

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

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**.

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 redisplays it if this administrator or a different one looks at this problem.

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.

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 t his 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.

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.

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.

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**.

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
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

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
- 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 3-1):

- 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 3-1 Composite Application Dashboard

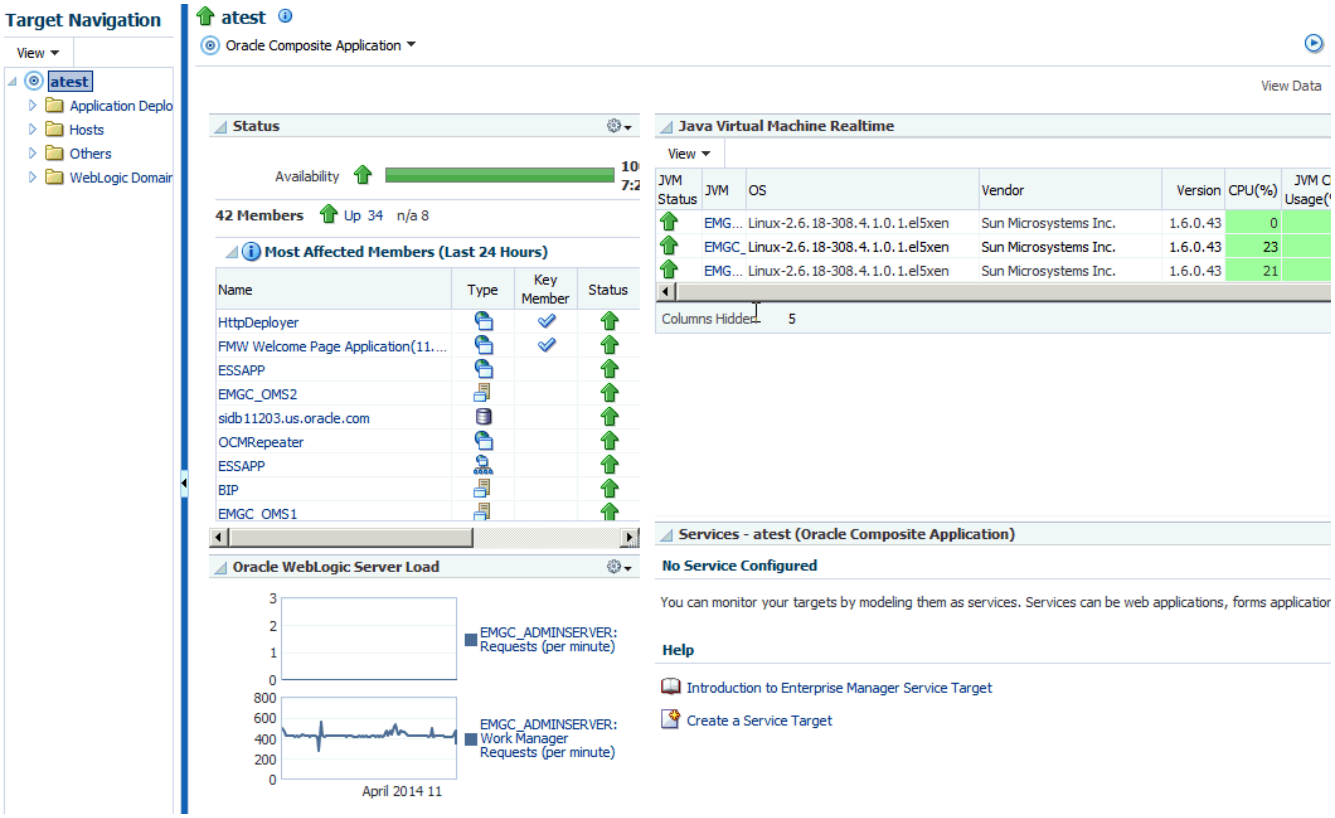

Target Navigation

View ▾

atest
- Application Deplo
- Hosts
- Others
- WebLogic Domair

atest

Oracle Composite Application ▾

View Data

Status

Availability 10 7:2

42 Members Up 34 n/a 8

Most Affected Members (Last 24 Hours)

Name	Type	Key Member	Status
HttpDeployer		✓	⬆
FMW Welcome Page Application(11....		✓	⬆
ESSAPP			⬆
EMGC_OMS2			⬆
sidb11203.us.oracle.com			⬆
OCMRepeater			⬆
ESSAPP			⬆
BIP			⬆
EMGC_OMS1			⬆

Java Virtual Machine Realtime

View ▾

JVM Status	JVM	OS	Vendor	Version	CPU(%)	JVM C Usage(
⬆	EMG...	Linux-2.6.18-308.4.1.0.1.el5xen	Sun Microsystems Inc.	1.6.0.43	0	
⬆	EMGC_	Linux-2.6.18-308.4.1.0.1.el5xen	Sun Microsystems Inc.	1.6.0.43	23	
⬆	EMG...	Linux-2.6.18-308.4.1.0.1.el5xen	Sun Microsystems Inc.	1.6.0.43	21	

Columns Hidden 5

Oracle WebLogic Server Load

3
2
1
0
EMGC_ADMINSERVER: Requests (per minute)
800
600
400
200
0
EMGC_ADMINSERVER: Work Manager Requests (per minute)
April 2014 11

Services - atest (Oracle Composite Application)

No Service Configured

You can monitor your targets by modeling them as services. Services can be web applications, forms applicatior

Help

Introduction to Enterprise Manager Service Target

Create a Service Target

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 3-2 First Page of the Create Composite Application Wizard

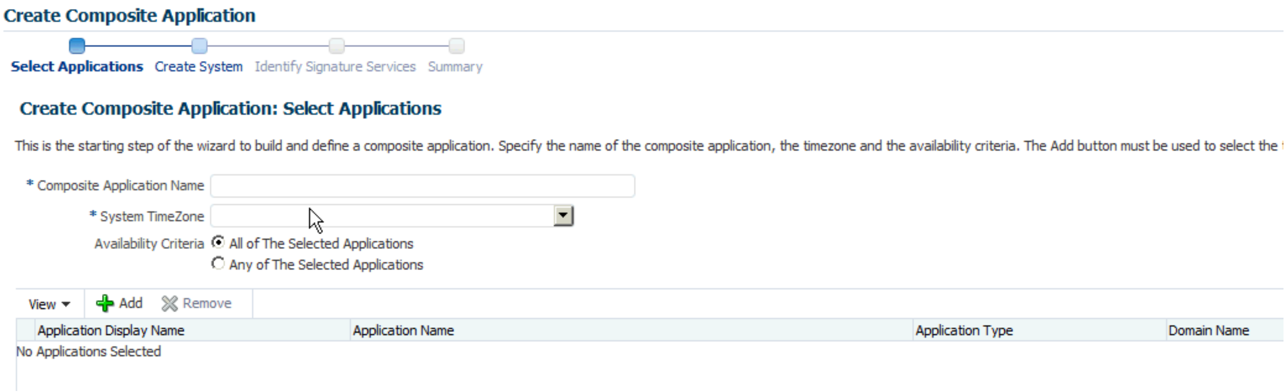

Create Composite Application

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

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 Select 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.

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.

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.

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

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* available from the Management page of the Enterprise Manager documentation library.

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.
- Instructions for configuring the Exalytics In-Memory System within Oracle Enterprise Manager.

Part III
Monitoring Oracle WebLogic Domains

This part describes how you can monitor Oracle WebLogic Domains.

- Monitoring WebLogic Domains

4 Monitoring WebLogic Domains

This section describes how to monitor WebLogic domains.
When using Enterprise Manager 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.

Updating the Agent Truststore

To update the Agent truststore (AgentTrust.jks), you use EMCTL. If the default demo certificate, or a self-signed certificate is being used on the WebLogic servers for t3s/iiops, then the Root CA certificate for this must be added to AgentTrust.jks in order for the Agent to be able to discover and monitor these WebLogic servers and J2EE applications using t3s. An EMCTL command is provided for this purpose.

```
emctl secure add_trust_cert_to_jks [-password <password> -trust_certs_loc <loc> -alias
<alias>]
```

Where:

- password = password to the `AgentTrust.jks` (if not specified, you will be prompted for the password at the command line)
- trust_certs_loc = location of the certificate file to import
- alias = alias for the certificate to import

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 4-1   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
```

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.

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
```

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.

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.

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:

5 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.

About Oracle SOA Management Pack Enterprise Edition

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

Table 5-1 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: • 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: • 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 *Enterprise Manager Administrator's Guide for Software and Server Provisioning and Patching*.
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.

Table 5-1 (Cont.) Highlights of Oracle SOA Management Pack Enterprise Edition

Feature	Benefit
Service level management	Enables you to monitor services from the end-user's perspective using service tests or synthetic transactions, model relationships between services and underlying IT components, and report on achieved service levels.
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.

6

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:

Supported Versions

For supported Service Bus versions, see the Enterprise Manager certification matrix on My Oracle Support:

To access the Enterprise Manager certification matrix, follow these steps:

1. Sign in to My Oracle Support: `http://support.oracle.com`
2. Click the **Certifications** tab.
3. In the **Certification Search**, from the **Product** list, select one of the following:
   - **Enterprise Manager Base Platform - OMS**, to view the certification for OMS.
   - **Enterprise Manager Base Platform - Agent**, to view the certification for Management Agent.
4. From the **Release** list, select release version, and then click **Search**.

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 the 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.

Understanding the Discovery Process

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

Table 6-1 Discovery Process

Step	Requirement	Description
1	Service Bus	Install the Service Bus software. **Note:** Before you launch the Service Bus Deployment Procedure, ensure that Sun JDK has been installed.
2	Enterprise Manager Cloud Control	Install Enterprise Manager. For information about installing the base release of Enterprise Manager Cloud Control, see the Enterprise Manager Basic Installation and Configuration Guide . Oracle recommends that you install the Enterprise Manager 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.
3	Oracle Management Agent (Management Agent)	Install Oracle Management Agent on the host where Service Bus is installed. If Service Bus and Enterprise Manager are on the same host, then you do not have to install a separate Management Agent. The Management Agent that comes with Enterprise Manager 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. You can install the Management Agent in one of the following ways: • Invoke the installer provided with Enterprise Manager, and select the installation type **Additional Management Agent**. Then apply the Agent patch on it. • Use the Agent Deploy application within the Enterprise Manager. For information about installing the Management Agent, see the Enterprise Manager Cloud Control Basic Installation and Configuration Guide.
4	Discovery in Enterprise Manager Cloud Control	Service Bus is automatically discovered when the Oracle WebLogic Domain to which it is deployed is discovered and added to Enterprise Manager.

Discovering Service Bus

The Service Bus deployed to Oracle WebLogic Managed Server is automatically discovered in Enterprise Manager 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

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 Enterprise Manager, first discover that Oracle WebLogic Manager Server in Enterprise Manager; 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 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 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. On 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 Enterprise Manager 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 Enterprise Manager.

   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, 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. 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".

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

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 Enterprise Manager, 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 displays the membership page that lists the Service Bus that is currently not being monitored. Click **OK**.

   Enterprise Manager 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 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.

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**.

Monitoring Service Bus in Enterprise Manager

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

For each Service Bus target being monitored, Enterprise Manager 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, Enterprise Manager 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.

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.
6. Click **Apply.**

Generating Service Bus Reports Using BI Publisher

> **Note:**
>
> Beginning with Enterprise Manager Cloud Control 13c Release 5, Oracle Analytics Publisher (formerly BI Publisher) must be accessed from a standalone Oracle Analytics Server. For more information, see Oracle Analytics Server in *Enterprise Manager Administrator's Guide*.

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.

System and Service

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

Table 6-2 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.

SOAP Test

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

Table 6-3 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 Configuation.`	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: `http://<host>:<port>/`. For example, `http://stade61.us.example.com:7001/`

7 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

List of Supported Versions

For the supported versions of the SOA Suite and the SOA Cloud Services (SOACS) for Enterprise Manager go to the Enterprise Manager certification matrix.

To access the Enterprise Manager certification matrix, follow these steps:

1. Sign in to My Oracle Support: `http://support.oracle.com`
2. Click the **Certifications** tab.
3. In the **Certification Search**, from the **Product** list, select one of the following:
   - **Enterprise Manager Base Platform - OMS**, to view the certification for OMS.

- **Enterprise Manager Base Platform - Agent**, to view the certification for Management Agent.

4. From the **Release** list, select release version, and then click **Search**.

Monitoring Templates

The following Oracle-certified default templates are being shipped for Enterprise Manager and Enterprise Manager agents. Table 7-1 describes the available templates, and the agents to which they apply:

Table 7-1 Monitoring Templates

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

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

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 7-1 New Domain Discovery

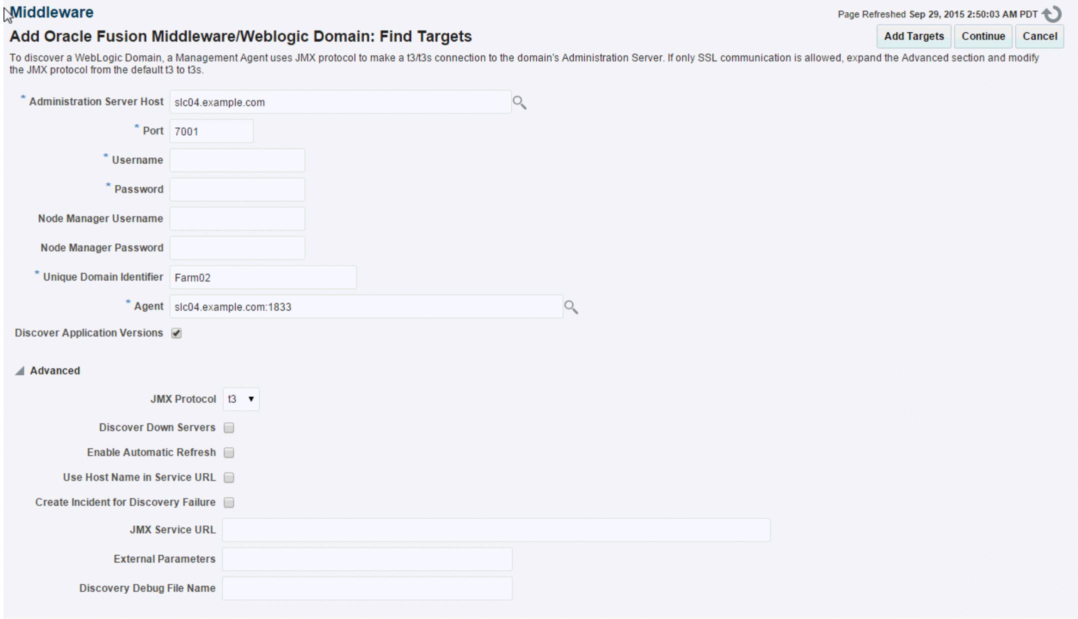

Middleware

Page Refreshed Sep 29, 2015 2:50:03 AM PDT

Add Oracle Fusion Middleware/Weblogic Domain: Find Targets

Add Targets | Continue | Cancel

To discover a WebLogic Domain, a Management Agent uses JMX protocol to make a t3/t3s connection to the domain's Administration Server. If only SSL communication is allowed, expand the Advanced section and modify the JMX protocol from the default t3 to t3s.

* Administration Server Host	slc04.example.com
* Port	7001
* Username	
* Password	
Node Manager Username	
Node Manager Password	
* Unique Domain Identifier	Farm02
* Agent	slc04.example.com:1833
Discover Application Versions	☑

Advanced

JMX Protocol	t3 ▾
Discover Down Servers	☐
Enable Automatic Refresh	☐
Use Host Name in Service URL	☐
Create Incident for Discovery Failure	☐
JMX Service URL	
External Parameters	
Discovery Debug File Name	

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 7-2 Upgrade Domain Discovery

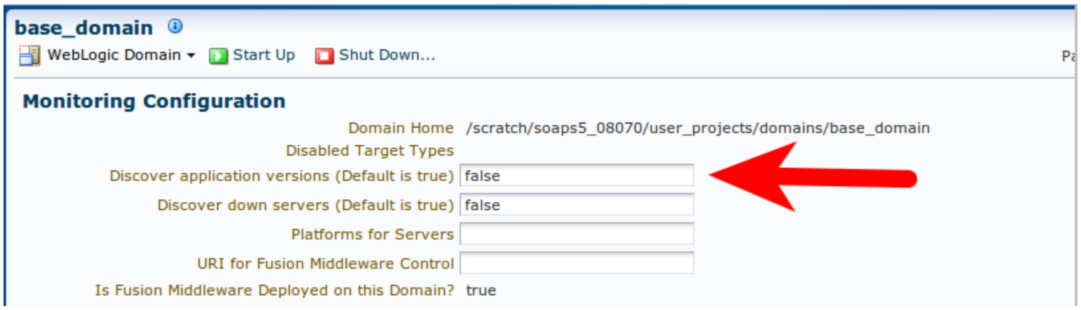

Monitoring Configuration

Domain Home	/scratch/soaps5_08070/user_projects/domains/base_domain
Disabled Target Types	
Discover application versions (Default is true)	false
Discover down servers (Default is true)	false
Platforms for Servers	
URI for Fusion Middleware Control	
Is Fusion Middleware Deployed on this Domain?	true

Note:

- If you have targets which were discovered with the **Discover Application Versions** box checked (which is the default, see Figure 7-1), 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 7-2.)
  - 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.

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 7-2 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 7-2   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`

Discovering the SOACS Instance Using the Hybrid Cloud Agent

You can discover a SOA instance on the Oracle Public Cloud in Enterprise Manager using the Hybrid Cloud Agent. To understand the architecture of the hybrid cloud see Enabling Hybrid Cloud Management in the *Oracle Enterprise Manager 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.

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.

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.

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.

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.

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 7-3 Metric and Collection Settings

Target Type	Metric Name	Collection Interval Update Steps
SOA Infrastructure	Response	Navigate to the associated weblogic server where SOA is deployed, to do so, follow these steps: 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. **Note:** This change is applicable to all the applications deployed in that WebLogic server.

Table 7-3 (Cont.) Metric and Collection Settings

Target Type	Metric Name	Collection Interval Update Steps
SOA Composite	Response	For SOA PS5 (11.1.1.6.0) or earlier, follow these steps: 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)** 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: 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.** **Note:** This change is applicable to all the SOA composites which are deployed in that SOA Infrastructure
SOA Composite	SOA Composite - Component Detail Metrics	Navigate to the associated SOA Infrastructure where soa composite is deployed. To do so, follow these steps: 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 7-3 (Cont.) Metric and Collection Settings

Target Type	Metric Name	Collection Interval Update Steps
SOA Composite	SOA Composite - Recoverable And Rejected Messages	Navigate to the associated SOA Infrastructure where soa composite is deployed. To do so, follow these steps: 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.**

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

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.

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.

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.**

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 11*g* Targets Only)
- Setting Search Criteria for Tracing an Instance
- Tracing an Instance Within a SOA Infrastructure
- Tracing Instance Across SOA Infrastructures

Configuring Instance Tracing (SOA 11*g* 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.

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

Instance Tracing for SOA 11g Targets

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

Table 7-4 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.

Instance Tracing for SOA 12c Targets

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

Table 7-5 Setting Search Criteria

Field	Description
Time	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. Additionally, you can add the following filters: • Instance Updated If you set this value to **None**, then it means that instance updated filter is not set at all. • Fault Occurred

Table 7-5 (Cont.) Setting Search Criteria

Field	Description
Composite	Use to restrict your search for business flows to a specific composite. 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: • **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	**Flow ID:** Use this to search for the flow ID of the business flow instance. **Flow Correlation ID**: Use this to search for the flow correlation ID of the business flow instance. **Initiating ECID**: Use this to search for the ECID of the business flow instance. **Flow Instance Name:** Use this to search for unique system and business identifiers that help you isolate a specific flow instance **Composite Instance Name:** Use this to specify the name or title of the composite instance name.
State	Select one of the following states: Select **Active** to search active instances. If you select a blank, then the filtering is ignored. • **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. Select **Inactive** to search inactive instances. If you select a blank, then the filtering is ignored. • **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.
Fault	Use to limit your search for business flows to only those with faults. If you leave this field blank, the Fault filter is ignored. To search for faults in any state, select **All.** To search for faults in a particular state, select one of the following: • **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.

Table 7-5   (Cont.) Setting Search Criteria

Field	Description
Fault Type	To search for all types of faults, select **All** To search for a particular type of fault, select one of the following: • **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	Use the **Name** field to enter a fault owner name. Ensure that the name entered is in the following format: `<partition>/<composite name>!<composite version>/<component name>` 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.** To drill down to a particular fault owner, select one of the following: • BPEL • BPMN • Mediator • Human Workflow • Decision • Spring • Case Management
Fault Details	You can fine grain your search parameters to drill down to granular result by providing all or some of the following details: • **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. Expand **Other** to display additional fields for filtering: • **HTTP Host** • **JNDI Name**
Restrict Search Rows	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. 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 this maximum limit set on the OMS, run the following command:`emctl set property -name oracle.sysman.core.uifwk.maxRows -value <max_limit_value>` **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.

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 7-5.
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.**

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 7-5.
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.**

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.

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.

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.

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.

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.

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

Generating SOA Reports Using BI Publisher

> **Note:**
>
> Beginning with Enterprise Manager Cloud Control 13c Release 5, Oracle Analytics Publisher (formerly BI Publisher) must be accessed from a standalone Oracle Analytics Server. For more information, see Oracle Analytics Server in *Enterprise Manager Cloud Control Administrator's Guide*.

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.

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.**

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.

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.**

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*.

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.

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)

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.

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 7-6 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.	• BPEL • BPMN • Mediator
Abort	Terminates the entire instance.	• BPEL • BPMN • Mediator
Continue	Ignores the fault and continues processing (marks the faulting activity as a success).	• BPEL • BPMN

Table 7-6 (Cont.) Overview of the Recovery Actions for Resolving Faults

Recovery Action	Description	Applicable To SOA Engine Type
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.	• BPEL • BPMN
Replay	Replays the entire scope again in which the fault occurred.	• BPEL • BPMN

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.

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`.

- To sort the fault details alphabetically, click the column header based on which you want to sort the details.
- To find out the number of rows to which the search results have been restricted, see the note below the table.

  For example, the following note appears if the rows were restricted to 20.

```
This table of search results is limited to 20 fault instances. Narrow the
results by using the search parameters.
```

Setting Search Criteria

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

Table 7-7 Setting Search Criteria

Field	Description
Time	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. 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.
Composite	Use to restrict your search for business flows to a specific 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.
Flow Instance	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.
Fault	Use to limit the search for business flows to only those with faults. If you leave this field blank, the Fault filter is ignored. To search for faults of any type, select **All or blank.** To search for faults in a particular type, select one of the following: • **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	To search for all types of faults, select **All.** To search for a particular type of fault, select one of the following: • **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 7-7 (Cont.) Setting Search Criteria

Field	Description
Fault Owner	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.** To drill down to a particular fault owner, select one of the following: • BPEL • BPMN • Mediator • Human Workflow • Decision • Spring • Case Management
Fault Details	You can fine grain your search parameters to drill down to granular result by providing all or some of the following details: • 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. Expand **Other** to display additional fields for filtering: • HTTP Host • JNDI Name
Restrict Search Rows	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. 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 this maximum limit set on the OMS, run the following command:`emctl set property -name oracle.sysman.core.uifwk.maxRows -value <max_limit_value>` **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.

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.

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.

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.

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.

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.

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`.

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.

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

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.

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.

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 esures 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 60mins/10mins = 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.
```

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.

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 7-7.
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 of 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.

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 7-9.
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.

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

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.
Job Failure Error	Provides details of the failed recovery jobs. The errors are shown from the last point of failure of the job. The details include: • **Failed Step Name,** the name of the step of the recovery job that failed. The job has two steps, mainly *pre-check* and *recover_faults.* The *recover_faults* 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. To verify if the faults were recovered, search for the faults for the given composite and time period again. 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.

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

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=<composite1 target name>, <composite 2 target name>
   variable.BatchDelay=300
   variable.BatchSize=10
   variable.EnableBatching=1
   variable.EngineType=BPEL
   variable.ErrorMsg=xxxx
   variable.FaultStartTime=01-01-2013 00:00:00 PST
   variable.FaultEndTime=01-02-2013 00:00:00 PST
   variable.FaultTimePeriod=Custom
   variable.RecoveryAction=Continue
   ```

   > **Note:**
   >
   > Currently, Oracle supports only one SOA-Infrastructure target to be entered in the *target_list* property.

4. Run the following command to submit a bulk recovery job with the updated property file as an input:

   ```
   emcli create_job -name=bulk522 -job_type=SOABulkRecovery -input_file=property
   file:/tmp/temp.properties
   ```

5. Set the preferred credentials or named credentials for the WebLogic Domain and SOA Server Host. By default, the job uses the preferred credentials, that is, WebLogic Administrator Credentials for the WebLogic domain and Normal Host Credentials for the SOA Server hosts.

   To set the preferred credentials, run the following commands:

   ```
   Setting WebLogic Domain Credentials:
   emcli set_preferred_credential -target_type=weblogic_domain -target
   name=<weblogic domain target name> -set_name=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
```

Viewing the Submitted Jobs and Outputs Using EMCLI

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

Table 7-8 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`	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.
	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**.
		Run the following command to get the Execution ID of the job:
		`emcli get_jobs -targets=<SOA-Infra target name>:oracle_soainfra -format=name:csv | grep BULK521`
		Use the Execution ID in the following command to view the details of the job submitted:
		`emcli get_job_execution_detail -execution=D4081BAB8942F246E040F00A5AA93E04 -xml -showOutput`

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.

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

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*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.**

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.**
    b. In Batch Time Period, enter **10 mins.** This would mean that, every batch would recover faults in 10mins time window. Since you have selected a time window of one hour (3:00 am to 4:00 am), there will be 60 / 10 = 6 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 **Later**, and provide a date and time to schedule the job.
   b. From repeat menu, select **Do not repeat.**
   c. Enter Frequency as **1 day.**
6. Click **Submit.**

Generating Error Hospital Reports

Use the Error Hospital page to view an aggregate count of errors that have occurred in all SOA Composites deployed in the SOA Infrastructure. This page does not list out individual faulted instances. To view the individual flows that have faults, go to the Faults and Rejected Messages tab on the SOA Infrastructure Home page.

The Error Hospital page is available at the SOA Infrastructure level, where system-wide faults data is aggregated. When accessed at the partition level, the Error Hospital page is limited to faults data associated only with that partition.

The Error Hospital page is arranged in the following sections:

- **Search Region:** You can update the necessary filters available in the Search section to drill down to a more granular result that meets your requirements. By default, the total faults that have occurred across all instances created in the last 24 hours is displayed. You must provide the **Instance Created From** and **Instance Created To** values as they are mandatory fields. In addition to these values, you may specify a time window for the fault to restrict your query to a specific time in the past.

  Additionally, you can select the fault attribute by which data is aggregated. For example, if you select Fault Code, each row in the first column represents a specific code and the remaining columns show the fault statistics aggregated for each code.
- **Error Hospital Report Table:** This table displays a report based on the filters specified in the search region. The data is always aggregated by one of the primary fault attributes selected from the list such as **Fault Name, Fault Code,** and so on. The default aggregation is by **Fault Name.** This report enables you to assess the error trends. For example, aggregate by Fault Code to see which code has the most faults. You can then select a single row which has maximum faults from the table, and perform a bulk recovery.
- **Charts Region:** The details of the Error Hospital Report are also available in a chart form. Essentially, the top faults aggregated by **Fault Name** are represented in a bar chart. The pie chart depicts the recovery required faults as against non-recoverable faults.

The major advantages are:

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 7-9   Setting Search Criteria for Error Hospital

Field	Description
Time	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. Additionally, you can add the following filters: • Instance Updated • Fault Occurred
Composite	Use to restrict your search for business flows to a specific composite. You can select the following option: • 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.
State	Select one of the following states: Select **Active** to search active instances. If you select a blank, then the filtering is ignored. • **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. Select **Inactive** to search inactive instances. If you select a blank, then the filtering is ignored. • **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.
Fault	Use to limit the search for business flows to only those with faults. If you leave this field blank, the Fault filter is ignored. To search for faults of any type, select **All or blank.** To search for faults in a particular type, select one of the following: • **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.

Table 7-9 (Cont.) Setting Search Criteria for Error Hospital

Field	Description
Fault Type	To search for all types of faults, select **All** To search for a particular type of fault, select one of the following: • **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	Use the **Name** field to enter a fault owner name. Ensure that the name entered is in the following format: `<partition>/<composite name>!<composite version>/<component name>` 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.** To drill down to a particular fault owner, select one of the following: • BPEL • BPMN • Mediator • Human Workflow • Decision • Spring • Case Management
Fault Details	You can fine grain your search parameters to drill down to granular result by providing all or some of the following details: • **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. Expand **Other** to display additional fields for filtering: • **HTTP Host** • **JNDI Name**
Restrict Search Rows	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. 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 this maximum limit set on the OMS, run the following command:`emctl set property -name oracle.sysman.core.uifwk.maxRows -value <max_limit_value>` **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.

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

- Generating an Error Hospital Report
- Customizing the Error Hospital Report

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 7-9.
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.

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.

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.**

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

Discovery

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

Table 7-10 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.

Monitoring

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

Table 7-11 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 `unsuspend_soametrics.pl` script available in the agent scripts directory under `%emd_root%`.

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 7-12 Error Message

Error Message	Workaround Steps
`Instance Tracing Job Fails`	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.

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 7-13 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.

Fault Management

This section contains the troubleshooting information for fault management:

- Bulk Recovery
- Fault Search and Recovery
- Fault Management and Instance Tracing Errors

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 7-14 Error Message

Error Message	Workaround Steps

`java.lang.IllegalArgumentExcepton: Invalid Job Identifier! The specified identifier does not match any valid fault recovery jobs.`
`java.lang.IllegalStateException The results job xxxx are not available because the processing has not yet completed.`
	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.
	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:
	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**.

`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]`	Ensure that the SOA Infrastructure is up and running, and submit another bulk recovery job with the same parameters.

Fault Search and Recovery

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

Table 7-15 Error Message

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

Fault Management and Instance Tracing Errors

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

Table 7-16 Error Message

Error Message	Workaround Steps
`Error occured when getting faults Java.rmi.RemoteException: EJB Exception: ; nested exception is: java.lang.RuntimeException: java.lang.RuntimeException: weblogic.jdbc.extensions.PoolDisabledSQLException: weblogic.common.resourcepool.ResourceDisabledException: Pool SOALocalTxDataSource is Suspended, cannot allocate resources to applications.`	Ensure the SOA Database is up and running.
`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]`	Ensure the SOA Database is up and running.
`Error occured 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]`	Ensure the SOA Database is up and running.

Information Publisher Reports

This section lists report related errors.

Table 7-17 Error Message

Error Message	Workaround Steps
`Report generation fails due to invalid database details.`	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.

BI Publisher Reports

This section lists BI Publisher report related errors.

Table 7-18 Error Message

Error Message	Workaround Steps
`Exception Encountered For One of SOA BIP Report If SOA Dehydration Is Not Configured`	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: `The report cannot be rendered because of an error, please contact the administrator. Parameter name: P_PARTITION_NAME Can not establish database connection(EMSOA)` 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: 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. Once you have created the EMSOA data source, the issue should be resolved.

Systems and Services

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

Table 7-19 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.

BPEL Recovery

The following error occurs when invalid credentials are provided.

Table 7-20 Error Message

Error Message	Workaround Steps
`Invalid Host and WebLogic Domain Credentials`	For the BPEL Recovery functionality to work, the host credentials and WebLogic Domain credentials must to be available in the preferred credential store. Set the valid credentials and try again.

SOA License Issue

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

Table 7-21 Error Message

Error Message	Workaround Steps
`The page requested is part of the SOA Management Pack EE.`	The SOA Management Pack EE must be enabled for the specific SOA Infrastructure target. To enable the license, follow these steps: 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.

Dehydration Store Issue

Data is not displayed on the Dehydration Store page.

Table 7-22 Error Message

Error Message	Workaround Steps
`Data is not displayed in the Dehydration Store page.`	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: 1. Compare the Database Host and SID value of the database target with the value collected for the WebLogic Server JDBC Datasource configuration. 2. If the values are different, select **Services** from the **Targets** menu. Select **DataSources**, then select **SOALocalTxtSource**, then click **Connection Pool** to update the Datasource Connection URL .

Part V
Managing Oracle Business Intelligence

The chapter in this part describes how you can discover, monitor, and administer Oracle Business Intelligence instance and Oracle Essbase targets in Enterprise Manager.

This part contains the following chapter:

8

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 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. 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

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

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.

A BI Instance target consists of a number of components, which can be monitored individually using Enterprise Manager.

Table 8-1 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, and graph generation. It also enables Oracle BI Server query access to Hyperion Financial Management and Oracle Online Analytical Processing (OLAP) data sources.

Oracle Essbase

Oracle Essbaseis 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.

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.

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

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.

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.

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.

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

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 8-2 describes the elements of the **Summary** section.

Table 8-2 Target General and Availability Summary

Element	Description
Up Since	*(Displayed only when the target is up)* Time the target was last started successfully.
Down Since	*(Displayed only when the target is down)* Time the target was last stopped.
Availability	Percentage availability of the target.
Version	Version of the target software.
Oracle Home	Location of the target binaries.
Oracle Instance	Location of the target content files, metadata, configuration files and log files.
Port	Port used by the target for communication.
Running Applications (Only for Essbase Server targets)	Number of Essbase applications currently up and running.
Unexposed Applications (Only for Essbase Server targets)	Number of Essbase applications currently not being accessed by any user.
Connected Users (Only for Essbase Server targets)	Number of users currently connected through one or more of the applications.
Storage Type (Only for Essbase application targets)	Type of data storage used by the application.

Table 8-2 (Cont.) Target General and Availability Summary

Element	Description
Cubes (Only for Essbase application targets)	Number of cubes contained in the application.
Query Tracking (Only for Essbase application targets)	Whether or not query tracking, that is, tracking data combinations having a large number of data values that require aggregation, is enabled.
Memory Usage (MB) (Only for Essbase application targets)	Memory used by the application in MB.
Threads (Only for Essbase application targets)	Number of application threads.

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 8-3 describes the elements of the **General** section.

Table 8-3 Target Status and Availability History

Element	Description
Current Status	Current status of the target, whether it is up and running, or down.
Up Since	*(Displayed only when the target is up)* Time the target was last started successfully.
Down Since	*(Displayed only when the target is down)* Time the target was last stopped.
Availability (%)	Percentage availability of the target.

Table 8-3 (Cont.) Target Status and Availability History

Element	Description
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.

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.

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.

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.

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.

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 8-4 describes the elements of the **Monitoring and Diagnostics** section.

Table 8-4 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.

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 8-5 describes the elements of the Alert History page.

Table 8-5 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.

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.*

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 8-6 describes the elements of the Log Messages page.

Table 8-6   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.

Table 8-6 (Cont.) Target Log Messages

Element	Description
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.
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 `oracle` 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.

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.

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.

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.*

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

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 8-7 describes the elements of the Dashboard Reports page.

Table 8-7 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.

Table 8-7 (Cont.) Oracle Business Intelligence Dashboard Reports

Element	Description
Compile Time	Time taken to convert all logical requests made for a particular dashboard. **Note:** In the Top Users by Resource Usage in Last 7 Days reports, this element represents the time taken to convert all logical requests made by a particular user, to physical requests.
Failed Logical Requests	Number of logical requests made for the dashboard that failed. **Note:** In the Top Users by Resource Usage in Last 7 Days reports, this element represents the number of logical requests made by a particular user that failed.
Total Logical Requests	Total number of logical requests made for the dashboard. **Note:** In the Top Users by Resource Usage in Last 7 Days reports, this element represents the total number of logical requests made by a particular user.

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 8-8 describes the elements of the Scheduler Reports page.

Table 8-8 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.

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 8-9 describes the elements of the **Metrics** section.

Table 8-9 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.

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 8-10 describes the elements of this section.

Table 8-10 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.

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 8-11 describes the elements of this section.

Table 8-11 Oracle Essbase Application Data Storage Details

Element	Description
Name	Name of the cube.
Dimensions	Number of dimensions the cube has.

Table 8-11 (Cont.) Oracle Essbase Application Data Storage Details

Element	Description
Connected Users	Number of users currently connected to the cube data.
Locks	Number of data block locks currently held on the cube.
Data Cache Size (KB)	Size, in KB, of the buffer in memory that holds uncompressed data blocks.

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

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

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.

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.

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.

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 8-12 describes the elements of the Monitoring Configuration page.

Table 8-12 Target Monitoring Configuration

Element	Description
Canonical Path	Component path of the form `instance_name/component_name`.
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.

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

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.*

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 8-13 describes the elements of the Monitoring Credentials page.

Table 8-13 Oracle Business Intelligence Instance Monitoring Credentials

Element	Description
Database Type	Type of the database.
Hostname	Name of the host on which the database is installed.
Port	Port used for communicating with the database.
Service Name	Name of the database service.
Username	User name used for database login.
Password	Password used for database login.

Part VI

Using JVM Diagnostics

The chapters in this part provide information regarding JVM Diagnostics.

- Introduction to JVM Diagnostics
- Using JVM Diagnostics
- Troubleshooting JVM Diagnostics

For more information about JVMD Hybrid Cloud, see Deploying JVMD for Hybrid Cloud in the *Enterprise Manager Administrator's Guide.*

9

Introduction to JVM Diagnostics

This section provides an overview of JVM Diagnostics.
It contains the following:

- Overview
- Supported Platforms and JVMs
- User Roles

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 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

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.

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.

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.

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.

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.

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.

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.

Supported Platforms and JVMs

To access the Enterprise Manager certification matrix, follow these steps:

1. Sign in to My Oracle Support: `http://support.oracle.com`
2. Click the **Certifications** tab.
3. In the **Certification Search**, from the **Product** list, select one of the following:
   - **Enterprise Manager Base Platform - OMS**, to view the certification for OMS.
   - **Enterprise Manager Base Platform - Agent**, to view the certification for Management Agent.
4. From the **Release** list, select release version, and then click **Search**.

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 Administrator's Guide*.

10 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

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 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 in 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.

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 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–5

     The 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 his 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.

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 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.

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 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_$LATCHNAME` fixed views in the target database.
     >
     > To grant select privileges to a user such as jvmadmin, enter a command as follows:
     >
     > ```
     > SQL> grant select on SYS.GV_$LATCHNAME to jvmadmin
     > ```
     >
     > 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.

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 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**.

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

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.

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

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).

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.

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.

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.

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.

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.

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.

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

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).

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 for details.
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.

> **Note:**
>
> Agent Diagnostics Charts (raw data), JVM and Host States (raw data), JVM Heap Memory Usage (raw data), and Thread (Active) States (raw data) metric groups are deprecated from JAVA Virtual Machine Pools.

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.

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 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.

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 10-1 Adding Cluster Database

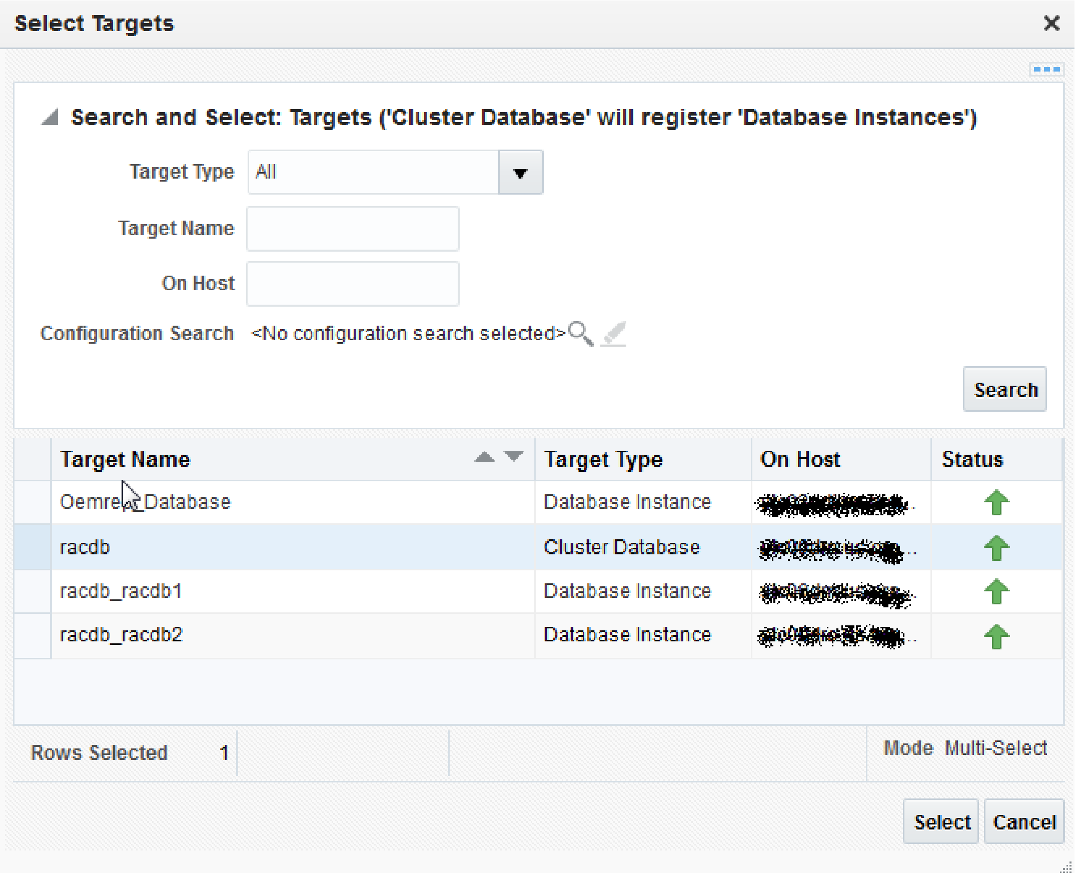

Select Targets

Search and Select: Targets ('Cluster Database' will register 'Database Instances')

Target Type: All

Target Name:

On Host:

Configuration Search: <No configuration search selected>

Target Name	Target Type	On Host	Status
Oemrep_Database	Database Instance	[redacted]	↑
racdb	Cluster Database	[redacted]	↑
racdb_racdb1	Database Instance	[redacted]	↑
racdb_racdb2	Database Instance	[redacted]	↑

Rows Selected 1

Mode Multi-Select

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 10-2 After registering cluster database

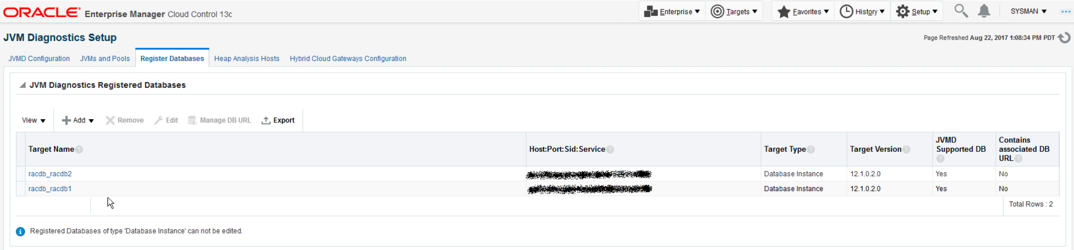

JVM Diagnostics Setup

Page Refreshed Aug 22, 2017 1:08:34 PM PDT

JVM Diagnostics Registered Databases

Target Name	Host:Port:Sid:Service	Target Type	Target Version	JVMD Supported DB	Contains associated DB URL
racdb_racdb2	[redacted]	Database Instance	12.1.0.2.0	Yes	No
racdb_racdb1	[redacted]	Database Instance	12.1.0.2.0	Yes	No

Total Rows : 2

Registered Databases of type 'Database Instance' can not be edited.

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 to connect to application via configured data source) in **Manage DB URL**.

Figure 10-3 Manage Database URL

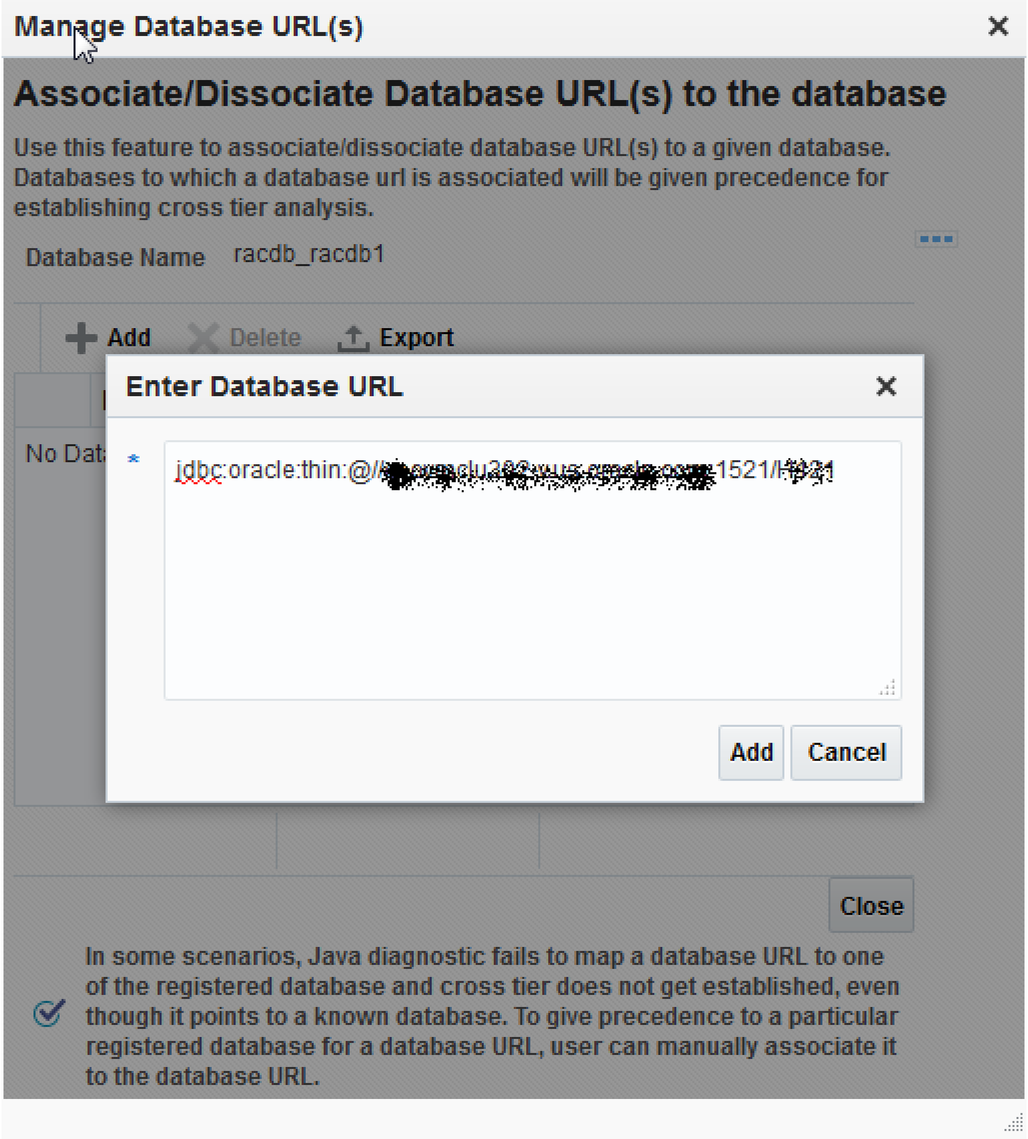

Manage Database URL(s)

Associate/Dissociate Database URL(s) to the database

Use this feature to associate/dissociate database URL(s) to a given database. Databases to which a database url is associated will be given precedence for establishing cross tier analysis.

Database Name racdb_racdb1

No Dat

Enter Database URL

* jdbc:oracle:thin:@//[redacted]1521/[redacted]

In some scenarios, Java diagnostic fails to map a database URL to one of the registered database and cross tier does not get established, even though it points to a known database. To give precedence to a particular registered database for a database URL, user can manually associate it to the 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.

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**.

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

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**.

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.

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.

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.

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.

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.

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` & `sqllplus` are present.
- There should be sufficient disk space in the system temp directory.
- A JVM Diagnostics DB User must have been created using the `create_jvm_diagnostic_db_user script`. 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.

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.

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**.

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.
- **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.

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.

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 help 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.

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.

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.

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.

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.

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.

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.

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.

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.**

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.

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.

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.

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

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.

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.

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.

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.

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.

Using Java Workload Explorer

Java Workload Explorer provides a detailed view of all performance statistics associated with the JVM and JVM Pool targets.

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.

Performance Analysis and Search Criteria

The Performance Analysis menu provides the following features:

Figure 10-4   Performance Analysis

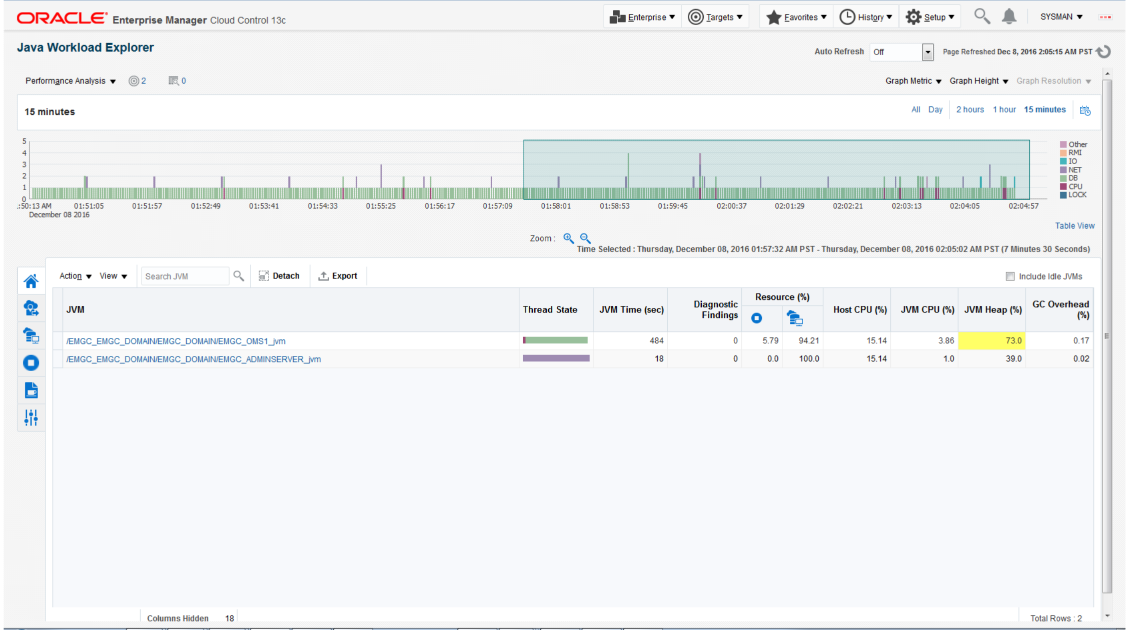

Auto Refresh: Off

Page Refreshed Dec 8, 2016 2:05:15 AM PST

15 minutes

Time Selected : Thursday, December 08, 2016 01:57:32 AM PST - Thursday, December 08, 2016 02:05:02 AM PST (7 Minutes 30 Seconds)

JVM	Thread State	JVM Time (sec)	Diagnostic Findings	Resource (%)		Host CPU (%)	JVM CPU (%)	JVM Heap (%)	GC Overhead (%)
				●					
/EMGC_EMGC_DOMAIN/EMGC_DOMAIN/EMGC_OMS1_jvm		484	0	5.79	94.21	15.14	3.86	73.0	0.17
/EMGC_EMGC_DOMAIN/EMGC_DOMAIN/EMGC_ADMINSERVER_jvm		18	0	0.0	100.0	15.14	1.0	39.0	0.02

Columns Hidden 18

Total Rows : 2

- 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 10-5 Search Criteria

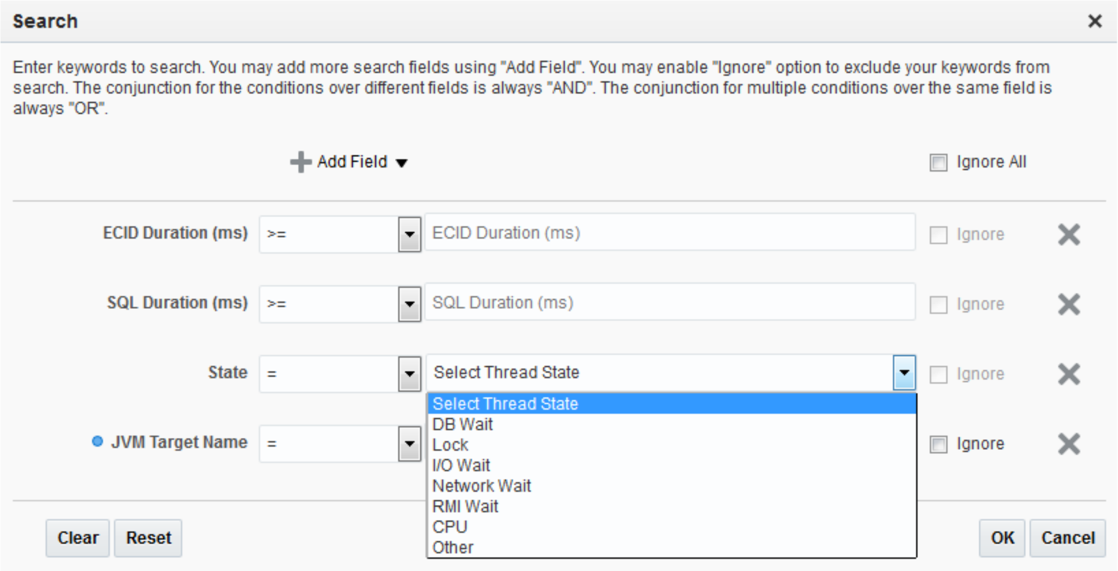

Search

Enter keywords to search. You may add more search fields using "Add Field". You may enable "Ignore" option to exclude your keywords from search. The conjunction for the conditions over different fields is always "AND". The conjunction for multiple conditions over the same field is always "OR".

ECID Duration (ms) >= ECID Duration (ms) □ Ignore

SQL Duration (ms) >= SQL Duration (ms) □ Ignore

State = Select Thread State □ Ignore

JVM Target Name = □ Ignore

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.

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).

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 10-6 Add to Search

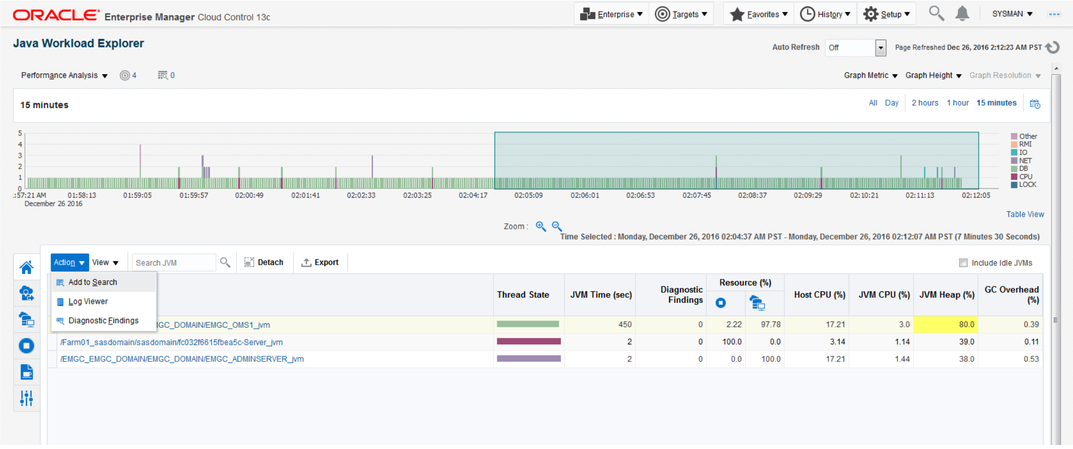

Auto Refresh: Off

Page Refreshed Dec 26, 2016 2:12:23 AM PST

Time Selected : Monday, December 26, 2016 02:04:37 AM PST - Monday, December 26, 2016 02:12:07 AM PST (7 Minutes 30 Seconds)

	Thread State	JVM Time (sec)	Diagnostic Findings	Resource (%)		Host CPU (%)	JVM CPU (%)	JVM Heap (%)	GC Overhead (%)
…GC_DOMAIN/EMGC_OMS1_jvm		450	0	2.22	97.78	17.21	3.0	80.0	0.39
/Farm01_sasdomain/sasdomain/fc032f6615fbea5c-Server_jvm		2	0	100.0	0.0	3.14	1.14	39.0	0.11
/EMGC_EMGC_DOMAIN/EMGC_DOMAIN/EMGC_ADMINSERVER_jvm		2	0	0.0	100.0	17.21	1.44	38.0	0.53

- Add to Set: It adds the element to the Set Criteria.

Figure 10-7 Add to set

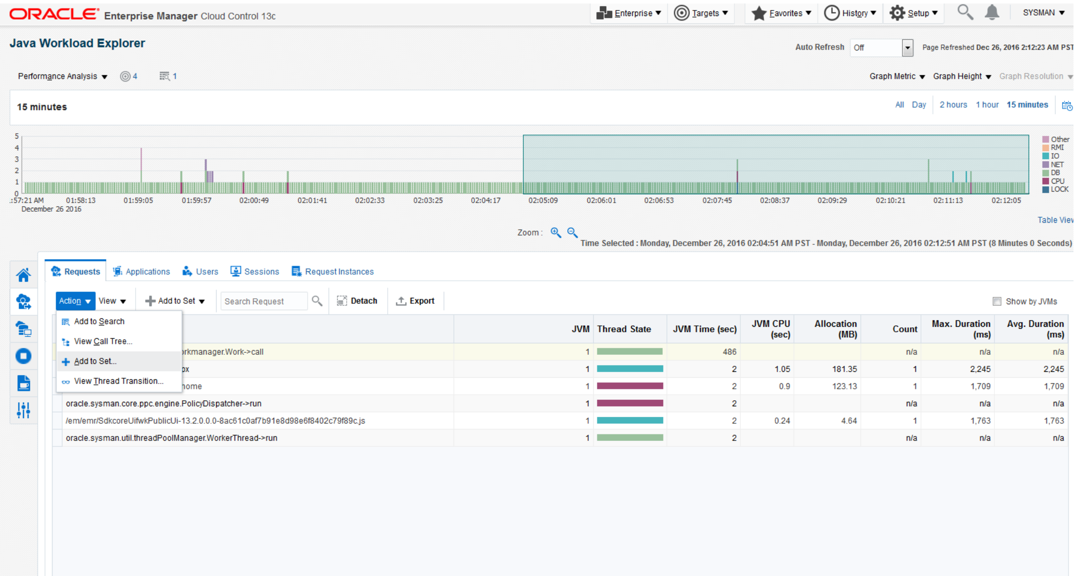

Auto Refresh: Off

Page Refreshed Dec 26, 2016 2:12:23 AM PST

Time Selected : Monday, December 26, 2016 02:04:51 AM PST - Monday, December 26, 2016 02:12:51 AM PST (8 Minutes 0 Seconds)

	JVM	Thread State	JVM Time (sec)	JVM CPU (sec)	Allocation (MB)	Count	Max. Duration (ms)	Avg. Duration (ms)
…rkmanager.Work->call	1		486			n/a	n/a	n/a
…px	1		2	1.05	181.35	1	2,245	2,245
…home	1		2	0.9	123.13	1	1,709	1,709
oracle.sysman.core.ppc.engine.PolicyDispatcher->run	1		2			n/a	n/a	n/a
/em/emr/SdkcoreUifwkPublicUi-13.2.0.0.0-8ac61c0af7b91e8d98e6f8402c79f89c.js	1		2	0.24	4.64	1	1,763	1,763
oracle.sysman.util.threadPoolManager.WorkerThread->run	1		2			n/a	n/a	n/a

- Log Viewer (Displays Log Messages).
- 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).

The tabs are:

- Overview:

Figure 10-8   Overview Tab

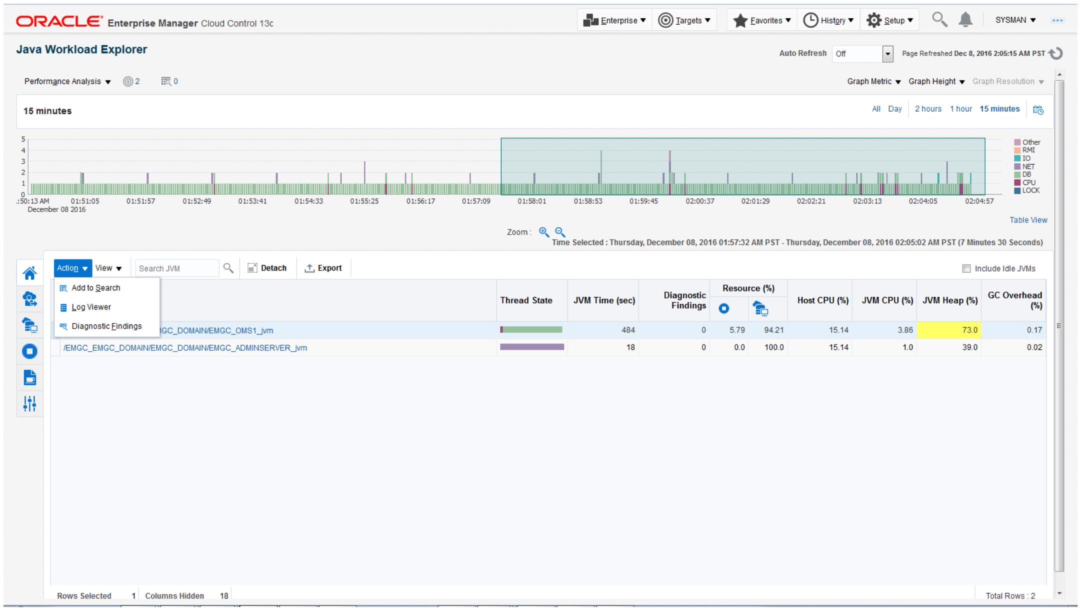

Time Selected : Thursday, December 08, 2016 01:57:32 AM PST - Thursday, December 08, 2016 02:05:02 AM PST (7 Minutes 30 Seconds)

	Thread State	JVM Time (sec)	Diagnostic Findings	Resource (%)		Host CPU (%)	JVM CPU (%)	JVM Heap (%)	GC Overhead (%)
GC_DOMAIN/EMGC_OMS1_jvm		484	0	5.79	94.21	15.14	3.86	73.0	0.17
/EMGC_EMGC_DOMAIN/EMGC_DOMAIN/EMGC_ADMINSERVER_jvm		18	0	0.0	100.0	15.14	1.0	39.0	0.02

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 10-9   Request Tab

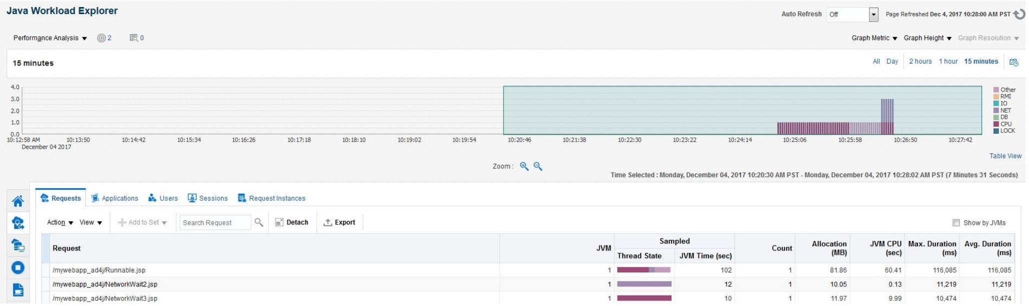

Time Selected : Monday, December 04, 2017 10:20:30 AM PST - Monday, December 04, 2017 10:28:02 AM PST (7 Minutes 31 Seconds)

Request	JVM	Sampled		Count	Allocation (MB)	JVM CPU (sec)	Max. Duration (ms)	Avg. Duration (ms)
		Thread State	JVM Time (sec)					
/mywebapp_ad4j/Runnable.jsp	1		102	1	81.86	60.41	116,085	116,085
/mywebapp_ad4j/NetworkWait2.jsp	1		12	1	10.05	0.13	11,219	11,219
/mywebapp_ad4j/NetworkWait3.jsp	1		10	1	11.97	9.99	10,474	10,474

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 10-10   External Resources Tab

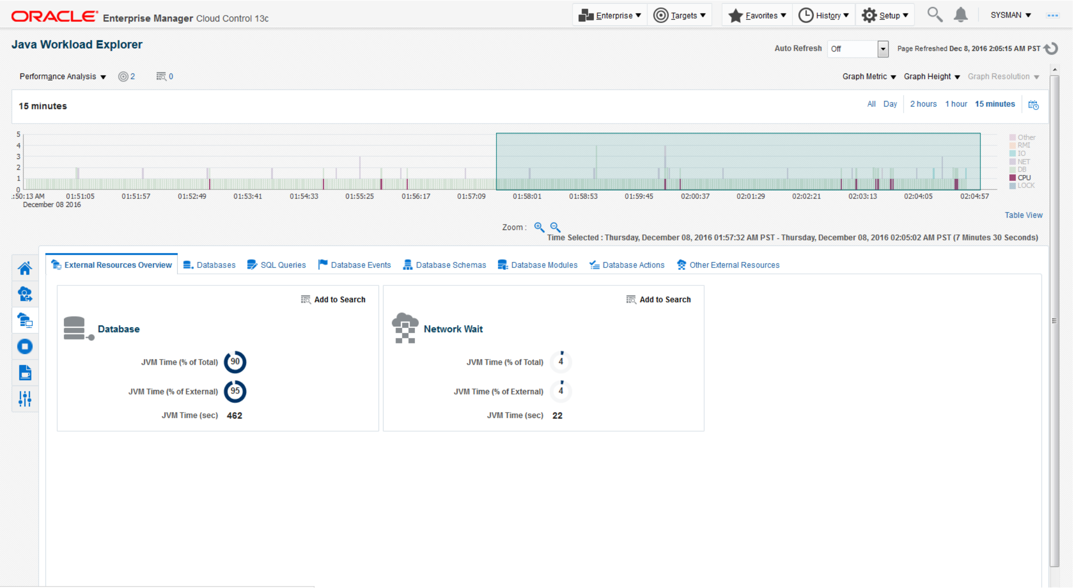

Time Selected : Thursday, December 08, 2016 01:57:32 AM PST - Thursday, December 08, 2016 02:05:02 AM PST (7 Minutes 30 Seconds)

Database
JVM Time (% of Total) 90
JVM Time (% of External) 95
JVM Time (sec) 462

Network Wait
JVM Time (% of Total) 4
JVM Time (% of External) 4
JVM Time (sec) 22

– 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 10-11   Internal Resources Tab

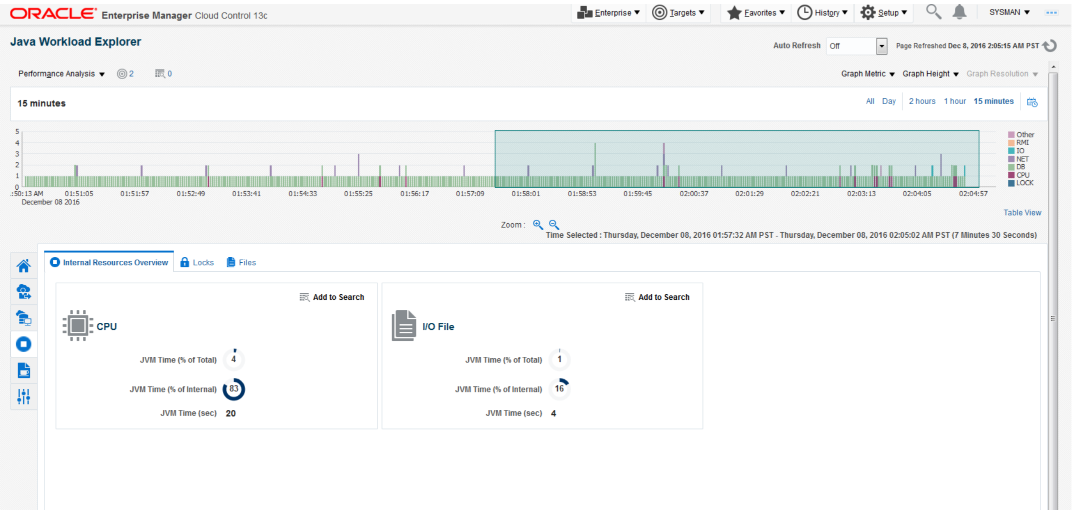

Auto Refresh Off

Page Refreshed Dec 8, 2016 2:05:15 AM PST

Time Selected : Thursday, December 08, 2016 01:57:32 AM PST - Thursday, December 08, 2016 02:05:02 AM PST (7 Minutes 30 Seconds)

CPU		I/O File	
JVM Time (% of Total)	4	JVM Time (% of Total)	1
JVM Time (% of Internal)	83	JVM Time (% of Internal)	16
JVM Time (sec)	20	JVM Time (sec)	4

  - 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 10-12 Code Tab

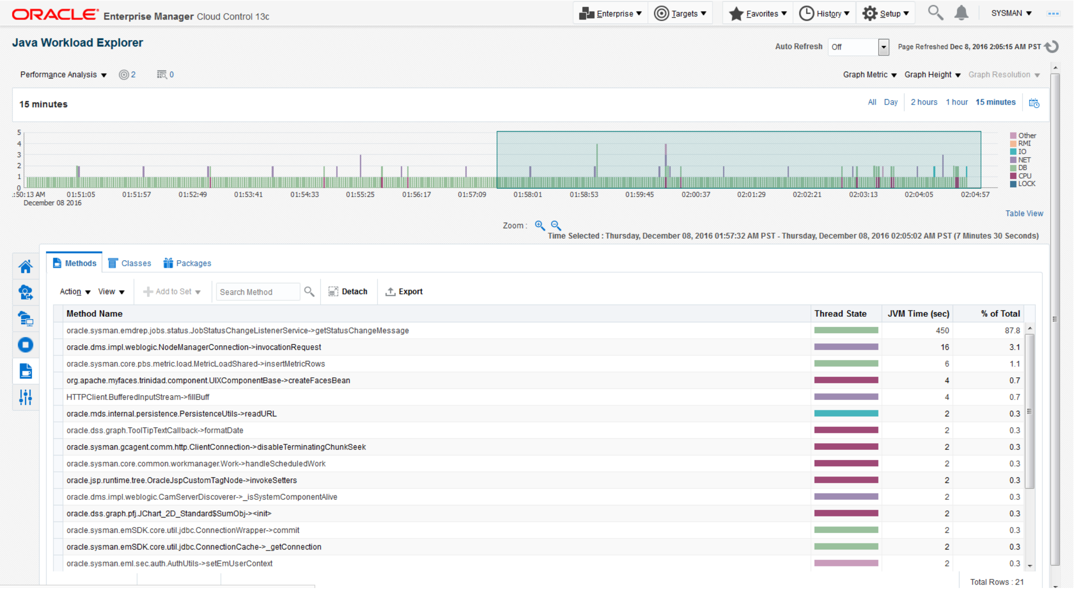

Time Selected : Thursday, December 08, 2016 01:57:32 AM PST - Thursday, December 08, 2016 02:05:02 AM PST (7 Minutes 30 Seconds)

Method Name	Thread State	JVM Time (sec)	% of Total
oracle.sysman.emdrep.jobs.status.JobStatusChangeListenerService->getStatusChangeMessage	🟩	450	87.8
oracle.dms.impl.weblogic.NodeManagerConnection->invocationRequest	🟪	16	3.1
oracle.sysman.core.pbs.metric.load.MetricLoadShared->insertMetricRows	🟩	6	1.1
org.apache.myfaces.trinidad.component.UIXComponentBase->createFacesBean	🟥	4	0.7
HTTPClient.BufferedInputStream->fillBuff	🟪	4	0.7
oracle.mds.internal.persistence.PersistenceUtils->readURL	🟦	2	0.3
oracle.dss.graph.ToolTipTextCallback->formatDate	🟥	2	0.3
oracle.sysman.gcagent.comm.http.ClientConnection->disableTerminatingChunkSeek	🟥	2	0.3
oracle.sysman.core.common.workmanager.Work->handleScheduledWork	🟥	2	0.3
oracle.jsp.runtime.tree.OracleJspCustomTagNode->invokeSetters	🟥	2	0.3
oracle.dms.impl.weblogic.CamServerDiscoverer->_isSystemComponentAlive	🟪	2	0.3
oracle.dss.graph.pfj.JChart_2D_Standard$SumObj-><init>	🟥	2	0.3
oracle.sysman.emSDK.core.util.jdbc.ConnectionWrapper->commit	🟩	2	0.3
oracle.sysman.emSDK.core.util.jdbc.ConnectionCache->_getConnection	🟩	2	0.3
oracle.sysman.eml.sec.auth.AuthUtils->setEmUserContext	▨	2	0.3

Total Rows : 21

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 10-13 Threads Tab

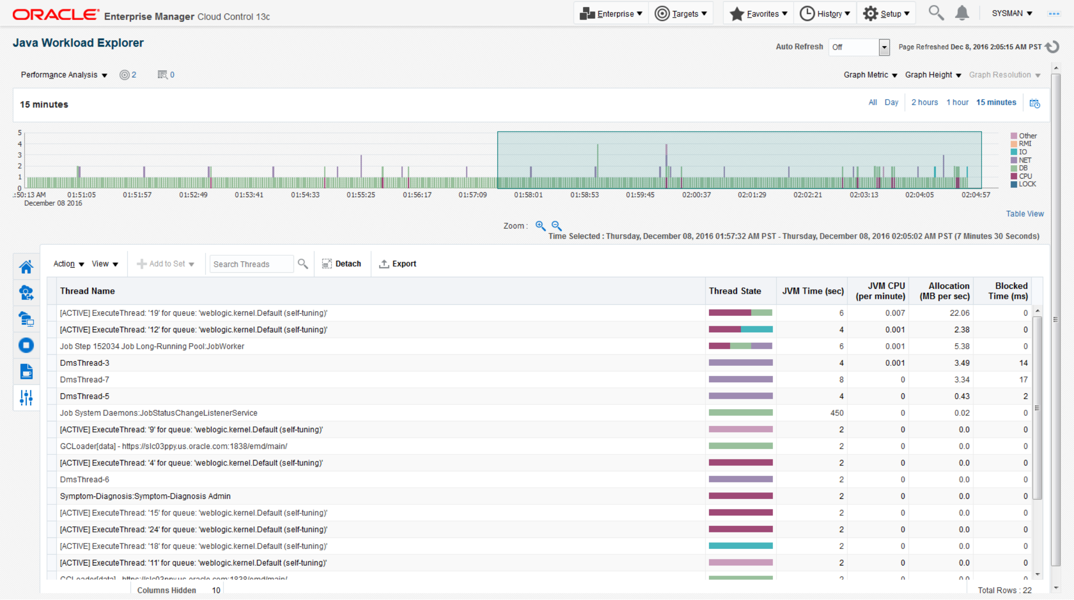

Time Selected : Thursday, December 08, 2016 01:57:32 AM PST - Thursday, December 08, 2016 02:05:02 AM PST (7 Minutes 30 Seconds)

Thread Name	Thread State	JVM Time (sec)	JVM CPU (per minute)	Allocation (MB per sec)	Blocked Time (ms)
[ACTIVE] ExecuteThread: '19' for queue: 'weblogic.kernel.Default (self-tuning)'	🟥🟩	6	0.007	22.06	0
[ACTIVE] ExecuteThread: '12' for queue: 'weblogic.kernel.Default (self-tuning)'	🟥🟦	4	0.001	2.38	0
Job Step 152034 Job Long-Running Pool:JobWorker	🟥🟩🟪	6	0.001	5.38	0
DmsThread-3	🟪	4	0.001	3.49	14
DmsThread-7	🟪	8	0	3.34	17
DmsThread-5	🟪	4	0	0.43	2
Job System Daemons:JobStatusChangeListenerService	🟩	450	0	0.02	0
[ACTIVE] ExecuteThread: '9' for queue: 'weblogic.kernel.Default (self-tuning)'	▨	2	0	0.0	0
GCLoader[data] - https://slc03ppy.us.oracle.com:1838/emd/main/	🟩	2	0	0.0	0
[ACTIVE] ExecuteThread: '4' for queue: 'weblogic.kernel.Default (self-tuning)'	🟥	2	0	0.0	0
DmsThread-6	🟪	2	0	0.0	0
Symptom-Diagnosis:Symptom-Diagnosis Admin	🟥	2	0	0.0	0
[ACTIVE] ExecuteThread: '15' for queue: 'weblogic.kernel.Default (self-tuning)'	🟥	2	0	0.0	0
[ACTIVE] ExecuteThread: '24' for queue: 'weblogic.kernel.Default (self-tuning)'	🟥	2	0	0.0	0
[ACTIVE] ExecuteThread: '18' for queue: 'weblogic.kernel.Default (self-tuning)'	🟦	2	0	0.0	0
[ACTIVE] ExecuteThread: '11' for queue: 'weblogic.kernel.Default (self-tuning)'	▨	2	0	0.0	0

Columns Hidden 10

Total Rows : 22

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

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 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.

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.

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"
```

11
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

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 11-1 Live Thread Analysis (Cross Tier)

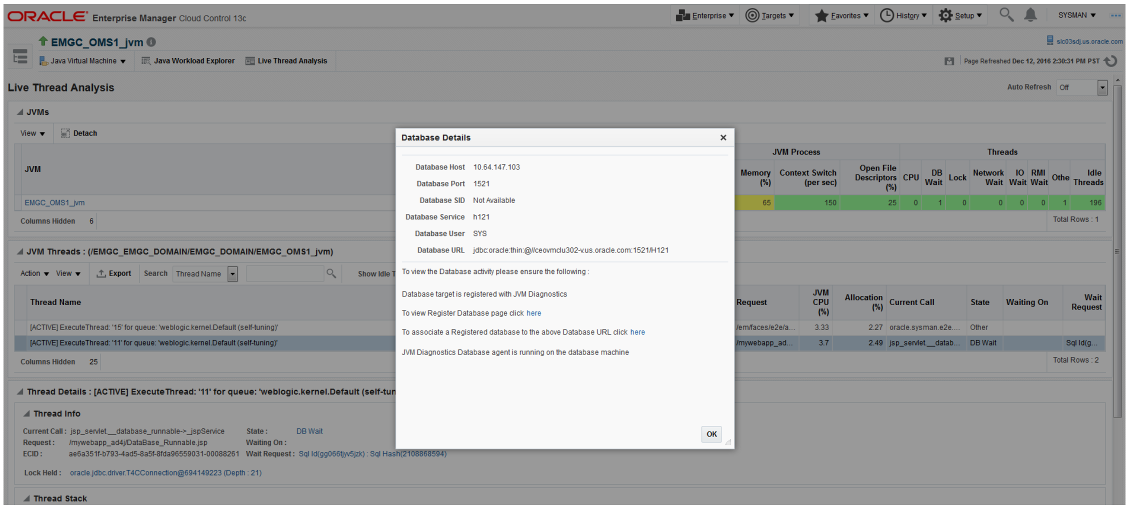

Database Details

Database Host 10.64.147.103

Database Port 1521

Database SID Not Available

Database Service h121

Database User SYS

Database URL jdbc:oracle:thin:@//ceovmclu302-v.us.oracle.com:1521/H121

To view the Database activity please ensure the following :

Database target is 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

OK

> **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 11-2   Live Thread Analysis: Associate / Disassociate a Registered Database

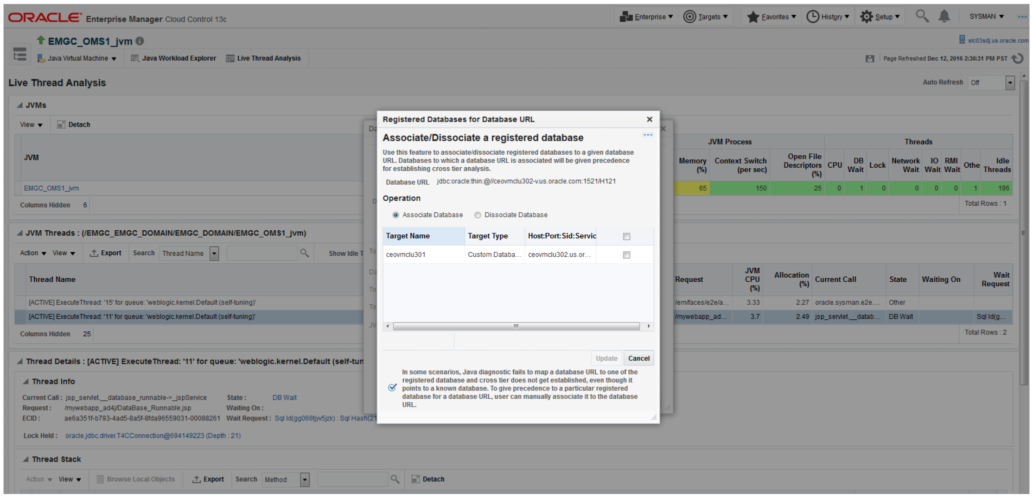

Registered Databases for Database URL

Associate/Dissociate a registered database

Use this feature to associate/dissociate registered databases to a given database URL. Databases to which a database URL is associated will be given precedence for establishing cross tier analysis.

Database URL   jdbc:oracle:thin:@//ceovmclu302-v.us.oracle.com:1521/H121

Operation

● Associate Database   ○ Dissociate Database

Target Name	Target Type	Host:Port:Sid:Servic	□
ceovmclu301	Custom Databa...	ceovmclu302.us.or...	□

Update   Cancel

In some scenarios, Java diagnostic fails to map a database URL to one of the registered database and cross tier does not get established, even though it points to a known database. To give precedence to a particular registered database for a database URL, user can manually associate it to the database URL.

- **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 11-3 Setup: Associate / Disassociate a Registered Database

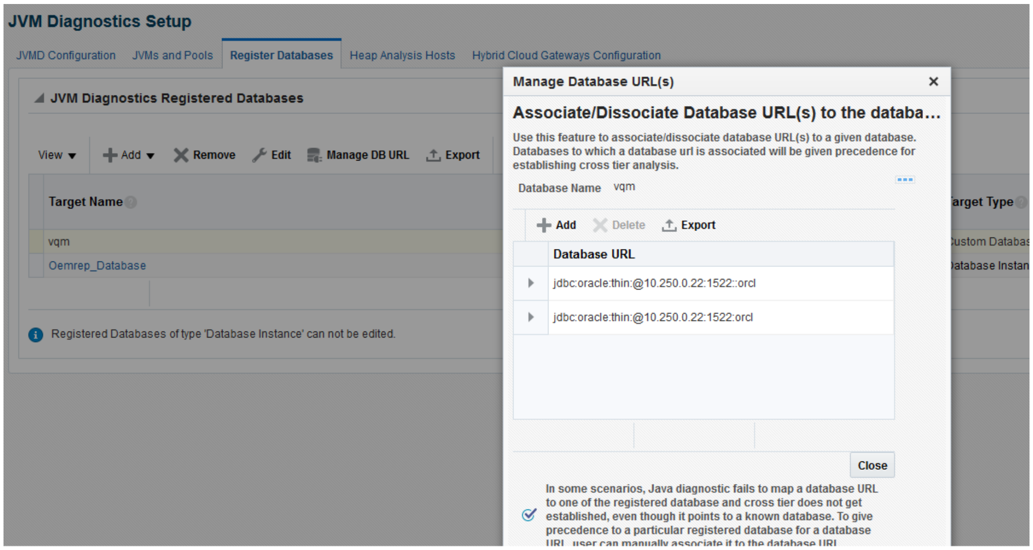

Target Name	arget Type
vqm	ustom Databas
Oemrep_Database	atabase Instan

Registered Databases of type 'Database Instance' can not be edited.

Associate/Dissociate Database URL(s) to the databa...

Use this feature to associate/dissociate database URL(s) to a given database. Databases to which a database url is associated will be given precedence for establishing cross tier analysis.

Database Name vqm

Database URL
jdbc:oracle:thin:@10.250.0.22:1522::orcl
jdbc:oracle:thin:@10.250.0.22:1522:orcl

In some scenarios, Java diagnostic fails to map a database URL to one of the registered database and cross tier does not get established, even though it points to a known database. To give precedence to a particular registered database for a database URL, user can manually associate it to the database URL

- If cross tier correlation cannot be established due to mismatch of the JVM Diagnostics Agent host name with the machine name stored in `V$ESSION` table of the database (for instance, inconsistent logical naming of machine), do the following:
  - Update the `v$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.

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.

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>:<webogic 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(WLInitialContextFacto
  ryDelegate.java:788)
  at
  weblogic.jndi.WLInitialContextFactoryDelegate.pushSubject(WLInitialContextFact
  oryDelegate.java:682)
  atweblogic.jndi.WLInitialContextFactoryDelegate.newContext(WLInitialContextFactoryDel
  egate.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**: `JVMD 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/aim
e/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
```

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.

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.

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.

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

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`

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.

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`.

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.

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 11-4 JVMD SLB Configuration

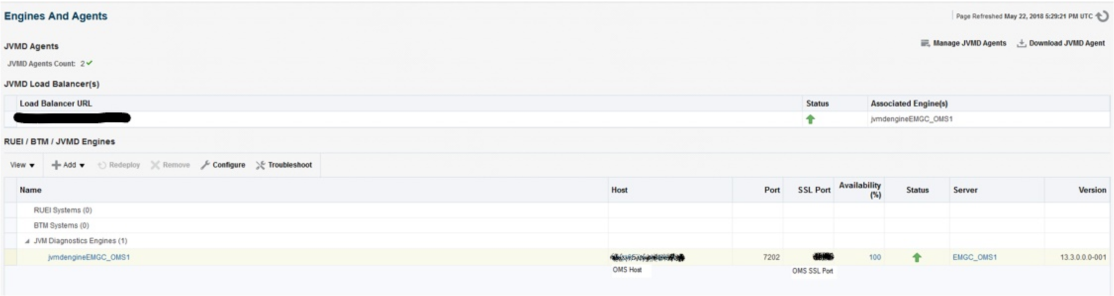

Engines And Agents

Page Refreshed May 22, 2018 5:29:21 PM UTC

JVMD Agents

Manage JVMD Agents   Download JVMD Agent

JVMD Agents Count: 2 ✓

JVMD Load Balancer(s)

Load Balancer URL	Status	Associated Engine(s)
[redacted]	↑	jvmdengineEMGC_OMS1

RUEI / BTM / JVMD Engines

View ▾   Add ▾   Redeploy   Remove   Configure   Troubleshoot

Name	Host	Port	SSL Port	Availability (%)	Status	Server	Version
RUEI Systems (0)							
BTM Systems (0)							
JVM Diagnostics Engines (1)							
jvmdengineEMGC_OMS1	[redacted] OMS Host	7202	[redacted] OMS SSL Port	100	↑	EMGC_OMS1	13.3.0.0.0-001

Virtual Server on SLB

SLB Virtual server port: `4901 aixcs2.us.example.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 11-5   Virtual Server on SLB

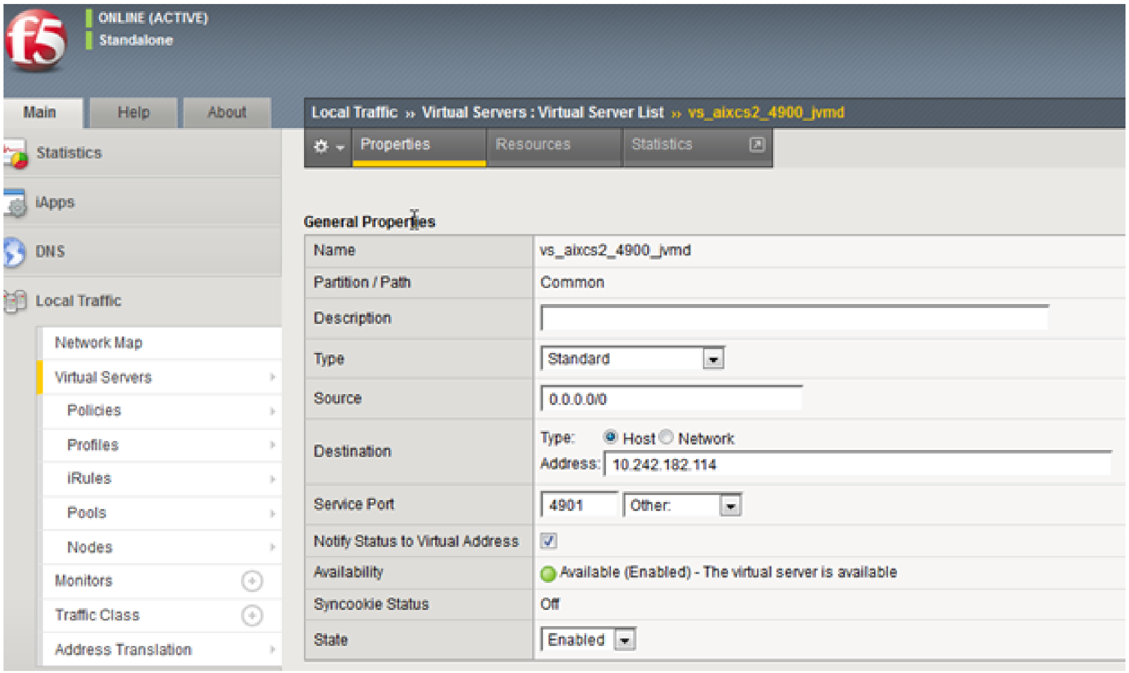

General Properties

Name	vs_aixcs2_4900_jvmd
Partition / Path	Common
Description	
Type	Standard
Source	0.0.0.0/0
Destination	Type: ● Host ○ Network Address: 10.242.182.114
Service Port	4901 Other:
Notify Status to Virtual Address	☑
Availability	Available (Enabled) - The virtual server is available
Syncookie Status	Off
State	Enabled

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.

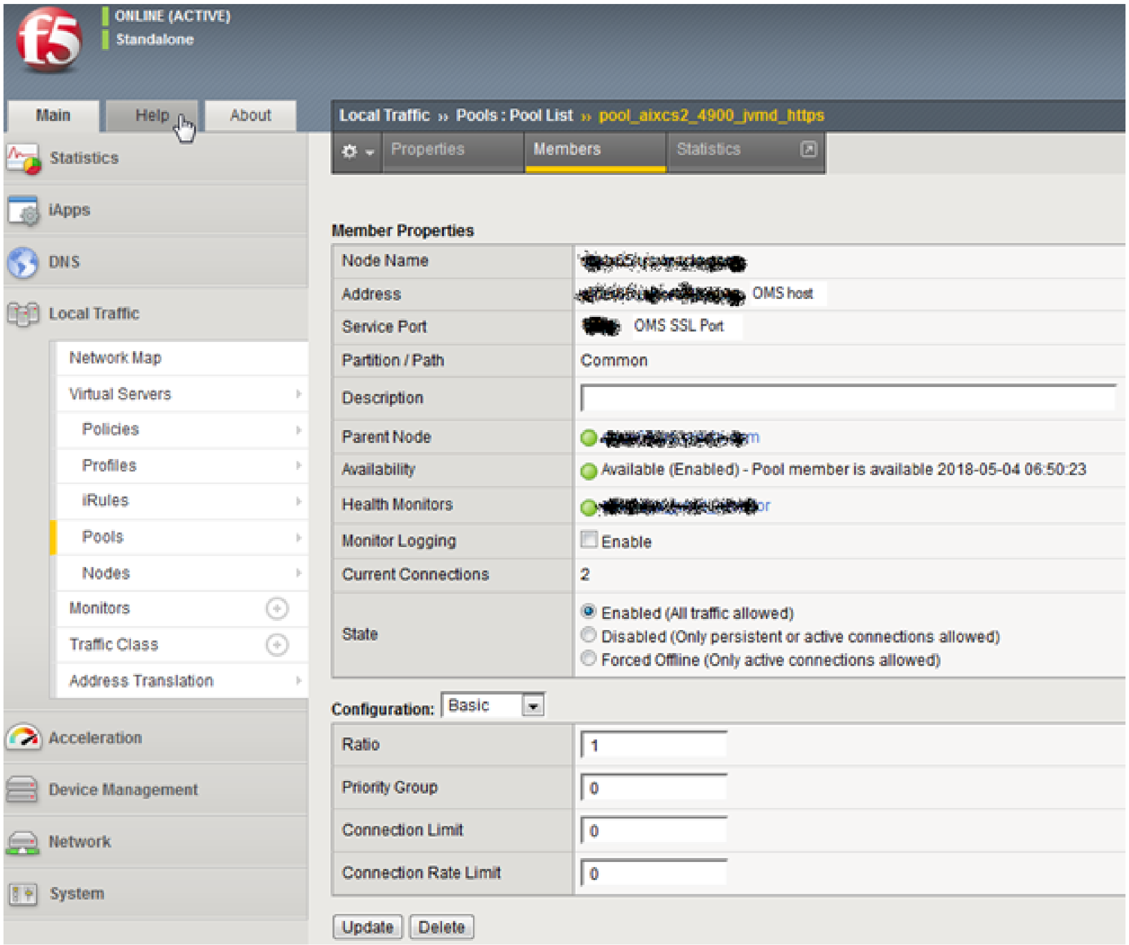

Member Properties

Node Name	[redacted]
Address	[redacted] OMS host
Service Port	[redacted] OMS SSL Port
Partition / Path	Common
Description	
Parent Node	● [redacted]
Availability	● Available (Enabled) - Pool member is available 2018-05-04 06:50:23
Health Monitors	● [redacted]
Monitor Logging	□ Enable
Current Connections	2
State	● Enabled (All traffic allowed) ○ Disabled (Only persistent or active connections allowed) ○ Forced Offline (Only active connections allowed)

Configuration: Basic

Ratio	1
Priority Group	0
Connection Limit	0
Connection Rate Limit	0

Update Delete

For more information about configuring the F5 SLB for EM and JVMD, see Configuring OMS High Availability with F5 BIG-IP Local Traffic Manager.

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.

Usage of the Try Changing Threads Parameter

This parameter should be used only when the JVM is highly active.

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.

Custom Provisioning Agent Deployment

You can customize the JVMD Agent deployment in the production environment by running custom provisioning scripts.

After the OMS has been installed, the `jvmd.zip` file can be found in the `plugins/oracle.sysman.emas.oms.plugin_12.1.0.0.0` directory in the Middleware installation directory. The zip file contains a set of scripts in the `customprov` directory. Details on using these scripts are described in the `README.TXT` present in the same directory. To use the custom provisioning scripts, follow these steps:

1. From the **Setup** menu, select **Middleware 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 JVMD Agent to all the specified servers.

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

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.

Part VII
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
- Troubleshooting and Best Practices
- Coherence Integration with JVM Diagnostics

12

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 13*c*. The following sections are covered in this chapter:

- About Coherence Management
- Configuring a Coherence Cluster
- Discovering Coherence Targets
- Enabling the Management Pack

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.

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 12-1 shows the configuration for monitoring standalone Coherence clusters using Enterprise Manager.

Figure 12-1 Coherence Cluster Configuration (Standalone Coherence Cluster)

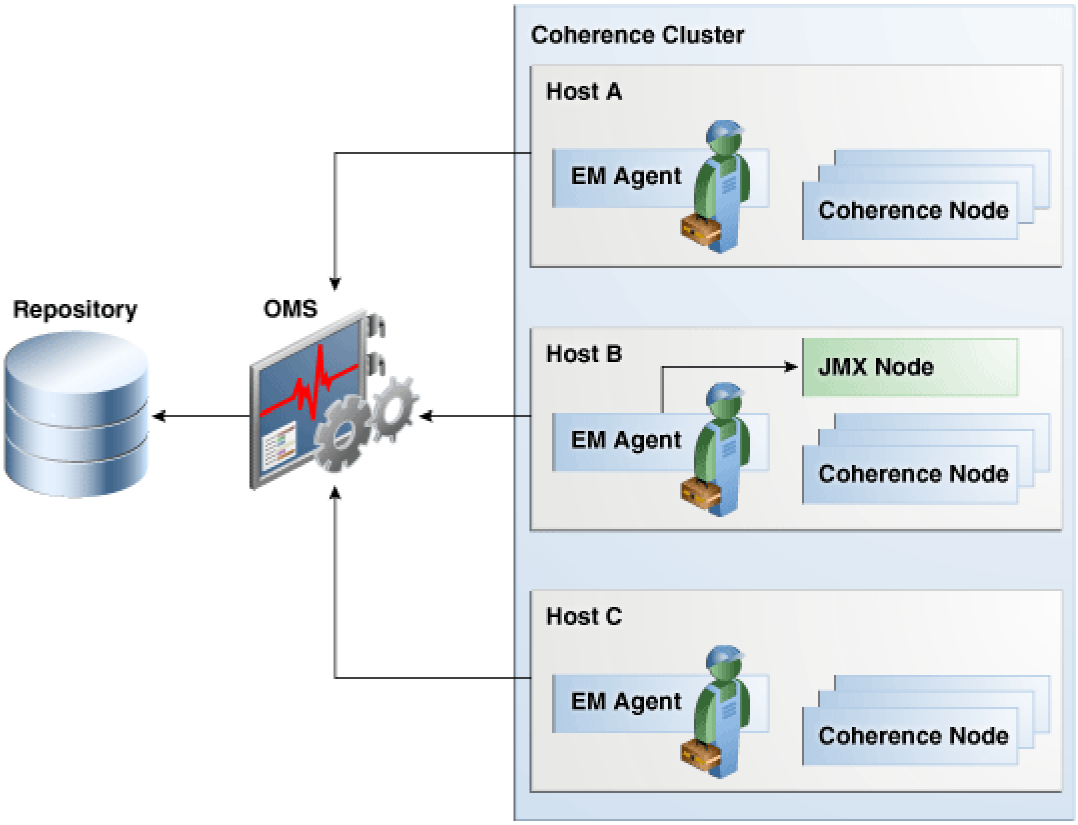

As shown in the figure, Coherence Management (JMX) node's MBean server will expose MBeans for entire Coherence cluster. Enterprise Manager will connect to this management node to discover and monitor Coherence cluster.

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

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**.

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.

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.

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
```

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.

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**.

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
```

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.

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 12-2 JConsole

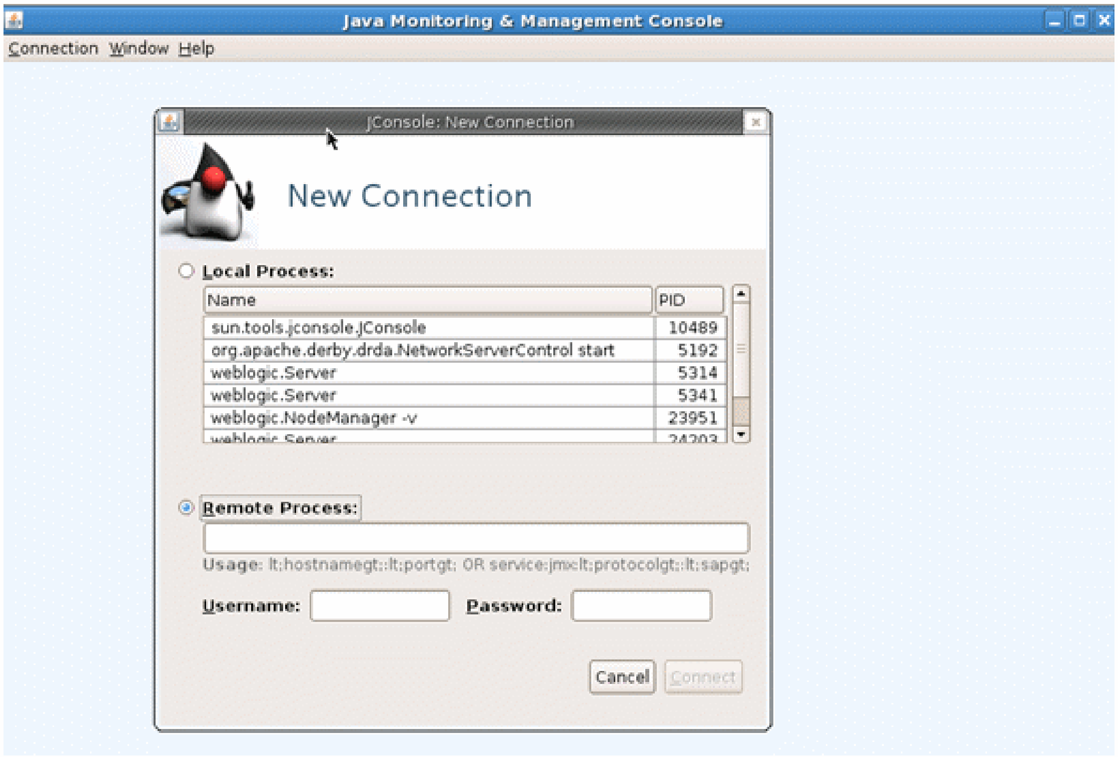

○ Local Process:

Name	PID
sun.tools.jconsole.JConsole	10489
org.apache.derby.drda.NetworkServerControl start	5192
weblogic.Server	5314
weblogic.Server	5341
weblogic.NodeManager -v	23951
weblogic.Server	24203

● Remote Process:

Usage: lt;hostnamegt::lt;portgt: OR service:jmx:lt;protocolgt::lt;sapgt:

Username: Password:

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 12-3 MBean Object Tree

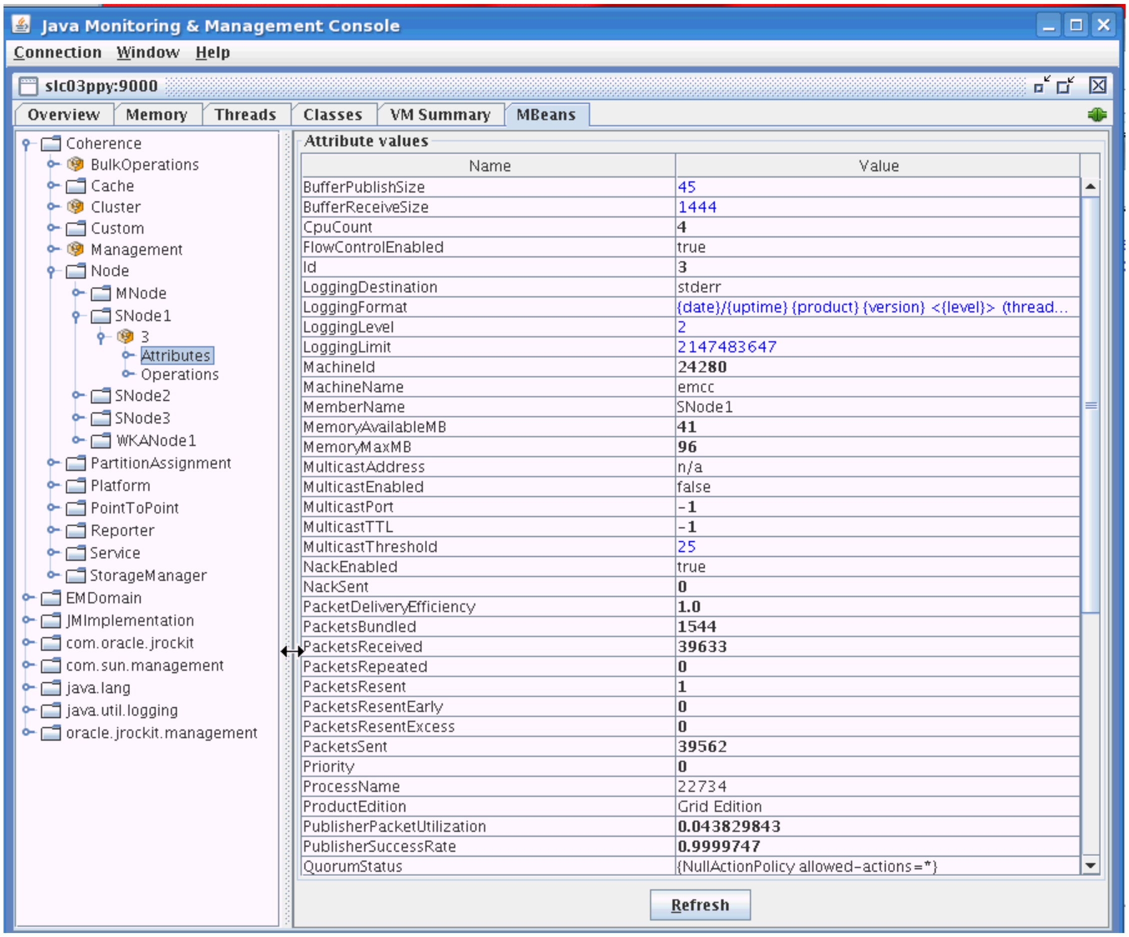

Name	Value
BufferPublishSize	45
BufferReceiveSize	1444
CpuCount	4
FlowControlEnabled	true
Id	3
LoggingDestination	stderr
LoggingFormat	{date}/{uptime} {product} {version} <{level}> (thread...
LoggingLevel	2
LoggingLimit	2147483647
MachineId	24280
MachineName	emcc
MemberName	SNode1
MemoryAvailableMB	41
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.

Discovering Coherence Targets

This section covers the following:

- Discovering a Standalone Coherence Cluster
- Discovering a Managed Coherence Cluster

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 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 12-4 Add Coherence Target

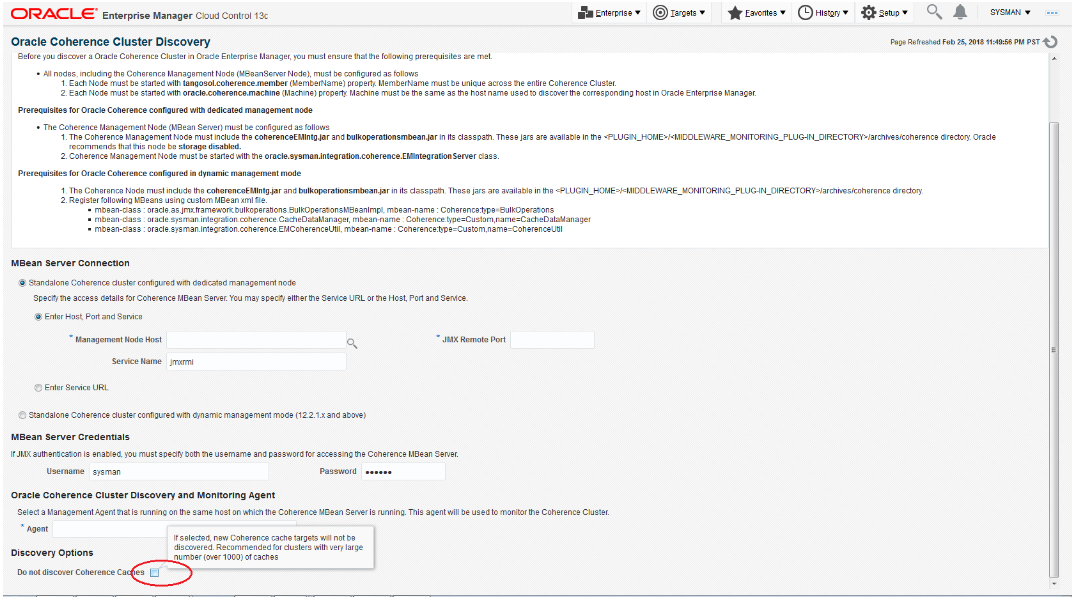

Oracle Coherence Cluster Discovery

Page Refreshed Feb 25, 2018 11:49:56 PM PST

Before you discover a Oracle 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
  1. Each Node must be started with **tangosol.coherence.member** (MemberName) property. MemberName must be unique across the entire Coherence Cluster.
  2. 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
  1. 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_PLUG-IN_DIRECTORY>/archives/coherence directory. Oracle recommends that this node be **storage disabled.**
  2. Coherence Management Node must be started with the **oracle.sysman.integration.coherence.EMIntegrationServer** class.

Prerequisites for Oracle Coherence configured in dynamic management mode

1. The Coherence Node must include the **coherenceEMIntg.jar** and **bulkoperationsmbean.jar** in its classpath. These jars are available in the <PLUGIN_HOME>/<MIDDLEWARE_MONITORING_PLUG-IN_DIRECTORY>/archives/coherence directory.
2. 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.EMCoherenceUtil, mbean-name : Coherence:type=Custom,name=CoherenceUtil

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 jmxrmi

○ 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 sysman

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 12-5   Add Coherence Target

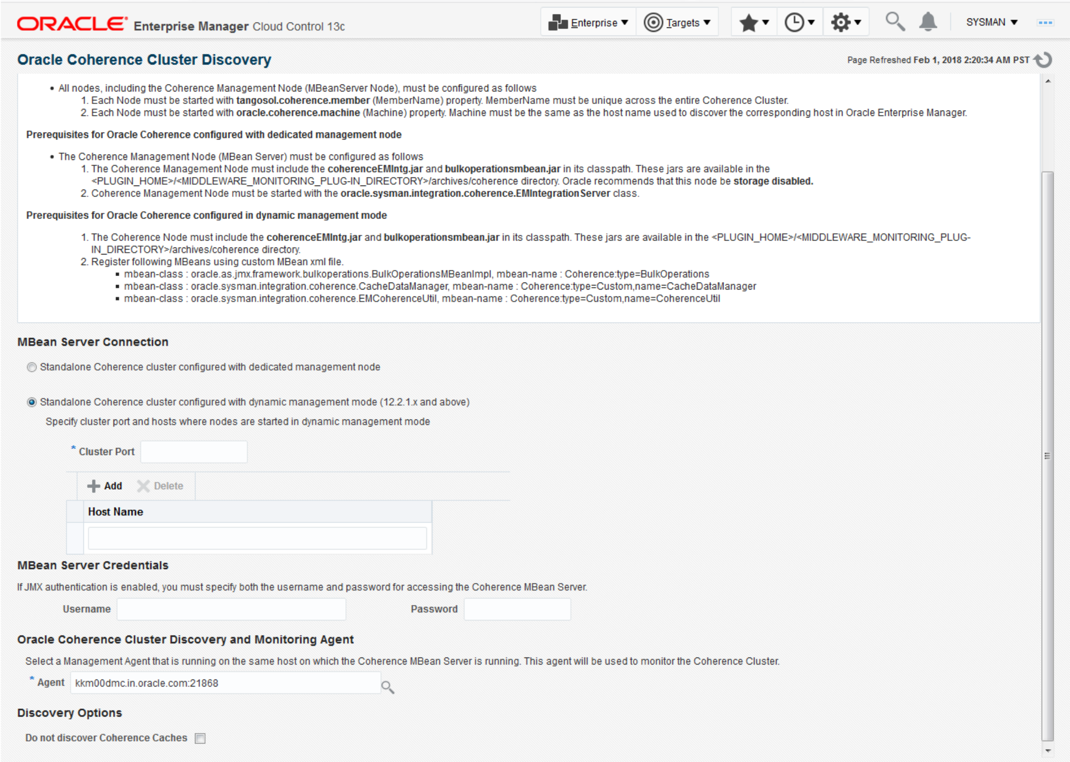

Oracle Coherence Cluster Discovery

Page Refreshed Feb 1, 2018 2:20:34 AM PST

- All nodes, including the Coherence Management Node (MBeanServer Node), must be configured as follows
  1. Each Node must be started with **tangosol.coherence.member** (MemberName) property. MemberName must be unique across the entire Coherence Cluster.
  2. 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
  1. 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_PLUG-IN_DIRECTORY>/archives/coherence directory. Oracle recommends that this node be **storage disabled.**
  2. Coherence Management Node must be started with the **oracle.sysman.integration.coherence.EMIntegrationServer** class.

Prerequisites for Oracle Coherence configured in dynamic management mode

1. The Coherence Node must include the **coherenceEMIntg.jar** and **bulkoperationsmbean.jar** in its classpath. These jars are available in the <PLUGIN_HOME>/<MIDDLEWARE_MONITORING_PLUG-IN_DIRECTORY>/archives/coherence directory.
2. 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.EMCoherenceUtil, mbean-name : Coherence:type=Custom,name=CoherenceUtil

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

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 kkm00dmc.in.oracle.com:21868

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*.

5. **MBean Server Credentials**: If JMX authentication is used, specify the username and password required to access the MBean Server.
6. Select the Management Agent that will be used to monitor the Coherence target.
7. 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.
8. 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.

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 12-6 Refresh Cluster

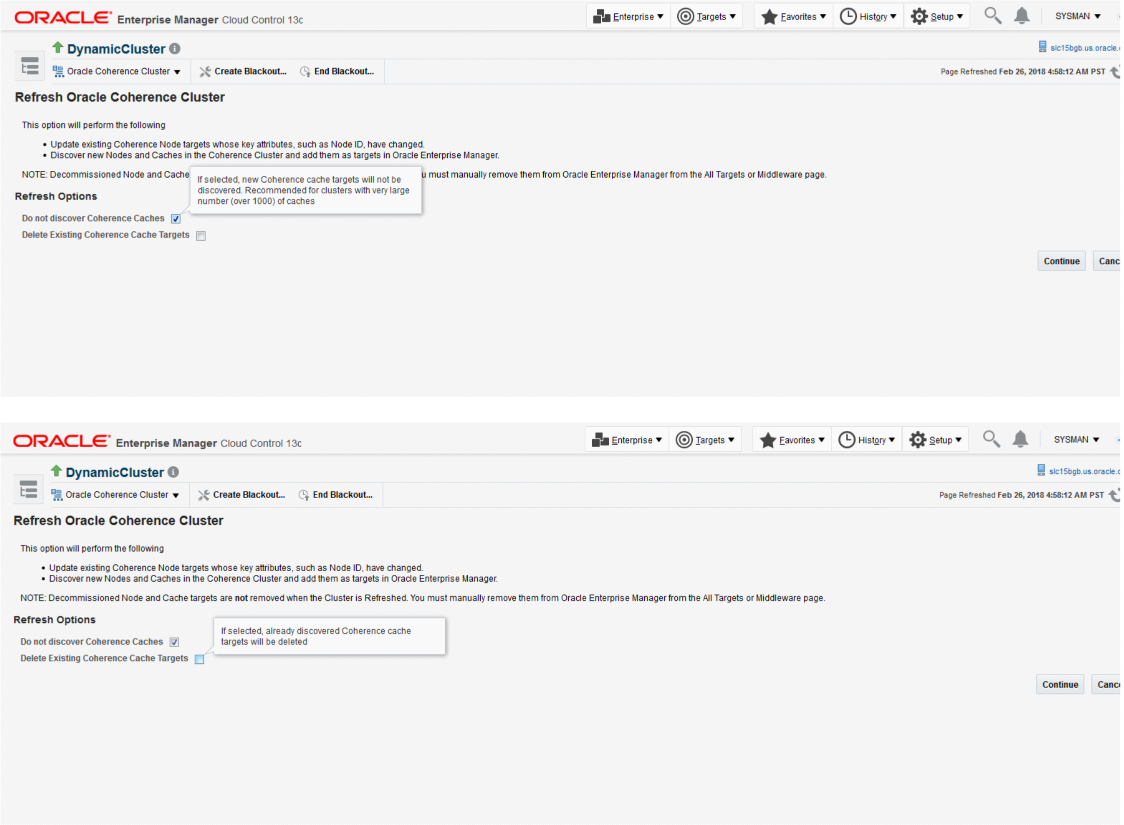

Refresh Oracle Coherence Cluster

This option will perform the following

- Update existing Coherence Node targets whose key attributes, such as Node ID, have changed.
- Discover new Nodes and Caches in the Coherence Cluster and add them as targets in Oracle Enterprise Manager.

NOTE: Decommissioned Node and Cache u must manually remove them from Oracle Enterprise Manager from the All Targets or Middleware page.

If selected, new Coherence cache targets will not be discovered. Recommended for clusters with very large number (over 1000) of caches

Refresh Options

Do not discover Coherence Caches ☑

Delete Existing Coherence Cache Targets □

Continue Canc

Refresh Oracle Coherence Cluster

This option will perform the following

- Update existing Coherence Node targets whose key attributes, such as Node ID, have changed.
- Discover new Nodes and Caches in the Coherence Cluster and add them as targets in Oracle Enterprise Manager.

NOTE: Decommissioned Node and Cache targets are **not** removed when the Cluster is Refreshed. You must manually remove them from Oracle Enterprise Manager from the All Targets or Middleware page.

Refresh Options

Do not discover Coherence Caches ☑

Delete Existing Coherence Cache Targets □

If selected, already discovered Coherence cache targets will be deleted

Continue Canc

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.

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 12-7 Mis-Configured Nodes

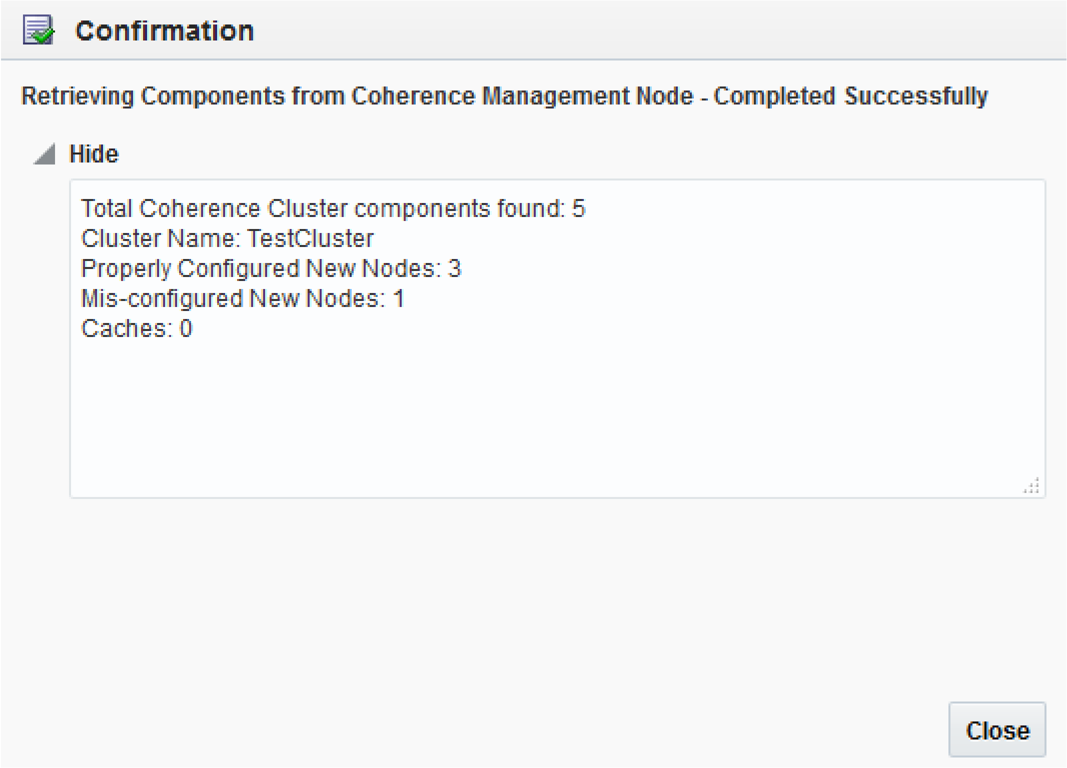

Retrieving Components from Coherence Management Node - Completed Successfully

Total Coherence Cluster components found: 5
Cluster Name: TestCluster
Properly Configured New Nodes: 3
Mis-configured New Nodes: 1
Caches: 0

This indicates that there are some improperly configured nodes in the cluster. Click **Close**. The following page is displayed.

Figure 12-8   Mis-configured Nodes II

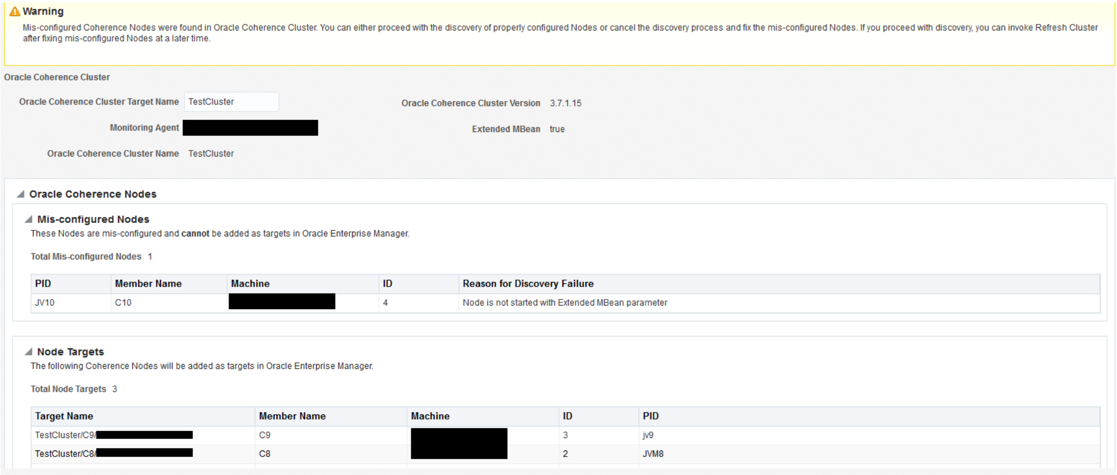

⚠ **Warning**

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
TestCluster/C9/[redacted]	C9	[redacted]	3	jv9
TestCluster/C8/[redacted]	C8		2	JVM8

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.

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 12-9   Add Oracle Fusion Middleware / WebLogic Domain: Find Targets

   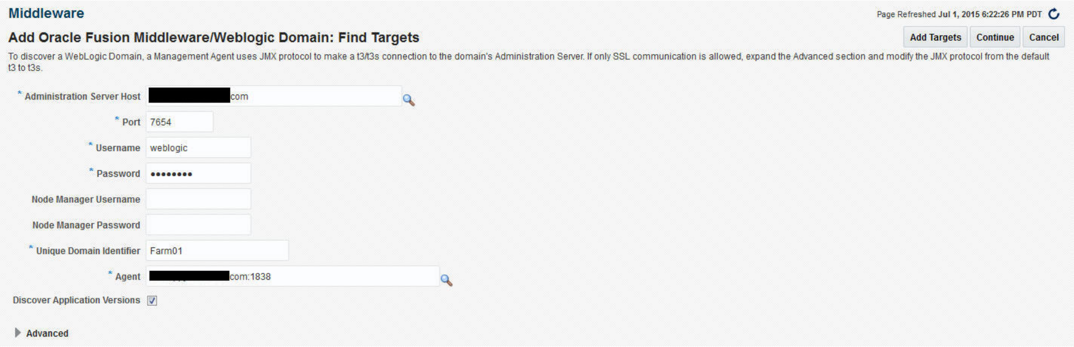

   Add Oracle Fusion Middleware/Weblogic Domain: Find Targets

   To discover a WebLogic Domain, a Management Agent uses JMX protocol to make a t3/t3s connection to the domain's Administration Server. If only SSL communication is allowed, expand the Advanced section and modify the JMX protocol from the default t3 to t3s.

* Administration Server Host	[redacted]com
* Port	7654
* Username	weblogic
* Password	••••••••
Node Manager Username	
Node Manager Password	
* Unique Domain Identifier	Farm01
* Agent	[redacted]com:1838
Discover Application Versions	☑

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 12-10   Targets and Agents Assignments

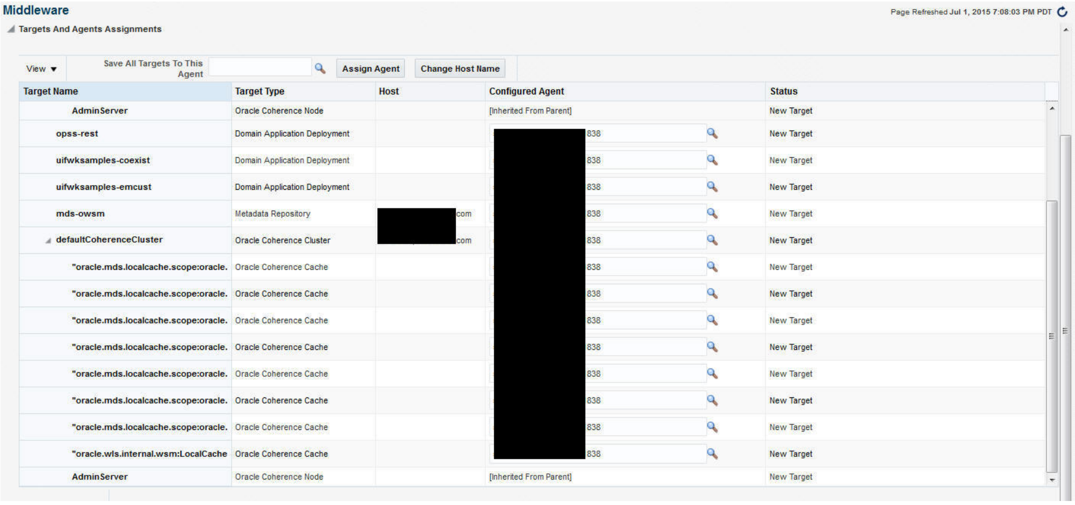
Middleware
Page Refreshed Jul 1, 2015 7:08:03 PM PDT
Targets And Agents Assignments
View ▾
Save All Targets To This Agent
Assign Agent
Change Host Name

Target Name	Target Type	Host	Configured Agent	Status
AdminServer	Oracle Coherence Node		[Inherited From Parent]	New Target
opss-rest	Domain Application Deployment		[redacted]838	New Target
uifwksamples-coexist	Domain Application Deployment		[redacted]838	New Target
uifwksamples-emcust	Domain Application Deployment		[redacted]838	New Target
mds-owsm	Metadata Repository	[redacted]com	[redacted]838	New Target
defaultCoherenceCluster	Oracle Coherence Cluster	[redacted]com	[redacted]838	New Target
"oracle.mds.localcache.scope:oracle.	Oracle Coherence Cache		[redacted]838	New Target
"oracle.mds.localcache.scope:oracle.	Oracle Coherence Cache		[redacted]838	New Target
"oracle.mds.localcache.scope:oracle.	Oracle Coherence Cache		[redacted]838	New Target
"oracle.mds.localcache.scope:oracle.	Oracle Coherence Cache		[redacted]838	New Target
"oracle.mds.localcache.scope:oracle.	Oracle Coherence Cache		[redacted]838	New Target
"oracle.mds.localcache.scope:oracle.	Oracle Coherence Cache		[redacted]838	New Target
"oracle.mds.localcache.scope:oracle.	Oracle Coherence Cache		[redacted]838	New Target
"oracle.wls.internal.wsm:LocalCache	Oracle Coherence Cache		[redacted]838	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.

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.

13 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 13.2.0.0.0 Management Agent on all hosts where Coherence nodes are running.
- Deploy the 13.2.0.0.0 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 13.2.0.0.0, you must ensure that the Fusion Middleware Plug-in is also upgraded to 13.2.0.0.0.

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

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 13-1 Navigation Tree

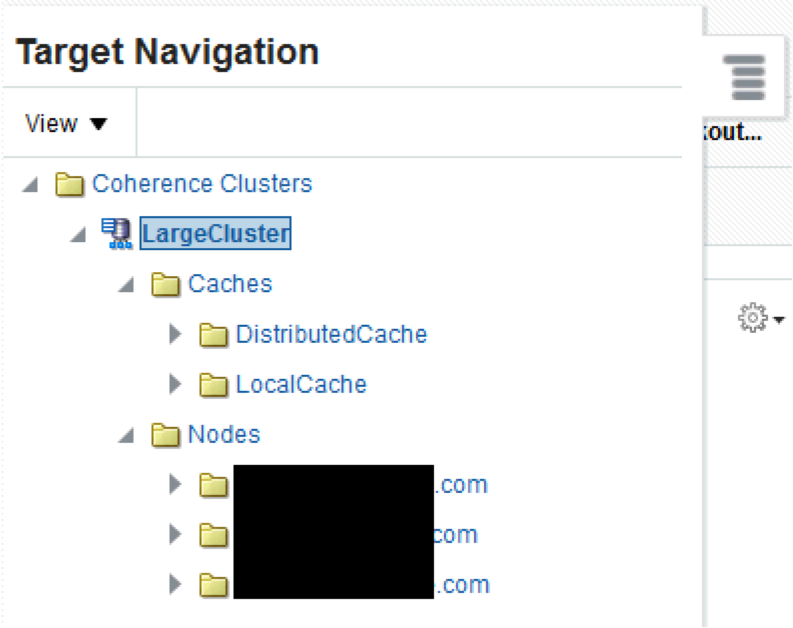

For a managed Coherence cluster, all the Coherence clusters in a domain are included in the Coherence Clusters folder. This folder appears at the same level as the WebLogic Domain folder. If multi-tenancy is supported for a target, you will also see the Domain Partition Coherence Caches folder.

Figure 13-2   Navigation Tree (Managed Clusters)

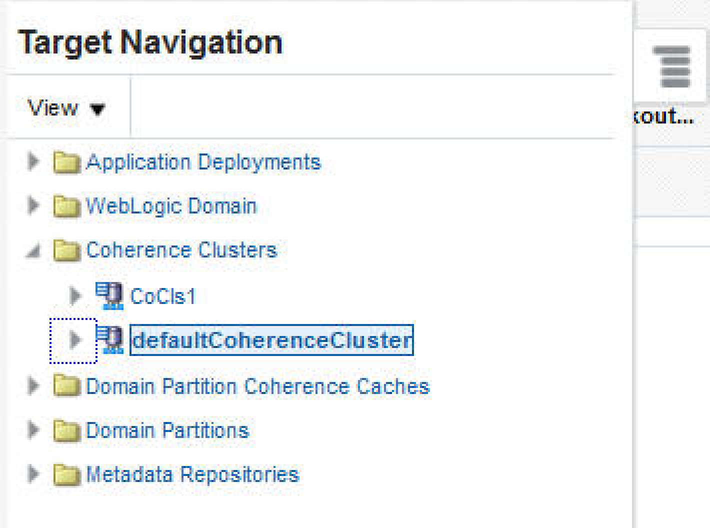

Target Navigation

View ▾

- Application Deployments
- WebLogic Domain
- Coherence Clusters
  - CoCls1
  - defaultCoherenceCluster
- Domain Partition Coherence Caches
- Domain Partitions
- Metadata Repositories

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 13-3   Cluster Home Page (Personalization Icon)

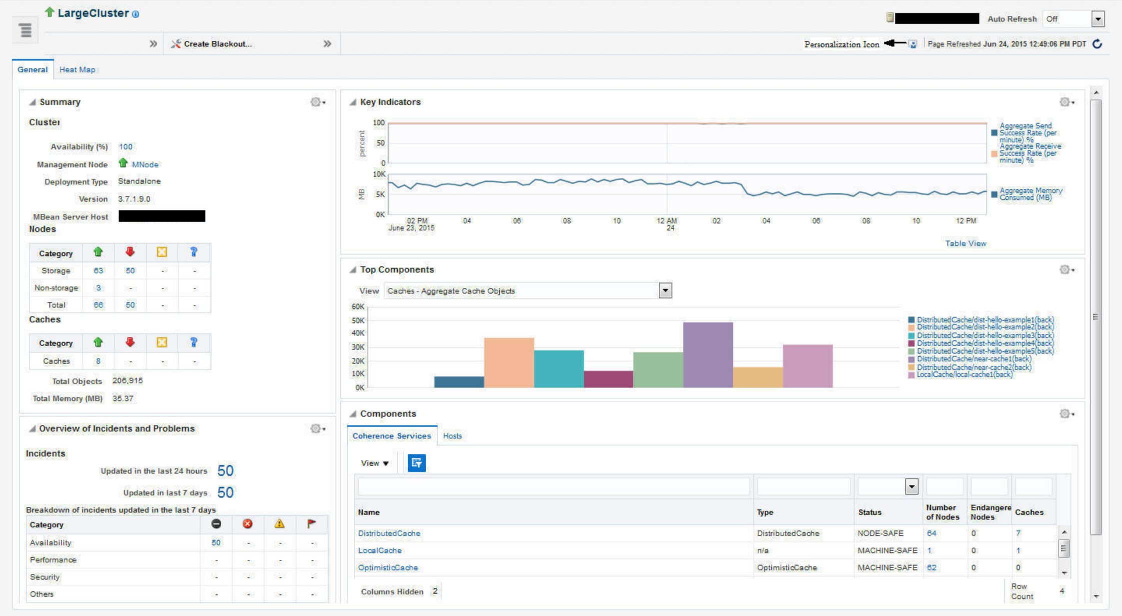

LargeCluster

[redacted] Auto Refresh Off

Create Blackout...

Personalization Icon ← | Page Refreshed Jun 24, 2015 12:49:06 PM PDT

General | Heat Map

Summary

Cluster

Availability (%) 100
Management Node ⬆ MNode
Deployment Type Standalone
Version 3.7.1.9.0
MBean Server Host [redacted]

Nodes

Category	⬆	⬇	☒	?
Storage	63	50	-	-
Non-storage	3	-	-	-
Total	66	50	-	-

Caches

Category	⬆	⬇	☒	?
Caches	8	-	-	-

Total Objects 206,915
Total Memory (MB) 35.37

Overview of Incidents and Problems

Incidents

Updated in the last 24 hours 50
Updated in last 7 days 50

Breakdown of incidents updated in the last 7 days

Category	⛔	⊗	⚠	⚑
Availability	50	-	-	-
Performance	-	-	-	-
Security	-	-	-	-
Others	-	-	-	-

Key Indicators

percent
100
50
0
MB
10K
5K
0K
02 PM
June 23, 2015
04
06
08
10
12 AM
24
02
04
06
08
10
12 PM
Aggregate Send Success Rate (per minute) %
Aggregate Receive Success Rate (per minute) %
Aggregate Memory Consumed (MB)
Table View

Top Components

View Caches - Aggregate Cache Objects

60K
50K
40K
30K
20K
10K
0K
DistributedCache/dist-hello-example1(back)
DistributedCache/dist-hello-example2(back)
DistributedCache/dist-hello-example3(back)
DistributedCache/dist-hello-example4(back)
DistributedCache/dist-hello-example5(back)
DistributedCache/near-cache1(back)
DistributedCache/near-cache2(back)
LocalCache/local-cache1(back)

Components

Coherence Services | Hosts

View ▾

Name	Type	Status	Number of Nodes	Endangere Nodes	Caches
DistributedCache	DistributedCache	NODE-SAFE	64	0	7
LocalCache	n/a	MACHINE-SAFE	1	0	1
OptimisticCache	OptimisticCache	MACHINE-SAFE	62	0	0

Columns Hidden 2   Row Count 4

You will see the page in Edit mode as shown below.

Figure 13-4   Cluster Home Page (Edit Mode)

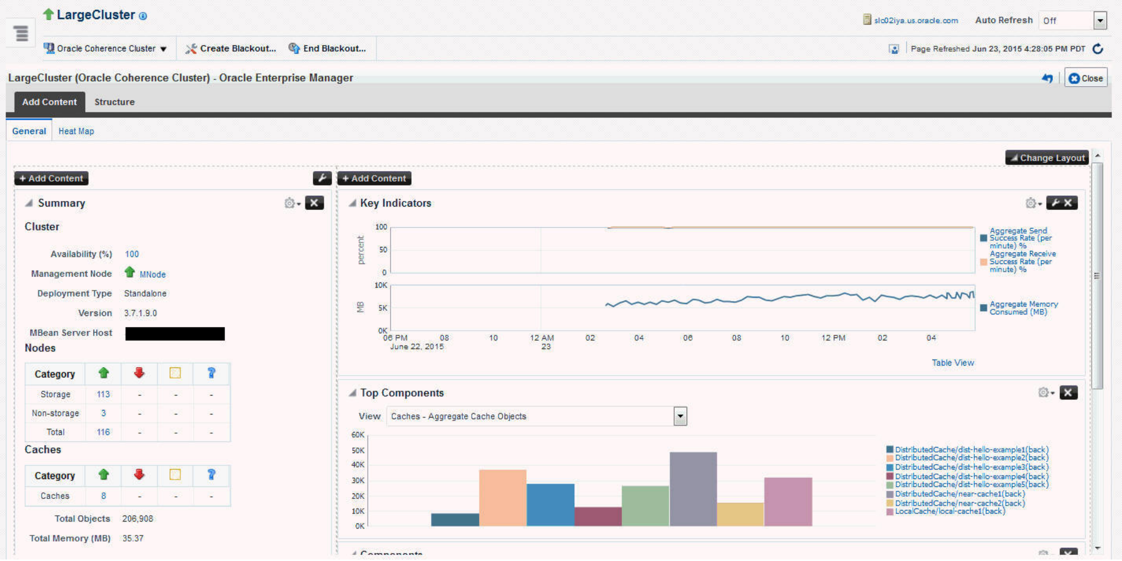
LargeCluster
Oracle Coherence Cluster ▾
Create Blackout...
End Blackout...
slc02iya.us.oracle.com
Auto Refresh
Off
Page Refreshed Jun 23, 2015 4:28:05 PM PDT
LargeCluster (Oracle Coherence Cluster) - Oracle Enterprise Manager
Close
Add Content
Structure
General
Heat Map
Change Layout
+ Add Content
+ Add Content
Summary
Key Indicators
100
50
0
percent
10K
5K
0K
MB
06 PM
June 22, 2015
08
10
12 AM
23
02
04
06
08
10
12 PM
02
04
Aggregate Send Success Rate (per minute) %
Aggregate Receive Success Rate (per minute) %
Aggregate Memory Consumed (MB)
Table View
Top Components
View
Caches - Aggregate Cache Objects
60K
50K
40K
30K
20K
10K
0K
DistributedCache/dist-hello-example1(back)
DistributedCache/dist-hello-example2(back)
DistributedCache/dist-hello-example3(back)
DistributedCache/dist-hello-example4(back)
DistributedCache/dist-hello-example5(back)
DistributedCache/near-cache1(back)
DistributedCache/near-cache2(back)
LocalCache/local-cache1(back)
Components

Cluster

Availability (%) 100
Management Node MNode
Deployment Type Standalone
Version 3.7.1.9.0
MBean Server Host [redacted]

Nodes

Category	⬆	⬇	□	?
Storage	113	-	-	-
Non-storage	3	-	-	-
Total	116	-	-	-

Caches

Category	⬆	⬇	□	?
Caches	8	-	-	-

Total Objects 206,908
Total Memory (MB) 35.37

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.

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.

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 13-5   Coherence Cluster Home Page

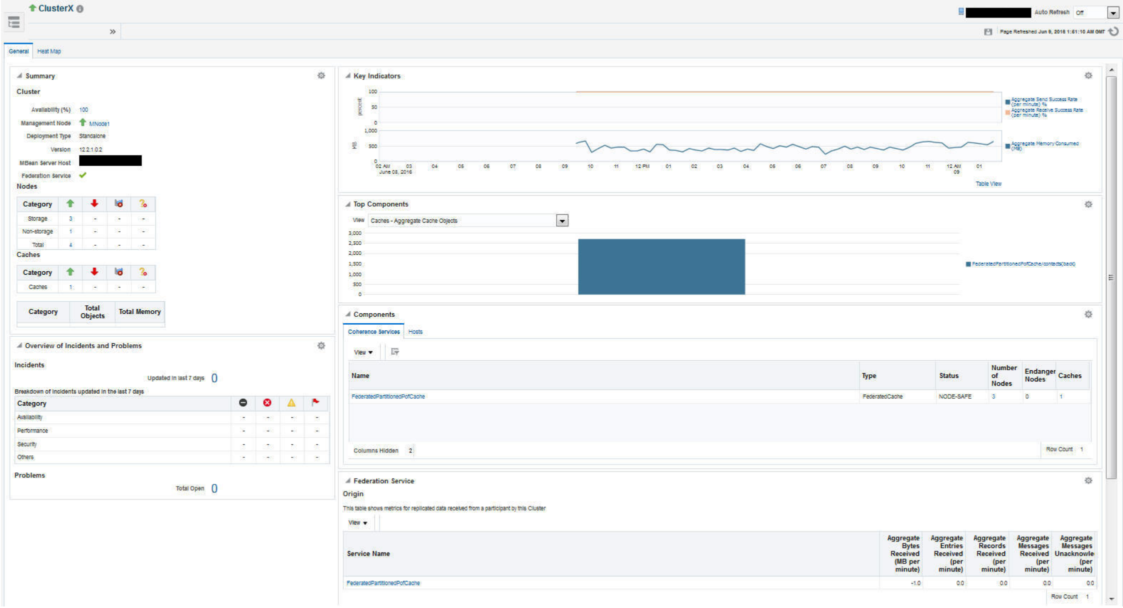

Summary

Cluster

Availability (%) 100
Management Node MNode1
Deployment Type Standalone
Version 12.2.1.0.2
MBean Server Host [redacted]
Federation Service ✓

Nodes

Category	↑	↓	◐	?
Storage	3	-	-	-
Non-storage	1	-	-	-
Total	4	-	-	-

Caches

Category	↑	↓	◐	?
Caches	1	-	-	-

Category	Total Objects	Total Memory

Incidents

Updated in last 7 days 0

Breakdown of incidents updated in the last 7 days

Category	⊖	⊗	⚠	⚑
Availability	-	-	-	-
Performance	-	-	-	-
Security	-	-	-	-
Others	-	-	-	-

Problems

Total Open 0

View Caches - Aggregate Cache Objects

Name	Type	Status	Number of Nodes	Endanger Nodes	Caches
FederatedPartitionedPofCache	FederatedCache	NODE-SAFE	3	0	1

Columns Hidden 2   Row Count 1

Origin

This table shows metrics for replicated data received from a participant by this Cluster

Service Name	Aggregate Bytes Received (MB per minute)	Aggregate Entries Received (per minute)	Aggregate Records Received (per minute)	Aggregate Messages Received (per minute)	Aggregate Messages Unacknowle (per minute)
FederatedPartitionedPofCache	-1.0	0.0	0.0	0.0	0.0

Row Count 1

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.

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.

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.

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.

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.
- **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.

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 13-6 Coherence Node Home Page

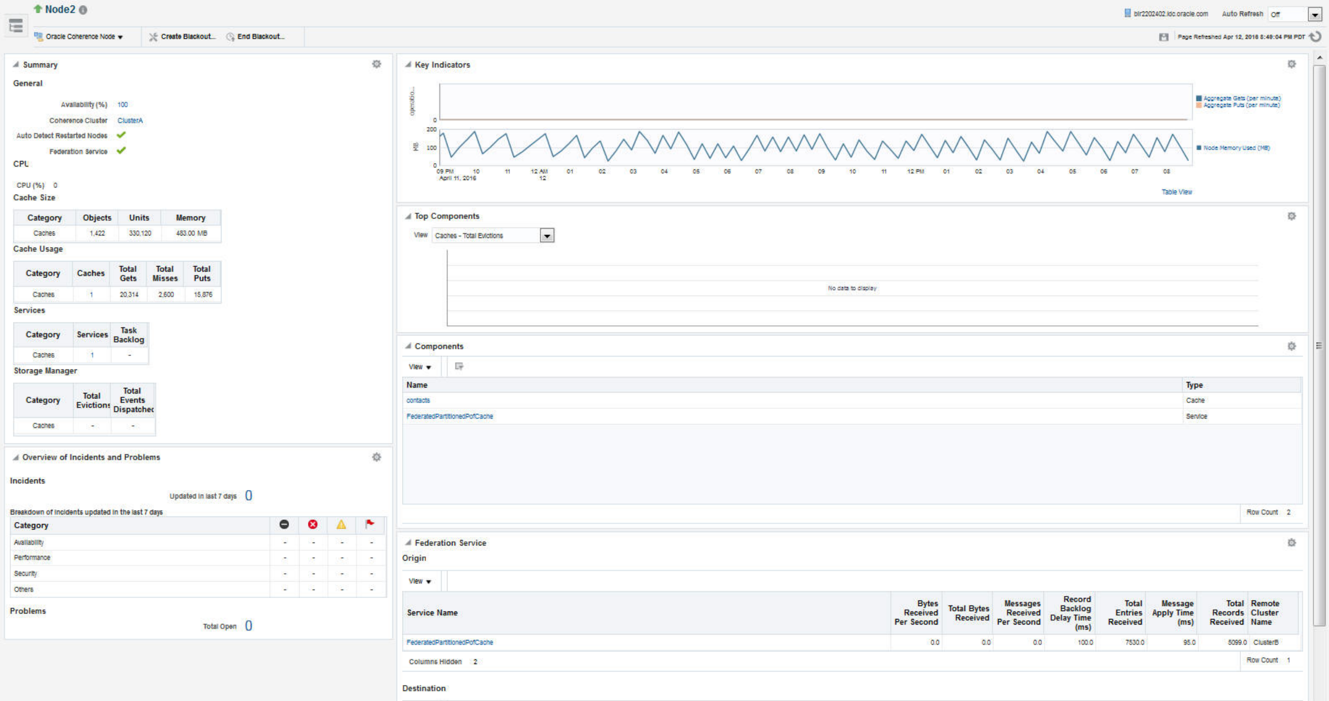

Summary

General

Availability (%) 100
Coherence Cluster ClusterA
Auto Detect Restarted Nodes ✓
Federation Service ✓

CPU

CPU (%) 0

Cache Size

Category	Objects	Units	Memory
Caches	1,422	330,120	483.00 MB

Cache Usage

Category	Caches	Total Gets	Total Misses	Total Puts
Caches	1	20,314	2,600	15,876

Services

Category	Services	Task Backlog
Caches	1	-

Storage Manager

Category	Total Evictions	Total Events Dispatched
Caches	-	-

Overview of Incidents and Problems

Incidents

Updated in last 7 days 0

Breakdown of incidents updated in the last 7 days

Category	⊖	⊗	⚠	⚑
Availability	-	-	-	-
Performance	-	-	-	-
Security	-	-	-	-
Others	-	-	-	-

Problems

Total Open 0

Top Components

View Caches - Total Evictions

Components

View ▾

Name	Type
contacts	Cache
FederatedPartitionedPofCache	Service

Row Count 2

Federation Service

Origin

View ▾

Service Name	Bytes Received Per Second	Total Bytes Received	Messages Received Per Second	Record Backlog Delay Time (ms)	Total Entries Received	Message Apply Time (ms)	Total Records Received	Remote Cluster Name
FederatedPartitionedPofCache	0.0	0.0	0.0	100.0	7530.0	95.0	5099.0	ClusterB

Columns Hidden 2

Row Count 1

Destination

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.

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**.

  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 Diagnosticsfor details.

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 13-7   Cache Home Page

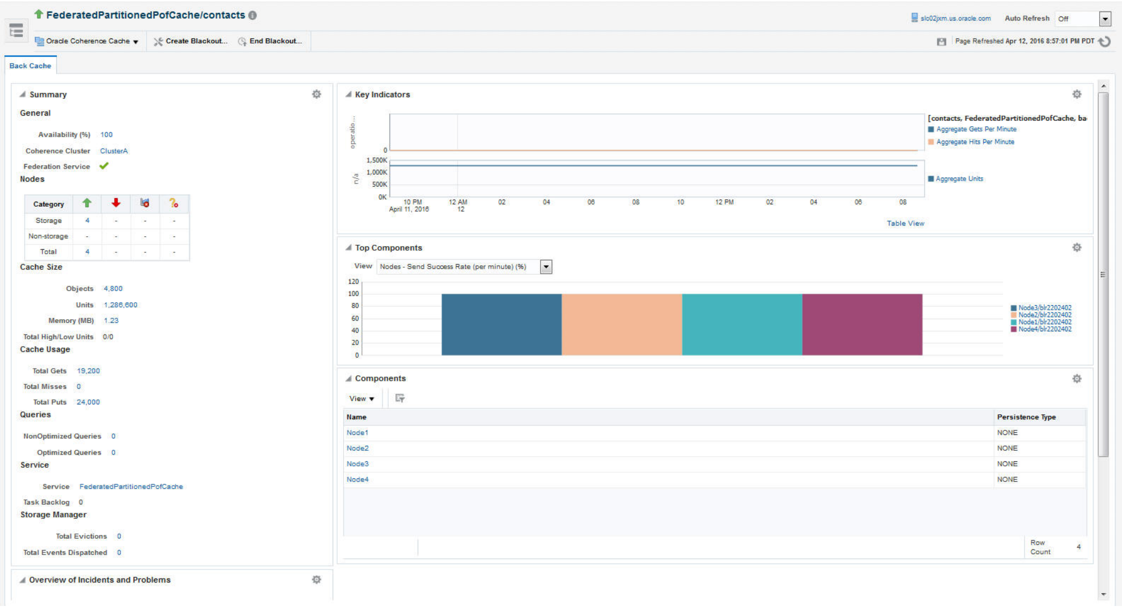

Summary

General

Availability (%) 100

Coherence Cluster ClusterA

Federation Service ✓

Nodes

Category	↑	↓		
Storage	4	-	-	-
Non-storage	-	-	-	-
Total	4	-	-	-

Cache Size

Objects 4,800

Units 1,286,600

Memory (MB) 1.23

Total High/Low Units 0/0

Cache Usage

Total Gets 19,200

Total Misses 0

Total Puts 24,000

Queries

NonOptimized Queries 0

Optimized Queries 0

Service

Service FederatedPartitionedPofCache

Task Backlog 0

Storage Manager

Total Evictions 0

Total Events Dispatched 0

View Nodes - Send Success Rate (per minute) (%)

Name	Persistence Type
Node1	NONE
Node2	NONE
Node3	NONE
Node4	NONE

Row Count 4

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.

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 13-8 Near Cache (Back Cache)

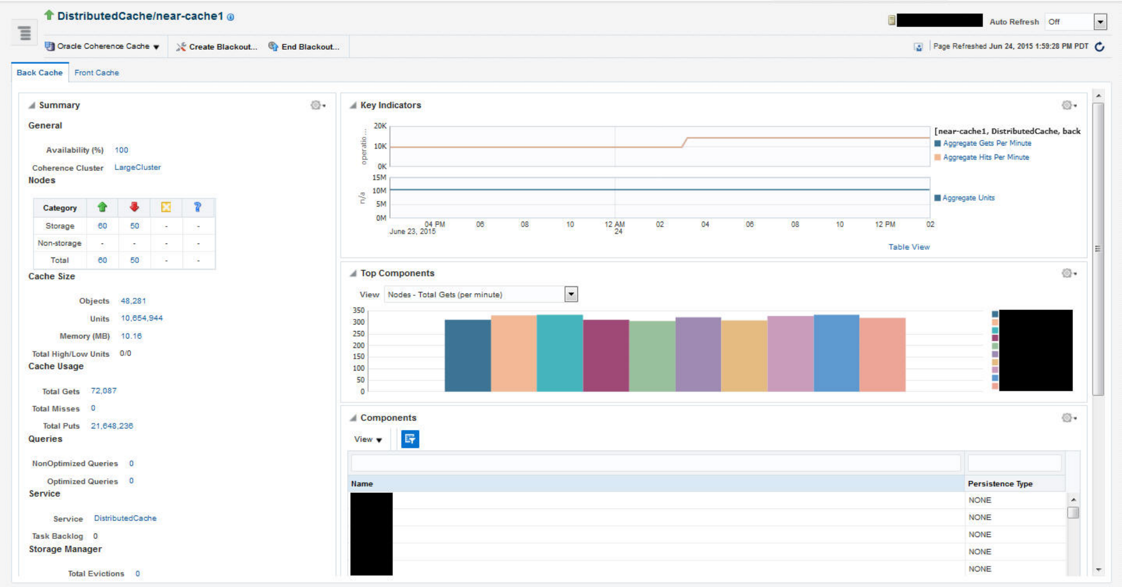

DistributedCache/near-cache1

[redacted] Auto Refresh Off

Oracle Coherence Cache ▾ | Create Blackout... | End Blackout...

Page Refreshed Jun 24, 2015 1:59:28 PM PDT

Back Cache | Front Cache

Summary

General

Availability (%) 100

Coherence Cluster LargeCluster

Nodes

Category	↑	↓	☒	?
Storage	60	50	-	-
Non-storage	-	-	-	-
Total	60	50	-	-

Cache Size

Objects 48,281

Units 10,654,944

Memory (MB) 10.16

Total High/Low Units 0/0

Cache Usage

Total Gets 72,087

Total Misses 0

Total Puts 21,648,236

Queries

NonOptimized Queries 0

Optimized Queries 0

Service

Service DistributedCache

Task Backlog 0

Storage Manager

Total Evictions 0

Key Indicators

Top Components

View Nodes - Total Gets (per minute)

[redacted]

Components

View ▾

Name	Persistence Type
[redacted]	NONE
[redacted]	NONE
[redacted]	NONE
[redacted]	NONE
[redacted]	NONE

Figure 13-9 Near Cache (Front Cache)

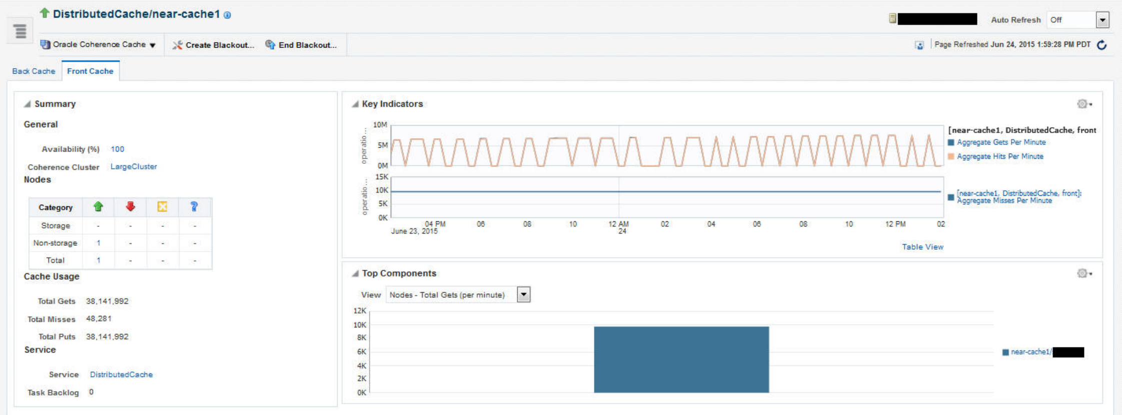

DistributedCache/near-cache1

[redacted] Auto Refresh Off

Oracle Coherence Cache ▾ | Create Blackout... | End Blackout...

Page Refreshed Jun 24, 2015 1:59:28 PM PDT

Back Cache | Front Cache

Summary

General

Availability (%) 100

Coherence Cluster LargeCluster

Nodes

Category	↑	↓	☒	?
Storage	-	-	-	-
Non-storage	1	-	-	-
Total	1	-	-	-

Cache Usage

Total Gets 38,141,992

Total Misses 48,281

Total Puts 38,141,992

Service

Service DistributedCache

Task Backlog 0

Key Indicators

operatio...
10M
5M
0M
operatio...
15K
10K
5K
0K
04 PM
June 23, 2015
06
08
10
12 AM
24
02
04
06
08
10
12 PM
02
[near-cache1, DistributedCache, front
Aggregate Gets Per Minute
Aggregate Hits Per Minute
[near-cache1, DistributedCache, front]: Aggregate Misses Per Minute
Table View

Top Components

View Nodes - Total Gets (per minute)

[redacted]

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**.
- **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.
- **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.

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.

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.

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.

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 13-10 Service Home Page

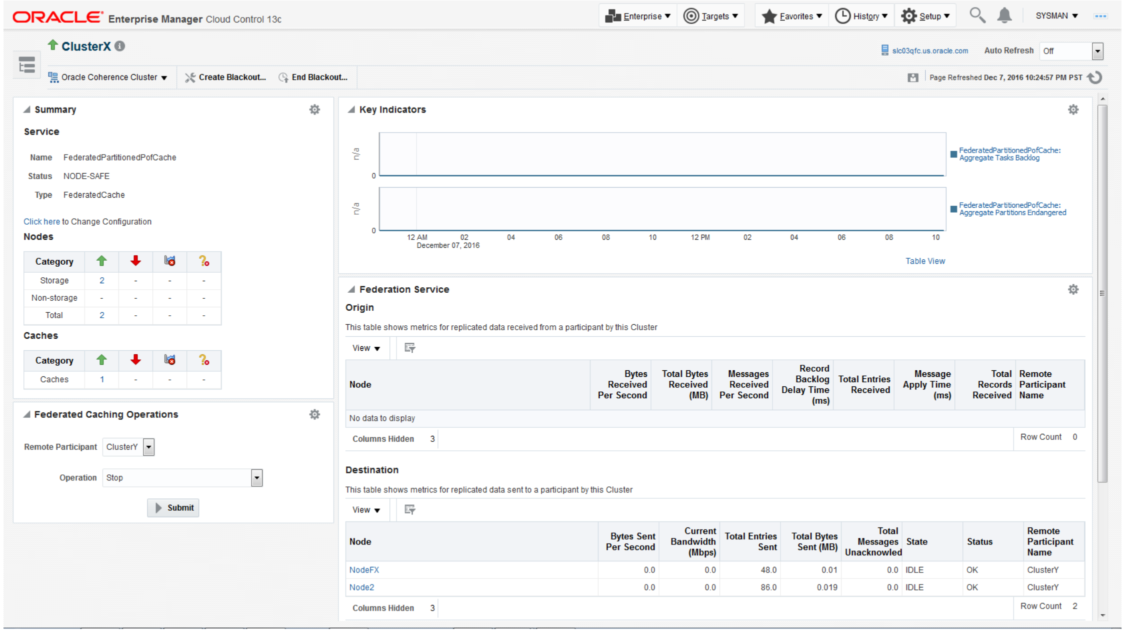

Auto Refresh: Off

Page Refreshed Dec 7, 2016 10:24:57 PM PST

Service

Name FederatedPartitionedPofCache

Status NODE-SAFE

Type FederatedCache

Click here to Change Configuration

Nodes

Category	↑	↓	☒	?
Storage	2	-	-	-
Non-storage	-	-	-	-
Total	2	-	-	-

Caches

Category	↑	↓	☒	?
Caches	1	-	-	-

Federated Caching Operations

Remote Participant ClusterY

Operation Stop

Federation Service

Origin

This table shows metrics for replicated data received from a participant by this Cluster

Node	Bytes Received Per Second	Total Bytes Received (MB)	Messages Received Per Second	Record Backlog Delay Time (ms)	Total Entries Received	Message Apply Time (ms)	Total Records Received	Remote Participant Name
No data to display								

Columns Hidden 3 Row Count 0

Destination

This table shows metrics for replicated data sent to a participant by this Cluster

Node	Bytes Sent Per Second	Current Bandwidth (Mbps)	Total Entries Sent	Total Bytes Sent (MB)	Total Messages Unacknowled	State	Status	Remote Participant Name
NodeFX	0.0	0.0	48.0	0.01	0.0	IDLE	OK	ClusterY
Node2	0.0	0.0	86.0	0.019	0.0	IDLE	OK	ClusterY

Columns Hidden 3 Row Count 2

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 13-11 Confirmation Pop Up

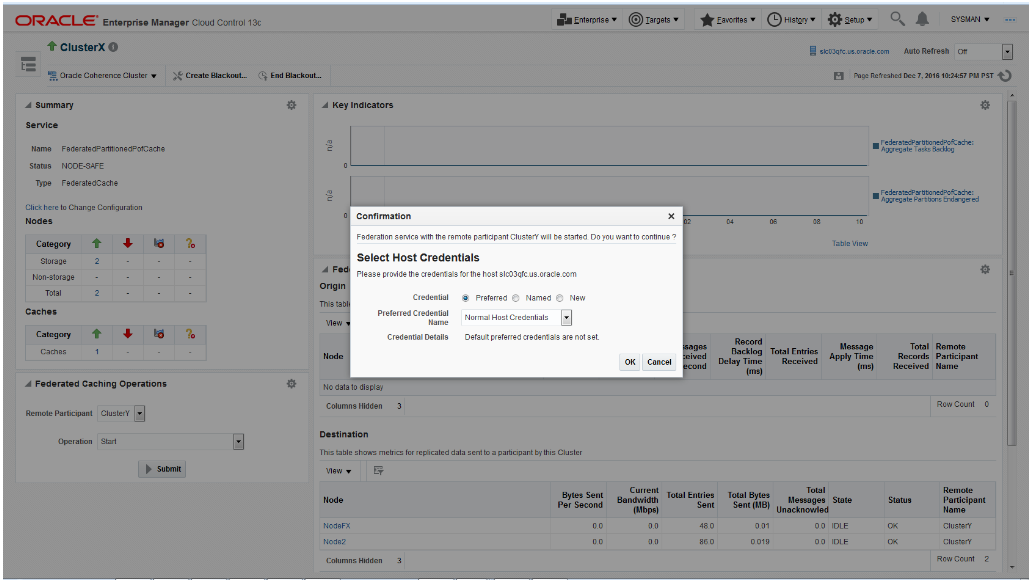
ORACLE Enterprise Manager Cloud Control 13c
Enterprise ▾
Targets ▾
Favorites ▾
History ▾
Setup ▾
SYSMAN ▾
ClusterX
Oracle Coherence Cluster ▾
Create Blackout...
End Blackout...
slc03qfc.us.oracle.com
Auto Refresh Off
Page Refreshed Dec 7, 2016 10:24:57 PM PST
Key Indicators
FederatedPartitionedPofCache: Aggregate Tasks Backlog
FederatedPartitionedPofCache: Aggregate Partitions Endangered
Table View
View ▾
Submit

Summary

Service

Name FederatedPartitionedPofCache

Status NODE-SAFE

Type FederatedCache

Click here to Change Configuration

Nodes

Category	↑	↓		?
Storage	2	-	-	-
Non-storage	-	-	-	-
Total	2	-	-	-

Caches

Category	↑	↓		?
Caches	1	-	-	-

Federated Caching Operations

Remote Participant ClusterY

Operation Start

Confirmation

Federation service with the remote participant ClusterY will be started. Do you want to continue ?

Select Host Credentials

Please provide the credentials for the host slc03qfc.us.oracle.com

Credential ● Preferred ○ Named ○ New

Preferred Credential Name Normal Host Credentials

Credential Details Default preferred credentials are not set.

OK Cancel

Origin

No data to display

Columns Hidden 3 Row Count 0

Destination

This table shows metrics for replicated data sent to a participant by this Cluster

Node	Bytes Sent Per Second	Current Bandwidth (Mbps)	Total Entries Sent	Total Bytes Sent (MB)	Total Messages Unacknowled	State	Status	Remote Participant Name
NodeFX	0.0	0.0	48.0	0.01	0.0	IDLE	OK	ClusterY
Node2	0.0	0.0	86.0	0.019	0.0	IDLE	OK	ClusterY

Columns Hidden 3 Row Count 2

Connection Manager Home Page

Use this page to view the Connection Manager details in the Coherence cluster.

Figure 13-12 Connection Manager Home Page

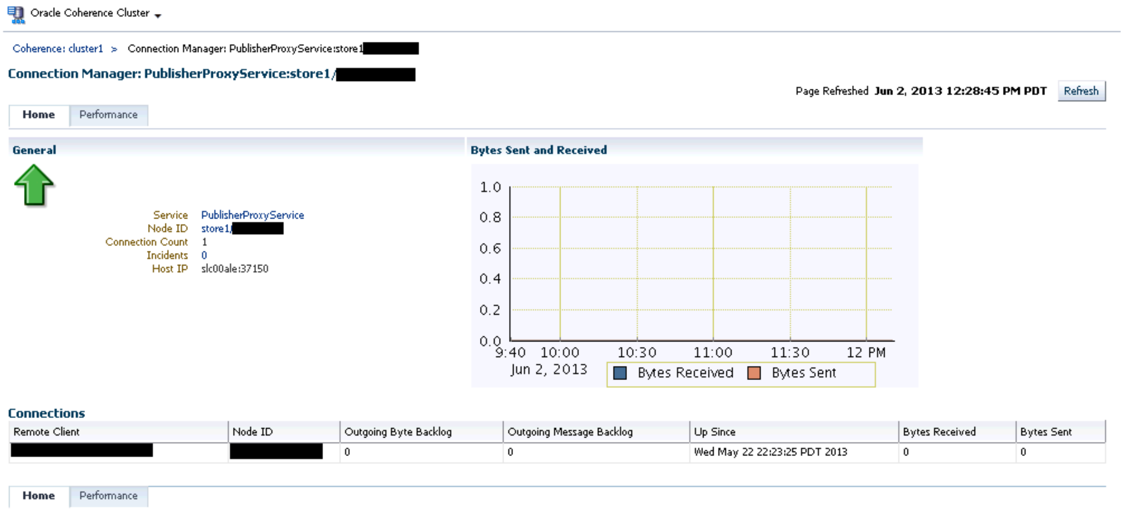

Coherence: cluster1 > Connection Manager: PublisherProxyService:store1[redacted]

Connection Manager: PublisherProxyService:store1/[redacted]

Page Refreshed Jun 2, 2013 12:28:45 PM PDT

General

Service PublisherProxyService

Node ID store1/[redacted]

Connection Count 1

Incidents 0

Host IP slc00ale:37150

Bytes Sent and Received

Connections

Remote Client	Node ID	Outgoing Byte Backlog	Outgoing Message Backlog	Up Since	Bytes Received	Bytes Sent
[redacted]	[redacted]	0	0	Wed May 22 22:23:25 PDT 2013	0	0

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.

Viewing the Summary Pages

These pages describe the target pages such as nodes, caches, services, and so on associated with the cluster.

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 13-13   All Nodes Page

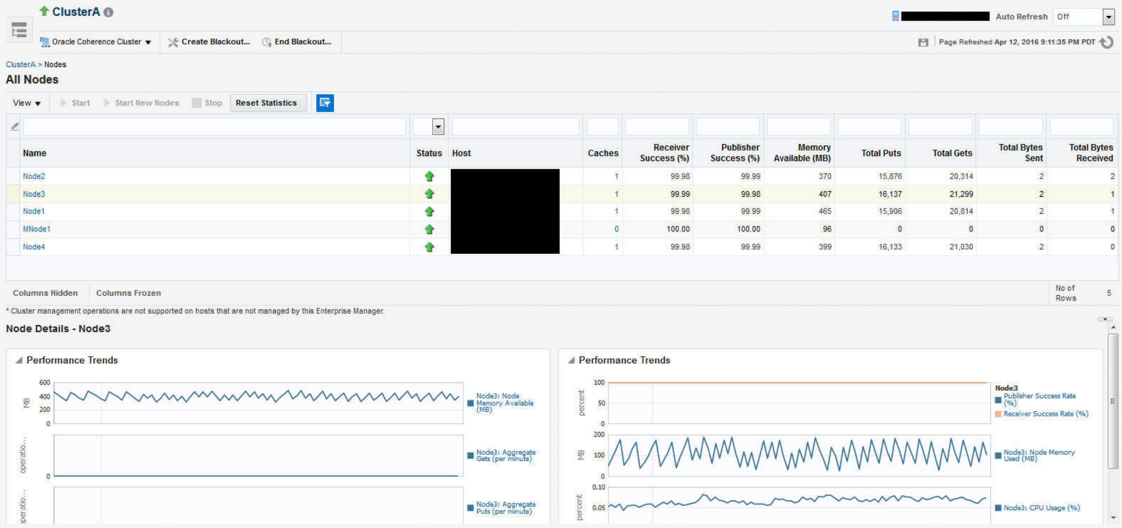

ClusterA

Oracle Coherence Cluster ▾ | Create Blackout... | End Blackout...

[redacted] Auto Refresh Off

Page Refreshed Apr 12, 2016 9:11:35 PM PDT

ClusterA > Nodes

All Nodes

View ▾ | Start | Start New Nodes | Stop | Reset Statistics

Name	Status	Host	Caches	Receiver Success (%)	Publisher Success (%)	Memory Available (MB)	Total Puts	Total Gets	Total Bytes Sent	Total Bytes Received
Node2	⬆	[redacted]	1	99.98	99.99	370	15,876	20,314	2	2
Node3	⬆	[redacted]	1	99.99	99.98	407	16,137	21,299	2	1
Node1	⬆	[redacted]	1	99.98	99.99	465	15,906	20,814	2	1
MNode1	⬆	[redacted]	0	100.00	100.00	96	0	0	0	0
Node4	⬆	[redacted]	1	99.98	99.99	399	16,133	21,030	2	0

Columns Hidden | Columns Frozen | No of Rows 5

* Cluster management operations are not supported on hosts that are not managed by this Enterprise Manager.

Node Details - Node3

Performance Trends
600
400
200
0
MB
Node3: Node Memory Available (MB)
operatio...
0
Node3: Aggregate Gets (per minute)
operatio...
Node3: Aggregate Puts (per minute)
Performance Trends
100
50
0
percent
Node3
Publisher Success Rate (%)
Receiver Success Rate (%)
200
100
0
MB
Node3: Node Memory Used (MB)
0.10
0.05
percent
Node3: CPU Usage (%)

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.

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 13-14   All Caches Page

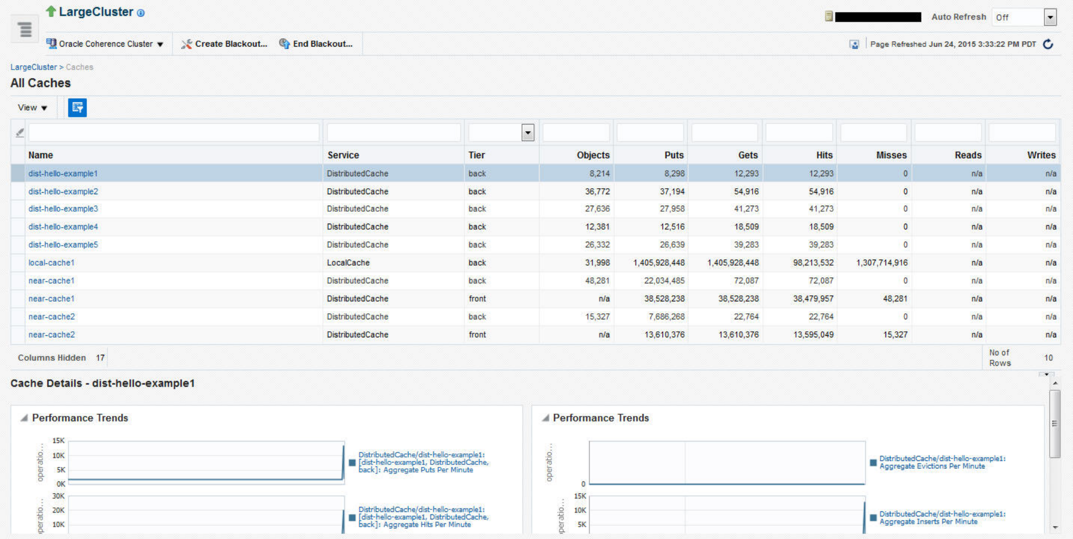

LargeCluster

Oracle Coherence Cluster ▾   Create Blackout...   End Blackout...

[redacted] Auto Refresh Off

Page Refreshed Jun 24, 2015 3:33:22 PM PDT

LargeCluster > Caches

All Caches

View ▾

Name	Service	Tier	Objects	Puts	Gets	Hits	Misses	Reads	Writes
dist-hello-example1	DistributedCache	back	8,214	8,298	12,293	12,293	0	n/a	n/a
dist-hello-example2	DistributedCache	back	36,772	37,194	54,916	54,916	0	n/a	n/a
dist-hello-example3	DistributedCache	back	27,636	27,958	41,273	41,273	0	n/a	n/a
dist-hello-example4	DistributedCache	back	12,381	12,516	18,509	18,509	0	n/a	n/a
dist-hello-example5	DistributedCache	back	26,332	26,639	39,283	39,283	0	n/a	n/a
local-cache1	LocalCache	back	31,998	1,405,928,448	1,405,928,448	98,213,532	1,307,714,916	n/a	n/a
near-cache1	DistributedCache	back	48,281	22,034,485	72,087	72,087	0	n/a	n/a
near-cache1	DistributedCache	front	n/a	38,528,238	38,528,238	38,479,957	48,281	n/a	n/a
near-cache2	DistributedCache	back	15,327	7,686,268	22,764	22,764	0	n/a	n/a
near-cache2	DistributedCache	front	n/a	13,610,376	13,610,376	13,595,049	15,327	n/a	n/a

Columns Hidden 17

No of Rows 10

Cache Details - dist-hello-example1

Performance Trends
operatio...
15K
10K
5K
0K
DistributedCache/dist-hello-example1: [dist-hello-example1, DistributedCache, back]: Aggregate Puts Per Minute
operatio...
30K
20K
10K
DistributedCache/dist-hello-example1: [dist-hello-example1, DistributedCache, back]: Aggregate Hits Per Minute
Performance Trends
operatio...
0
DistributedCache/dist-hello-example1: Aggregate Evictions Per Minute
operatio...
15K
10K
5K
DistributedCache/dist-hello-example1: Aggregate Inserts Per Minute

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.

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 13-15 Services Page

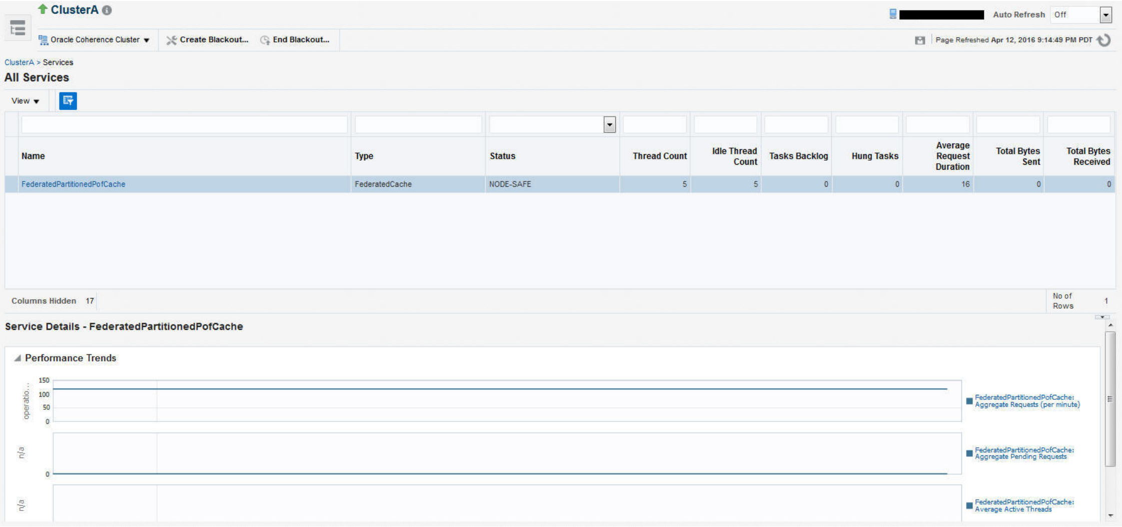

ClusterA

[redacted] Auto Refresh Off

Oracle Coherence Cluster ▾ | Create Blackout... | End Blackout...

Page Refreshed Apr 12, 2016 9:14:49 PM PDT

ClusterA > Services

All Services

View ▾

Name	Type	Status	Thread Count	Idle Thread Count	Tasks Backlog	Hung Tasks	Average Request Duration	Total Bytes Sent	Total Bytes Received
FederatedPartitionedPofCache	FederatedCache	NODE-SAFE	5	5	0	0	16	0	0

Columns Hidden 17

No of Rows 1

Service Details - FederatedPartitionedPofCache

Performance Trends

operatio...
150
100
50
0
FederatedPartitionedPofCache: Aggregate Requests (per minute)
n/a
0
FederatedPartitionedPofCache: Aggregate Pending Requests
n/a
FederatedPartitionedPofCache: Average Active Threads

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.

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.

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.

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

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 13-16   Configure Log Location Settings

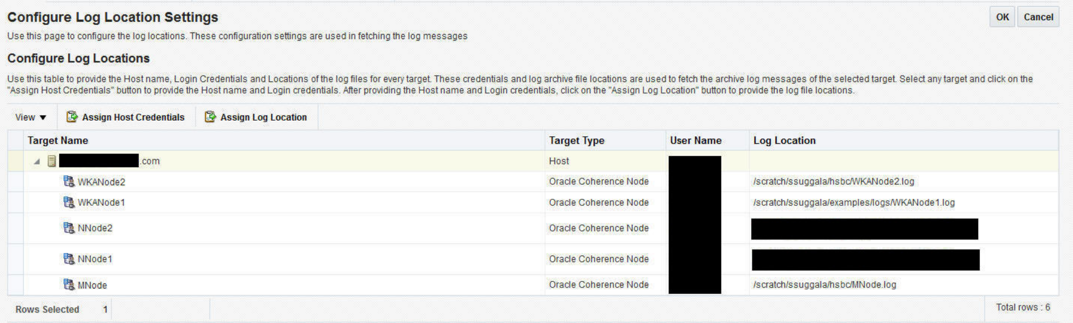

Configure Log Location Settings

Use this page to configure the log locations. These configuration settings are used in fetching the log messages

Configure Log Locations

Use this table to provide the Host name, Login Credentials and Locations of the log files for every target. These credentials and log archive file locations are used to fetch the archive log messages of the selected target. Select any target and click on the "Assign Host Credentials" button to provide the Host name and Login credentials. After providing the Host name and Login credentials, click on the "Assign Log Location" button to provide the log file locations.

Target Name	Target Type	User Name	Log Location
[redacted].com	Host	[redacted]	
WKANode2	Oracle Coherence Node	[redacted]	/scratch/ssuggala/hsbc/WKANode2.log
WKANode1	Oracle Coherence Node	[redacted]	/scratch/ssuggala/examples/logs/WKANode1.log
NNode2	Oracle Coherence Node	[redacted]	[redacted]
NNode1	Oracle Coherence Node	[redacted]	[redacted]
MNode	Oracle Coherence Node	[redacted]	/scratch/ssuggala/hsbc/MNode.log

Rows Selected 1   Total rows : 6

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.

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 13-17   Log Messages

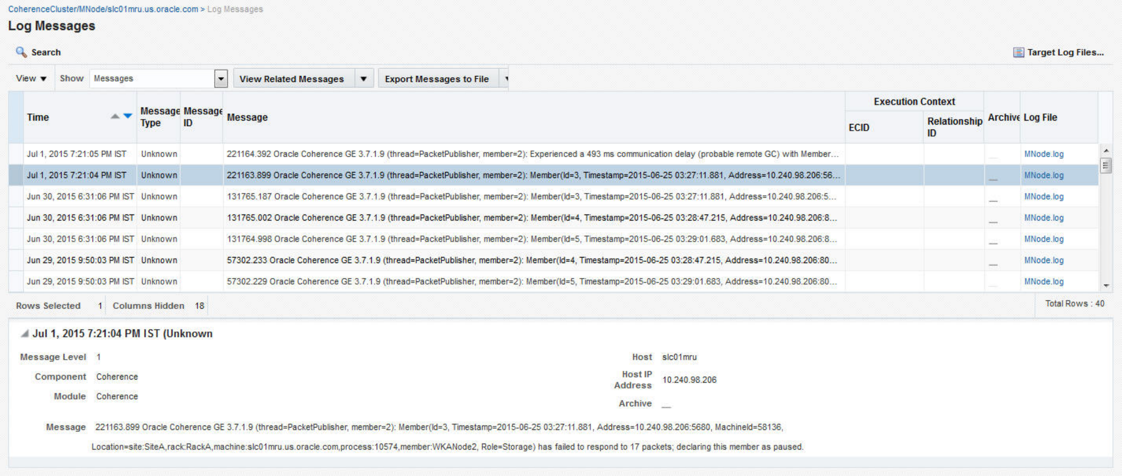

Time	Message Type	Message ID	Message	Execution Context		Archive	Log File
				ECID	Relationship ID		
Jul 1, 2015 7:21:05 PM IST	Unknown		221164.392 Oracle Coherence GE 3.7.1.9 (thread=PacketPublisher, member=2): Experienced a 493 ms communication delay (probable remote GC) with Member...			—	MNode.log
Jul 1, 2015 7:21:04 PM IST	Unknown		221163.899 Oracle Coherence GE 3.7.1.9 (thread=PacketPublisher, member=2): Member(Id=3, Timestamp=2015-06-25 03:27:11.881, Address=10.240.98.206:56...			—	MNode.log
Jun 30, 2015 6:31:06 PM IST	Unknown		131765.187 Oracle Coherence GE 3.7.1.9 (thread=PacketPublisher, member=2): Member(Id=3, Timestamp=2015-06-25 03:27:11.881, Address=10.240.98.206:5...			—	MNode.log
Jun 30, 2015 6:31:06 PM IST	Unknown		131765.002 Oracle Coherence GE 3.7.1.9 (thread=PacketPublisher, member=2): Member(Id=4, Timestamp=2015-06-25 03:28:47.215, Address=10.240.98.206:8...			—	MNode.log
Jun 30, 2015 6:31:06 PM IST	Unknown		131764.998 Oracle Coherence GE 3.7.1.9 (thread=PacketPublisher, member=2): Member(Id=5, Timestamp=2015-06-25 03:29:01.683, Address=10.240.98.206:8...			—	MNode.log
Jun 29, 2015 9:50:03 PM IST	Unknown		57302.233 Oracle Coherence GE 3.7.1.9 (thread=PacketPublisher, member=2): Member(Id=4, Timestamp=2015-06-25 03:28:47.215, Address=10.240.98.206:80...			—	MNode.log
Jun 29, 2015 9:50:03 PM IST	Unknown		57302.229 Oracle Coherence GE 3.7.1.9 (thread=PacketPublisher, member=2): Member(Id=5, Timestamp=2015-06-25 03:29:01.683, Address=10.240.98.206:80...			—	MNode.log

Rows Selected 1   Columns Hidden 18   Total Rows : 40

Jul 1, 2015 7:21:04 PM IST (Unknown

Message Level 1
Component Coherence
Module Coherence
Host slc01mru
Host IP Address 10.240.98.206
Archive —

Message 221163.899 Oracle Coherence GE 3.7.1.9 (thread=PacketPublisher, member=2): Member(Id=3, Timestamp=2015-06-25 03:27:11.881, Address=10.240.98.206:5680, MachineId=58136, Location=site:SiteA,rack:RackA,machine:slc01mru.us.oracle.com,process:10574,member:WKANode2, Role=Storage) has failed to respond to 17 packets; declaring this member as paused.

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.

Viewing the Performance Pages

This section describes the Performance Summary page, and the connection manager performance pages.

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.

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.

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.

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.

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 13-18 Coherence Cluster Topology

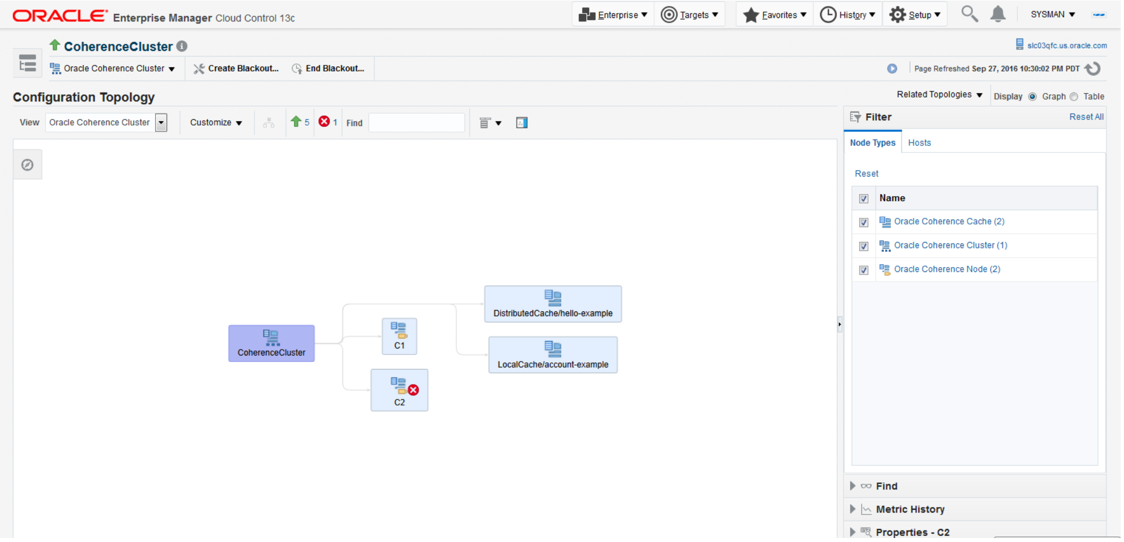

The topology is organized into 3 tiers, cluster, hosts, and services. All nodes running on the host are grouped under host group. All caches of a service are grouped under the service folder. When you select a node, links to all the caches it supports are displayed. If you select a cache, links to all 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.

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.

14

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

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.

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

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.

15

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

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.

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 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 Administrator's Guide*.

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 jamconsport=<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
```

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 jamconsport=<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
```

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.

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.

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.

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.

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 VIII

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

16 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.

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

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.

  - 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.

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.

Monitoring Oracle Identity Management Components in Enterprise Manager

You can use Enterprise Manager to monitor the following Identity Management 11*g* components:

Table 16-1 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	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. 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: • Oracle Adaptive Access Manager Cluster • Oracle Access Manager Cluster • Oracle Identity Manager Cluster Each cluster target is a logically related group of components that are managed as a unit. Every target is part of a WebLogic domain.
Oracle Directory Server Enterprise Edition	The following types of targets will be created for each Oracle Directory Server Enterprise Edition deployment: • 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. Each target provides an interface in Enterprise Manager with access to target overview, customizable performance summary, process control, and configuration management.

Table 16-2 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 11*g* providing access to metrics, alerts, performance charts, and configuration management of the SOA infrastructure instance and its service engines.

Table 16-3 Targets Associated with Identity Management 12c Targets

Enterprise Manager Target Type	Purpose
Oracle Identity Federation Oracle Identity Manager Oracle Unified Directory Oracle Directory Integration Platform Oracle Internet Directory	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. 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: • Oracle Adaptive Access Manager Cluster • Oracle Access Manager Cluster • Oracle Identity Manager Cluster Each cluster target is a logically related group of components that are managed as a unit. Every target is part of a WebLogic domain.

17

Prerequisites for Discovering Oracle Identity Management Targets

This section lists the system requirements and prerequisites needed to discover identity management targets.

System Requirements

Table 17-1 lists the supported Oracle Identity Management products in the Management Pack Plus for Identity Management in Enterprise Manager Cloud Control 13*c*.

Note: For the most up-to-date list of supported platforms, check My Oracle Support Certification Matrix on My Oracle Support (`https://support.oracle.com`).

Table 17-1 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 17-2 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 17-2 (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);
	11*g* 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 11*g* PS 1 (11.1.1.2.0); 11*g* 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 10*g*	Oracle Database
Oracle Identity Management Suite - Single Sign-On Server	10.1.4.2; 10.1.4.3.0	Oracle Application Server 10*g*	Oracle Database
Oracle Identity Management Suite - Delegated Administration Services	10.1.4.2; 10.1.4.3.0	Oracle Application Server 10*g*	Oracle Database
Oracle Identity Management Suite - Directory Integration Platform	10.1.4.2; 10.1.4.3.0	Oracle Application Server 10*g*	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);
	11*g* 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);
	11*g* 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	11*g* (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	11*g* (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 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; 11*g* (11.1.1.3.0); 11.1.1.3.0; 11.1.1.5.0	Not Applicable	Not Applicable

Table 17-2 (Cont.) Supported Identity Management Products and Platforms in Enterprise Manager Cloud Control Release 3

Product	Version	Application Server	Directory Server/ Database
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

Installing Oracle Enterprise Manager

Before you begin configuring Enterprise Manager to manage your Identity Management components, you must install and configure Enterprise Manager on at least one host computer on your network. Oracle recommends that you install Enterprise Manager 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 Enterprise Manager Management Agent on every host that includes the components you want to manage with Enterprise Manager.

For more information, see *Oracle Enterprise Manager Basic Installation Guide*.

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 13*c* 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 13*c* 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 17-1.

Figure 17-1 Plug-Ins Deploy On Options

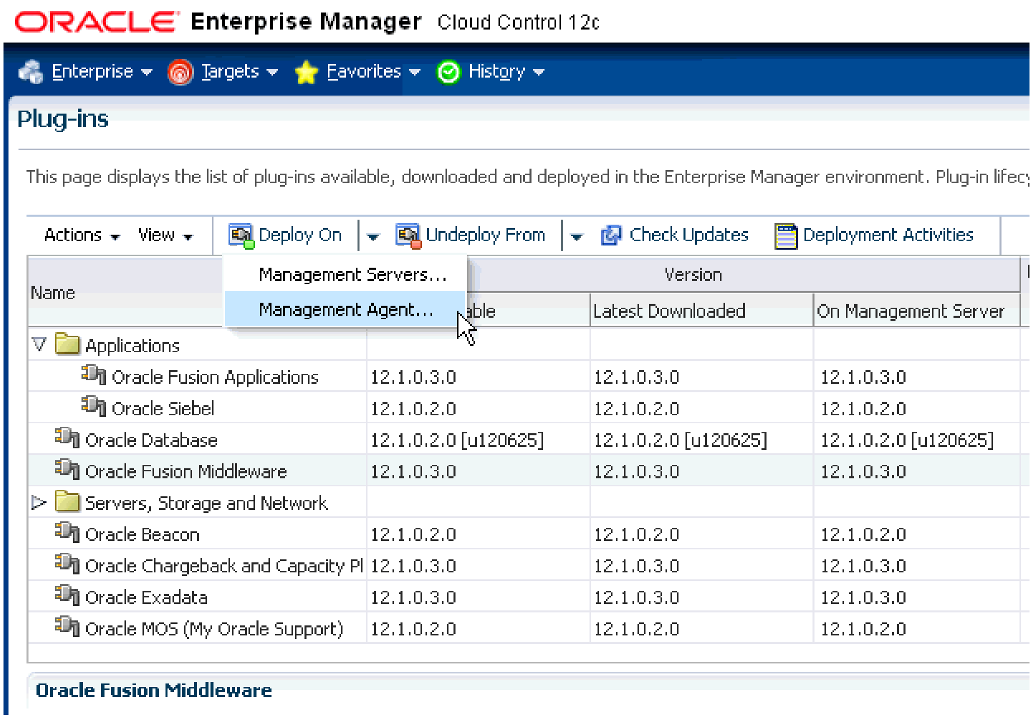

This page displays the list of plug-ins available, downloaded and deployed in the Enterprise Manager environment. Plug-in lifecy

Name	Version		
	able	Latest Downloaded	On Management Server
Applications			
Oracle Fusion Applications	12.1.0.3.0	12.1.0.3.0	12.1.0.3.0
Oracle Siebel	12.1.0.2.0	12.1.0.2.0	12.1.0.2.0
Oracle Database	12.1.0.2.0 [u120625]	12.1.0.2.0 [u120625]	12.1.0.2.0 [u120625]
Oracle Fusion Middleware	12.1.0.3.0	12.1.0.3.0	12.1.0.3.0
Servers, Storage and Network			
Oracle Beacon	12.1.0.2.0	12.1.0.2.0	12.1.0.2.0
Oracle Chargeback and Capacity Pl	12.1.0.3.0	12.1.0.3.0	12.1.0.3.0
Oracle Exadata	12.1.0.3.0	12.1.0.3.0	12.1.0.3.0
Oracle MOS (My Oracle Support)	12.1.0.2.0	12.1.0.2.0	12.1.0.2.0

Oracle Fusion Middleware

18

Discovering and Configuring Oracle Identity Management Targets

This section provides the information needed to discover and configure Oracle Identity Management targets.

Discovering Identity Management Targets

This section describes how to discover Identity Management targets.

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 18-1.

Figure 18-1 Server Properties - Statistics Tab

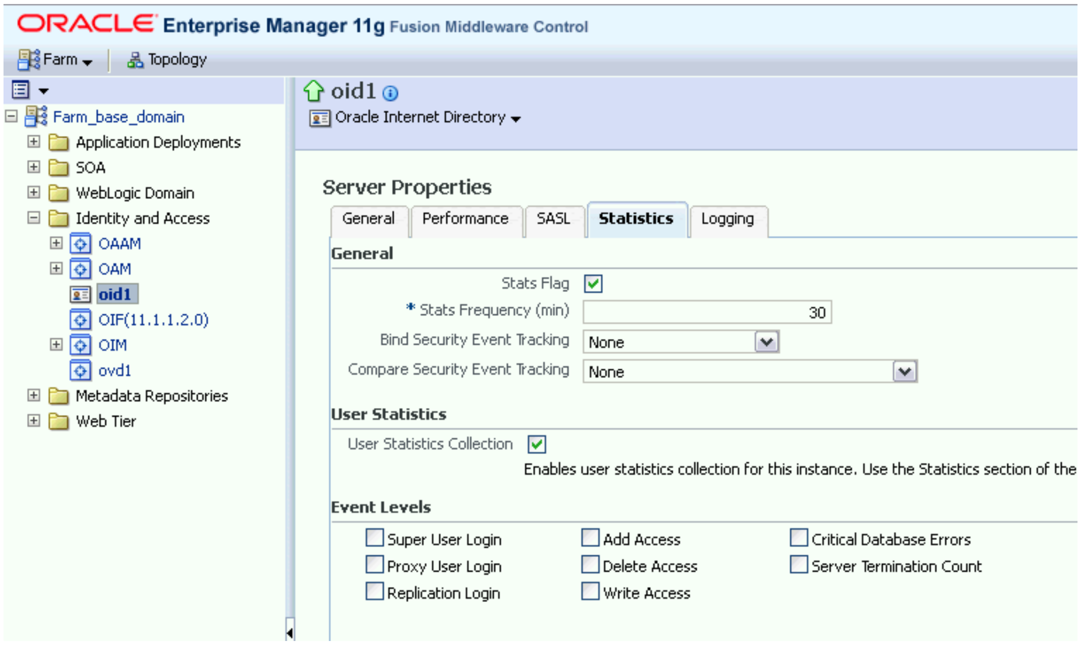

General

Stats Flag ☑

* Stats Frequency (min) 30

Bind Security Event Tracking None

Compare Security Event Tracking None

User Statistics

User Statistics Collection ☑

Enables user statistics collection for this instance. Use the Statistics section of the

Event Levels

□ Super User Login	□ Add Access	□ Critical Database Errors
□ Proxy User Login	□ Delete Access	□ Server Termination Count
□ Replication Login	□ Write Access	

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 18-2.

Figure 18-2 Shared Properties - General Tab

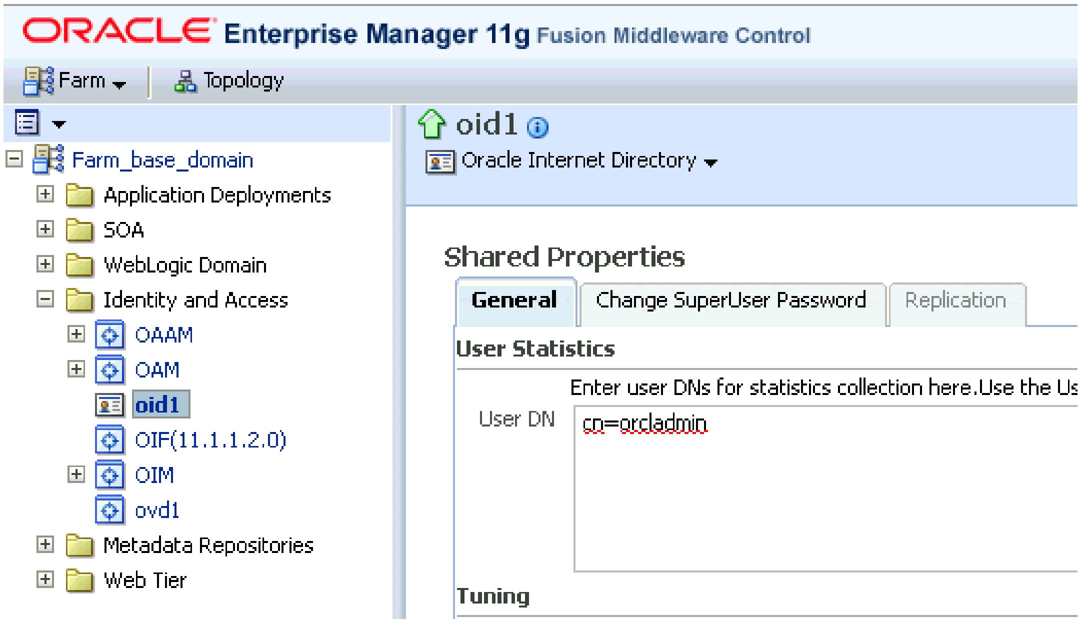

User Statistics

Enter user DNs for statistics collection here.Use the Us

User DN cn=orcladmin

Tuning

Creating Identity Management Elements

This section describes how to create Identity Management elements.

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.

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.

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 IX

Discovering and Monitoring Non-Oracle Middleware

This chapter describes how to discover and monitor non-Oracle middleware components.

19

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

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.

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.

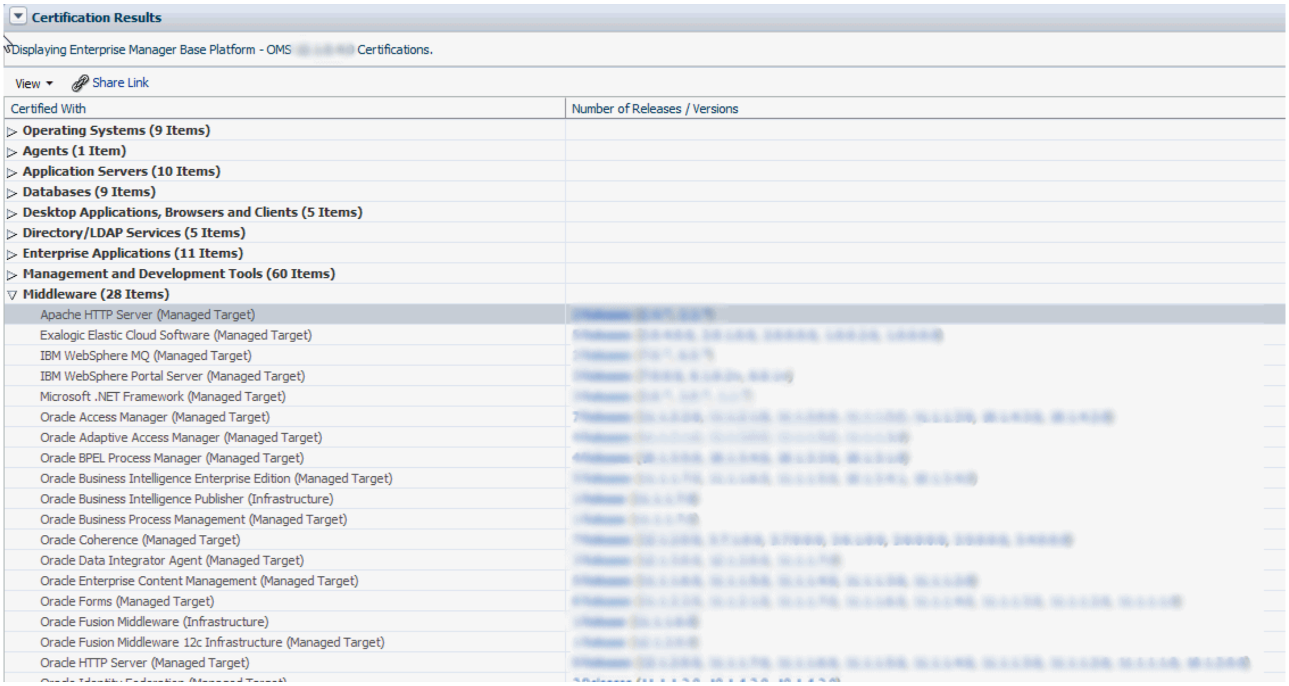

Certification Results

Displaying Enterprise Manager Base Platform - OMS [illegible] Certifications.

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)	[illegible]
Exalogic Elastic Cloud Software (Managed Target)	[illegible]
IBM WebSphere MQ (Managed Target)	[illegible]
IBM WebSphere Portal Server (Managed Target)	[illegible]
Microsoft .NET Framework (Managed Target)	[illegible]
Oracle Access Manager (Managed Target)	[illegible]
Oracle Adaptive Access Manager (Managed Target)	[illegible]
Oracle BPEL Process Manager (Managed Target)	[illegible]
Oracle Business Intelligence Enterprise Edition (Managed Target)	[illegible]
Oracle Business Intelligence Publisher (Infrastructure)	[illegible]
Oracle Business Process Management (Managed Target)	[illegible]
Oracle Coherence (Managed Target)	[illegible]
Oracle Data Integrator Agent (Managed Target)	[illegible]
Oracle Enterprise Content Management (Managed Target)	[illegible]
Oracle Forms (Managed Target)	[illegible]
Oracle Fusion Middleware (Infrastructure)	[illegible]
Oracle Fusion Middleware 12c Infrastructure (Managed Target)	[illegible]
Oracle HTTP Server (Managed Target)	[illegible]

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.

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**.

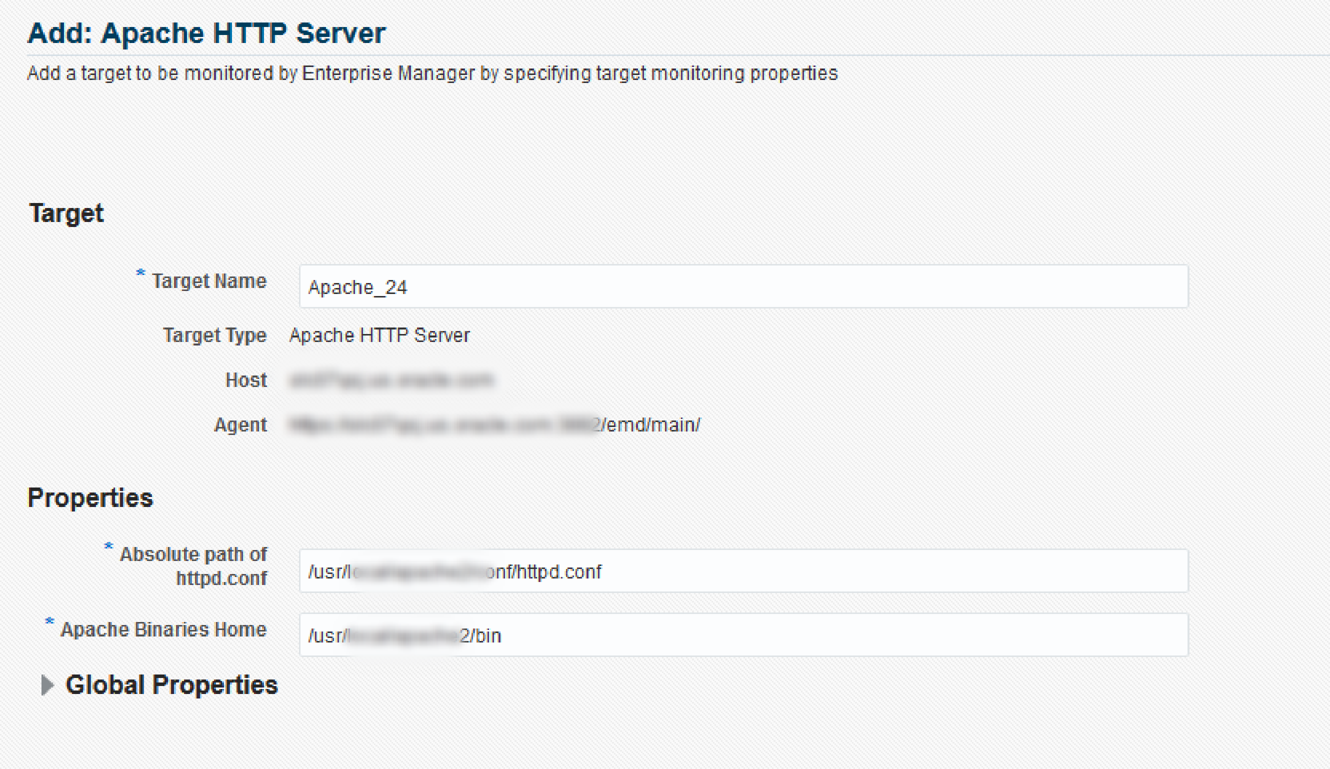

Add: Apache HTTP Server

Add a target to be monitored by Enterprise Manager by specifying target monitoring properties

Target

* Target Name	Apache_24
Target Type	Apache HTTP Server
Host	[redacted]
Agent	[redacted]/emd/main/

Properties

* Absolute path of httpd.conf	/usr/l[redacted]onf/httpd.conf
* Apache Binaries Home	/usr/[redacted]2/bin

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.

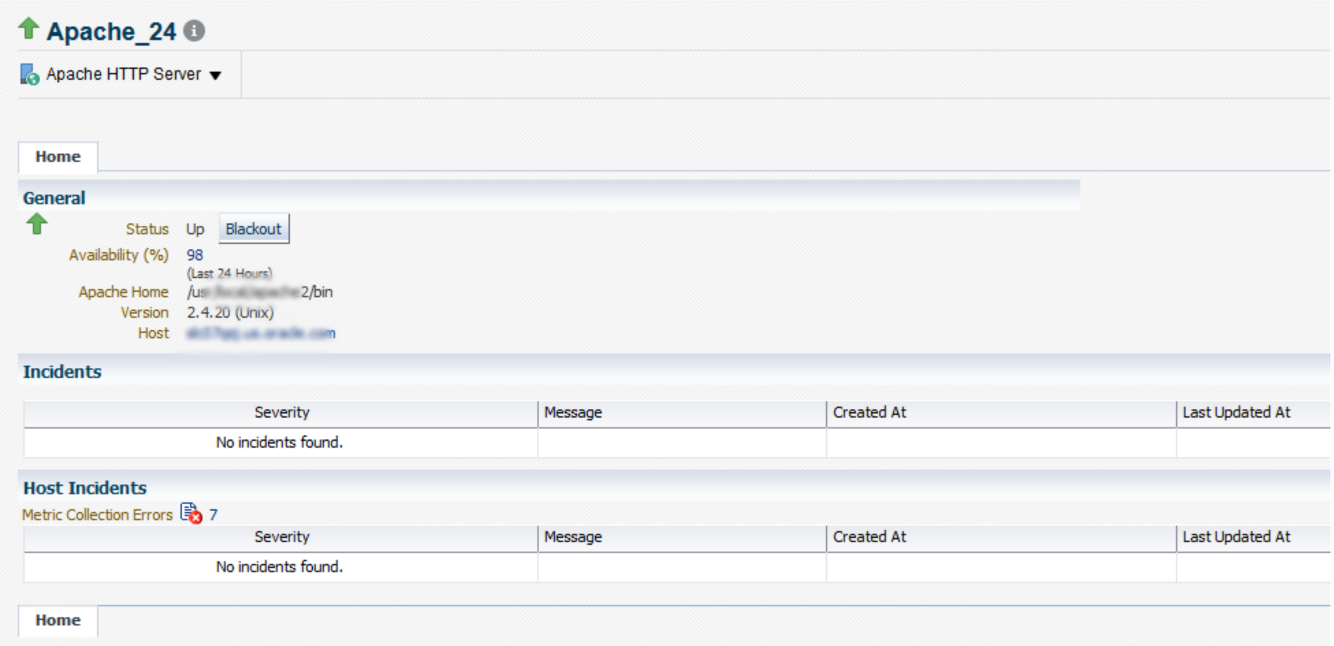

General

Status Up
Availability (%) 98 (Last 24 Hours)
Apache Home /us[redacted]2/bin
Version 2.4.20 (Unix)
Host [redacted]

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.			

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.

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.**

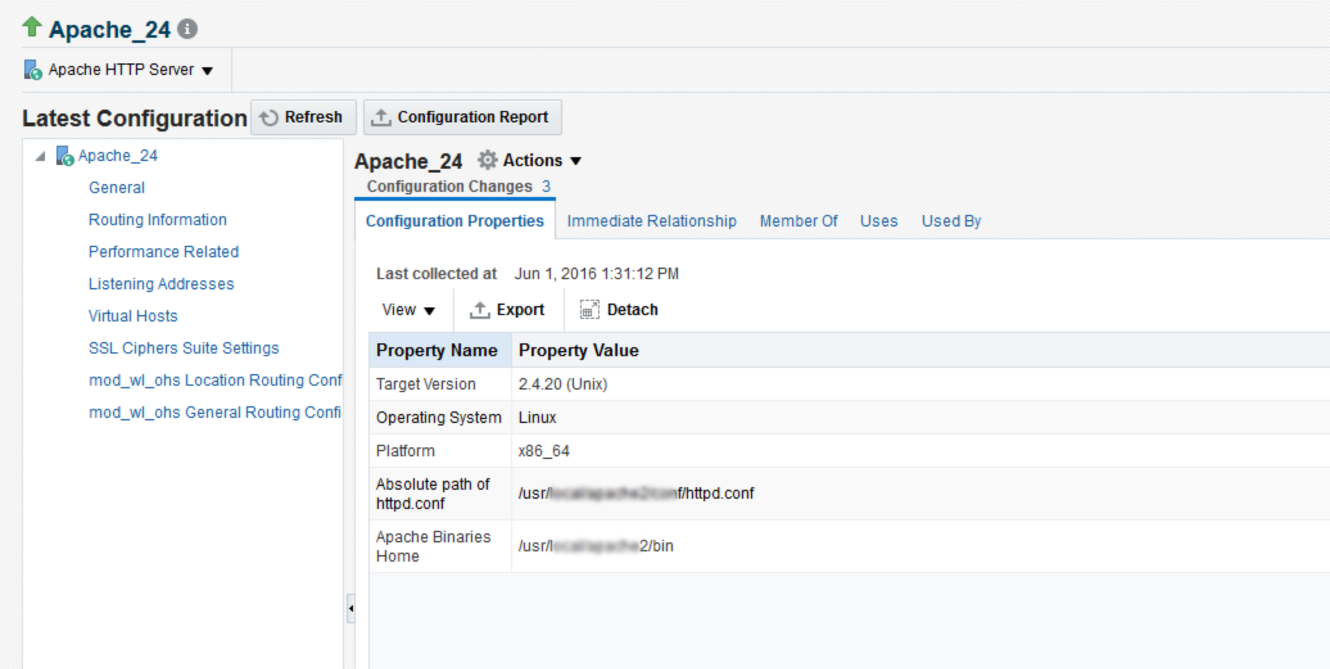

Last collected at Jun 1, 2016 1:31:12 PM

Property Name	Property Value
Target Version	2.4.20 (Unix)
Operating System	Linux
Platform	x86_64
Absolute path of httpd.conf	/usr/[redacted]f/httpd.conf
Apache Binaries Home	/usr/l[redacted]2/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.

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 X
Managing Oracle Data Integrator

The chapter in this part describes how you can configure and monitor Oracle Data Integrator.

This part contains the following chapter:

20

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

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.

Monitoring Oracle Data Integrator

This section describes the following:

- Monitoring Oracle Data Integrator
- Monitoring ODI Agents
- Monitoring Repositories
- Monitoring Load Plan Executions and Sessions

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:

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.

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.

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.

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.

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.

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:

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.

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.

Element	Description
View Performance	Click the eye glass icon to view the performance data of the Agent. The metrics include: • 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.

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:

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.

Repositories

Use this region to view details of the work repositories.

Element	Description
Name	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. • 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.
Status	Displays the status of the Work Repository database. • 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.

Element	Description
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. • 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.

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. • Up (green arrow): on • Down (red arrow): off
Host	Displays the name of the host on which the database resides.
SID	Database SID.

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

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.

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.

Search Sessions/LPEs

Use this region to search for sessions and load plan executions for all master and work repositories.

The latest specified search criteria are always retained. Specify a new criteria and click **Search** to see the updated results. Or, click **Reset** to reset the search form (you must still click **Search** to see the updated results). Note that the search criteria are reset each time you log out or navigate away from all the tabbed pages of the Oracle Data Integrator Cloud Control application.

Element	Description
Master Repository	Select the Master Repository containing the session information.
Work Repository	Select the Work Repository containing the session information.
Execution Agent	Select the Agent used to execute the session.
Context	Select the session's execution context.
Execution Type	Select Sessions, Load Plan Executions, or All.
Begin Date	Use the calendar icon to select a date at which to start the search for sessions. Only session started after this date will be returned.
End Date	Use the calendar icon to select a date at which to end the search for load plan executions and sessions. Only load plan executions and sessions ended before this date will be returned.
User Name	Name of the ODI user who started the execution.
Status	Select All or narrow the search to display specific statuses: Error, Running, Done, Warning, or Waiting. For example, you can select to view only Running and Warning statuses.
Message	Error message of the Load Plan Execution/Session run.
Keywords	Type keywords to narrow the search. When using multiple keywords, use a comma to separate each keyword, do not include spaces. For example use: lpe1,lpe2.
Execution Name	Type the name of the load plan execution.
Error Records	Select All or narrow the search to display load plan executions and sessions With Error Records or Without Error Records.
Execution ID	Specific Load Plan Execution or Session identifier.

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: • 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.
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.

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.

Element	Description
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.

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

Starting Up, Shutting Down, and Restarting Oracle Data Integrator Agents

> **Note:**
>
> - Oracle Process Manager and Notification (OPMN) is used for release 11*g* Standalone Agents. WebLogic Management Framework is used for release 12*c* 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 11*g* Standalone Agents managed by OPMN instances, and not for 12*c* 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*.

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.

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.**

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.

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.

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 12*c* 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.

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

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.

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.

Target Incidents

The Target Incidents region displays notifications raised by the Agents attached to this Repository.

Element	Description
Severity	Seriousness of the incident. • 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.

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. • Fatal - Corresponding service is no longer available. For example, a monitored target is down (target down event). A Fatal severity is the highest level severity and only applies to the Target Availability event type. • Critical - Immediate action is required in a particular area. The area is either not functional or indicative of imminent problems. • Warning - Attention is required in a particular area, but the area is still functional. • Advisory - While the particular area does not require immediate attention, caution is recommended regarding the area's current state. • Clear - Conditions that raised the incident have been resolved.
ID	Incident ID.
Summary	Summary description of the incident.
Category	Classification of an incident, for example, Error.

Load Balancing Agents

The Load Balancing Agents region displays (if using ODI Load Balancing) the status and session metrics for the Agents declared as child Agents of the current Agent.

Element	Description
Name	Displays the name of the agent. This is the name you specified when you created the Agent in Oracle Data Integrator. Select an Agent to display the corresponding Agent Home page.
Status	Displays the status of the Agent. • 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.

Element	Description
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. • 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.

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.

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).

Element	Description
Port	Displays the HTTP listener port number. If using SSO, enter the port of the machine where Oracle HTTP Server 10*g* or 11*g* Webgate is installed.
Context Root	Displays the Web application's context root.
Protocol	Displays the protocol of the connection

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