Oracle® Enterprise Manager

Cloud Control Administrator's Guide

12*c* Release 1 (12.1.0.1)

E24473-01

October 2011

ORACLE®

Oracle Enterprise Manager Cloud Control Administrator's Guide, 12*c* Release 1 (12.1.0.1)

E24473-01

Contents

3 Notifications

6 Group Management

7 Job System and Corrective Actions

8 Compliance

12 Sizing Your Enterprise Manager Deployment

16 Enterprise Manager Diagnosability

17 Updating Cloud Control

20 Monitoring Using Web Services and JMX

21 Configuring Services

22 Identity Management

23 Lifecycle Management

24 Starting and Stopping Enterprise Manager Components

25 Locating and Configuring Enterprise Manager Log Files

26 Maintaining and Troubleshooting the Management Repository

31 Installation to Support Real-time Configuration Change Monitoring

32 Setting Up Enterprise Manager High Availability

Index

Preface

This guide describes how to set up a Private Cloud, manage and deploy virtualization targets with Oracle Enterprise Manager 12c Release 1.

The preface covers the following:

Audience

This document is intended for cloud administrators who want to setup and manage the cloud infrastructure. It is also intended for administrators and users of the Self Service Portal.

Documentation Accessibility

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

Access to Oracle Support

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

Related Documents

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

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

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

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
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

Administration

This section covers Enterprise Manager administrative functionality.

1 Adding Targets

Oracle components that are managed and monitored by Cloud Control, such as an Oracle Database or an Oracle WebLogic Server domain, are known as “managed targets”.

Before a target can be managed, a Management Agent must first be installed on the host machine the target is running on. The target itself must then be assigned to a Management Agent, thereby promoting it to managed target status.

Cloud Control offers two modes for installing Management Agents to monitor and manage potential targets: Manually or using automatic target discovery.

The manual process is exactly that: You explicitly add a specific host or Oracle component as a target to bring under management.

In automatic discovery, you create an Enterprise Manager job that searches host machines for possible targets on a regularly-scheduled basis. A key benefit is that as new Oracle components are added to your infrastructure, they can be found and brought under management.

See the following sections for more information about discovering and adding new targets:

- Configuring Automatic Discovery
- Manually Adding Targets

1.1 Configuring Automatic Discovery

Automatic discovery refers to the process of scanning host machines for Oracle components that can be managed and monitored by Enterprise Manager Cloud Control. You can decide upon a schedule for discovery, the target types to be discovered and the hosts to scan for targets. Discovered targets can then be promoted to managed target status, enabling Enterprise Manager to manage the targets.

Once automatic discovery has been configured, you can check the Auto Discovery Results page on a periodic basis to see what new targets have been discovered.

See the following sections for instructions on using the various self-discovery options:

- Configuring Automatic Discovery of Un-Managed Host Machines Using IP Scan
- Configuring Automatic Discovery of Targets On Managed Hosts
- Checking For and Promoting Discovered Targets

1.1.1 Configuring Automatic Discovery of Un-Managed Host Machines Using IP Scan

In automatic host discovery, a single Management Agent is tasked to scan the entire network based on IP address ranges that you specify. It then returns a list of "un-managed" host machines - that is, host machines that do not yet have a Management Agent installed - with a list of ports in use that match the ranges you specified. The name of the service using each port is also returned.

By looking at the list of services and ports, you should be able to determine what types of Oracle components have been discovered. For example, if a host is returned with port 7001 in use, you can assume that this port is associated with an Oracle WebLogic Server domain that can be promoted to managed target status.

The next step is to deploy Management Agents to the hosts you want to promote to managed status. Once a Management Agent is deployed to the host, any Oracle components running on the host will be discovered and reported as potential targets. These components can then be promoted to managed target status, enabling them to be managed and monitored by Cloud Control.

To be most effective, automatic discovery should ideally be run by a network administrator with an overview of what Oracle components are running on what ports.

Note that because the network will be scanned, the Sudo Privilege Delegation must be set on the Management Agent host that will perform the scan. Typically you will use the Management Agent installed by default on the Oracle Management Service host that Cloud Control is running, which means you can set this privilege on the Cloud Control machine. Privileges are managed via the Manage Privilege Delegation Settings page, which is accessed by selecting **Security>Privilege Delegation** from the **Setup** menu. XXX LINK TO SECURITY CHAPTER

To discover and configure hosts using IP scan, follow these steps:

1. Log into Enterprise Manager Cloud Control.
2. Click **Setup** and then click **Add Target** followed by **Configure Auto Discovery** from the drop-down menus.
3. Click the **Configure** icon in the Hosts and Virtual Server Discovery Using IP Scan in the Configure Auto Discovery table.
4. Click **Create**. You will not create the discover job. By default, the Name field will be populated with a title including that date and time the job was created. Note that you can edit the discovery jobs and schedule discovery to run immediately or later.
5. Click **Add**. You will now select the Management Agent that will perform the network scan. You can select the Management Agent that is installed by default on the Oracle Management Service host, or can select another Agent if desired.

   Note that because the entire network will be scanned, the Sudo Privilege Delegation must be set on the Management Agent host that will perform the scan.
6. Select the agent in the IP Ranges for scan table, and enter the IP ranges to scan. You can specify any or even all of the following:
   - One or more absolute hostnames, each separated by a space; for example: `host1.acme.com host3.acme.com`
   - One or more IP addresses, each separated by a space

- A range of addresses; for example: `10.0.0-255.1-250`. Note that IP addresses and IP ranges must be separated by a comma; for example: `10.0.0-255.1-250`
- Classless Inter-Domain Routing (CIDR) notations; for example: `128.16.10.0/24`

Separate each value with a space; for example:

```
host1.acme.com 192.168.0.1 128.16.10.0/24
10.0.0-255.1-250,254
```

7. A default list of ports to scan within the IP ranges you specified is listed in the Configure Ports table. These are default ports typically used by the listed Oracle components.

   To modify the port values for a component, select the component in the table and change the values accordingly. Up to 10 ports and/or port ranges can be specified.

8. If you want to add more component ports to the list, click **Add**. Enter the name of the service to include, and specify the port(s) or port range to scan.

9. Click **Save and Submit IP Scan** when finished. The Job Details panel now opens. This is where you can specify:

   - The schedule at which the discovery job will run. Note that you can start the job immediately.
   - The credentials set on the Management Agent that will perform the scan.

     As noted, the Sudo Privilege Delegation must be set on the Management Agent host that will perform the scan. The named credential that will be used must be configured to run as root.

10. After the discovery job executes, you can check for discovered hosts that may contain potential targets. You can do this two ways:

    - Select the job in the Host Discovery page, then click **View Discovered Targets**;

      or:

    - Select **Add Target>Auto Discovery Results** from the **Setup** menu.

11. Click the **Host Targets** tab. All discovered hosts are listed, with the open ports and identifiable service names shown. Based on your understanding of the Oracle components deployed on your network, you should be able to determine the types of potential targets that have been discovered.

12. Select a host in the table, then click **Promote** to promote the host to managed target status. The Add Host Targets wizard opens. You will use this wizard to install a Management Agent on the host.

    For instructions on installing a Management Agent, see "Installing Oracle Management Agent" in the *Enterprise Manager Cloud Control Basic Installation Guide*.

13. After the Management Agent has been successfully installed on the host, targets running on the host will be discovered as potential targets. See Section 1.1.3, "Checking For and Promoting Discovered Targets" for details on promoting targets.

1.1.2 Configuring Automatic Discovery of Targets On Managed Hosts

By default, automatic discovery of new Oracle targets on hosts that are already managed targets - that is, host machines that have a Management Agent installed - is

pre-configured to run on each host. Automatic target discovery is the most efficient way to discover potential targets on managed hosts, as Cloud Control can search one or more hosts for multiple types of targets at the same time.

Automatic discovery is enabled for all supported target types by default except for Oracle Fusion Middleware, which requires that a search paramter be provided.

The automatic discovery configuration is defined within a"discovery module", which you can modify to suit your requirements. You can schedule discovery to run on all hosts in the discovery module at the same interval, or can configure separate schedules for each host.

To configure automatic discovery on one or more managed hosts, follow these steps:

1. Login into Enterprise Manager Cloud Control.
2. Select **Add Target>Configure Auto Discovery** from the **Setup** menu.
3. Click the **Configure** icon in the **Multiple Target-type Discovery on Single Host** row in the Discovery table.
4. Expand **Search**, then enter the hostname for the host you want to check for targets in the Agent Host Name field. The host must have a Management Agent installed on it.
5. Click **Search**. The host will be added to the table below.
6. Select the host in the table and click **Configure**.
7. Set the schedule at which the discovery job will be run, in days. This schedule will be applied to all selected hosts. By default the job will run every 24 hours.
8. Select the Oracle component types you want to search the host for. Note that you must supply search parameters for some target types. To specify a parameter, select the target type in the Discovery Module column and click **Edit Parameters**.
   - Oracle Cluster and High Availability Service: No parameters required
   - Oracle Database, Listener and Automatic Storage Management: Specify the path to the Clusterware Home.
   - Oracle Home Discovery: No parameters required.
   - Oracle Secure Backup: No parameters required.
   - Oracle Fusion Middleware: Specify * (the "star" character) to search all Middleware Homes, or specify the path to one or more Middleware Homes on the host, each separated by a comma.
9. Click **OK** when finished. The host has been added to the discovery module.
10. Repeat these steps for each additional host you want to add to the discovery module.
11. Click **Run Discovery Now** to discover targets immediately.

    Note that the discovery job will also run at the scheduled daily interval.
12. Once targets have been discovered, you can promote them to managed status. See Section 1.1.3, "Checking For and Promoting Discovered Targets" for details on promoting targets.

1.1.3 Checking For and Promoting Discovered Targets

Once automatic discovery has been configured, you should check the Auto Discovery Results page on a regular basis to see what targets have been discovered. You can then promote targets to managed status.

To promote discovered targets to managed status, follow these steps:

1. Log in into Enterprise Manager.
2. After the discovery job executes, you can check for discovered hosts that may contain potential targets. You can do this two ways:
   - Select the job in the Host Discovery page, then click **View Discovered Targets**; or
   - Select **Add Target>Auto Discovery Results** from the **Setup** menu.
3. Select a target to promote, then click **Promote**. A wizard specific to the target type you are promoting opens. Supply the required values.
4. Click the **Non-Host Targets**.You can choose one or several targets to promote.
5. Note that you can optionally click **Ignore** for a discovered target, essentially marking it to be processed at a later time.

   Ignored targets will be displayed in the Ignored Targets tab, and will remain in Cloud Control as un-managed targets until you decide to either promote or remove them.
6. Check the target type home page to verify that the target is promoted as an Enterprise Manager target. Once a target is successfully promoted, the Management Agent installed on the target host will begin collecting metric data on the target.

> **Note:** Enterprise Manager does not support simultaneous promotion of multiple targets. Additionally, multiple selection of database targets has been disabled to avoid a user selecting RAC databases across clusters. This is similar to the user-guided discovery feature where a user cannot discover targets across a cluster in the same session.

1.2 Manually Adding Targets

In addition to automatic discovery, Cloud Control allows you to manually add hosts as well as a wide variety of Oracle software and components as managed targets. When you add a target manually, you do not need to go through the process of discovery by adding the target directly. Discovering targets in this way eliminates the need to consume resources on the agent to perform discovery when it is not needed.

You must be able to specify the properties of a target to be managed and create an Enterprise Manager managed target.

Not all target types can be manually added. During registration with the discovery framework, the target type owner indicates whether a target type can be manually added or not.

See the following sections for instructions:

- Manually Adding Host Targets
- Manually Adding Non-Host Targets

1.2.1 Manually Adding Host Targets

A wizard guides you through the process of manually deploying a Management Agent to a new host target.

For instructions on installing a Management Agent, see "Installing Oracle Management Agent" in the *Enterprise Manager Cloud Control Basic Installation Guide*.

1.2.2 Manually Adding Non-Host Targets

A configuration page or wizard based on target type metadata listing all the instance properties required to manage target is displayed.

You can specify a name for the target and provide the required configuration information.

To add targets manually to Enterprise Manager, follow these steps:

1. Log in into Enterprise Manager.
2. Click **Setup,** then click **Add Target** followed by **Add Targets Manually** from the drop-down menus. Enterprise Manager displays the Add Targets Manually page.
3. Under the Add Targets Manually page, go to the Add Targets Manually sub-section and choose an option:
   - Add Non-Host Targets Using Guided Process

     Choose one of the target types to add, such as **Oracle Cluster and High Availability Service**, **Oracle Database Machine**, or **WebLogic Domain Discovery**. This process will also add related targets.
   - Add Non-Host Targets by Specifying Target Monitoring Properties

     Choose one of the target types to add, such as **Fusion J2EE Application**, **Applications Utilities**, or **Supplier Portal**.
4. After you select the target type, you will follow a wizard specific to the target type to add the target.

   Upon confirmation, the target becomes a managed target in Enterprise Manager. Enterprise Manager simply accepts the information, performs validation of the supplied data where possible and starts monitoring the target.

2

Incident Management

Incident and problem management allows you to monitor and resolve service disruptions quickly and efficiently.

This chapter covers the following topics:

2.1 Overview: Monitoring and Managing Via Incidents

Enterprise Manager Cloud Control 12c greatly expands target monitoring and management capability beyond previous releases by letting you focus on what is important from a broader monitoring/management perspective rather than focusing on discrete events that may or may not be relevant to a particular situation.

> **Note:** Also available is the mobile application for managing the incidents and problems, on the go. For more information, see Chapter 9, "Cloud Control Mobile".

What is an event?

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: database target going down,

performance threshold violation, change in application configuration files, successful completion of job execution, or job failure.

Previous versions of Enterprise Manger generated alerts for exception conditions--metric alerts. For Enterprise Manager 12c,metric alerts are a type of event , one of many different event types. This event model significantly raises the number of of conditions in an IT infrastructure for which Enterprise Manager can detect and raises events.

What is an incident?

An incident is the situation or issue you need to act on. By definition, an incident is 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.

When you create an incident, you define a macroscopic set of conditions that you want to monitor and/or manage. In general, you should not need to modify an incident once it is defined, but only the events that make up the incident. An incident may consist of a single event, as might be the case when you are only interested in whether a single database is up or down, or something more complex with multiple events as would be the case when monitoring host resources where you are interested in a variety of metric alert conditions such as disk space utilization, CPU load.

Managing your environment via incidents is carried out through Incident Manger, Enterprise Manager's new console which provides you with a centralized location from which to view, manage, diagnose and resolve incidents as well as identify, resolve and eliminate the root cause of disruptions. See "Incident Manager Console" on page 2-11 for more information.

2.2 Events

By definition, an event is a significant occurrence within your IT infrastructure that Enterprise Manager can detect and subsequently notify interested parties or take action on. An event has very specific attributes that allow Enterprise Manager (and ultimately an Enterprise Manager administrator) to identify, categorize, and manage the event. All events have the following attributes

- Type
- Severity
- Entity on which the event is raised
- Message
- Timestamp
- Category

Event Types

Previous versions of Enterprise Manager let you to monitor and manage by discrete signals and notified you by raising a metric alert as a result of threshold violations. For Enterprise Manager 12c, a metric alert is now just one of several categories of event conditions for which Enterprise Manager can monitor. These event conditions are called Event Types. As shown in the following list, the range of events types greatly expands Enterprise Manager's monitoring flexibility:

- Target Availability (up/down)
- Metric Alert

- Job Status Change
- Compliance Standard Violation Event
- Metric Evaluation Error
- SLA Alert
- High Availability
- User reported event

Incidents allow you to manage many discrete event types by providing an intuitive way to combine them into meaningful issues that you can act upon.

Event Severity

Another event attribute is severity. Just as previous versions of Enterprise Manager utilized metric alert severity levels, this concept has been extended to all event types. Events can have the following severity levels:

- **Fatal**

  The monitored target is down.

- **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. This severity is used primarily for compliance standards.

- **Informational**

  A specific area condition has just occurred.

2.3 Incidents

A previously mentioned, you monitor and manage your Enterprise Manager environment via incidents and not discrete events (even though an incident can conceivably consist of a single event). Managing by incident means rather than managing discrete events for your system (for example, a single target metric warning threshold being exceeding), you now manage an incident that may consist of one significant event (for example, a target down event) or combination of related events (for example, multiple metric thresholds being exceeded). Incidents add an intuitive layer of abstraction that allows you to manage your monitored systems more efficiently. Once you define an incident, it becomes Enterprise Manager's job to monitor for the specified events of interest. This allows you to quickly identify, resolve, and eliminate root causes of monitored system disruptions.

When an incident is created, Enterprise Manager makes available a variety of functions that covers the incident management workflow allowing you to manage and track the incident through its complete lifecycle. Incident management functions allow you to:

- Assign incident ownership

- Track the incident resolution status
- Set incident priority
- Set incident escalation level
- Access My Oracle Support directly for in-context access to the knowledge base.
- Access direct in-context diagnostic/action links to relevant Enterprise Manager functionality allowing you to quick diagnose or resolve the incident

All incident management/tracking operations are carried out from the Incident Manager console. Incident management can be automated using incident rules.

The following examples illustrate how incidents are constructed and how attributes map to various stages of the incident lifecycle.

2.3.1 Incident Composed of a Single Event

The simplest incident is composed of a single event. Since the incident serves as a wrapper for related events, in this case the incident and the event can be considered one in the same.

In the following example, you are concerned about the availability of a single, very important database. Hence, you are only interested in a single event, which in this case is a database availability event. You create an incident whereby if the database encounters problems, Enterprise Manager will raise an event and open the incident. Once open, you will have available all incident management functionality required to track and manage its resolution.

Figure 2–1 Incident with a Single Event

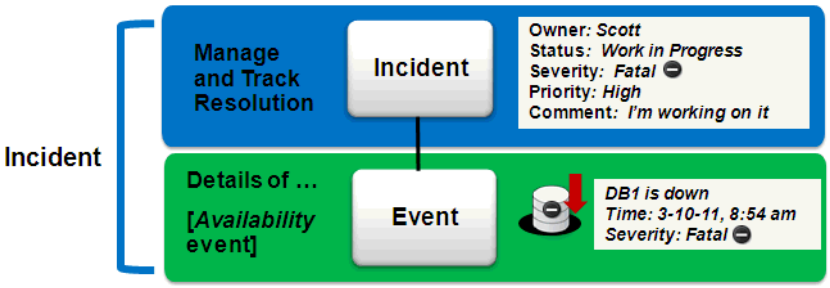

The figure shows how both the incident and event attributes are used to help you manage the incident. From the figure, we see that the database DB1 has gone down and an event of Fatal severity has been raised. An incident is opened and the owner/administrator Scott is currently working to resolve the issue. The incident severity is currently Fatal as the incident inherits the worst severity of all the events within incident. In this case there is only one event associated with the incident so the severity is Fatal.

As an open incident, you can use the Incident Manager console to track its ownership, its resolution status, set the priority, if necessary, add annotations to the incident to share information with others when working in a collaborative environment. In addition, you have direct access to pertinent information from MOS and links to other areas of Enterprise Manager that will help you resolve the database problem.

2.3.2 Incident Composed of Multiple Events

Situations of interest may involve more than a single event. It is the incident's ability to monitor for multiple events that demonstrates the power and flexibility of monitoring and managing via incidents rather than discrete events.

Figure 2–2 Incident with Multiple Events

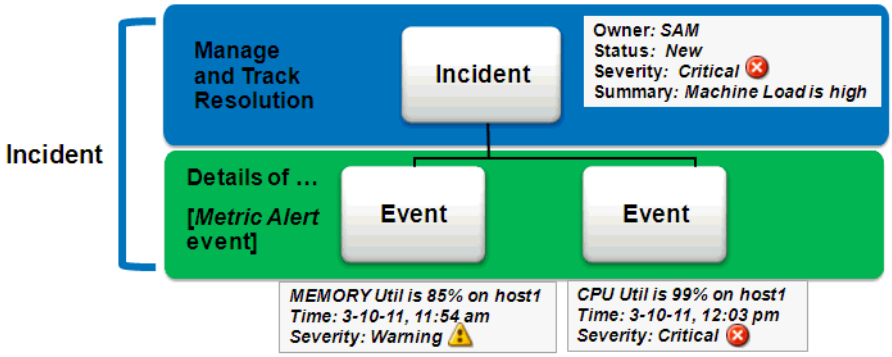

As shown in the figure, the incident is newly opened and has not yet been acknowledged. The incident severity is Critical even though one of the events (Memory Utilization) is only at a Warning severity level. Incidents inherit the worst severity of all the events within incident. The incident summary indicates why this incident should be of interest, in this case, "Machine Load is high". This message is an intuitive indicator for all administrators looking at this incident. By default, the incident summary is pulled from the message of the event, however, this message can be changed by any administrator working on the incident.

Because administrators are interested in overall machine load, administrator Sam has created an incident for these two metric events because they are related--together they represent a host overload situation. An administrator needs to take action because memory is filling up and consumed CPU resource is too high. In its current state, this condition will impact any applications running on the host.

Helpdesk Incident Resolution

If your IT group relies on a helpdesk to resolve this host overload issue, you will want the incident to file a helpdesk ticket in order to have you helpdesk team manage the incident. Here, you can use the ticketing connector to integrate the incident with a helpdesk ticket. It will automatically open a ticket when the incident is created in addition to tracking the ticket ID, and status of the ticket. This provides administrators with a way, from within Enterprise Manager, to view some ticket attributes and not have to accesses third-party helpdesk console. Enterprise Manager also allows you to link out to a Web-based third-part console directly from the ticket so that you can launch the console in context directly from the ticket.

2.4 Incident Rule Sets

With previous versions of Enterprise Manager, you used notification rules to choose the individual targets and conditions for which you want to perform actions and/or receive notifications (send e-mail, page, open a trouble ticket) from Enterprise Manager. For Enterprise Manager 12c, the concept and function of notification rules has been replaced with a two-tier system consisting of Incident Rules and Incident Rule Sets.

- **Incident Rules**: Operate at the lowest level granularity (on discrete events) and performs the same role as notification rules from earlier releases.
- **Incident Rule Set**: Operate at a higher level and consist of one or more rules and manages the automation of business processes relating to incidents and events. An incident rule set instructs Enterprise Manager to take specific actions when incidents, events, or problems occur.

An incident rule set is a set of one or more rules that apply to a common set of objects such as targets (hosts, databases, groups), jobs, or Web applications. The set of objects to which the rule set applies do not have to be of the same type. A rule set allows you to logically combine different rules relating to the common set of objects (such as jobs, targets, applications) into a single manageable unit. Operationally, rules within a rule set are executed in a specified order as are the rule sets themselves. The following figure shows typical rule set structure and now the individual rules are applied to a heterogeneous group of targets.

Figure 2–3 Incident Rule Set Application

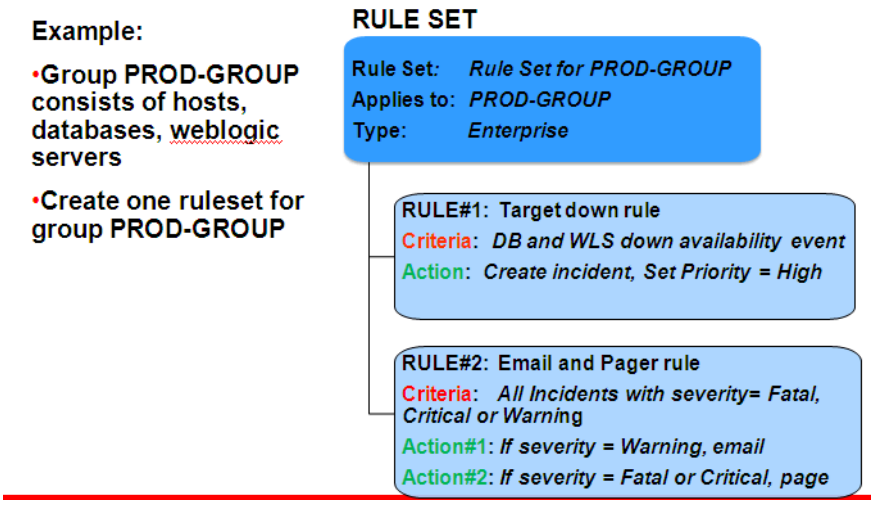

2.4.1 Out-of-Box Rule Sets

Enterprise Manager provides out-of-box rule sets for incident creation, event deletion based on typical scenarios. The following rule sets are immediately available upon installation.

Incident Management Rule Sets for All Targets

- Incident creation Rule for target down.
- Incident creation Rule for target unreachable (for Agents and hosts).
- Incident creation Rule for metric alerts (for critical severity only).
- Out-of-box Incident creation rule for Service Level Agreement Alerts.
- Incident creation rule for compliance score violation
- Incident creation rule for high-availability events.
- Auto-clear Rule for metric alerts older than 7 days.

- Auto clear Rule for job status change terminal status events older than 7 days.
- Clear Application Dependency and Performance (ADP) alerts after without incidents after 7 days.

Event Management Rule Set for Self-Update

- Notification Rule for new updates

> **Note:** Out-of-Box rule sets cannot be deleted. They can only be disabled or updated.

Some examples of the types of actions that a rule can perform are:

- Create an incident based on an event
- Perform notification actions such as generating a helpdesk ticket
- Perform actions to manage incident workflow notification via e-mail/PL/SQL methods/ SNMP traps. For example, if a problem occurs on a database, send e-mail to administrator Joe. If the incident remains unacknowledged for more than 2 days, escalate the incident.

2.4.2 Incident Rule Set Types

There are two types of Incident Rule Sets:

- **Enterprise**: Used to implement all operational practices within your IT organization. All supported actions are available for this type of rule set. However, because this type of rule set can perform all actions, there are restrictions as to who can create an enterprise rule set.

  In order to create or edit an enterprise rule set, an administrator must have been granted the "Create Enterprise Rule Set " privilege on the "Enterprise Rule Set" resource. An enterprise rule set can have multiple authors, however, if the originator of the rule set wants other administrators to edit the rule set, he will need to share access in order to work collaboratively. Incident rule sets are visible to all administrators.
- **Private**: If an administrator does not have the Create Enterprise Rule Set resource privilege and consequently cannot create an enterprise rule set, but wants to be notified about something he is monitoring, he can create a private rule set. The only action a private rule set can perform is to send e-mail to the rule set owner. Any administrator can create a private rule set.

2.4.3 Rules

Rules are instructions within a rule set that automate actions on incoming events or incidents or problems (specialized incidents for Oracle software). Because rules operate on *incoming* incidents/events/problems, if you create a new rule, it will not act retroactively on incidents/events/problems that have already occurred.

Every rule is composed of two parts:

- **Criteria**: The events/incidents/problems on which rule applies.
- **Action**(s): The ordered set of one or more operations on the specified events/ incidents/ problems. Each action can be executed based on additional conditions.

The following table illustrates how rule criteria and actions determines rule application:

Table 2–1 Rule Operation

Rule Criteria	Rule Action	
	Condition	**Actions**
CPU Util(%), Tablespace Used(%) metric alert events of warning or critical severity		Create incident.
Incidents of warning or critical severity	If severity = critical	Notify by page
	If severity =warning	Notify by e-mail
Incidents open for more than 7 days		Set escalation level to 1

From the rule operation example shown in the table, the rule applies to two metric alert events: CPU Utilization and Tablespace Used. Whenever these events reach either Warning or Critical severity threshold levels, an incident is created. Additional conditions have been added to the rule criteria that determines what actions are to be taken. When the incident severity level (the incident severity is inherited from the worst event severity) reaches Warning, Enterprise Manager sends an e-mail to the administrator. If the incident severity level reaches Critical, Enterprise Manager sends a page to the administrator.

2.4.4 Incident Rule Set Guidelines

When creating incident rule sets, adhering to the following guideline will result in efficient use of system resource as well as operational efficiency.

- For rule sets that operate on targets (for example, hosts and databases), use groups to consolidate all targets into a single target for the rule set.
- Consolidate all rules that apply to the group members within the same rule set.
- Leverage the execution order of rules within the rule set

When deciding how to use different rule types within the rule set, adhere to the following rule usage guidelines:

Table 2–2 Incident Rule Usage Guidelines

Rule Usage	Application
Rules on Events…	To create incidents for the alerts/events managed in Enterprise Manager
	To create tickets for incidents managed by helpdesk analysts
	Create Incidents based on event, then create ticket for the incident
	Send events to third party management systems
	To send notifications on events (no incident created)
Rules on Incidents	Automate management of incident workflow (assign owner, set priority, escalation levels..) and send notifications
	Create tickets based on incident conditions. For example, create a ticket if the incident is escalated to level 2.

Table 2–2 (Cont.) Incident Rule Usage Guidelines

Rule Usage	Application
Rules on Problems	Automate management of problem workflow (assign owner, set priority, escalation levels..) and send notifications

Incident Rule Set Example

The following example illustrates many of the implementation guidelines just discussed. All targets have been consolidated into a single group, all rules that apply to group members are part of the same rule set, and the execution order of the rules has been

All rules in the rule set perform three types of actions: incident creation, notification, and escalation.

- Example *Rule Set*:
  - Rule set applies to target: Group target G
  - Rules in the *rule set*:
    1. Rule(s) to create incidents for specified events
    2. Rule(s) that send notifications on the incidents
    3. Rule(s) that escalate incidents based on some condition (e.g. length of time incident is open)

In a more detailed view of the rule set, we can see how the guidelines have been followed.

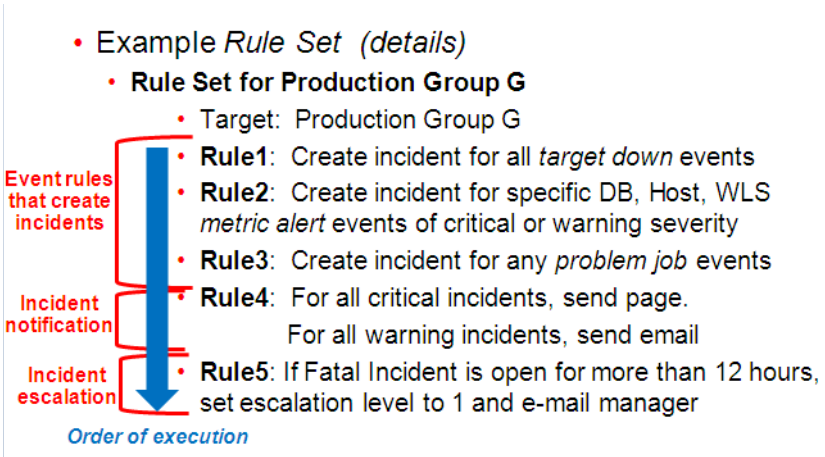

In this detailed view, there are five rules that apply to all group members. The execution sequence of the rules (rule 1 - rule 5) has been leveraged to correspond to the three types of rule actions in the rule set: Rules 1-3

- Rules 1-3: Incident Creation
- Rule 4: Notification
- Rule 5: Escalation

By synchronizing rule execution order with the progression of rule action categories, maximum efficiency is achieved.

2.4.5 Event Prioritization

When working in a large enterprise it is conceivable that when systems are under heavy load, an extraordinarily large number of incidents and events will be generated. All of these need to be processed in a timely and efficient manner in accordance with your business priorities. To have them processed sequentially can result in long waits before incidents can be resolved: High priority events/incidents need to be addressed before those of low priority.

In order to determine which event/incidents are high priority, Enterprise Manager uses a prioritization protocol based on two incident/event attributes: Lifecycle Status of the target and the Incident/Event Type. A target's Lifecycle Status is set when it is added to Enterprise Manager for monitoring. At that time, you determine where in the prioritization hierarchy that target belongs; the highest level being "mission critical" and the lowest being "development."

Target Lifecycle Status

- Mission Critical (highest priority)
- Production
- Sage
- Test
- Development (lowest priority)

Incident/Event Type

- Availability (highest priority)
- All events/incidents (Fatal severity)
- All events/incidents (Warning and Critical severities)
- All events/incidents (Informational) (lowest priority)

2.5 Problems

A problem represents the underlying root cause of the incident requires further analysis beyond the immediate resolution of the incident. For Enterprise Manager 12c, problems focus on the diagnostic incidents and problems diagnostic incidents/problems generated by "Advanced Diagnostic Repository (ADR)". Because the Support Workbench problems and diagnostic incidents are propagated to Incident Manager, you can perform additional tracking such as viewing problems across different databases. A problem represents the root cause of all the Oracle software incidents.

When a problem is raised for Oracle software, Oracle has determined that the only recourse is to open an SR, send support the diagnostic logs, and eventually provide a patch. As an incident, Enterprise Manager makes available all tracking, diagnostic, and reporting functions for problem management. Whenever you view all open incidents and problems, whether you are using the Incident Manager console, or in context of a target/group home page, you can easily determine what issues are actually affecting your database.

The following figure shows the tracking and diagnostic functionality available for problems from the Incident Manger console.

Figure 2–4

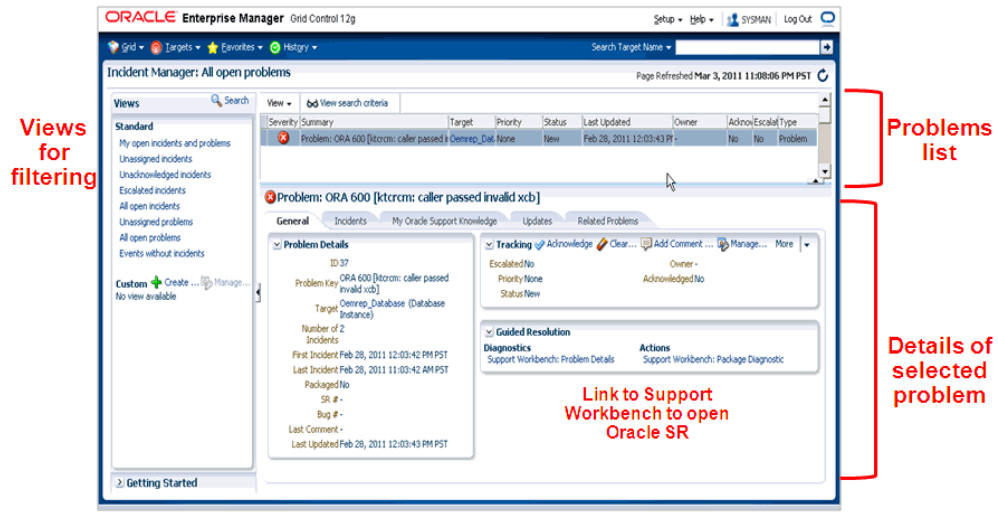

2.6 Incident Manager Console

Incident Manager provides, in one location, the ability to search, view, manage, and resolve incidents and problems impacting your environment. Use Incident Manager to perform the following tasks:

- Respond and Work on an Incident
- Filter Incidents, Problems, and Events by Using Views
- Manage and Automate Incident Workflow
- Suppress Incidents and Problems

Figure 2–5 incident Manager Console

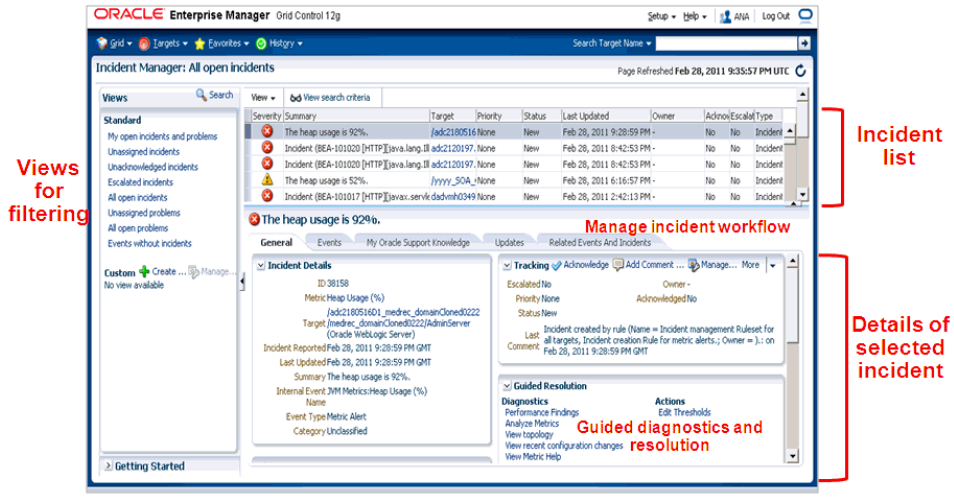

The advantages of using Incident Manager include:

- Ability to manage events, incidents, and problems, for example, the number of unmonitored thresholds exceeds the acceptable percentage.
- Ability to assign incidents to specific personnel; prioritize, escalate, and track incidents through various states of resolution
- From the person working the incidents and problems:

  See what incidents and problems are assigned to me, acknowledge that I am working the incidents and problems, and provide information to the user community regarding the progress of the resolution.
- Integration with My Oracle Support
- In-context diagnostics and resolution links

2.7 Moving from Enterprise Manager 10/11g to 12c

Enterprise Manager 12c incident management/monitoring functionality leverages your existing pre-12c monitoring setup out-of-box. Migration is seamless and transparent. For example, if your Enterprise Manager 10/11g monitoring system sends you e-mails based on specific monitoring conditions, you will continue to receive those e-mails without interruption. To take advantage of 12c features, however, you may need to perform additional migration tasks.

Important: Alerts that were generated pre-12c will still be available.

Incident Rules

When you migrate to Enterprise Manger 12c, all of your existing notification rules are automatically converted to incident rules. Technically, they are converted to event rules first with incidents automatically being created for each event rule.

In general, event rules allow you to define which events should become incidents. However, they also allow you to take advantage of the Enterprise Manager's increased monitoring flexibility: While an incident is open, if you want to monitor the status of individual events within the incident, you can utilize individual event rules as a way to obtain those notifications.

For more information on incident rule migration, see the following documents:

- Appendix A, " Overview of Notification in Enterprise Manager Cloud Control" section "Migrating Notification Rules to Incident Rulesets" in the upgrade guide
- Chapter 29 "Updating Incident Rules" in the Enterprise Manger Upgrade guide.

Privilege Requirements

The 'Rule Set' resource privilege is now required in order to edit/create enterprise rule sets and rules contined within. The exception to this is migrated notification rules. When pre-12c notification rules are migrated to event rules, the original notification rule owners will still be able to edit their own rules without having been granted the Create Enterprise Rule Set resource privilege. However, they must be granted the 'Rule Set' resource privilege if they wish to create new rules. Enterprise Manager Super Administrators, by default, can edit and create rule sets.

- Privileges on events are calculated based on the privilege on the underlying source objects. For example, the user will have VIEW privilege on an event if he can view the target for the event.
- Privileges on an incident are calculated based on the privileges on participating events.
- Similarly, problem privileges are calculated based on privileges on underlying incidents.

2.8 Before Working with Incidents

Before using Incident Manager, ensure all relevant Enterprise Manager administrator accounts have been granted the appropriate privileges to manage incidents. Also, ensure that the notification system is properly configured to allow automated notification for incidents.

Granting User Privileges for Events, Incidents and Problems

Users are granted privileges for events, incidents, and problems in the following situations:

For events, two privileges are defined:

- The View Event privilege allows you to view an event and add comments to the event.
- The Manage Event privilege allows you to take update actions on an event such as closing a manually-clear event, creating an incident for an event, and creating a ticket for an event. You can associate an event with an incident.

For incidents, two privileges are defined:

- The View Incident privilege allows you to view an incident, and add comments to the incident.
- The Manage Incident privilege allows you to take update actions on an incident. The update actions supported for an incident includes incident assignment and prioritization, resolution management, manually closing manually-clear events, manually create a problem for an incident, and create a ticket for an incident.

For problems, two privileges are defined:

- The View Problem privilege allows you to view a problem and add comments to the problem.
- The Manage Problem privilege allows you to take update actions on the problem. The update actions supported for a problem include problem assignment and prioritization, resolution management, manually closing the problem, and update customer-defined attribute values. It also includes ability to create a Service Request and gather diagnostics using Support Workbench.

2.9 Working with Incidents

You can perform the following tasks using Incident Manager:

- Setting Up Incident Views
- Using Views to Filter Incidents, Problems, and Events
- Responding and Working on a Simple Incident
- Suppressing Incidents and Problems

2.9.1 Setting Up Incident Views

Incident views allow you to save commonly used incident search criteria for repeated use. You can set up a filter in Incident Manager to view incidents for targets managed by, for example, targets managed by a specific group of administrators.

By specifying preferences to view the following for each of the incidents in the list: incident severity, incident message, acknowledgement flag, date the incident triggered, administrator assigned to it, resolution status, priority, escalated flag, ID, and category, you can filter extraneous incidents. Once the view preference is saved, Enterprise Manager will display only the list of matching incidents.

You can then search the incidents for only the ones with specific attributes, such as priority P1. While reviewing the incidents, you can specify one-click access to this list so that it can be easily accessed for daily triaging activity. Accordingly, you can save the search criteria as a filter named "All P1 incidents for my targets". The filter becomes available in the UI for immediate use. The filter will show up anytime you log in to access the specific incidents quickly.

Perform the following steps:

1. Navigate to the Incident Manager page.

   From the **Enterprise** menu on the Enterprise Manager home page, select **Monitoring**, then select **Incident Manager**.

2. In the **Views** region located on the left, click **Search**.

   a. In the **Search** region, search for Incidents using the **Type** list and select **Incidents**.

   b. In the Criteria region, choose all the criteria that are appropriate. To add fields to the criteria, click **Add Fields...** and select the appropriate fields.

   c. After you have provided the appropriate criteria, click **Get Results**.

   d. To view all the columns associated with this table, in the **View** menu, select **Columns**, then select **Show All**.

      Validate that the list of incidents match what you are looking for. If not, change the search criteria as needed.

   e. Click the **Create View...** button.

2.9.2 Using Views to Filter Incidents, Problems, and Events

A view is a set of search criteria for filtering incidents and problems in the system. You can define views to help you gain quick access to the incidents and problems on which you need to focus. For example, you may define a view to display all the incidents associated with the production databases that you own.

2.9.3 Responding and Working on a Simple Incident

Before you begin working on resolving an incident, ensure your Enterprise Manager account has been granted the appropriate privileges to manage incidents from your managed system.

- Privileges on events are calculated based on the privilege on the underlying source objects. For example, the user will have VIEW privilege on an event if he can view the target for the event.
- Privileges on an incident are calculated based on the privileges on participating events.

- Similarly, problem privileges are calculated based on privileges on underlying incidents.

Perform the following steps:

1. Navigate to Incident Manager.

   From the **Enterprise** menu on the Enterprise Manager home page, select **Monitoring**, then select **Incident Manager**.

2. To view incidents assigned to you, in the table look at the Owner column. If the Owner column is not displayed, in the **View** menu select **Columns**, then select Owner.

   Work on the incident with the highest priority. Be aware that as you are working on an individual incident, new incidents might be coming in. Update the list of incidents by clicking the Refresh icon.

3. To work on an incident, click the incident. In the General section, click **Manage** and change the fields as appropriate. For example, set the status to **Work in Progress** and in the **Owner** field, type your name.

4. If the solution for the incident is unknown, use one or all of the following methods made available in the Incident page:

   - Use the **Guided Resolution** region and access any recommendations, diagnostic and resolution links available.
   - Check My Oracle Support Knowledge base for known solutions for the incident.
   - Study related incidents available through the Related Events and Incidents tab.

5. Once the solution is known and can be resolved right away, resolve the incident by using tools provided by the system, if possible.

6. In most cases, once the underlying cause has been fixed, the incident is cleared in the next evaluation cycle. However, in cases like "log based" incidents, clear the event.

2.9.4 Suppressing Incidents and Problems

There are times when it is convenient to hide an incident or problem from the list in the All Open Incidents page or the All Open Problems page. For example, you may want to suppress an incident is while the incident is being actively worked on and you do not need to be notified.

To suppress an incident or problem:

1. Navigate to the Incident Manager page.

   From the **Enterprise** menu on the Enterprise Manager home page, select **Monitoring**, then select **Incident Manager**.

2. Select either the All Open Incidents view or the All Open Problems view. Choose the appropriate incident or problem. Click the **General** tab.

3. In the resulting Details region, click **More**, then select **Suppress**.

4. On the resulting Suppress pop-up, choose the appropriate suppression type. Add a comment if desired. Click **OK**.

2.9.5 Searching My Oracle Support Knowledge

To access My Oracle Support Knowledge base entries from within Incident Manager, perform the following steps:

1. Navigate to the Incident Manager page.

   From the **Enterprise** menu on the Enterprise Manager home page, select **Monitoring**, then select **Incident Manager**.

2. Select one of the standard views. Choose the appropriate incident or problem in the View table.

3. In the resulting Details region, click My Oracle Support Knowledge. Sign in to My Oracle Support.

4. On the My Oracle Support page, click the **Knowledge** tab to browse the knowledge base and knowledge alerts.

2.9.6 Open Service Request

There are times when you may need assistance from Oracle Support to resolve a problem. To submit a service request (SR), perform the following steps:

1. Navigate to the Incident Manager page.

   From the **Enterprise** menu on the Enterprise Manager home page, select **Monitoring**, then select **Incident Manager**.

2. Select one of the standard views. Choose the appropriate problem from table.

3. In the resulting Details region, click **My Oracle Support Knowledge**. Sign in to My Oracle Support if you are not already signed in.

4. On the My Oracle Support page, click the **Service Requests** tab.

5. Click **Create SR** button. Click **Help** to learn how to create a new SR.

2.10 Incident Manager - Advanced Tasks

You can perform the following advanced tasks using Incident Manager:

- Creating an Incident Manually
- Managing Workload Distribution of Incidents
- Managing and Automating Incident Workflow

2.10.1 Creating an Incident Manually

To create an incident manually, perform the following:

1. Create incident in context of an event.

2. Enter details and save the incident.

3. Set yourself as owner of the incident and update status to Work in Progress.

Example Scenario

As per the operations policy, the DBA manager has setup rules to create incidents for all critical issues for his databases. The remainder of the issues are triaged at the event level by one of the DBAs.

One of the DBA receives e-mail for an "SQL Response" event (not associated with an incident) on the production database. He accesses the details of the event by clicking on the link in the e-mail. He reviews the details of the event. This is an issue that needs to be tracked and resolved, so he opens an incident to track the resolution of the issue. He marks the status of the incident as "Work in progress".

2.10.2 Managing Workload Distribution of Incidents

Incident Manager enables you to manage incidents and problems to be addressed by your team

Perform the following tasks:

1. Navigate to Incident Manager.

   From the **Enterprise** menu on the Enterprise Manager home page, select **Monitoring**, then select **Incident Manager**.

2. Use the standard or custom views to identify the incidents for which your team is responsible. You may want to focus on unassigned and unacknowledged incidents and problems.

3. Review the list of incidents. This includes: determining person assigned to the incident, checking its status, progress made, and actions taken by the incident owner.

4. Add comments, change priority, re-assign the incident as needed by clicking on the Manage button in the Incident Details region.

Example Scenario

The DBA manager uses Incident Manager to view all the incidents owned by his team. He ensures all of them are correctly assigned; if not, he re-assigns and prioritizes them appropriately. He monitors the escalated events for their status and progress, adds comments as needed for the owner of the incident. In the console, he can view how long each of the incidents has been open. He also reviews the list of unassigned incidents and assigns them appropriately.

2.10.3 Managing and Automating Incident Workflow

Data centers follow operational practices that enable them to manage events and incidents by business priority and in a collaborative manner. Enterprise Manager provides the following features to enable this management and automation:

- Sending notifications to the appropriate administrators
- Assigning initial ownership of an incident and perhaps transferring ownership based on shift assignments or expertise
- Tracking its resolution status
- Assigning priorities based on the component affected and nature of the incident
- Escalating incidents in order to meet service level agreements (SLA)
- Accessing My Oracle Support knowledge articles
- Opening Oracle Service Requests to request assistance with problems with Oracle software
- Generating management and operational reports to track the status of incidents

You can manage an incident by doing the following:

1. In the **All Open Incidents** view, click the incident.
2. In the resulting Details page, click the **General** tab, then click **Manage**.

   You can then adjust the priority, escalate the incident, and assign it to a specific engineer.

2.11 Set Up Tasks to Perform Before Using Incident Rules

Before you use incident rules, ensure the following prerequisites have been set up:

- User's Enterprise Manager account has notification preferences (e-mail and schedule).
- Automated notification for incidents.
- Subscription to incident rule.
- User's Enterprise Manager account has been granted the appropriate privileges to manage incidents from his managed system.
- If you decide to use connectors, tickets, or advanced notifications, you need to configure them before using them in the actions page.

To perform these tasks, click **Setup** on the Enterprise Manager home page, select **Security**, then select **Administrators** to access the Administrators page.

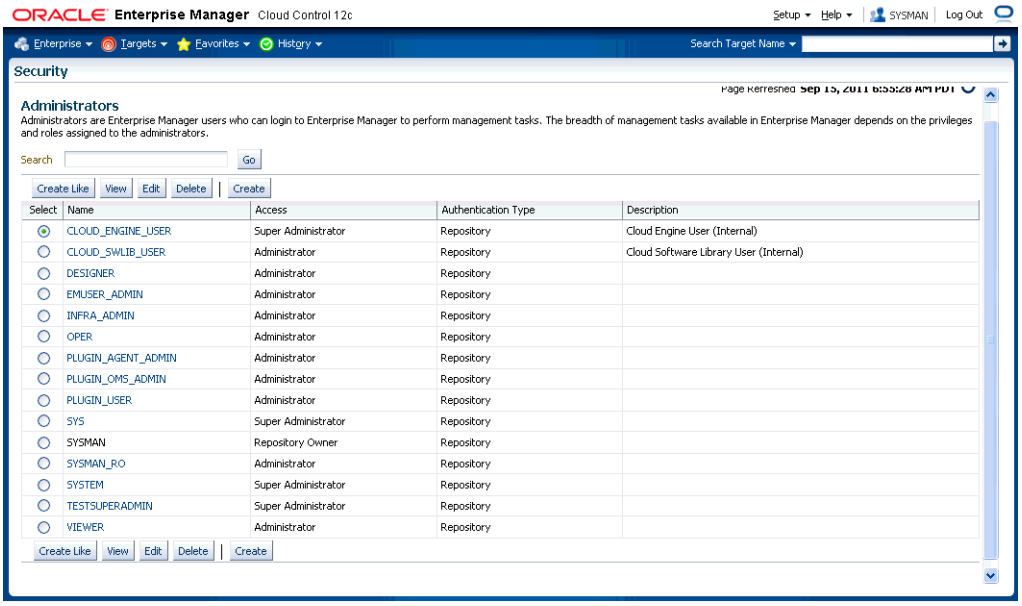

Privileges Required for Enterprise Rule Sets

As the owner of the rule set, an administrator can perform the following:

- Update or delete the rule set, and add, modify, or delete the rules in the rule set.
- Assign co-authors of the rule set. Co-authors can edit the rule set the same as the author. However, they cannot delete rule sets nor can they add additional co-authors.

- When a rule action is to update an event, incident, or problem (for example, change priority or clear an event), the action succeeds only if the owner has the privilege to take that action on the respective event, incident, or problem.
- Additionally, user must be granted privilege to create an enterprise rule set.

If an incident or problem rule has an update action (for example, change priority), it will take the action only if the owner of the respective rule set has manage privilege on the matching incident or problem.

To acquire privileges, click **Setup** on the Enterprise Manager home page, select **Security**, then select **Administrators** to access the Administrators page. Select an administrator from the list and click **Edit** to access the Administrator properties wizard as shown in the following graphic.

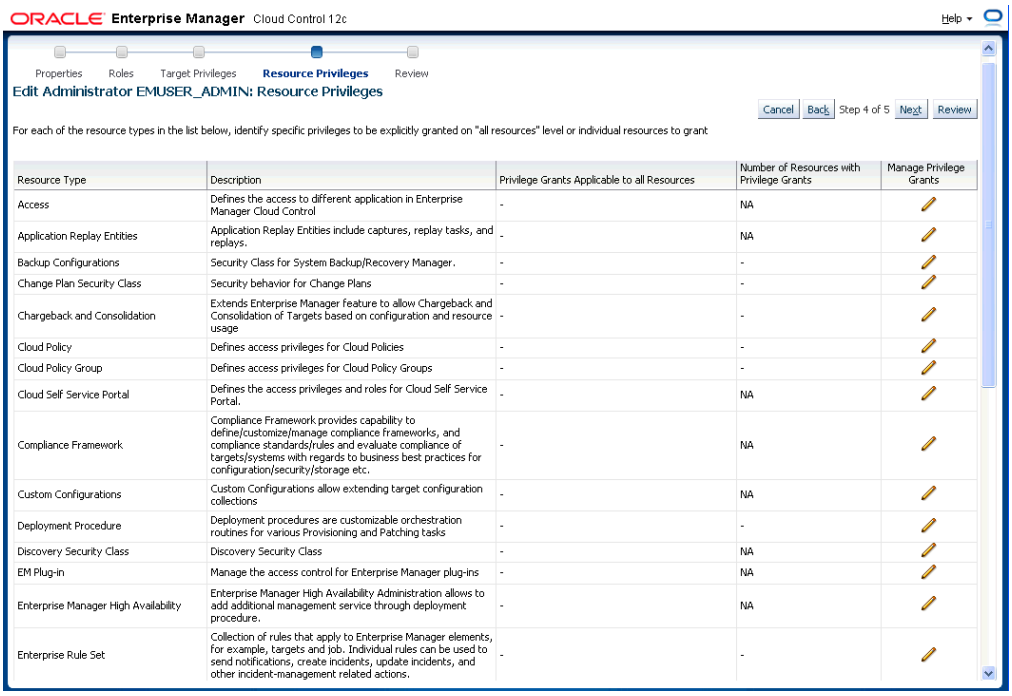

2.12 Working with Incident Rules

You can perform the following tasks using Incident Rules:

- Set Up the Monitoring Environment by Defining Incident Rules
- Create an Incident Rule
- Set Up a Rule to Create Incident in Response to an Event
- Create a Rule to Receive Notification Regarding Informational Events
- Set Up a Rule to Escalate a Problem

2.12.1 Setting Up the Monitoring Environment by Defining Incident Rules

One way to set up your monitoring environment is by defining incident rules. You can set up rules to:

1. Create incident in response to an event

2. Send notifications to different users
3. Manage escalation of incidents and problems
4. Create ticket for incidents
5. Notify different administrators for different classes of events
6. Create notification subscription to existing Enterprise Rules

2.12.2 Creating an Incident Rule

To create an incident rule, perform the following steps:

1. Navigate to the Incident Rules page.

   From the **Setup** menu located at the top-right of the Enterprise Manager home page, select **Incidents**, then select **Incident Rules**.

2. On the Incident Rules - All Enterprise Rules page, edit the existing rule set (highlight the rule set and click **Edit...**) or create a new rule set. Rules are created in the context of a rule set!
3. In the Rules tab of the Edit Rule Set page, click **Create...** and select the type of rule to create (Event, Incident, Problem) on the Select Type of Rule to Create page. Click **Continue**.
4. In the Create New Rule wizard, provide the required information. Click **Help** for information regarding the wizard pages.
5. Once you have finished defining the rule, click **Continue** to add the rule to the rule set. Click **Save** to save the changes made to the rule set.

2.12.3 Creating a Rule to Create an Incident

To create a rule that creates an incident, perform the following steps:

1. Navigate to the Incident Rules page.

   From the **Setup** menu located at the top-right of the Enterprise Manager home page, select **Incidents**, then select **Incident Rules**.

2. Determine whether there is an existing rule set that contains a rule that manages the event. In the Incident Rules page, use the Search option to find the events for the target and the associated rule set.

   Note: In the case where there is no existing rule set, create a rule set by clicking **Create Rule Set...** You then create the rule as part of creating the rule set.

3. Select the rule set that will contain the new rule. Click **Edit...** In the Rules tab of the Edit Rule Set page,

   1. Click **Create ...**
   2. Select "Incoming events and updates to events"
   3. Click **Continue**.

   select **Actions**, then select **Add rule for event...** (or click **Create...**, select **Event Rule,** and click **Continue**).

   Provide the rule details using the Create New Rule wizard.

   a. Select the Event Type the rule will apply to, for example, Metric Alert. (Metric Alert is available for rule sets of the type Targets.) You can then specify metric alerts by selecting **Specific Metrics**. The table for selecting metric alerts

displays. Click the **+Add** button to launch the metric selector. On the Select Specific Metric Alert page, select the target type, for example, Database Instance. A list of relevant metrics display. Select the ones in which you are interested. Click **OK**.

You also have the option to select the severity and corrective action status.

b. Once you have provided the initial information, click **Next**. Click **+Add** to add the actions to occur when the event is triggered. One of the actions is to **Create Incident**.

As part of creating an incident, you can assign the incident to a particular user, set the priority, and create a ticket. Once you have added all the conditional actions, click **Continue**.

c. After you have provided all the information on the Add Actions page, click **Next** to specify the name and description for the rule. Once on the Review page, verify that all the information is correct. Click **Back** to make corrections; click **Continue** to return to the Edit (Create) Rule Set page.

d. Click **Save** to ensure that the changes to the rule set and rules are saved to the database.

4. Test the rule by generating a metric alert event on the metrics chosen in the previous steps.

2.12.4 Creating a Rule to Manage Escalation of Incidents

Before you set up a rule to manage escalations, ensure the following prerequisite task has been performed:

- DBA has setup appropriate thresholds for the metric so that critical metric alert is generated as expected.

Perform the following steps:

1. Navigate to the Incident Rules page.

   From the **Setup** menu located at the top-right of the Enterprise Manager home page, select **Incidents**, then select **Incident Rules**.

2. Determine whether there is an existing rule set that contains a rule that manages the incident. In the Incident Rules page, use the Search option to find the incident and the associated rule set.

   Note: In the case where there is no existing rule set, create a rule set by clicking **Create Rule Set...** You then create the rule as part of creating the rule set.

3. Select the rule set that will contain the new rule. Click **Edit...** In the Rules tab of the Edit Rule Set page, and then:

   1. Click **Create ...**
   2. Select "Incoming events and updates to events"
   3. Click **Continue**.

4. For demonstration purposes, the escalation is in regards to a production database.

   As per the organization's policy, the DBA manager is notified for escalation level 1 incidents. Similarly, the DBA director and operations VP are paged for incidents escalated to levels 2 and 3 respectively.

   Provide the rule details using the Create New Rule wizard.

a. Select Specific Incidents where the Target Attribute has a value of Database Instance.

b. In the Conditions for Actions region located on the Add Actions page, select **Execute the actions on the conditions specified**.

   Select **How long the incident is open and in a particular state (select time and optional expressions)**

   Select the Time to be 30 minutes and the Attribute Name to be **Escalation Level** with a value of 1. Click **Continue**.

c. In the Basic Notification region, type the name of the administrator to be notified by e-mail or page.

d. Repeat steps b and c to notify the DBA director when escalation level is 2 and the Operations VP when the escalation level is 3.

e. Review the summary and save the rule.

f. Click **Next** until you get to the Summary screen. Verify that the information is correct and click **Save**.

5. Review the sequence of existing enterprise rules and position the newly created rule in the sequence.

   On the Edit Rule Set page, select **Actions**, then select **Reorder Rules**. Click **Save** to save the change to the sequence.

Example Scenario

In many companies, the operations team handles incidents at different escalation levels. An incident is escalated to a higher level based on how long the incident remains unresolved.

To facilitate this process, the administration manager creates a rule to escalate unresolved incidents based on their age:

- To level 1 if the incident is open for 30 minutes
- To level 2 if the incident is open for 1 hour
- To level 3 if the incident is open for 90 minutes

As per the organization's policy, the DBA manager is notified for escalation level 1. Similarly, the DBA director and operations VP are paged for incidents escalated to levels "2" and "3" respectively.

Accordingly, the administration manager inputs the above logic and the respective Enterprise Manager administrator IDs in a separate rule to achieve the above notification requirement. Enterprise Manager administrator IDs represents the respective users with required target privileges and notification preferences (that is, e-mail addresses and schedule).

2.12.5 Creating a Rule to Escalate a Problem

Before you create a rule to escalate a problem, ensure the following prerequisites are met:

- Incident rule has been setup to generate appropriate incidents as to when a critical issue occurs.
- Administrator attaches incidents to the problem if they feel the underlying issue is the one being tracked by the problem.

Perform the following steps:

1. Navigate to the Incident Rules page.

   From the **Setup** menu located at the top-right of the Enterprise Manager home page, select **Incidents**, then select **Incident Rules**.

2. On the Incident Rules - All Enterprise Rules page, either create a new rule set (click **Create Rule Set...**) or edit an existing rule set (highlight the rule set and click **Edit...**). (Rules are created in the context of a rule set.)

3. In the Rules section of the Edit Rule Set page, select **Create...** to create an enterprise rule to automate actions on the problem. Select Problem Rule on the **Select Type of Rule to Create** page. Click **Continue**.

4. On the Create New Rule page, select **Specific problems** and add the following criteria:

   The Attribute Name is **Incident Count**, the Operator is **Greater than or equals** and the Values is **20**. Click **Next**.

5. In the Conditions for Actions region on the Add Actions page select **Always execute the action**. As the actions to take when the rule matches the condition:

   - In the Notifications region, send e-mail to the owner of the problem and to the Operations Manager.
   - In the Update Problem region, select **Escalate** and choose **1** as the appropriate level.

   Click **Continue**.

6. Review the rules summary. Make corrections as needed. Click **Save**.

Example Scenario

In an organization, whenever an unresolved problem has more than 20 occurrences of associated incidents, the problem should be escalated to prioritize the resolution. Accordingly, a problem rule is created to observe the count of incidents attached to the problem and escalate the problem when the count reaches the limit.

The problem owner and the Operations manager are notified by way of e-mail.

2.12.6 Setting Up Automated Notification for Private Rule

A DBA has setup a backup job on the database that he is administering. As part of the job, the DBA has subscribed to e-mail notification for "completed" job status.

Before you create the rule, ensure the following prerequisites are met:

- You have the privilege to create jobs.
- You have created a database backup job.

Perform the following steps:

1. Navigate to the Incident Rules page.

   From the **Setup** menu located at the top-right of the Enterprise Manager home page, select **Incidents**, then select **Incident Rules**.

2. On the Incident Rules - All Enterprise Rules page, either edit an existing rule set (highlight the rule set and click **Edit...**) or create a new rule set.

   Note: The rule set must be defined as a Private rule set.

3. In the Rules tab of the Edit Rule Set page, select **Create...** and select **Event Rule**. Click **Continue**.
4. On the Select Events page, select **Job Status Change** as the Event Type. Select the job in which you are interested either by selecting a specific job or selecting a job by providing a pattern, for example, Backup Management.

   Add additional criteria by adding an attribute: Target Type as Database Instance.
5. Add conditional actions: Event matches the following criteria (Severity is Informational) and E-mail Me for notifications.
6. Review the rules summary. Make corrections as needed. Click **Save**.
7. Create a database backup job and subscribe for e-mail notification when the job completes.

When the job completes, Enterprise Manager publishes the informational event for "Job Complete" state of the job. The newly created rule matches the rule and e-mail is sent out to the DBA.

The DBA receives the e-mail and clicks the link to access the details section in Enterprise Manager console for the event.

2.12.7 Creating a Rule to Receive Notification Regarding Events

To create a rule to receive notification on events, perform the following steps:

1. Navigate to Incident Rules-All Enterprise Rules page.

   From the Enterprise Manager home page, select **Setup** located at the top-right of the page, select **Incidents**, then select **Incident Rules**.
2. Edit an existing enterprise rule set.

   Highlight the rule set and click **Edit...**
3. In the Rules section of the Edit Rules Set page:
   1. Click **Create ...**
   2. Select "Incoming events and updates to events"
   3. Click **Continue**.

   Select an event type, for example, Target Availability. Add Specific Target Availability Events, for example, Host and select the specific availability events in which you are interested. Additional Criteria can include Severity of Critical. Click **Next**.
4. To be notified of the event, define additional actions using the Add Actions page. For the conditions under which the event occurs, select Always Execute the Actions. For the notifications, provide information in the Basic Notifications region.
5. When you receive the e-mail regarding the event, click on the link to access the details section in Enterprise Manager console for the event.

2.12.8 Setting Up Escalations

In an organization, there are times where incidents need to be escalated. To escalate an incident to another person, perform the following steps:

1. Navigate to the Incident Rules page.

Click **Setup** located at the top-right of the page. Select **Incidents**, then select **Incident Rules**.

2. On the Incident Rules - All Enterprise Rules page, highlight the rule set of which the rule is a member. Click **Edit**.
3. On the **Edit Rule Set** page, scroll down to the Rules section and highlight the rule you want to edit. Click **Edit...**
4. In the Edit Rule Set wizard, select the incidents to which you want this escalation to apply by selecting **Specific Incidents** and selecting **Escalation level** in the Attribute Name. Provide an escalation value. Click **Next**.
5. In the Add Actions page, click **+Add** to add an action.

   In the Update Incident section, check **Escalate to** and choose the option to which to escalate the incident. For example, choose 1 in the associated list. Click **Continue**.
6. Click **Next** to specify the name and description. Click **Next** again to access the Review page.
7. Review the rules summary. Make corrections as needed. Click **Next**, then click **Save**.

2.13 Incident Rules - Advanced Tasks

You can perform the following advanced tasks using Incident Rules:

- Setting Up a Rule to Send Different Notifications for Different Severity States of an Event
- Creating a Rule to Create a Ticket for Incidents
- Creating a Rule to Notify Different Administrators Based on the Event Type
- Creating Notification Subscription to Existing Enterprise Rules
- Manually Ensuring That There Are No Events That Should Be Incidents

2.13.1 Setting Up a Rule to Send Different Notifications for Different Severity States of an Event

Before you perform this task, ensure the following prerequisite is met:

- DBA has setup appropriate thresholds for the metric so that a critical metric alert is generated as expected.

Perform the following tasks:

1. Navigate to the Incident Rules page.

   From the **Setup** menu located at the top-right of the Enterprise Manager home page, select **Incidents**, then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, highlight a rule set and click **Edit...**. (Rules are created in the context of a rule set. If there is no existing rule set to manage the newly added target, create a rule set.)
3. In the Edit Rule Set page, locate the Rules section. From the **Actions** menu, select **Add Rule for Event**.
4. Provide the rule details as follows:

   a. For Type, select **Metric Alerts** as the Type.

b. In the Additional Criteria section, select **Severity** as the Attribute Name and **Critical** as the Value. Click **Next**.

c. On the Add Actions page, click **+Add**. In the Notifications section, provide the contact information for the DBA to be paged. Click **Continue** until you reach the Edit Rule Set page.

d. Highlight the rule again. For Event Type Specific Criteria, select **Metric Alert** as the Type.

e. In the Additional Criteria section, select **Severity** as the Attribute Name and **Warning** as the Value. Click **Next**.

f. On the Add Actions page, click **+Add**. In the Notifications section, provide the contact information for the DBA to be e-mailed. Click **Continue** until you reach the Edit Rule Set page.

g. Click **Next** until you get to the Summary screen. Verify that the information is correct and click **Save**.

Example Scenario

The Administration Manager sets up a rule to page the specific DBA when a critical metric alert event occurs for a database in a production database group and to e-mail the DBA when a warning metric alert event occurs for the same targets. This task occurs when a new group of databases is deployed and DBAs request to create appropriate rules to manage such databases.

2.13.2 Creating a Rule to Create a Ticket for Incidents

According to the operations policy of an organization, all critical incidents from a production database should be tracked by way of Remedy tickets. An incident rule is created to invoke the Remedy ticket connector to generate a ticket when a critical incident occurs for the database. When such an incident occurs, the ticket is generated by the incident rule, the incident is associated with the ticket, and the operation is logged for future reference to the updates of the incident. While viewing the details of the incident, the DBA can view the ticket ID and, using the attached URL link, access the Remedy to get the details about the ticket.

Before you perform this task, ensure the following prerequisites are met:

- Monitoring support has been set up.
- Remedy ticket support has been setup.

Perform the following steps:

1. Navigate to the Incident Rules page.

   From the **Setup** menu located at the top-right of the Enterprise Manager home page, select **Incidents**, then select **Incident Rules**.

2. On the Incident Rules - All Enterprise Rules page, highlight a rule set and click **Edit...**. (Rules are created in the context of a rule set. If there is no existing rule set , create a rule set.)

3. Select the appropriate rule that covers the incident conditions for which tickets should be generated and click **Edit...**.

   a. Specify that a ticket should be generated for incidents covered by the rule.

   b. Specify the ticket template to be used.

4. Repeat step 3 until all appropriate rules have been edited.

5. Click **Save**.

2.13.3 Creating a Rule to Notify Different Administrators Based on the Event Type

As per operations policy for production databases, the alerts that relate to application issues should go to the application DBAs and the alerts that relate to system parameters should go to the system DBAs. Accordingly, the respective incidents will be assigned to the appropriate DBAs and they should be notified by way of e-mail.

Before you set up rules, ensure the following prerequisites are met:

- DBA has setup appropriate thresholds for the metric so that critical metric alert is generated as expected.
- Incident rule has been setup to create incident for all such events.
- Respective notification setup is complete, for example, global SMTP gateway, e-mail address, and schedule for individual DBAs.

Perform the following steps:

1. Navigate to the Incident Rules page.

   From the **Setup** menu located at the top-right of the Enterprise Manager home page, select **Incidents**, then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, highlight a rule set and click **Edit...**. (Rules are created in the context of a rule set. If there is no existing rule set , create a rule set.)
3. Search the list of enterprise rules matching the events from the production database.
4. Select the rule which creates the incidents for the metric alert events for the database. Click **Edit**.
5. Enter the specific metrics that identify *application* issues, as condition to match the incidents.
6. Enter the specific metrics, which identifies issues with *system parameters*, as condition to match the incidents.
7. Type a summary message, for example: Assign the incident to Cindy (Enterprise Manager administrator handling the system parameter issues). For the action, select to e-mail her.
8. Review the rules summary. Make corrections as needed. Click **Save**.

2.13.4 Creating Notification Subscription to Existing Enterprise Rules

A DBA is aware of an enterprise rule that will escalate incidents managed by him when not resolved in 2 hours. The DBA wants to be notified when the rule escalates the Incident. The DBA can subscribe to the Rule, which escalates the Incident and will be notified whenever the rule escalates the Incident.

Before you set up a notification subscription, ensure the following prerequisites are met:

- There exists an open incident for a database.
- There exists a rule that escalates High Priority Incidents for databases that have not been resolved in hours.

Perform the following steps:

1. Navigate to the Incident Rules page.

   From the **Setup** menu located at the top-right of the Enterprise Manager home page, select **Incidents**, then select **Incident Rules**.

2. On the Incident Rules - All Enterprise Rules page, edit the existing rule set Highlight the rule set and click **Edit...** (Rules are created in the context of a rule set. If there is no existing rule set , create a rule set.)

3. In the Rules section of the Edit Rule Set page, highlight the rule you want to change and click **Edit...**.

4. Edit the rule associated with e-mail notification. Subscribe yourself to receive e-mail notifications.

5. Review the rules summary. Make corrections as needed. Click **Save**.

As a result of the edit to the enterprise rule, when an incident stays unresolved for 2 hours, the rule marks it to escalation level 1. An e-mail is sent out to the DBA notifying him about the escalation of the incident.

The DBA receives the e-mail notification and views the details pertaining to the database down Incident. The DBA clicks on the e-mail link that takes him to the Incident details page after successful login to Enterprise Manager.

2.13.5 Manually Ensuring That There Are No Events That Should Be Incidents

Perform the following steps:

1. Navigate to the Incident Rules page.

   From the **Setup** menu located at the top-right of the Enterprise Manager home page, select **Incidents**, then select **Incident Rules**.

2. On the Incident Rules - All Enterprise Rules page, edit the existing rule set Highlight the rule set and click **Edit...** (Rules are created in the context of a rule set. If there is no existing rule set , create a rule set.)

3. In the Rules section of the Edit Rule Set page, choose **Events** in the Create Rule For list. Click **Go**.

4. Provide the following parameters:

   - Specific event type: select Metric Alerts.
   - Target type for the event: select Database Instance.
   - Further filtering of events based on event lifecyle conditions like severity, indicate that the severity should be **Critical**.
   - Action that should take place when such an event occurs, enter that a specific DBA should be paged.
   - When prompted if any other actions should be taken, answer **yes**.
   - When prompted for specific lifecycle condition, enter that the severity is **Warning**.
   - When prompted for what action to take when this event occurs, enter that a specific DBA should be e-mailed.
   - When asked whether another action should take place, enter that there are no other actions for this rule.

Example Scenario

During the initial phase of Enterprise Manager uptake, every day the DBA manager reviews the un-acknowledged events on the databases his team is responsible for and filters them to view only the ones which are not tracked by ticket or incident. He browses such events to ensure that none of them requires incidents to track the issue. If he feels that one such event requires an incident to track the issue, he creates an incident directly for this event. He also creates incident rules to create an incident when similar events occur in future.

If there are certain events he triages and feels nobody else has to follow-up on the event, he marks it as acknowledged. Enterprise Manager filters out events from the Incident Manager that have been acknowledged.

3 Notifications

The notification system allows you to notify Enterprise Manager administrators when specific incidents, events, or problems arise.

> **Note:** This chapter assumes that you are familiar with incident management. For information about monitoring and managing your IT infrastructure via incident management, see Chapter 2, "Incident Management".

As an integral part of the management framework, notifications can also perform actions such as executing operating system commands (including scripts) and PL/SQL procedures when specific incidents, events, or problems occur. This capability allows you to automate IT practices. For example, if an incident (such as monitoring of the operational (up/down) status of a database) arises, you may want the notification system to automatically open an in-house trouble-ticket using an OS script so that the appropriate IT staff can respond in a timely manner.

By using Simple Network Management Protocol (SNMP) traps, the Enterprise Manager notification system also allows you to send traps to SNMP-enabled third-party applications such as HP OpenView for events published in Enterprise Manager. Some administrators may want to send third-party applications a notification when a certain metric has exceeded a threshold.

This chapter covers the following:

- Setting Up Notifications
- Extending Notification Beyond E-mail
- Passing Corrective Action Status Change Information
- Passing Job Execution Status Information
- Passing User-Defined Target Properties to Notification Methods
- Management Information Base (MIB)
- Troubleshooting Notifications

3.1 Setting Up Notifications

All Enterprise Manager administrators can set up e-mail notifications for themselves. Super Administrators also have the ability to set up notifications for other Enterprise Manager administrators.

3.1.1 Setting Up a Mail Server for Notifications

Before Enterprise Manager can send e-mail notifications, you must first specify the Outgoing Mail (SMTP) servers to be used by the notification system. Once set, you can then define e-mail notifications for yourself or, if you have Super Administrator privileges, other Enterprise Manager administrators.

You specify the Outgoing Mail (SMTP) server on the Notification Methods page (Figure 3–1). To display the Notification Methods page, from the **Setup** menu choose **Notifications-->Notification Methods**.

Note: You must have Super Administrator privileges in order to set up SMTP servers.

Specify one or more outgoing mail server names, the mail server authentication credentials (User Name, Password, and Confirm Password), if required, the name you want to appear as the sender of the notification messages, and the e-mail address you want to use to send your e-mail notifications. This address, called the Sender's Mail Address, must be a valid address on each mail server that you specify. A message will be sent to this e-mail address if any problem is encountered during the sending of an e-mail notification. Example 3–1 shows sample notification method entries.

Example 3–1 Mail Server Settings

- **Outgoing Mail (SMTP) Server -** smtp01.mycorp.com:587, smtp02.mycorp.com
- **User Name -** myadmin
- **Password -** ******
- **Confirm Password -** ******
- **Identify Sender As -** Enterprise Manager
- **Sender's E-mail Address -** mgmt_rep@mycorp.com
- **Use Secure Connection** - *No*: E-mail is not encrypted. *SSL*: E-mail is encrypted using the Secure Sockets Layer protocol. *TLS, if available*: E-mail is encrypted using the Transport Layer Security protocol if the mail server supports TLS. If the server does not support TLS, the e-mail is automatically sent as plain text.

Figure 3–1 Defining a Mail Server

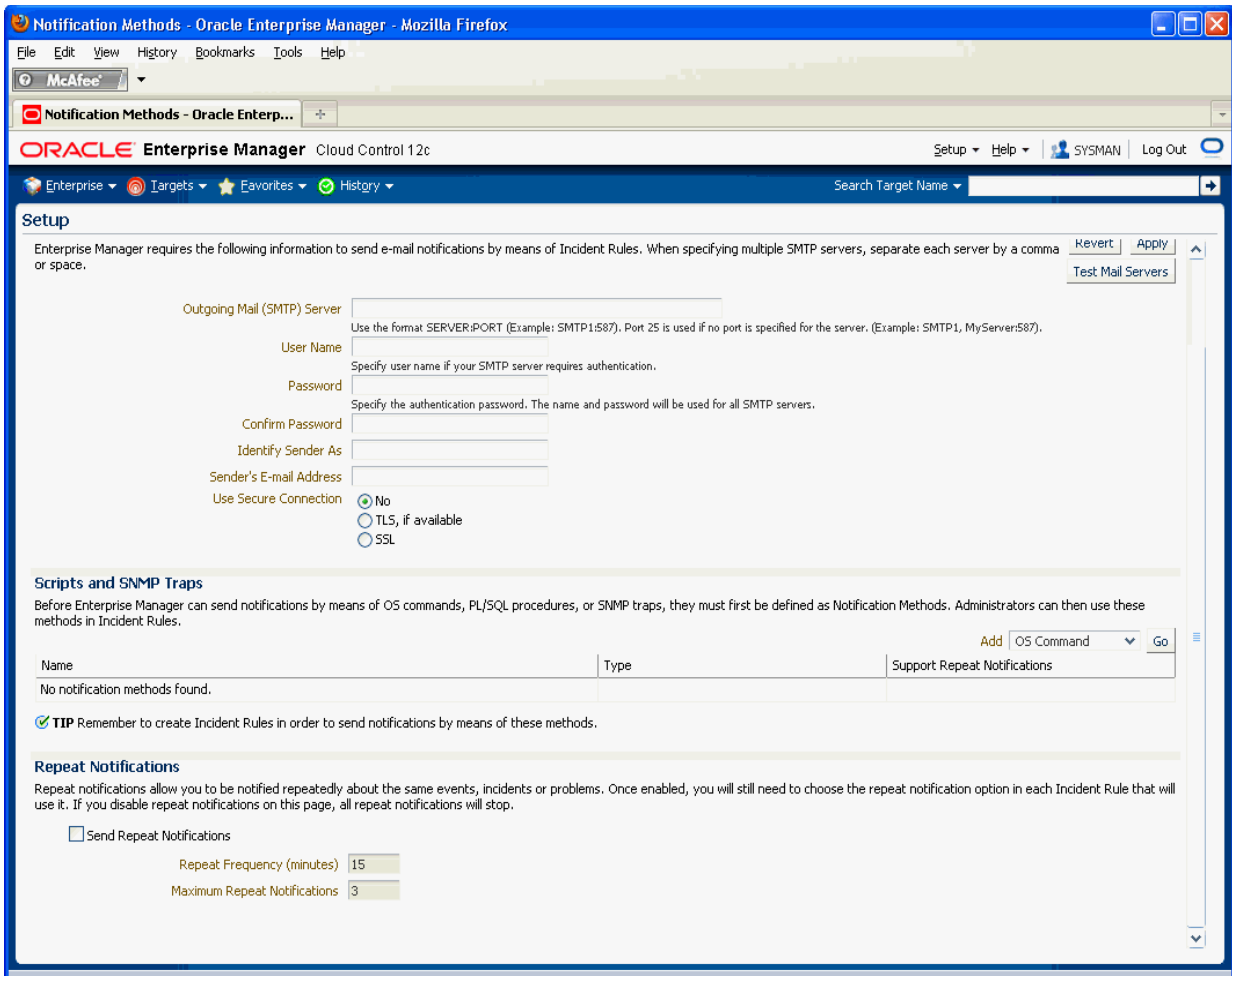

Note: The e-mail address you specify on this page is not the e-mail address to which the notification is sent. You will have to specify the e-mail address (where notifications will be sent) from the Password and E-mail page.
Setup-->MyPreferences-->Enterprise Manager Password & E-mail.

After configuring the e-mail server, click **Test Mail Servers** to verify your e-mail setup. You should verify that an e-mail message was received by the e-mail account specified in the **Sender's E-mail Address** field.

Defining multiple mail servers will improve the reliability of e-mail notification delivery. E-mail notifications will be delivered if at least one e-mail server is up. The notification load is balanced across multiple e-mail servers by the OMS, which switches through them (servers are allocated according to availability) after 20 e-mails have been sent. Switching is controlled by the *em.notification.emails_per_connection* emoms property.

3.1.1.1 Setting Up Repeat Notifications

Repeat notifications allow administrators to be notified repeatedly until an incident is either acknowledged or the number of **Maximum Repeat Notifications** has been

reached. Enterprise Manager supports repeat notification for all notification methods (e-mail, OS command, PL/SQL procedure, and SNMP trap). To enable this feature for a notification method, select the **Send Repeat Notifications** option. In addition to setting the maximum number of repeat notifications, you can also set the time interval at which the notifications are sent.

Important: For Oracle database versions 10 and higher, it is recommend that no modification be made to *aq_tm_processes* init.ora parameter. If, however, this parameter must be modified, its value should be at least one for repeat notification functionality. If the Enterprise Manager Repository database version is 9.2, the *aq_tm_processes* init.ora parameter must be set to at least one to enable repeat notification functionality.

Repeat Notifications for Incident Rules

Setting repeat notifications globally at the notification method level may not be provide sufficient flexibility. For example, you may want to have different repeat notification settings based on event type. Enterprise Manager accomplishes this by allowing you to set repeat notifications for individual incident rule sets or individual rules within a rule set. Repeat notifications set at the rule level take precedence over those defined at the notification method level.

Important: Repeat notifications for rules will only be sent if the **Send Repeat Notifications** option is enabled in the Notification Methods page.

For PL/SQL, OS command, and SNMP trap notification methods, you must enable each method to support repeat notifications. You can select **Supports Repeat Notifications** option when adding a new notification method or by editing an existing method.

Figure 3–2 Enabling Repeat Notification for an OS Command Notification Method

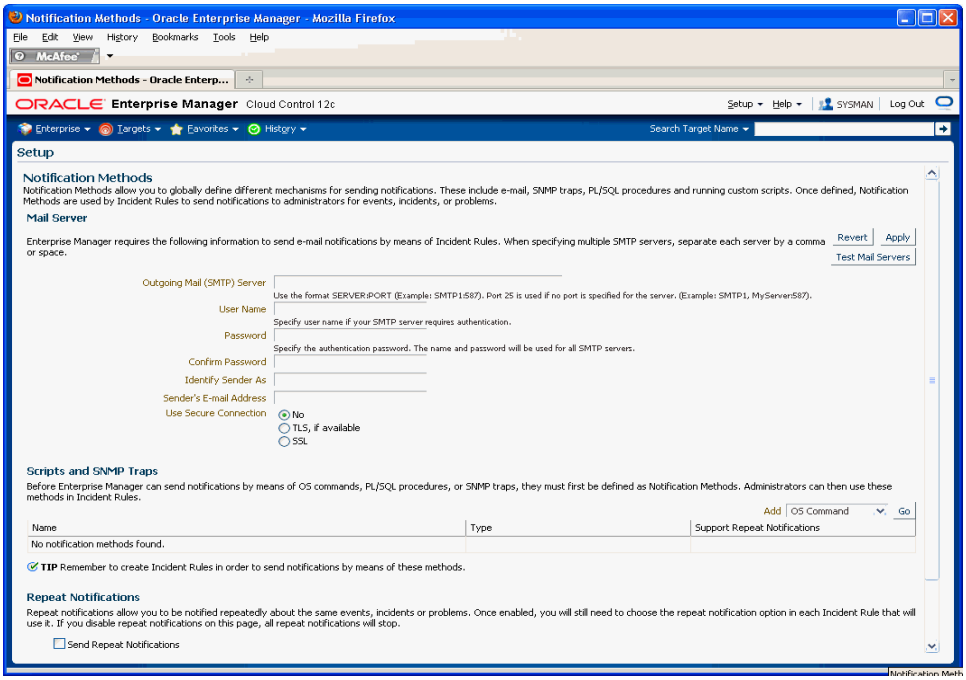

3.1.2 Setting Up E-mail for Yourself

If you want to receive notifications by e-mail, you will need to specify your e-mail address(s) in the Password & E-mail page (**Setup-->MyPreferences-->Enterprise Manager Password & E-mail**). In addition to defining notification e-mail addresses, you associate the notification message format (long, short, pager) to be used for your e-mail address.

Setting up e-mail involves three steps:

Step 1: Define e-mail addresses.

Step 2: Set up a Notification Schedule.

Step 3: Subscribe to incident rules in order to receive e-mails.

3.1.2.1 Defining E-mail Addresses

An e-mail address can have up to 128 characters. There is no upper limit with the number of e-mail addresses.

To add an e-mail address:

1. From the **Setup** menu, choose **MyPreferences-->Enterprise Manager Password & E-mail**.
2. Click **Add Another Row** to create a new e-mail entry field in the **E-mail Addresses** table.
3. Specify the e-mail associated with your Enterprise Manager account. All e-mail notifications you receive from Enterprise Manager will be sent to the e-mail addresses you specify.

   For example, `user1@oracle.com`

Select the *E-mail Type* (message format) for your e-mail address. *E-mail (Long)* sends a HTML formatted e-mail that contains detailed information. Example 3–2 shows a typical notification that uses the long format.

E-mail (Short) and *Pager(Short)* (Example 3–3) send a concise, text e-mail that is limited to a configurable number of characters, thereby allowing the e-mail be received as an SMS message or page. The content of the message can be sent entirely in the subject, entirely in the body or split across the subject and body. For example, in the last case, the subject could contain the severity type (for example, Critical) and the target name. The body could contain the time the severity occurred and the severity message. Since the message length is limited, some of this information may be truncated. If truncation has occurred there will be an ellipsis end of the message. *Pager(Short)* addresses are used for supporting the paging feature in incident rules. Note that the incident rules allow the rule author to designate some users to receive a page for critical issues.

4. Click **Apply** to save your e-mail address.

Example 3–2 Long E-mail Notification for Metric Alerts

```
Target type=Host
Target name=adc6140830.us.oracle.com
Message=Filesystem / has 54.39% available space, fallen below warning (60) or
critical (30) threshold.
Severity=Warning
Event reported time=Apr 28, 2011 2:33:55 PM PDT
Event Type=Metric Alert
Event name=Filesystems:Filesystem Space Available (%)
Metric Group=Filesystems
Metric=Filesystem Space Available (%)
Metric value=54.39
Key Value=/
Key Column 1=Mount Point
Rule Name=NotifRuleSet1,Event rule1
Rule Owner=SYSMAN
```

Example 3–3 Short E-mail Notification for Alerts

```
Subject is :
EM:Unreachable Start:myhost
Body is :
Nov 16, 2006 2:02:19 PM EST:Agent is Unreachable (REASON = Connection refused)
but the host is UP
```

More about E-mail(Short) and Pager(Short) Formats

Enterprise Manager does not directly support message services such as paging or SMS, but instead relies on external gateways to, for example, perform the conversion from e-mail to page. Beginning with Enterprise Manager 12c, the notification system allows you to tag e-mail addresses explicitly as 'page' or 'e-mail'. Explicit system differentiation between these two notification methods allows you to take advantage of the multiple action capability of incident rules. For example, the e-mail versus page distinction is required in order to send you an e-mail if an event severity is 'warning' or page you if the severity is 'critical'. To support this capability, a Pager format has been made available that sends an abbreviated version of the short format e-mail.

EMOMS properties can be used for controlling the size and format of the short e-mail. The following table lists emoms properties for Notification System.

Table 3–1 EMOMS Properties for Notifications

Property Name	Default Value	Description
em.notification.emails_per_minute	250	Email delivery limits per minute. The Notification system uses this value to throttle number of Email delivery per minutes. Customer should set the value lower if doesn't want to over flow the Email server, or set the value higher if the Email server can handle high volume of Emails.
em.notification.cmds_per_minute		OS Command delivery limits per minute. The Notification system uses this value to throttle number of OS Command delivery per minutes.
em.notification.os_cmd_timeout	30	OS Command delivery timeout in seconds. This value indicates how long to allow OS process to execute the OS Command delivery. Set this value higher if the OS command script requires longer time to complete execution.
em.notification.plsql_per_minute	250	PL/SQL delivery limits per minute. The Notification system uses this value to throttle number of PL/SQL delivery per minutes.
em.notification.java_per_minute	500	JAVA delivery limits per minute. The Notification system uses this value to throttle number of Java delivery per minutes.
em.notification.ticket_per_minute	250	Ticket delivery limits per minute. The Notification system uses this value to throttle number of Ticket delivery per minutes.
em.notification.traps_per_minute	250	SNMP delivery limits per minute. The Notification system uses this value to throttle number of SNMP Trap per minutes.
em.notification.locale.plsql	OMS Locale	This property specifies the Locale delivered by advanced PL/SQL notification. Customer can define this property to overwrite the default Locale where OMS installed.
em.notification.locale.email	OMS Locale	This property specifies the Locale delivered by Email. Customer can define this property to overwrite the default Locale where OMS installed.
em.notification.locale.osmcd	OMS Locale	This property specifies the Locale delivered by OS Command. Customer can define this property to overwrite the default Locale where OMS installed.
em.notification.locale.snmp	OMS Locale	This property specifies the Locale delivered by SNMP trap. Customer can define this property to overwrite the default Locale where OMS installed.
em.notification.oscmd.max_env_var_length	512	The maximum length of OS Common environment variable value.
em.notification.snmp.max_oid_length	2560	The maximum length of SNMP OID value.

Table 3–1 (Cont.) EMOMS Properties for Notifications

Property Name	Default Value	Description
em.notification.min_delivery_threads	6	The minimum number of active threads in the thread pool initially and number of active threads are running when system is in low activities. Set the value higher will use more system resources and delivery and deliver more notifications.
em.notification.max_delivery_threads	24	The maximum number of active threads in the thread pool when the system is in the high activities. This vale should greater than em.notification.min_delivery_threads. Set the value higher will use more system resources and deliver more notifications.
em.notification.short_format_length	>=1 (155)	The size limit of the total number of characters in short email format. The customers should modify this property value to fit their Email or Pager limit of content size.
em.notification.snmp_packet_length	>=1 (5120)	The maximum size of SNMP Protocol Data unit.
em.notification.email_content_transfer_encoding	8-bit, 7-bit(QP), 7-bit(BASE64) (8-bit)	The character set that can encode the Email. Oracle supports three character sets : 8-bit, 7-bit(QP), and7-bit(BASE64).
em.notification.emails_per_connection	>=1 (20)	The maximum number of emails delivery to same email gateway before rotate to next available email gateway if customers configure multiple email gateways. It is used for email gateway load balance.
em.notification.short_format	both, subject, body (both)	Use short format on both subject and body, subject only, or body only.

You must establish the maximum size your device can support and whether the message is sent in subject, body or both.

You can modify the EMOMS properties by using the Enterprise Manager command line control `emclt get/set/delete/list` property command.

Get Property Command

```
emctl get [-sysman_pwd "sysman password"]-name em.notification.short_format_length
```

Set Property Command

```
emctl set  property -name em.notification.short_format_length -value 155
```

Emoms Properties Entries for a Short E-mail Format

```
emctl set  property -name em.notification.short_format_length -value 155
emctl set property -name em.notification.short_format -value both
```

3.1.2.2 Setting Up a Notification Schedule

Once you have defined your e-mail notification addresses, you will need to define a notification schedule. For example, if your e-mail addresses are user1@oracle.com, user2@oracle.com, user3@oracle.com, you can choose to use one or more of these e-mail addresses for each time period in your notification schedule.

> **Note:** When you enter e-mail addresses for the first time, a 24x7 weekly notification schedule is set automatically. You can then review and modify the schedule to suit your monitoring needs.

A notification schedule is a repeating schedule used to specify your on-call schedule—the days and time periods and e-mail addresses that should be used by Enterprise Manager to send notifications to you. Each administrator has exactly one notification schedule. When a notification needs to be sent to an administrator, Enterprise Manager consults that administrator's notification schedule to determine the e-mail address to be used. Depending on whether you are Super Administrator or a regular Enterprise Manager administrator, the process of defining a notification schedule differs slightly.

If you are a regular Enterprise Manager administrator and are defining your own notification schedule:

1. From **Setup** menu, choose **Notifications-->My Notification Schedule**.
2. Follow the directions on the Notification Schedule page to specify when you want to receive e-mails.

3.1.2.3 Subscribe to Receive E-mail for Incident Rules

An incident rule is a user-defined rule that specifies the criteria by which notifications should be sent for specific events that make up the incident. An incident rule set, as the name implies, consists of one or more rules associated with the same incident.

When creating an incident rule, you specify criteria such as the targets you are interested in, the types of events to which you want the rule to apply. Specifically, for a given rule, you can specify the criteria you are interested in and the notification methods (such as e-mail) that should be used for sending these notifications. For example, you can set up a rule that when any database goes down or any database backup job fails, e-mail should be sent and the "log trouble ticket" notification method should be called. Or you can define another rule such that when the CPU or Memory Utilization of any host reach critical severities, SNMP traps should be sent to another management console.

Notification flexibility is further enhanced by the fact that with a single rule, you can perform multiple actions based on specific conditions. Example: When monitoring a condition such as machine memory utilization, for an incident severity of 'warning' (memory utilization at 80%), send the administrator an e-mail, if the severity is 'critical' (memory utilization at 99%), page the administrator immediately.

You can subscribe to a rule you have already created.

1. From the **Setup** menu, choose **Incidents-->Incident Rules**.
2. On the Incident Rules page, select the desired rule.
3. From the **Actions** menu, choose **Notifications** and then **Basic Notifications**.

Out-of-Box Incident Rules

Enterprise Manager comes with two incident rule sets that cover the most common monitoring conditions, they are:

- Incident Management Ruleset for All Targets
- Event Management Ruleset for Self Update

If the conditions defined in the out-of-box incident rules meet your requirements, you can simply subscribe to receive e-mail notifications for the conditions defined in the rule using the subscribe procedure shown in the previous section.

Creating Your Own Incident Rules

You can define your own custom rules. The following procedure documents the process of incident rule creation for non-Super Administrators.

To create your own incident rule:

1. From the **Setup** menu, choose **Incidents-->Incident Rules**.

   The Incident Rules page displays. From this page you can create a new rule set, to which you can add new rules. Alternatively, if you have the requisite permissions, you can add new rules to existing

2. Click **Create Rule Set...**

   The create rule set page displays.

3. Specify the **Name**, **Description**, and the **Targets** to which the rules set should apply.
4. Click the **Rules** tab and then click **Create**.
5. Choose the incoming incident, event or problem to which you want the rule to apply. See "Working with Incident Rules" on page 2-19 for more information.
6. Click **Continue**.

   Enterprise Manager displays the Create Incident Rule pages. Enter the requisite information on each page to create your incident rule.

7. Follow the wizard instructions to create your rule.

   Once you have completed defining your rule, the wizard returns you to the create rule set page.

8. Click **Save** to save the incident rule set.

3.1.3 Setting Up E-mail for Other Administrators

If you have Super Administrator privileges, you can set up e-mail notifications for other Enterprise Manager administrators. To set up e-mail notifications for other Enterprise Manager administrators, you need to:

Step 1: Ensure Each Administrator Account has an Associated E-mail Address

Each administrator to which you want to send e-mail notifications must have a valid e-mail address.

1. From the **Setup** menu, choose **Security-->Administrators**.
2. For each administrator, define an e-mail address. This sets up a 24x7 notification schedule for this user that uses all the e-mail addresses specified.

Enterprise Manager also allows you to specify an administrator address when editing an administrator's notification schedule.

Step 2: Define Administrators' Notification Schedules

Once you have defined e-mail notification addresses for each administrator, you will need to define their respective notification schedules. Although a default 24x7 notification schedule is created when you specify an e-mail address for the first time, you should review and edit the notification schedule as needed.

1. From the **Setup** menu, choose **Notifications-->Notification Schedule**.

   From the vertical navigation bar, click Schedules (under Notification). The **Notification Schedule** page appears.
2. Specify the administrator who's notification schedule you wish to edit and click **Change**.
3. Click **Edit Schedule Definition**. The **Edit Schedule Definition: Time Period** page appears. If necessary, modify the rotation schedule.
4. Click **Continue**. The **Edit Schedule Definition: E-mail Addresses** page appears.
5. Follow the directions on the **Edit Schedule Definition: E-mail Addresses** page to modify the notification schedule.
6. Click **Finish** when you are done.
7. Repeat steps three through seven for each administrator.

Step 3: Assign Incident Rules to Administrators

With the notification schedules set, you now need to assign the appropriate incident rules for each designated administrator.

1. From the **Setup** menu, choose **Incidents** and then Incident Rules.
2. Select the desired **Ruleset** and click **Edit**.
3. Click on the Rules tab., select the desired rule, and click **Edit**.
4. Click **Add Actions**, select desire action, and click **Edit**.
5. Enter the **Administrator** name on either **E-mail To** or **E-mail Cc** field in the **Basic Notification** region.
6. Click **Continue**, click **Next**, click **Next**, click **Continue**, and finally click **Save**.

3.1.4 E-mail Customization

Enterprise Manager allows Super Administrators to customize global e-mail notifications for the following types: All events, incidents, problems, and specific event types installed. You can alter the default behavior for all events by customizing *Default Event Email Template*. In addition, you can further customize the behavior for a specific event type by customizing the template for the event type. For instance, you can customize the *Metric Alert Events* template for the metric alert event type. Using predefined building blocks (called attributes and labels) contained within a simple script, Super Administrators can customize alert e-mails by selecting from a wide variety of information content.

To customize an e-mail:

1. From the **Setup** menu, choose **Notifications and then Customize Email Formats**.
2. Choose the **Type** and **Format**.

3. Click **Customize**. The Customize E-mail Template page is displayed.

From the Customize E-mail Template page, you can modify the content of the e-mail template Enterprise Manager uses to generate e-mail notifications. Extensive information on script formatting, syntax, and options is available from the Edit E-mail Template page via imbedded assistance and online help.

Figure 3–3 E-mail Customization

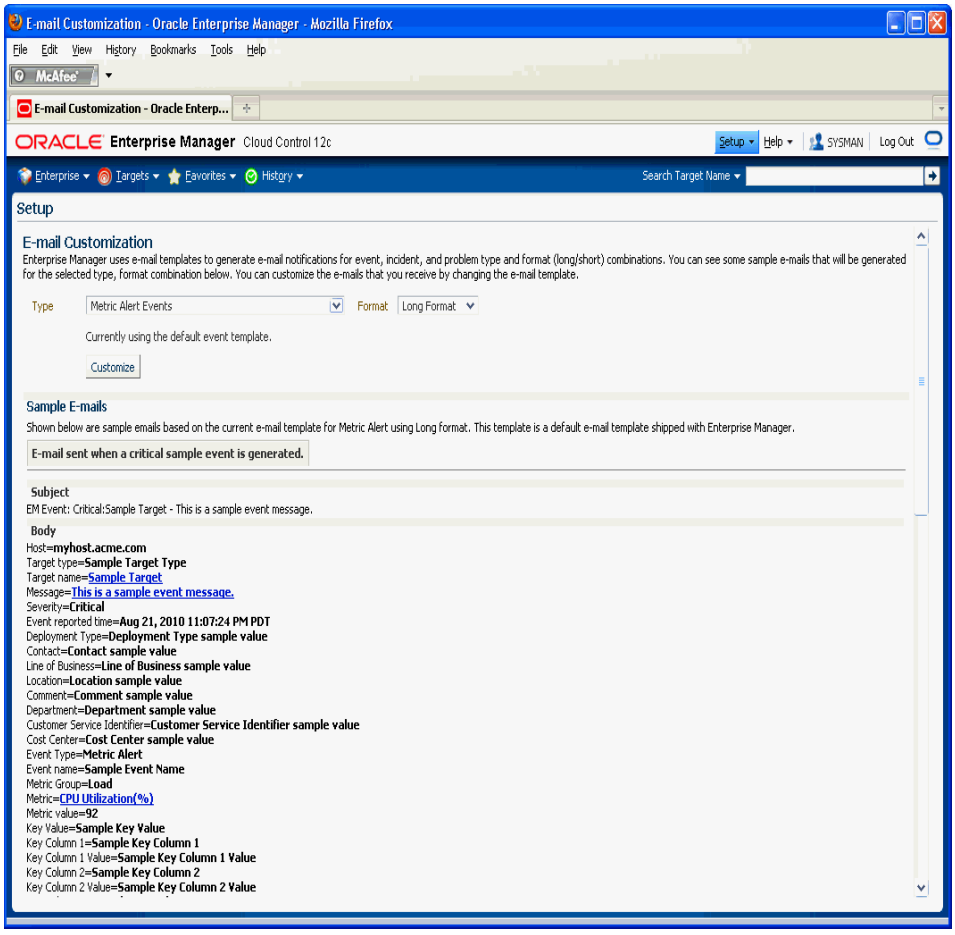

3.1.4.1 E-mail Customization Reference

The following reference summarizes the semantics and component syntax of the pseudo-language used to define e-mails. The pseudo-language provides you with a simple, yet flexible way to customize e-mail notifications. The following is a summary of pseudo-language conventions/limitations:

- You can add comments (or any free-form text) using separate lines beginning with "--" or at end of lines.
- You can use attributes.
- You can use IF & ELSE & ENDIF control structures. You can also use multiple conditions using "AND" or "OR". Nested IF statements are not supported.
- You can insert spaces for formatting purposes. Spaces at the beginning of a line will be ignored in the actual e-mail. To insert spaces at the beginning of a line, use the [SP] attribute.

- Use "/" to escape and "[" or "]" if you want to add attribute names, operators, or IF clauses to the actual e-mail.
- HTML is not supported.

Reserved Words and Operators

The following table lists all reserved words and operators used when modifying e-mail scripts.

Table 3–2 Reserved Words and Operators

Reserved Word/Operator	Description
IF, ELSIF, ENDIF, ELSE	Used in IF-ELSE constructs.
AND, OR	Boolean operators – used in IF-ELSE constructs only.
NULL	To check NULL value for attributes - used in IF-ELSE constructs only.
|	Pipe operator – used to show the first non-NULL value in a list of attributes. For example: `METRIC_NAME|SEVERITY`
EQ, NEQ	Equal and Not-Equal operators – applicable to NULL, STRING and NUMERIC values.
/	Escape character – used to escape reserved words and operators. Escape characters signify that what follows the escape character takes an alternative interpretation.
[,]	Delimiters used to demarcate attribute names and IF clauses.

Syntax Elements

Literal Text

You can specify any text as part of the e-mail content. The text will be displayed in the e-mail and will not be translated if the Oracle Management Services (OMS) language setting is changed. For example, 'my Oracle Home' appears as 'my Oracle Home' in the generated e-mail.

Predefined Attributes

Predefined attributes/labels will be substituted with actual values in a specific context. To specify a predefined attribute/label, use the following syntax:

```
[PREDEFINED_ATTR]
```

Attribute names can be in either UPPER or LOWER case. The parsing process is case-insensitive.

A pair of square brackets is used to distinguish predefined attributes from literal text. For example, for a job e-mail notification, the actual job name will be substituted for `[EXECUTION_STATUS]`. For a metric alert notification, the actual metric column name will be substituted for `[METIRC_COLUMN]`.

You can use the escape character "/" to specify words and not have them interpreted as predefined labels/attributes. For example, "`/[NEW/]`" will not be considered as the predefined attribute `[NEW]` when parsed.

Operators

`EQ, NEQ` – for text and numeric values

`NULL`- for text and numeric values

`GT, LT, GE, LE` – for numeric values

Control Structures

The following table lists acceptable script control structures.

Table 3–3 Control Structures

Control Structure	Description
Pipe "|"	Two or more attributes can be separated by '|' character. For example, `[METRIC_NAME|SEVERITY]` In this example, only the applicable attribute within the current alert context will be used (replaced by the actual value) in the e-mail. If more than one attributes are applicable, only the left-most attribute is used.

Table 3–3 (Cont.) Control Structures

Control Structure	Description

IF	Allows you to make a block of text conditional. Only one level of IF and ELSIF is supported. Nested IF constructs are not supported.
	
	All attributes can be used in IF or ELSIF evaluation using EQ/NEQ operators on NULL values. Other operators are allowed for "SEVERITY" and "REPEAT_COUNT" only.
	
	Inside the IF block, the values need to be contained within quotation marks " ". Enterprise Manager will extract the attribute name and its value based on the position of "EQ" and other key words such as "and", "or". For example,
	
	`[IF REPEAT_COUNT EQ "1" AND SEVERITY EQ "CRITICAL" THEN]`
	
	The statement above will be true when the attributes of the alert match the following condition:
	
	■ Attribute Name: REPEAT_COUNT
	■ Attribute Value: 1
	■ Attribute Name: SEVERITY
	■ Attribute Value: CRITICAL
	
	Example IF Block:
	
	`[IF EXECUTION_STATUS NEQ NULL]`
	`  [JOB_NAME_LABEL]=[EXECUTION_STATUS]`
	`  [JOB_OWNER_LABEL]=[JOB_OWNER]`
	`[ENDIF]`
	
	`[IF SEVERITY_CODE EQ CRITICAL ]`
	`  [MTRIC_NAME_LABEL]=[METRIC_GROUP]`
	`  [METRIC_VALUE_LABEL]=[METRIC_VALUE]`
	`  [TARGET_NAME_LABEL]=[TARGET_NAME]`
	`  [KEY_VALUES]`
	`[ENDIF]`
	
	Example IF and ELSEIF Block:
	
	`[IF SEVERITY_CODE EQ CRITICAL]`
	`  statement1`
	
	`[ELSIF SEVERITY_CODE EQ WARNING]`
	`  statement2`
	
	`[ELSIF SEVERITY_CODE EQ CLEAR]`
	`  statement3`
	
	`[ELSE]`
	`  statement4`
	
	`[ENDIF]`

Comments

You can add comments to your script by prefacing a single line of text with two hyphens "--". For example,

```
-- Code added on 8/3/2009
    [IF REPEAT_COUNT NEQ NULL]
    . . .
```

Comments may also be placed at the end of a line of text.

```
[IF SEVERITY_SHORT EQ W] -- for Warning alert
```

HTML Tags in Customization Content

Use of HTML tags is not supported.

When Enterprise Manager parses the e-mail script, it will convert the “<” and “>” characters of HTML tags into encoded format (< and >). This ensures that the HTML tag is not treated as HTML by the destination system.

Examples

E-mail customization template scripts support three main operators.

- Comparison operators: EQ/NEQ/GT/LT/GE/LE
- Logic operators: AND/OR
- Pipeline operator: |

3.2 Extending Notification Beyond E-mail

Notification Methods are the mechanisms by which notifications are sent. Enterprise Manager Super Administrators can set up e-mail notifications by configuring the 'e-mail' notification method. Most likely this would already have been set up as part of the Oracle Management Service installation.

Enterprise Manager Super Administrators can also define other custom notification methods. For example, event notifications may need to be forwarded to a 3rd party trouble-ticketing system. Assuming APIs to the third-party trouble-ticketing system are available, a custom notification method can be created to call a custom OS script that has the appropriate APIs. The custom notification method can be named in a user-friendly fashion, for example, "Log trouble ticket". Once the custom method is defined, whenever an administrator needs to send alerts to the trouble-ticketing system, he simply needs to invoke the now globally available notification method called "Log trouble ticket".

Custom notification methods can be defined based on any custom OS script, any custom PL/SQL procedure, or by sending SNMP traps. A fourth type of notification method (Java Callback) exists to support Oracle internal functionality and cannot be created or edited by Enterprise Manager administrators.

Only Super Administrators can define OS Command, PL/SQL, and SNMP Trap notification methods. However, any Enterprise Manager administrator can add these notification methods (once defined by the Super Administrator) as actions to their incident rules.

Through the Notification Methods page, you can:

- Set up the outgoing mail servers if you plan to send e-mail notifications through incident rules
- Create other custom notification methods using OS and PL/SQL scripts and SNMP traps.
- Set global repeat notifications.

3.2.1 Custom Notification Methods Using Scripts and SNMP Traps

You can create other custom notification methods based on OS scripts, PL/SQL procedures, or SNMP traps. Any administrator can then use these methods in incident rules.

The length of the SNMP OID value is limited to 2560 bytes by default. Configure emoms property *em.notification.snmp.max_oid_length* to change the default limit.

For Enterprise Manager 12c, SNMP traps are delivered for event notifications only. SNMP trap notifications are not supported for incidents or problems.

Note: SNMP advanced notification methods defined using previous versions of Enterprise Manager (pre-12c) will continue to function without modification. Traps will conform to the older Enterprise Manager MIB definition.

3.2.1.1 Adding a Notification Method based on an OS Command or Script

Notification system invokes the custom script when an incident rule matches the OS Command advanced notification action. Custom script receives notifications for matching events, incidents and problem through environment variables.

The length of any environment variable's value is limited to 512 characters by default. Configure emoms property named em.notification.oscmd.max_env_var_length for changing the default limit.

Note: Notification methods based on OS commands must be configured by an administrator with Super Administrator privileges.

Step 1: Define your OS command or script.

You can specify an OS command or script that will be called by the notification system when an incident rule matches the OS Command advanced notification action. You can use incident, event, or problem context information, corrective action execution status and job execution status within the body of the script. Passing this contextual information to OS commands/scripts allows you to customize automated responses specific event conditions. For example, if an OS script opens a trouble ticket for an in-house support trouble ticket system, you will want to pass severity levels (critical, warning, and so on) to the script to open a trouble ticket with the appropriate details and escalate the problem. For more information on passing specific types of information to OS Command or Scripts, see:

- "Passing Event, Incident, Problem Information to an OS Command or Script" on page 3-19
- "Passing Corrective Action Execution Status to an OS Command or Script" on page 3-50
- "Passing Job Execution Status to an OS Command or Script" on page 3-50

Step 2: Deploy the script on each Management Service host.

You must deploy the OS Command or Script on each Management Service host machine that connects to the Management Repository. The OS Command is run as the user who started the Management Service.

The OS Command or Script should be deployed on the same location on each Management Service host machine. The OS Command should be an absolute path, for example, /u1/bin/logSeverity.sh. The command is run by the user who started the Management Service. If an error is encountered during the running of the OS Command, the Notification System can be instructed to retry the sending of the notification to the OS Command by returning an exit code of 100. The procedure is initially retried after one minute, then two minutes, then three minutes and so on, until the notification is a day old, at which point it will be purged.

Example 3–4 shows the parameter in emoms.properties that controls how long the OS Command can execute without being killed by the Management Service. This is to prevent OS Commands from running for an inordinate length of time and blocking the delivery of other notifications. By default the command is allowed to run for 30 seconds before it is killed. The *em.notification.os_cmd_timeout* emoms property can be configured to change the default timeout value.

Example 3–4 Changing the em.notification.os_cmd_timeout EMOMS Property

```
emctl set property -name em.notification.os_cmd_timeout value 30
```

Step 3: Register your OS Command or Script as a new Notification Method.

Add this OS command as a notification method that can be called in incident rules. Log in as a Super Administrator. From the **Setup** menu, choose **Notifications-->Notification Methods**. From this page, you can define a new notification based on the 'OS Command' type. See "Adding a Notification Method based on an OS Command or Script" on page 3-17.

The following information is required for each OS command notification method:

- Name
- Description

  Both Name and Description should be clear and intuitive so that the function of the method is clear to other administrators.

- OS Command

You must enter the full path of the OS command or script in the OS command field (for example, `/u1/bin/myscript.sh`). For environments with multiple Management Services, the path must be exactly the same on each machine that has a Management Service. Command line parameters can be included after the full path (for example, /u1/bin/myscript.sh arg1 arg2).

Example 3–5 shows information required for the notification method.

Example 3–5 OS Command Notification Method

```
Name Trouble Ticketing
Description Notification method to log trouble ticket for a severity occurrence
OS Command /private/mozart/bin/logTicket.sh
```

> **Note:** There can be more than one OS Command configured per system.

Step 4: Assign the notification method to an instance rule.

You can edit an existing rule (or create a new instance rule), then go to the Methods page. From the **Setup** menu, choose **Incidents-->Incident Rules**. The Incident Rules page provides access to all available rule sets.

Passing Event, Incident, Problem Information to an OS Command or Script

The notification system passes information to an OS script or executable using system environment variables.

Conventions used to access environmental variables vary depending on the operating system:

- UNIX: $ENV_VARIABLE
- Windows:%ENV_VARIABLE%

The notification system sets the following environment variables before calling the script. The script can then use any or all of these variables within the logic of the script.

Environment Variables Common to Event, Incident and Problem

Table 3–4 Generic Environment Variables

Environment Variable	Description
NOTIF_TYPE	Type of notification and possible values NOTIF_NORMAL, NOTIF_RETRY, NOTIF_DURATION, NOTIF_REPEAT, NOTIF_CA, NOTIF_RCA
REPEAT_COUNT	How many times the notification has been sent out before this notification.
RULESET_NAME	The name of the ruleset that triggered this notification.
RULE_NAME	The name of the rule that triggered this notification.
RULE_OWNER	The owner of the ruleset that triggered this notification.
MESSAGE	The message of the event, incident, or problem.
MESSAGE_URL	EM console URL for this message.

Table 3–5 Category-Related Environment Variables

Environment Variable	Description
CATEGORIES_COUNT	Number of categories in this notification. This value is equal to1 if one category is associated with event, incident or problem. It is equal to 0 if no category associated with event, incident or problem.
CATEGORY_CODES_ COUNT	Number of category codes in this notification.
CATEGORY_n	Category is translated based on locale defined in OMS server. Valid values for the suffix "_n" are between 1.. $CATEGORIES_ COUNT
CATEGORY_CODE_n	Codes for the categories. Valid values for the suffix "_n" are between 1..$CATEGORY_CODES_COUNT

The following lists the common environment variables for User Defined Target Properties. They will be populated under the following cases: (a) When an event has a related target, (b) When an incident or a problem have single event source and have a related target.

Table 3–6 User-Defined Target Property Environment Variables

Environment Variable	Description
ORCL_GTP_COMMENT	Comment
ORCL_GTP_CONTACT	Contact
ORCL_GTP_COST_ CENTER	Cost Center
ORCL_GTP_ DEPARTMENT	Department
ORCL_GTP_ DEPLOYMENT_TYPE	Deployment type
ORCL_GTP_LINE_OF_ BUS	Line of Business
ORCL_GTP_LOCATION	Location

Event Notification-Specific Environment Variables

Table 3–7 Event Notification-Specific Environment Variables

Environment Variable	Description
EVENT_NAME	Event Name.
EVENT_REPORTED_ TIME	Event reported date.
EVENT_SOURCE_ COUNT	Number of Sources associated with this event.
EVENT_TYPE	Event type.
EVENT_OCCURRENCE_ TIME	Event occurrence time.
EVENT_TYPE_ATTRS	The list of event type specific attributes.
EVENT_CONTEXT_ ATTRS	Event context data.
LAST_UPDATED_TIME	Last updated time
SEQUENCE_ID	Event sequence global unique identifier.
SEVERITY	Severity of event, it is translated.
SEVERITY_CODE	Code for event severity. Possible values are the following. FATAL, CRITICAL, WARNING, MINOR_WARNING, INFORMATIONAL, and CLEAR
ACTION_MSG	Message describing the action to take for resolving the event.
TOTAL_OCCURRENCE_ COUNT	Total number of duplicate occurrences
SEQUENCE_ID	Event sequence ID
RCA_DETAILS	If RCA is associated with this events.

The following tables lists the environment variables for the incident associated with an event. They are populated when the event is associated with an incident.

Table 3–8 Associated Incident Environment Variables

Environment Variable	Description
ASSOC_INCIDENT_ ACKNOWLEDGED_BY_ OWNER	Set to yes, if associated incident was acknowledged by owner
ASSOC_INCIDENT_ ACKNOWLEDGED_ DETAILS	The details of associated incident acknowledgement. For example: No - if not acknowledged Yes By userName - if acknowledged
ASSOC_INCIDENT_ STATUS	Associated Incident Status
ASSOC_INCIDENT_ID	Associated Incident ID
ASSOC_INCIDENT_ PRIORITY	Associated Incident priority. Supported value are Urgent, Very High, High, Medium,Low, None.
ASSOC_INCIDENT_ OWNER	Associated Incident Owner if it is existed.
ASSOC_INCIDENT_ ESCALATION_LEVEL	Escalation level of the associated incident has a value between 0 to 5.

Following lists the common environment variables related to the Source Object. They are populated when $SOURCE_OBJ_TYPE is not TARGET.

Table 3–9 Source Object-Related Environment Variables

Environment Variable	Description
SOURCE_OBJ_TYPE	Type of the Source object. For example, JOB, TEMPLATE.
SOURCE_OBJ_NAME	Source Object Name.
SOURCE_OBJ_NAME_ URL	Source's event console URL.
SOURCE_OBJ_SUB_TYPE	Sub-type of the Source object. For example, it provides the underlying job type for job status change events.
SOURCE_OBJ_OWNER	Owner of the Source object.

Following lists the common environment variables for the target, associated with the given issue. They are populated when the issue is related to a target.

Table 3–10 Target-Related Environment Variables

Environment Variable	Description
TARGET_NAME	Name of Target
TARGET_TYPE	Type of Target
TARGET_OWNER	Owner of Target
HOST_NAME	The name of the host on which the target is deployed upon.
TARGET_URL	Target's Enterprise Manager Console URL.

Table 3–10 (Cont.) Target-Related Environment Variables

Environment Variable	Description
TARGET_LIFECYCLE_STATUS	Life Cycle Status of the target. Possible values: Production, MissionCritical, Stage, Test, and Development. It is null if not defined.
TARGET_VERSION	Target Version of the target

Events are classified into multiple types. For example, the mertc_alert event type is used for modeling metric alerts. Following SQL query lists the environment variables corresponding to the event type specific attributes.

```
Select event_class as event_type, upper(name) as env_var_name
from em_event_class_attrs
where notif_order != 0
and event_class is not null
union
select event_class as event_type, upper(name) || '_NLS' as env_var_name
from em_event_class_attrs
where notif_order != 0
and event_class is not null
and is_translated = 1
order by event_type, env_var_name;
```

There is environment variable payload specific to each event type which can be accessed from the OS scripts.The following tables list notification attributes for the most critical event types.

Table 3–11 Environment variables specific to Metric Alert event type

Environment Variable	Description
COLL_NAME	The name of the collection collecting the metric.
COLL_NAME_NLS	The translated name of the collection collecting the metric
KEY_COLUMN_X	Internal name of Key Column X where X is a number between 1 and 7.
KEY_COLUMN_X_NLS	Translated name of Key Column X where X is a number between 1 and 7.
KEY_COLUMN_X_VALUE	Value of Key Column X where X is a number between 1 and 7.
KEY_VALUE	Monitored object for the metric corresponding to the Metric Alert event.
METRIC_COLUMN	The name of the metric column
METRIC_COLUMN_NLS	The translated name of the metric column.
METRIC_DESCRIPTION	Brief description of the metric.
METRIC_DESCRIPTION_NLS	Translated brief description of the metric.
METRIC_GROUP	The name of the metric.
METRIC_GROUP_NLS	The translated name of the metric
NUM_KEYS	The number of key metric columns in the metric.

Table 3–11 (Cont.) Environment variables specific to Metric Alert event type

Environment Variable	Description
SEVERITY_GUID	The guid of the severity record associated with this metric alert.
VALUE	Value of the metric when the event triggered.

Table 3–12 Environment variables specific to Target Availability event type

Environment Variable	Description
AVAIL_SEVERITY	The transition severity that resulted in that status of the target to change to the current availability status..
AVAIL_SUB_STATE	The sub-status of a target for the current status.
CYCLE_GUID	The guid of the first severity record in this availability cycle.
METRIC_GUID	Metric Guid of response metric.
SEVERITY_GUID	The guid of the severity record associated with this availability status.
TARGET_STATUS	The current availability status of the target.
TARGET_STATUS_NLS	The translated current availability status of the target.

Table 3–13 Environment variables specific to Job Status Change event type

Environment Variable	Description
EXECUTION_ID	Unique ID of the job execution..
EXECUTION_LOG	The job output of the last step executed.
EXECUTION_STATUS	The internal status of the job execution.
EXECUTION_STATUS_ NLS	The translated status of the job execution.
EXEC_STATUS_CODE	Execution status code of job execution.
STATE_CHANGE_GUID	Unique ID of last status change

You can use SQL queries to list the deployed event types in your deployment and the payload specific to each one of them. The following SQL can be used to list all internal event type names which are registered in the Enterprise Manager.

```
select class_name as event_type_name from em_event_class;
```

Following SQL lists environment variables specific to metric_alert event type.

```
select env_var_name
  from
    ( Select event_class as event_type, upper(name) as env_var_name
     from em_event_class_attrs
    where notif_order != 0
    and event_class is not null
    union
    select event_class as event_type, upper(name) || '_NLS' as env_var_name
    from em_event_class_attrs
    where notif_order != 0
```

```
and event_class is not null
and is_translated = 1)
where event_type = 'metric_alert';
```

You can also obtain the description of notification attributes specific to an event type directly from the Enterprsie Manager console:

1. From the **Setup** menu, choose **Notifications** and then **Customize Email Formats**.
2. Select the event type.
3. Click **Customize**.
4. Click **Show Predefined Attributes**.

Environment variables, ending with the suffix _NLS, provide the translated value for given attribute. For example, METRIC_COLUMN_NLS environment variable will provide the translated value for the metric_column attribute. Translated values will be in the locale of the OMS.

Environment Variables Specific to Incident Notifications

Table 3–14 Incident-Specific Environment Variables

Environment Variable	Description
SEVERITY	Incident Severity, it is translated. Possible Values: Fatal, Critical, Warning, Informational, Clear
SEVERITY_CODE	Code for Severity. Possible values are the FATAL, CRITICAL, WARNING, MINOR_WARNING, INFORMATIONAL, and CLEAR
INCIDENT_REPORTED_TIME	Incident reported time
INCIDENT_ACKNOWLEDGED_BY_OWNER	Set yes, if incident is acknowledged by owner.
INCIDENT_ID	Incident ID
INCIDENT_OWNER	Incident Owner
ASSOC_EVENT_COUNT	The number events associated with this incident.
INCIDENT_STATUS	Incident status. There are two internal fixed resolution status. NEW CLOSED Users can define additional statuses.
ESCALATED	Is Incident escalated
ESCALATED_LEVEL	The escalated level of incident.
PRIORITY	Incident priority. It is the translated priority name. Possible Values: Urgent, Very High, Hight, Medium, Low,None

Table 3–14 (Cont.) Incident-Specific Environment Variables

Environment Variable	Description
PRIOTITY_CODE	Incident priority code It is the internal value defined in EM. PRIORITY_URGENT PRIORITY_VERY_HIGH PRIORITY_HIGH PRIORITY_MEDIUM PRIORITY_LOW PRIORITY_NONE
TICKET_STATUS	Status of external ticket, if it exists.
TICKET_ID	ID of external ticket, if it exists.
LAST_UPDATED_TIME	Incident last update time

Following lists the associated problem's environment variables, when the incident is associated with a problem.

Table 3–15 Associated Problem Environment Variables

Environment Variable	Description
ASSOC_PROBLEM_ ACKNOWLEDGED_BY_ OWNER	Set to yes, if this problem was acknowledged by owner
ASSOC_PROBLEM_ STATUS	Associated Problem Status
ASSOC_PROBLEM_ID	Associated Problem ID
ASSOC_PROBLEM_ PRIORITY	Associated Problem priority
ASSOC_PROBLEM_ OWNER	Associated Problem Owner if it is existed.
ASSOC_PROBLEM_ ESCALATION_LEVEL	Escalation level of the associated Problem has a value between 0 to 5.

Environment Variables Specific to Problem Notifications

Table 3–16 Problem-Specific Environment Variables

Environment Variable	Description
SEVERITY	Problem Severity, it is translated.
SEVERITY_CODE	Code for Severity. Possible values are the FATAL, CRITICAL, WARNING, MINOR_WARNING, INFORMATIONAL, and CLEAR

Table 3–16 (Cont.) Problem-Specific Environment Variables

Environment Variable	Description
PROBLEM_REPORTED_TIME	Problem reported time
PROBLEM_ACKNOWLEDGED_BY_OWNER	Set yes, if problem is acknowledged by owner.
PROBLEM_ID	Problem ID
PROBLEM_KEY	Problem Key
PROBLEM_OWNER	Problem Owner
ASSOC_INCIDENT_COUNT	The number incident associated with this problem.
PROBLEM_STATUS	Incident status. There are STATUS_NEW STATUS_CLOSED Any other user defined status
ESCALATED	Is Incident escalated. Yes if it is escalated, otherwise No.
ESCALATED_LEVEL	The escalated level of incident.
PRIORITY	Incident priority. It is the translated priority name.
PRIOTITY_CODE	Incident priority code It is the internal value defined in Enterprise Manager. PRIORITY_URGENT PRIORITY_VERY_HIGH PRIORITY_HIGH PRIORITY_MEDIUM PRIORITY_LOW PRIORITY_NONE
LAST_UPDATED_TIME	Last updated time
SR_ID	Oracle Service Request Id, if it exists.
BUD_ID	Oracle Bug ID, if an associated bug exists.

Environment Variables Common to Incident and Problem Notifications

An incident or problem may be associated with multiple event sources. An event source can be a Target, a Source Object, or both.

3.2.1.1.1 Environment Variables Related to Event Sources

Number of event sources will be set in EVENT_SOURCE_COUNT environment variable. Using the EVENT_SOURCE_COUNT information, script can be written to loop through the relevant environment variables to fetch the information about multiple event sources. Environment variables for all event sources are prefixed with EVENT_SOURCE_. Environment variables for source objects are suffixed with SOURCE_<attribute_name> . For example, EVENT_SOURCE_1_SOURCE_TYPE provides the source object type of first event source. Environment variables for a target

are suffixed with TARGET_<attribute_name>. For example, EVENT_SOURCE_1_TARGET_NAME provides the target name of first event source.

The following table lists the environment variables for source object of x-th Event Source.

Table 3–17 Source Object of the x-th Event Source

Environment Variable	Description
EVENT_SOURCE_x_SOURCE_GUID	Source Object GUID.
EVENT_SOURCE_x_SOURCE_TYPE	Source Object Type
EVENT_SOURCE_x_SOURCE_NAME	Source Object Name.
EVENT_SOURCE_x_SOURCE_OWNER	Source Object Owner.
EVENT_SOURCE_x_SOURCE_SUB_TYPE	Source Object Sub-Type.
EVENT_SOURCE_x_SOURCE_URL	Source Object URL to EM console.

The following table lists the environment variables for target of x-th Event Source.

Table 3–18 Target of x-th Event Source

Environment Variable	Description
EVENT_SOURCE_x_TARGET_GUID	Target GUID
EVENT_SOURCE_x_TARGET_NAME	Target name
EVENT_SOURCE_x_TARGET_OWNER	Target Owner
EVENT_SOURCE_x_TARGET_VERSION	Target version
EVENT_SOURCE_x_TARGET_LIFE_CYCLE_STATUS	Target life cycle status
EVENT_SOURCE_x_TARGET_TYPE	Target Type
EVENT_SOURCE_x_HOST_NAME	Target Host Name
EVENT_SOURCE_x_TARGET_URL	Target URL to EM Console.

3.2.1.1.2 Script Examples

The sample OS script shown in Example 3–6 appends environment variable entries to a log file. In this example, the script logs a severity occurrence to a file server. If the file

server is unreachable then an exit code of 100 is returned to force the Oracle Management Service Notification System to retry the notification

Example 3–6 Sample OS Command Script

```
#!/bin/ksh

LOG_FILE=/net/myhost/logs/event.log
if test -f $LOG_FILE
then
echo $TARGET_NAME $MESSAGE $EVENT_REPORTE_TIME >> $LOG_FILE
else
   exit 100
fi
```

Example 3–7 shows an OS script that logs alert information for both incidents and events to the file 'oscmdNotify.log'. The file is saved to the /net/myhost/logs directory.

Example 3–7 Alert Logging Scripts

```
#!/bin/sh
#
  LOG_FILE=/net/myhost/logs/oscmdNotify.log

  echo '-------------' >> $LOG_FILE

  echo 'issue_type=' $ISSUE_TYPE >> $LOG_FILE
  echo 'notif_type=' $NOTIF_TYPE >> $LOG_FILE
  echo 'message=' $MESSAGE >> $LOG_FILE
  echo 'message_url'  = $MESSAGE_URL >>$LOG_FILE
  echo 'severity=' $SEVERITY >> $LOG_FILE
  echo 'severity_code'  = $SEVERITY_CODE >>$LOG_FILE
  echo 'ruleset_name=' $RULESET_NAME >> $LOG_FILE
  echo 'rule_name=' $RULE_NAME >> $LOG_FILE
  echo 'rule_owner=' $RULE_OWNER >> $LOG_FILE
  echo 'repeat_count=' $REPEAT_COUNT >> $LOG_FILE
  echo 'categories_count'  = $CATEGORIES_COUNT >>$LOG_FILE
  echo 'category_1'  = $CATEGORY_1 >>$LOG_FILE
  echo 'category_2'  = $CATEGORY_2 >>$LOG_FILE
  echo 'category_code_1'  = $CATEGORY_CODE_1 >>$LOG_FILE
  echo 'category_code_2'  = $CATEGORY_CODE_2 >>$LOG_FILE
  echo 'category_codes_count'  = $CATEGORY_CODES_COUNT >>$LOG_FILE

# event
if [ $ISSUE_TYPE -eq 1 ]
then
  echo 'host_name=' $HOST_NAME >> $LOG_FILE
  echo 'event_type=' $EVENT_TYPE >> $LOG_FILE
  echo 'event_name=' $EVENT_NAME >> $LOG_FILE
  echo 'event_occurrence_time=' $EVENT_OCCURRENCE_TIME >> $LOG_FILE
  echo 'event_reported_time=' $EVENT_REPORTED_TIME >> $LOG_FILE
  echo 'sequence_id=' $SEQUENCE_ID >> $LOG_FILE
  echo 'event_type_attrs=' $EVENT_TYPE_ATTRS >> $LOG_FILE
  echo 'source_obj_name=' $SOURCE_OBJ_NAME >> $LOG_FILE
  echo 'source_obj_type=' $SOURCE_OBJ_TYPE >> $LOG_FILE
  echo 'source_obj_owner=' $SOURCE_OBJ_OWNER >> $LOG_FILE
  echo 'target_name'  = $TARGET_NAME >>$LOG_FILE
  echo 'target_url'  = $TARGET_URL >>$LOG_FILE
  echo 'target_owner=' $TARGET_OWNER >> $LOG_FILE
```

```
  echo 'target_type=' $TARGET_TYPE >> $LOG_FILE
  echo 'target_version=' $TARGET_VERSION >> $LOG_FILE
  echo 'lifecycle_status=' $TARGET_LIFECYCLE_STATUS >> $LOG_FILE
  echo 'assoc_incident_escalation_level'  = $ASSOC_INCIDENT_ESCALATION_LEVEL
>>$LOG_FILE
  echo 'assoc_incident_id'  = $ASSOC_INCIDENT_ID >>$LOG_FILE
  echo 'assoc_incident_owner'  = $ASSOC_INCIDENT_OWNER >>$LOG_FILE
  echo 'assoc_incident_acknowledged_by_owner'  = $ASSOC_INCIDENT_ACKNOWLEDGED_BY_
OWNER >>$LOG_FILE
  echo 'assoc_incident_acknowledged_details'  = $ASSOC_INCIDENT_ACKNOWLEDGED_
DETAILS >>$LOG_FILE
  echo 'assoc_incident_priority'  = $ASSOC_INCIDENT_PRIORITY >>$LOG_FILE
  echo 'assoc_incident_status'  = $ASSOC_INCIDENT_STATUS >>$LOG_FILE
  echo 'ca_job_status'  = $CA_JOB_STATUS >>$LOG_FILE
  echo 'event_context_attrs'  = $EVENT_CONTEXT_ATTRS >>$LOG_FILE
  echo 'last_updated_time'  = $LAST_UPDATED_TIME >>$LOG_FILE
  echo 'sequence_id'  = $SEQUENCE_ID >>$LOG_FILE
  echo 'test_date_attr_noref'  = $TEST_DATE_ATTR_NOREF >>$LOG_FILE
  echo 'test_raw_attr_noref'  = $TEST_RAW_ATTR_NOREF >>$LOG_FILE
  echo 'test_str_attr1'  = $TEST_STR_ATTR1 >>$LOG_FILE
  echo 'test_str_attr2'  = $TEST_STR_ATTR2 >>$LOG_FILE
  echo 'test_str_attr3'  = $TEST_STR_ATTR3 >>$LOG_FILE
  echo 'test_str_attr4'  = $TEST_STR_ATTR4 >>$LOG_FILE
  echo 'test_str_attr5'  = $TEST_STR_ATTR5 >>$LOG_FILE
  echo 'test_str_attr_ref'  = $TEST_STR_ATTR_REF >>$LOG_FILE
  echo 'total_occurrence_count'  = $TOTAL_OCCURRENCE_COUNT >>$LOG_FILE
fi

# incident
if [ $ISSUE_TYPE -eq 2 ]
then
  echo 'action_msg=' $ACTION_MSG >> $LOG_FILE
  echo 'incident_id=' $INCIDENT_ID >> $LOG_FILE
  echo 'incident_creation_time=' $INCIDENT_CREATION_TIME >> $LOG_FILE
  echo 'incident_owner=' $INCIDENT_OWNER >> $LOG_FILE
echo 'incident_acknowledged_by_owner'  = $INCIDENT_ACKNOWLEDGED_BY_OWNER >>$LOG_
FILE
  echo 'incident_status'  = $INCIDENT_STATUS >>$LOG_FILE
  echo 'last_modified_by=' $LAST_MODIFIED_BY >> $LOG_FILE
  echo 'last_updated_time=' $LAST_UPDATED_TIME >> $LOG_FILE
  echo 'assoc_event_count=' $ASSOC_EVENT_COUNT >> $LOG_FILE
  echo 'adr_incident_id=' $ADR_INCIDENT_ID >> $LOG_FILE
  echo 'occurrence_count=' $OCCURRENCE_COUNT >> $LOG_FILE
  echo 'escalated=' $ESCALATED >> $LOG_FILE
  echo 'escalated_level=' $ESCALATED_LEVEL >> $LOG_FILE
  echo 'priority=' $PRIORITY >> $LOG_FILE
  echo 'priority_code'  = $PRIORITY_CODE >>$LOG_FILE
  echo 'ticket_id=' $TICKET_ID >> $LOG_FILE
  echo 'ticket_status=' $TICKET_STATUS >> $LOG_FILE
  echo 'ticket_url=' $TICKET_ID_URL >> $LOG_FILE
  echo 'total_duplicate_count=' $TOTAL_DUPLICATE_COUNT >> $LOG_FILE
  echo 'source_count=' $EVENT_SOURCE_COUNT >> $LOG_FILE
  echo 'event_source_1_host_name'  = $EVENT_SOURCE_1_HOST_NAME >>$LOG_FILE
  echo 'event_source_1_target_guid'  = $EVENT_SOURCE_1_TARGET_GUID >>$LOG_FILE
  echo 'event_source_1_target_name'  = $EVENT_SOURCE_1_TARGET_NAME >>$LOG_FILE
  echo 'event_source_1_target_owner'  = $EVENT_SOURCE_1_TARGET_OWNER >>$LOG_FILE
  echo 'event_source_1_target_type'  = $EVENT_SOURCE_1_TARGET_TYPE >>$LOG_FILE
  echo 'event_source_1_target_url'  = $EVENT_SOURCE_1_TARGET_URL >>$LOG_FILE
  echo 'event_source_1_target_lifecycle_status'  = $EVENT_SOURCE_1_TARGET_
LIFECYCLE_STATUS >>$LOG_FILE
```

```
  echo 'event_source_1_target_version'  = $EVENT_SOURCE_1_TARGET_VERSION >>$LOG_
FILE
fi
exit 0
```

Example 3–8 shows a script that sends an alert to an HP OpenView console from Enterprise Manager Cloud Control. When a metric alert is triggered, the Enterprise Manager Cloud Control displays the alert. The HP OpenView script is then called, invoking opcmsg and forwarding the information to the HP OpenView management server.

Example 3–8 HP OpenView Script

```
/opt/OV/bin/OpC/opcmsg severity="$SEVERITY" app=OEM msg_grp=Oracle msg_
text="$MESSAGE" object="$TARGET_NAME"
```

3.2.1.1.3 Migrating pre-12c OS Command Scripts This section describes how to map pre-12c OS Command notification shell environment variables to 12c OS Command shell environment variables.

Please note that Policy Violations are no longer supported beginning with the 12c release.

Migrating Metric Alert Event Types

Following table is the mapping for the OS Command shell environment variables when the event_type is 'metric_alert'.

Table 3–19 pre-12c/12c metric_alert Environment Variable Mapping

Pre-12c Environment Variable	Corresponding 12c Environment Variables
ACKNOWLEDGED	ASSOC_INCIDENT_ACKNOWLEDGED_BY_OWNER
ACKNOWLEDGED_BY	ASSOC_INCIDENT_OWNER
CYCLE_GUID	CYCLE_GUID
HOST	HOST_NAME
KEY_VALUE	Note: See detail description below.
KEY_VALUE_NAME	Note: See detail description below
MESSAGE	MESSAGE
METRIC	METRIC_COLUMN
NOTIF_TYPE	NOTIF_TYPE; use the map in section 2.3.5
REPEAT_COUNT	REPEAT_COUNT
RULE_NAME	RULESET_NAME
RULE_OWNER	RULE_OWNER
SEVERITY	SEVERITY
TARGET_NAME	TAGER_NAME
TARGET_TYPE	TARGET_TYPE
TIMESTAMP	EVENT_REPORTED_TIME
VIOLATION_CONTEXT	EVENT_CONTEXT_ATTRS
VIOLATION_GUID	SEVERITY_GUID

Table 3–19 (Cont.) pre-12c/12c metric_alert Environment Variable Mapping

Pre-12c Environment Variable	Corresponding 12c Environment Variables
POLICY_RULE	No mapping, obsolete in NG release.

To get KEY_VALUE_NAME and KEY_VALUE from NG, perform the following steps.

- If $NUM_KEYS variable is null, then $KEY_VALUE_NAME and $KEY_VALUE are null.
- If $NUM_KEYS equals 1

  KEY_VALUE_NAME=$KEY_COLUMN_1

  KEY_VALUE=$KEY_VALUE_1_VALUE
- If $NUM_KEYS is greater than 1

  KEY_VALUE_NAME="$KEY_COLUMN_1;$KEY_COLUMN_2;..;KEY_COLUMN_x"

  KEY_VALUE="$KEY_COLUMN_1_VALUE;$KEY_COLUMN_2_VALUE;..;KEY_COLUMN_x_VALUE "

  Where x is the value of $NUM_KEYS and ";" is the separator.

Migrating Target Availability Event Types

Following table is the mapping for the OS Command shell environment variables when the event_type is 'target_availability'.

Table 3–20 pre-12c/12c target_availability Environment Variable Mappings

Pre-12c Environment Variable	Corresponding 12c Environment Variables
TARGET_NAME	TARGET_NAME
TARGET_TYPE	TARGET_TYPE
METRIC	Status
CYCLE_GUID	CYCLE_GUID
VIOLATION_CONTEXT	EVENT_CONTEXT_ATTRS
SEVERITY	TARGET_STATUS
HOST	HOST_NAME
MESSAGE	MESSAGE
NOTIF_TYPE	NOTIF_TYPE; use the map in section 2.3.5
TIMESTAMP	EVENT_REPORTED_TIME
RULE_NAME	RULESET_NAME
RULE_OWNER	RULE_OWNER
REPEAT_COUNT	REPEAT_COUNT
KEY_VALUE	""
KEY_VALUE_NAME	""

Migrating Job Status Change Event Types

Following table is the mapping for the OS Command shell environment variables when the event_type is 'jos_status_change'.

Table 3–21 pre-12c/12c job_status_change Environment Variable Mappings

Pre-12c Environment Variable	Corresponding 12c Environment Variables
JOB_NAME	SOURCE_OBJ_NAME
JOB_OWNER	SOURCE_OBJ_OWNER
JOB_TYPE	SOURCE_OBJ_SUB_TYPE
JOB_STATUS	EXECUTION_STATUS
NUM_TARGETS	1 if $ TARGET_NAME is not null, 0 otherwise
TARGET_NAME1	TARGET_NAME
TARGET_TYPE1	TARGET_TYPE
TIMESTAMP	EVENT_REPORTED_TIME
RULE_NAME	RULESET_NAME
RULE_OWNER	RULE_OWNER

Migrating Corrective Action-Related OS Scripts

Refer to section "Migrating Metric Alert Event Types" for mapping the following environment variables while receiving notifications for corrective actions.

KEY_VALUE

KEY_VALUE_NAME

METRIC

METRIC_VALUE

RULE_NAME

RULE_OWNER

SEVERITY

TIMESTAMP

VIOLATION_CONTEXT

Use the map below for mapping other environment variables.

Table 3–22 pre-12c/12c Corrective Action Environment Variable Mappings

Pre-12c Environment Variable	Corresponding 12c Environment Variables
NUM_TARGETS	1
TARGET_NAME1	TAGER_NAME
TARGET_TYPE1	TARGET_TYPE
JOB_NAME	CA_JOB_NAME
JOB_OWNER	CA_JOB_OWNER

Table 3–22 (Cont.) pre-12c/12c Corrective Action Environment Variable Mappings

Pre-12c Environment Variable	Corresponding 12c Environment Variables
JOB_STATUS	CA_JOB_STATUS
JOB_TYPE	CA_JOB_TYPE

Notification Type Mapping

Table 3–23 pre-12c/12c notif_type Mappings

notif_type (12c)	notif_type (Pre-12c)
NOTIF_NORMAL	1
NOTIF_REPEAT	4
NOTIF_DURATION	9
NOTIF_RETRY	2

3.2.1.2 Adding a Notification Method Based on a PL/SQL Procedure

A user-defined PL/SQL procedure can receive notifications for matching events, incidents and problem.

> **Note:** PL/SQL procedures used with pre-12c versions of Enterprise Manager will continue to work without modification. However, you should update the procedures to use the new signatures.

Complete the following four steps to define a notification method based on a PL/SQL procedure.

Step 1: Define the PL/SQL procedure.

The procedure must have one of the following signatures depending on the type of notification that will be received.

For Events:

PROCEDURE event_proc(event_msg IN gc$notif_event_msg)

For Incidents:

PROCEDURE incident_proc(incident_msg IN gc$notif_incident_msg)

For Problems:

PROCEDURE problem_proc(problem_msg IN gc$notif_problem_msg)

> **Note:** The notification method based on a PL/SQL procedure must be configured by an administrator with Super Administrator privileges before a user can select it while creating/editing a incident rule.

For more information on passing specific types of information to scripts or PL/SQL procedures, see the following sections:

"Passing Information to a PL/SQL Procedure" on page 3-35

"Passing Corrective Action Status Change Information" on page 3-50

"Passing Job Execution Status Information" on page 3-51

Step 2: Create the PL/SQL procedure on the Management Repository.

Create the PL/SQL procedure on the repository database using one of the following procedure specifications:

PROCEDURE event_proc(event_msg IN gc$notif_event_msg)

PROCEDURE incident_proc(incident_msg IN gc$notif_incident_msg)

PROCEDURE problem_proc(problem_msg IN gc$notif_problem_msg)

The PL/SQL procedure must be created on the repository database using the database account of the repository owner (such as SYSMAN)

If an error is encountered during the running of the procedure, the Notification System can be instructed to retry the sending of the notification to the procedure by raising a user-defined exception that uses the error code -20000. The procedure initially retried after one minute, then two minutes, then three minutes and so on, until the notification is a day old, at which point it will be purged.

Step 3: Register your PL/SQL procedure as a new notification method.

Log in as a Super Administrator. From the **Setup** menu, choose **Notifications-->Notification Methods** to access the Notification Methods page. From this page, you can define a new notification based on 'PL/SQL Procedure'. See "Adding a Notification Method Based on a PL/SQL Procedure" on page 3-33.

Make sure to use a fully qualified name that includes the schema owner, package name and procedure name. The procedure will be executed by the repository owner and so the repository owner must have execute permission on the procedure.

Create a notification method based on your PL/SQL procedure. The following information is required when defining the method:

- Name
- Description
- PL/SQL Procedure

You must enter a fully qualified procedure name (for example, OWNER.PKGNAME.PROCNAME) and ensure that the owner of the Management Repository has execute privilege on the procedure.

An example of the required information is shown in Example 3–9.

Example 3–9 PL/SQL Procedure Required Information

```
Name Open trouble ticket
Description Notification method to open a trouble ticket in the event
PLSQL Procedure ticket_sys.ticket_ops.open_ticket
```

Step 4: Assign the notification method to an incident rule.

You can edit an existing rule (or create a new incident rule). From the **Setup** menu, choose **Incidents-->Incident Rules**. The Incident Rules page displays. From here, you can add an action to a rule specifying the new PL/SQL procedure found under **Advanced Notification Method.**

There can be more than one PL/SQL-based method configured for your Enterprise Manager environment.

Information about how incident, event, and problem information is passed to the PLSQL procedure is covered in the next section.

Passing Information to a PL/SQL Procedure

Passing event, incident, and problem information (payload) to PL/SQL procedures allows you to customize automated responses to these conditions. All 3 types of notification payloads have a common element - gc$notif_msg_info. It provides generic information that applies to all types of notifications. In addition, each of the 3 payloads have one specific element that provides the payload specific to the given issue type.

gc$notif_event_msg (payload for event notifications)

gc$notif_event_msg contains two objects - event payload object and message information object.

Table 3–24 Event Notification Payload

Attribute	Datatype	Additional Information
event_payload	gc$notif_event_payload	Event notification payload. See gc$notif_event_payload type definition for detail.
msg_info	gc$notif_msg_info	Notification message. See gc$notif_msg_info definition for detail.

gc$notif_incident_msg (payload for incident notifications)

gc$notif_incident_msg type contains two objects - incident payload and message information. This object represents the delivery payload for Incident notification message, contains all data associated with Incident notification, and can be accessed by user's custom PL/SQL procedures.

Table 3–25 Incident Notification Payload

Attribute	Datatype	Additional Information
incident_payload	gc$notif_incident_payload	Incident notification payload. See gc$notif_incident_payload type definition for detail.
msg_info	gc$notif_msg_info	Envelope level notification information. See gc$notif_msg_info type definition for detail.

gc$notif_problem_msg (payload for problem notifications)

This object represents the delivery payload for Problem notification message, contains all data associated with problem notification, and can be accessed by a user's custom PL/SQL procedures.

Table 3–26 Problem Notification Payload

Attribute	Datatype	Additional Information
problem_payload	gc$notif_problem_payload	Problem notification payload. See gc$notif_problem_payload type definition for detail.

Table 3–26 (Cont.) Problem Notification Payload

Attribute	Datatype	Additional Information
msg_info	gc$notif_msg_info	Notification message. See gc$notif_msg_info type definition for detail.

gc$notif_msg_info (common for event/incident/problem payloads)

This object contains the generic notification information including notification_type, rule set and rule name, etc. for Event, Incident or Problem delivery payload.

Table 3–27 Event, Incident, Problem Common Payload

Attribute	Datatype	Description
notification_type	VARCHAR2(32)	Type of notification, can be one of the following values. GC$NOTIFICATION.NOTIF_NORMAL GC$NOTIFICATION.NOTIF_RETRY GC$NOTIFICATION.NOTIF_REPEAT GC$NOTIFICATION.NOTIF_DURATION GC$NOTIFICATION.NOTIF_CA GC$NOTIFICATION.NOTIF_RCA
repeat_count	NUMBER	Repeat notification count
ruleset_name	VARCHAR2(256)	Name of the rule set that triggered the notification
rule_name	VARCHAR2(256)	Name of the rule that triggered the notification
rule_owner	VARCAH2(256)	EM User who owns the rule set
message	VARCHAR2(4000)	Message about event/incident/problem.
message_url	VARCHAR2(4000)	Link to the Enterprise Manager console page that provides the details of the event/incident/problem.

gc$notif_event_payload (payload specific to event notifications)

This object represents the payload specific to event notifications.

Table 3–28 Common Payloads for Events, Incidents, and Problems

Attribute	Datatype	Additional Information
event_instance_guid	RAW(16)	Event instance global unique identifier.
event_sequence_guid	RAW(16)	Event sequence global unique identifier.
Target	gc$notif_target	Related Target Information object. See gc$notif_target type definition for detail.
Source	gc$notif_source	Related Source Information object, that is not a target. See gc$notif_source type definition for detail.
event_attrs	gc$notif_event_attr_array	The list of event specified attributes. See gc$notif_event_attr type definition for detail.

Table 3–28 (Cont.) Common Payloads for Events, Incidents, and Problems

Attribute	Datatype	Additional Information
corrective_action	gc$notif_corrective_action_job	Corrective action information, optionally populated when corrective action job execution has completed.
event_type	VARCHAR2(20)	Event type - example: Metric Alert.
event_name	VARCHAR2(512)	Event name.
event_msg	VARCHAR2(4000)	Event message.
reported_date	DATE	Event reported date.
Occurrence_date	DATE	Event occurrence date.
Severity	VARCHAR2(128)	Event Severity. It is the translated severity name.
severity_code	VARCHAR2(32)	Event Severity code. It is the internal severity name used in Enterprise Manager.
assoc_incident	gc$notif_issue_summary	Summary of associated incident. It is populated if the event is associated with an incident. See gc$notif_issue_summary type definition for detail
action_msg	VARCHAR2(4000)	Message describing the action to take for resolving the event.
rca_detail	VARCHAR2(4000)	Root cause analysis detail. The size of RCA details output is limited to 4000 characters long.
event_context_data	gc$notif_event_context_array	Event context data. See gc$notif_event_context type definition for detail.
categories	gc$category_string_array	List of categories that the event belongs to. Category is translated based on locale defined in OMS server. Notification system sends up to 10 categories.
category_codes	gc$category_string_array	Codes for the categories. The size of array is up to 10.

gc$notif_incident_payload (payload specific to incident notifications)

It contains the incident specific attributes, associated problem and ticket information.

Table 3–29 Incident Notification Payloads

Attribute	Datatype	Additional Information
incident_attrs	gc$notif_issue_attrs	Incident specific attributes. See gc$notif_issue_attrs type definition for detail.
assoc_event_count	NUMBER	The total number of events associated with this incident.
ticket_status	VARCHAR2(64)	The status of external Ticket, if it exists.
ticket_id	VARCHAR2(128)	The ID of external Ticket, if it exists.
ticket_url	VARCHAR2(4000)	The URL for external Ticket, if it exists.
assoc_problem	gc$notif_issue_summary	Summary of the problem, if it has an associated problem. See gc$notif_issue_summary type definition for detail.

Table 3–29 (Cont.) Incident Notification Payloads

Attribute	Datatype	Additional Information

gc$notif_problem_payload (payload specific to problems)

It contains problem specific attributes, key, Service Request(SR) and Bug information.

Table 3–30 Problem Payload

Attribute	Datatype	Additional Information
problem_attrs	gc$notif_issue_attrs	Problem specific attributes. See gc$notif_ issue_attrs type definition for detail.
problem_key	VARCHAR2(850)	Problem key if it is generated.
assoc_incident_count	NUMBER	Number of incidents associated with this problem.
sr_id	VARCHAR2(64)	Oracle Service Request Id, if it exists.
sr_url	VARCHAR2(4000)	URL for Oracle Service Request, if it exists.
bug_id	VARCHAR2(64)	Oracle Bug ID, if an associated bug exists.

gc$notif_issue_attrs (payload common to incidents and problems)

It provides common details for incident and problem. It contains the details such as id, severity, priority, status, categories, acknowledged by owner, and source information associated with.

Table 3–31 Payload Common to Incidents and Problems

Attribute	Datatype	Additional Information
Id	NUMBER(16)	ID of the incident or problem.
Severity	VARCHAR2(128)	Issue Severity. It is the translated.
severity_code	VARCHAR2(32)	Issue Severity Code.The possible values are defined in descending order of severity: GC$EVENT.FATAL GC$EVENT.CRITICAL GC$EVENT.WARNING GC$EVENT.MINOR_WARNING GC$EVENT.INFORMATIONAL GC$EVENT.CLEAR
priority	VARCHAR2(128)	Issue Priority. It is the translated priority name.

Table 3–31 (Cont.) Payload Common to Incidents and Problems

Attribute	Datatype	Additional Information
priority_code	VARCHAR2(32)	Issue Priority. It is the internal value defined in EM. The possible values are defined in descending order of priority: GC$EVENT.PRIORITY_URGENT GC$EVENT.PRIORITY_VERY_HIGH GC$EVENT.PRIORITY_HIGH GC$EVENT.PRIORITY_MEDIUM GC$EVENT.PRIORITY_LOW GC$EVENT.PRIORITY_NONE
status	VARCHAR2(32)	Status of Issue. The possible values are GC$EVENT.STATUS_NEW GC$EVENT.STATUS_CLOSED Any other user defined status.
escalation_level	NUMBER(1)	Escalation level of the issue, has a value between 0 to 5.
owner	VARCHAR(256)	Issue Owner. Set to NULL if no owner exists.
acknowledged_by_owner	NUMBER(1)	Set to 1, if this issue was acknowledged by owner.
creation_date	DATE	Issue creation date.
closed_date	DATE	Issue closed date, null if not closed.
categories	gc$category_string_array	List of categories that the event belongs to. Category is translated based on locale defined in OMS server. Notification system sends up to 10 categories.
category_codes	gc$category_string_array	Codes for the categories. Notification system sends up to 10 category codes.
source_info_arr	gc$notif_source_info_array	Array of source information associated with this issue. See $gcnotif_source_info type definition for detail.
last_modified_by	VARCHAR2(256)	Last modified by user.
last_updated_date	DATE	Last updated date.

gc$notif_issue_summary (common to event and incident payloads)

This object represents the associated incident summary in Event payload, or associated problem summary in Incident payload, respectively.

Table 3–32 Payload

Attribute	Datatype	Additional Information
id	NUMBER	Issue Id, either Incident Id or Problem Id.
severity	VARCHAR(128)	The severity level of an issue. It is translated severity name.

Table 3–32 (Cont.) Payload

Attribute	Datatype	Additional Information
severity_code	VARCHAR2(32)	Issue Severity Code, has one of the following values. GC$EVENT.FATAL GC$EVENT.CRITICAL GC$EVENT.WARNING GC$EVENT.MINOR_WARNING GC$EVENT.INFORMATIONAL GC$EVENT.CLEAR
priority	VARCHAR2(128)	Current priority. It is the translated priority name.
priority_code	VARCHAR2(32)	Issue priority code, has one of the following values. GC$EVENT.PRIORITY_URGENT GC$EVENT.PRIORITY_VERY_HIGH GC$EVENT.PRIORITY_HIGH GC$EVENT.PRIORITY_MEDIUM GC$EVENT.PRIORITY_LOW GC$EVENT.PRIORITY_NONE
status	VARCHAR2(64)	Status of issue. The possible values are GC$EVENT.STATUS_NEW GC$EVENT.STATUS_CLOSED GC$EVENT.WIP (work in progress) GC$EVENT.RESOLVED any other user defined status
escalation_level	VARCHAR2(2)	Issue escalation level range from 0 to 5, default 0.
owner	VARCHAR2(256)	Issue Owner. Set to NULL if no owner exists.
acknowledged_by_owner	NUMBER(1)	Set to 1, if this issue was acknowledged by owner.

gc$category_string_array

gc$category_string_array is an array of string containing the categories which event, incident or problem is associated with. Note that notification system delivers up to 10 categories.

gc$notif_event_context_array

gc$notif_event_context_array provides information about the additional diagnostic data that was captured at event detection time. Note that notification system delivers up to 200 elements from the captured event context.

CREATE OR REPLACE TYPE gc$notif_event_context_array AS VARRAY(200) OF gc$notif_event_context;

gc$notif_event_context

This object represents the detail of event context data which is additional contextual information that is captured by the source system at the point of event generation that

could be useful for diagnosis. The context for an event should consist of set of keys and values along with data type (Number or String only).

Table 3–33 Event Context Type

Attribute	Datatype	Additional Information
Name	VARCHAR2(256)	The event context name.
Type	NUMBER(1)	The data type of the value, which is stored (0) - for numeric data (1) - for string data.
Value	NUMBER	The numerical value.
string_value	VARCHAR2(4000)	The string value.

gc$notif_corrective_action_job

This object provides information about the execution of a corrective action job. Note that the corrective actions are supported for metric alert and target availability events only.

Table 3–34 Corrective Action Job-Specific Attributes

Attribute	Datatype	Additional Information
job_guid	RAW(16)	Corrective Action Job global unique identifier.
job_name	VARCHAR2(128)	The value will be the name of the Corrective Action. It applies to Metric Alert and Target Availability Events.
job_owner	VARCHAR2(256)	Corrective action job owner.
job_type	VARCHAR2(256)	Corrective action job type.
job_status	VARCHAR2(64)	Corrective action job execution status.
job_status_code	NUMBER	Corrective action job execution status code. It is the internal value defined in EM.
job_step_output	VARCHAR2(4000)	The value will be the text output from the Corrective Action execution. This will be truncated to last 4000 characters.
job_execution_guid	RAW(16)	Corrective Action Job execution global unique identifier.
job_state_change_guid	RAW(16)	Corrective Action Job change global unique identifier.
occurred_date	DATE	Corrective action job occurred date.

gc$notif_source_info_array

It is used providing access to the multiple sources that an incident or a problem could be related to. NOTE: The notification system delivers up to 200 sources associated with an incident or a problem.

CREATE OR REPLACE TYPE gc$notif_source_info_array AS VARRAY(200) OF gc$notif_source_info;

gc$notif_source_info

Notification Source Information which is used for referencing Source information containing either target or source, or both.

Table 3–35 Source Information Type

Attribute	Datatype	Additional Information
target	gc$notif_target	It is populated when the event is related to a target. See gc$notif_target type definition for detail.
Source	gc$notif_source	It is populated when the event is related to a (non-target) source. See gc$notif_source type definition for detail.

gc$notif_source

Source object is used for referencing source objects other than a job target.

Table 3–36 Payload

Attribute	Datatype	Additional Information
source_guid	RAW(16)	Source's global unique identifier.
source_type	VARCHAR2(120)	Type of the Source object, e.g., TARGET, JOB, TEMPLATE, etc.
source_name	VARCHAR2(256)	Source Object Name.
source_owner	VARCHAR2(256)	Owner of the Source object.
source_sub_type	VARCHAR2(256)	Sub-type of the Source object, for example, within the TARGET these would be the target types like Host, Database etc.
source_url	VARCHAR2(4000)	Source's event console URL.

gc$notif_target

Target information object is used for providing target information.

Table 3–37 Target Information

Attribute	Datatype	Additional Information
target_guid	RAW(16)	Target's global unique identifier.
target_name	VARCHAR2(256)	Name of target.
target_owner	VARCHAR2(256)	Owner of target.
target_lifecycle_status	VARCHAR2(1024)	Life Cycle Status of the target.
target_version	VARCHAR2(64)	Target Version of the target.
target_type	VARCHAR2(128)	Type of a target.
target_timezone	VARCHAR2(64)	Target's regional time zone.
host_name	VARCHAR2(256)	The name of the host on which the target is deployed upon.
target_url	VARCHAR2(4000)	Target's EM Console url.

Table 3–37 (Cont.) Target Information

Attribute	Datatype	Additional Information
udtp_array	gc$notif_udtp_array	The list of user defined target properties. It is populated for events that are associated with a target. It is populated for incidents and problems, when they are associated with a single source (gc$notif_source_info).

gc$notif_udtp_array

It is array of gc$notif_udtp type and size is up to 20.

CREATE OR REPLACE TYPE gc$notif_udtp_array AS VARRAY(20) OF gc$notif_udtp;

gc$notif_udtp

User defined Target Properties (UDTP) is used for referencing User defined target properties. UDTP should consist of set of property key name and property value.

Table 3–38 Payload

Attribute	Datatype	Additional Information
name	VARCHAR2(64),	The name of property.
value	VARCHAR2(1024)	Property value.
label	VARCHAR(256)	Property label.
nls_id	VARCHAR(64)	Property nls id

gc$notif_event_attr_array

Array of gc$notif_event_attr is used for referencing Event specific attributes, and its size is up to 35.

CREATE OR REPLACE TYPE gc$notif_event_attr_array AS VARRAY(35) OF gc$notif_event_attr;

gc$notif_event_attr

It is used for referencing Event type specific attributes.

Table 3–39 Event Attribute Type

Attribute	Datatype	Additional Information
name	VARCHAR2(64)	The internal name of event type specific attribute.
value	VARCHAR2(4000)	value.
nls_value	VARCHAR2(4000)	Translated value for the attribute.

3.2.1.2.1 Migrating Pre-12c PL/SQL Advanced Notification Methods

Pre-12c notifications map to event notifications in Enterprise Manager 12c. Event types metric_alert, target_availability and job_status_alert correspond to the pre-12c notification functionality. Note that policy violations functionality is removed for this release.

This section describes the mapping between Enterprise Manager 12c PL/SQL notification payload to the pre-12c PL/SQL notification payload. You can use this information for updating the existing pre-12c PL/SQL user callback procedures to use the 12c PL/SQL notification payload.

Please note that Policy Violations are no longer supported in the 12c release.

Mapping for MGMT_NOTIFY_SEVERITY

When event type is metric_alert

Use the following map when gc$notif_event_payload .event_type='metric_alert'.

Table 3–40 Metric Alert Mapping

MGMT_NOTIFY_SEVERITY	12c Notification Payload
TARGET_NAME	gc$notif_target.target_name
TARGET_TYPE	gc$notif_target.target_type
TIMEZONE	gc$notif_target.target_timezone
HOST_NAME	gc$notif_target.host_name
MERTIC_NAME	gc$notif_event_attr.value where its name=' metric_group' in gc$notif_event_attr_array.
METRIC_DESCRIPTION	gc$notif_event_attr.value where its name=' metric_description' in gc$notif_event_attr_ array.
METRIC_COLUMN	gc$notif_event_attr.value where its name=' metric_column' in gc$notif_event_attr_array.
METRIC_VALUE	gc$notif_event_attr.value where its name=' value' in gc$notif_event_attr_array.
KEY_VALUE	It is applied for multiple keys based metric when value of gc$notif_event_ attr.name='num_keys' is not null and is greater than 0 in gc$notif_event_attr_array. See detail descriptions below.
KEY_VALUE_NAME	It is applied for multiple keys based metric when value of gc$notif_event_ attr.name='num_keys' is not null and is greater than 0 in gc$notif_event_attr_array. See detail descriptions below.
KEY_VALUE_GUID	gc$notif_event_attr.value where its name='key_ value' in gc$notif_event_attr_ array.
CTXT_LIST	gc$notif_event_context_array
COLLECTION_ TIMESTAMP	gc$notif_event_payload. reported_date
SEVERITY_CODE	Derive from gc$notif_event_ payload.severity_code, see section 1.2.1.1.2 for the mapping over the value.
MESSAGE	gc$notif_msg_info.message
SEVERITY_GUID	gc$notif_event_attr.value where its name=' severity_guid' in gc$notif_event_attr_array.
METRIC_GUID	gc$notif_event_attr.value where its name=' metric_guid_id' in gc$notif_event_attr_array.

Table 3–40 (Cont.) Metric Alert Mapping

MGMT_NOTIFY_SEVERITY	12c Notification Payload
TARGET_GUID	gc$notif_target.target_guid
RULE_OWNER	gc$notif_msg_info.rule_owner
RULE_NAME	gc$notif_msg_info.ruleset_name

The following example shows you how to obtain similar pre-12c KEY_VALUE and KEY_VALUE_NAME from an Enterprise Manager 12c notification payload.

First, check its gc$notif_event_attr.value where its name='num_keys' in gc$notif_event_attr_array.

If it is null or its value is equal to 0, then KEY_VALUE and KEY_VALUE_NAME are null for this event.

If it is not null and value is equal to 1, then it is single key metric.

KEY_VALUE_NAME= value of gc$notif_event_attr where name='key_column_1' in gc$notif_event_attr_array.

KEY_VALUE = value of gc$notif_event_attr where name='key_value' in gc$notif_event_attr_array.

For example: METRIC= Filesystem Space Available (%) num_key=1

KEY_VALUE_NAME= Mount Point

KEY_VALUE= /

If it is not null and value is greater than 1, then it multiple keys metric.

KEY_VALUE_NAME = value of gc$notif_event_attr where name='key_column_1' + ";" +

value of gc$notif_event_attr where name='key_column_2' + ";" +

. …. (up to where name ='key_column_num_key')

KEY_VALUE = value of gc$notif_event_attr where name='key_column_1_value' + ";" +

value of gc$notif_event_attr where name='key_column_2_value' + ";" +

. …. (up to where name='key_column_num_key_value')

The ";" is separator between names or values.

For example: METRIC= Program's Max CPU Utilization (%) which num_key=2

KEY_VALUE_NAME= Program Name;Owner

KEY_VALUE= loadcpu;userId

Severity Code mapping from 12c to pre-12c when the event type is metric_alert

Table 3–41 Severity Code Mapping

12c Severity Code	Pre-12c Severity Code
GC_EVENT_RECEIVER.FATAL 32	MGMT_GLOBAL.G_SEVERITY_CRITICAL 25
GC_EVENT_RECEIVER.CRITICAL 16	MGMT_GLOBAL.G_SEVERITY_CRITICAL 25

Table 3–41 (Cont.) Severity Code Mapping

12c Severity Code	Pre-12c Severity Code
GC_EVENT_RECEIVER.WARNING 8	MGMT_GLOBAL.G_SEVERITY_WARNING 20
GC_EVENT_RECEIVER.CLEAR 0	MGMT_GLOBAL.G_SEVERITY_CLEAR 15

When event type is target_availability

Use the following map when gc$notif_event_payload .event_type='target_availability'.

Table 3–42 Target Availability Mapping

MGMT_NOTIFY_SEVERITY	12c Notification Payload
TARGET_NAME	gc$notif_target.target_name
TARGET_TYPE	gc$notif_target.target_type
TIMEZONE	gc$notif_target.target_timezone
HOST_NAME	gc$notif_target.host_name
MERTIC_NAME	Use fixed value "Response".
METRIC_DESCRIPTION	NULL
METRIC_COLUMN	Use fixed value "Status".
METRIC_VALUE	gc$notif_event_attr.value where its name='target_status' in gc$notif_event_attr_ array.
KEY_VALUE	NULL
KEY_VALUE_NAME	NULL
KEY_VALUE_GUID	NULL
CTXT_LIST	gc$notif_event_context_array
COLLECTION_TIMESTAMP	gc$notif_event_payload. reported_date
SEVERITY_CODE	gc$notif_event_attr.value where its name=' avail_severity' in gc$notif_event_attr_array.
MESSAGE	gc$notif_msg_info.message
SEVERITY_GUID	gc$notif_event_attr.value where its name=' severity_guid' in gc$notif_event_attr_array.
METRIC_GUID	gc$notif_event_attr.value where its name=' metric_guid_id' in gc$notif_event_attr_array.
TARGET_GUID	gc$notif_target.target_guid
RULE_OWNER	gc$notif_msg_info.rule_owner
RULE_NAME	gc$notif_msg_info.ruleset_name

Mapping for MGMT_NOTIFY_JOB

Use the following map when gc$notif_event_payload .event_type=job_status_change'.

Table 3–43 Job Status Change Mapping

MGMT_NOTIFY_JOB	12c Notification Payload
JOB_NAME	gc$notif_source.source_name
JOB_OWNER	gc$notif_source.source_owner
JOB_TYPE	gc$notif_source.source_sub_type
JOB_STATUS	gc$notif_event_attr.value where its name=' execution_status_code' in gc$notif_event_ attr_array.
STATE_CHANGE_GUID	gc$notif_event_attr.value where its name=' state_change_guid' in gc$notif_event_attr_ array.
JOB_GUID	gc$notif_source.source_guid
EXECUTION_ID	gc$notif_event_attr.value where its name=' execution_id' in gc$notif_event_attr_array.
TARGETS	gc$notif_target.target_name, gc$notif_ target.target_type
RULE_OWNER	gc$notif_msg_info.rule_owner
RULE_NAME	gc$notif_msg_info.ruleset_name
OCCURRED_DATE	gc$notif_event_payload. reported_date

Mapping for MGMT_NOTIFY_CORRECTIVE_ACTION

Note that corrective action related payload is populated when gc$notif_msg_ info.notification_type is set to NOTIF_CA.

For mapping the following attributes, use the mapping information provided for MGMT_NOTIFY_SEVERITY object Table 3–40, " Metric Alert Mapping"

MERTIC_NAME

METRIC_COLUMN

METRIC_VALUE

KEY_VALUE

KEY_VALUE_NAME

KEY_VALUE_GUID

CTXT_LIST

RULE_OWNER

RULE_NAME

OCCURRED_DATE

For mapping the job related attributes in MGMT_NOTIFY_CORRECTIVE_ACTION object, use the following map.

Table 3–44 Corrective Action Mapping

MGMT_NOTIFY_CORRECTIVE_ACTION	12c Notification Payload
JOB_NAME	gc$ notif_corrective_action_job.job_name

Table 3–44 (Cont.) Corrective Action Mapping

MGMT_NOTIFY_CORRECTIVE_ACTION	12c Notification Payload
JOB_OWNER	gc$ notif_corrective_action_job.job_owner
JOB_TYPE	gc$ notif_corrective_action_job.job_type
JOB_STATUS	gc$ notif_corrective_action_job.status_code
STATE_CHANGE_GUID	gc$ notif_corrective_action_job. job_state_ change_guid
JOB_GUID	gc$ notif_corrective_action_job. job _guid
EXECUTION_ID	gc$ notif_corrective_action_job. job_ execution_guid
OCCURRED_DATE	gc$ notif_corrective_action_job.occurred_ date
TARGETS	There can be at most one target. Use the values from gc$notif_target.target_name, gc$notif_target.target_type for the associated target.

3.2.1.3 Adding a Notification Method Based on an SNMP Trap

Enterprise Manager supports integration with third-party management tools through the SNMP. For example, you can use SNMP to notify a third-party application that a selected metric has exceeded its threshold.

The trap is an SNMP Version 1 trap and is described by the MIB definition shown at the end of this chapter. See "Management Information Base (MIB)" on page 3-55.

For more comprehensive configuration information, see the documentation specific to your platform; SNMP configuration differs from platform to platform.

> **Note:** Notification methods based on SNMP traps must be configured by an administrator with Super Administrator privileges before any user can then choose to select one or more of these SNMP trap methods while creating/editing a incident rule.

Step 1: Define a new notification method based on an SNMP trap.

Log in to Enterprise Manager as a Super Administrator. Click Setup and then click Notification Method from the vertical navigation bar to access the Notification Methods page. From this page you can add a new method based on an SNMP trap.

You must provide the name of the host (machine) on which the SNMP master agent is running and other details as shown in the following example. In Example 3–10, the SNMP host will receive your SNMP traps.

Example 3–10 SNMP Trap Required Information

```
Name HP OpenView Console
Description Notification method to send trap to HP OpenView console
SNMP Trap Host Name machine1.us.oracle.com
SNMP Host Port 162
SNMP Community public
```

```
This SNMP host will receive your SNMP traps.
```

Note: A Test Trap button exists for you to test your setup.

Metric severity information will be passed as a series of variables in the SNMP trap.

An example SNMP Trap is shown in Example 3–11. Each piece of information is sent as a variable embedded in the SNMP Trap.

Example 3–11 SNMP Trap

```
*************V1 TRAP***[3]*****************
Community : public
Enterprise :1.3.6.1.4.1.111.15.2
Generic :6
Specific :3
TimeStamp :48960
Agent adress :10.228.163.210
1.3.6.1.4.1.111.15.3.1.1.2.1: NOTIF_NORMAL
1.3.6.1.4.1.111.15.3.1.1.3.1: Memory Utilization is 98.343%, crossed warning (90)
or critical (95) threshold.
1.3.6.1.4.1.111.15.3.1.1.4.1:
https://adc6140830.us.oracle.com:15430/em/redirect?pageType=sdk-core-event-console
-detailEvent&
issueID=AADFAE01AD3A95E3E040E40AD2A32041
1.3.6.1.4.1.111.15.3.1.1.5.1: Critical
1.3.6.1.4.1.111.15.3.1.1.6.1: CRITICAL
1.3.6.1.4.1.111.15.3.1.1.7.1: 0
1.3.6.1.4.1.111.15.3.1.1.10.1: Aug 19, 2011 4:50:18 PM PDT
1.3.6.1.4.1.111.15.3.1.1.11.1: Capacity
1.3.6.1.4.1.111.15.3.1.1.12.1: Capacity
1.3.6.1.4.1.111.15.3.1.1.13.1: Metric Alert
1.3.6.1.4.1.111.15.3.1.1.14.1: Load:memUsedPct
1.3.6.1.4.1.111.15.3.1.1.15.1: 15
1.3.6.1.4.1.111.15.3.1.1.16.1:
1.3.6.1.4.1.111.15.3.1.1.17.1: No
1.3.6.1.4.1.111.15.3.1.1.18.1: New
1.3.6.1.4.1.111.15.3.1.1.19.1: None
1.3.6.1.4.1.111.15.3.1.1.20.1: 0
1.3.6.1.4.1.111.15.3.1.1.21.1: adc6140830.us.oracle.com
1.3.6.1.4.1.111.15.3.1.1.22.1:
https://adc6140830.us.oracle.com:15430/em/redirect?pageType=TARGET_
HOMEPAGE&targetName=adc6140
830.us.oracle.com&targetType=host
1.3.6.1.4.1.111.15.3.1.1.23.1: Host
1.3.6.1.4.1.111.15.3.1.1.24.1: adc6140830.us.oracle.com
1.3.6.1.4.1.111.15.3.1.1.25.1: SYSMAN
1.3.6.1.4.1.111.15.3.1.1.27.1: 4.8.0.0.0
1.3.6.1.4.1.111.15.3.1.1.28.1:
1.3.6.1.4.1.111.15.3.1.1.39.1: snmp ruleset
1.3.6.1.4.1.111.15.3.1.1.40.1: snmp ruleset,snmp rule
1.3.6.1.4.1.111.15.3.1.1.41.1: SYSMAN
1.3.6.1.4.1.111.15.3.1.1.42.1: AADFAE01AD3A95E3E040E40AD2A32041
1.3.6.1.4.1.111.15.3.1.1.61.1: METRIC_GUID=86821B5F0CE858D6E4A7F7390E88B73C
1.3.6.1.4.1.111.15.3.1.1.62.1: SEVERITY_GUID=AADFADD572DE08E2E040E40AD2A3202C
1.3.6.1.4.1.111.15.3.1.1.63.1: CYCLE_GUID=AADFAE01AD3695E3E040E40AD2A32041
1.3.6.1.4.1.111.15.3.1.1.64.1: COLL_NAME=LoadLinux
1.3.6.1.4.1.111.15.3.1.1.65.1: METRIC_GROUP=Load
1.3.6.1.4.1.111.15.3.1.1.66.1: METRIC_COLUMN=Memory Utilization (%)
```

```
1.3.6.1.4.1.111.15.3.1.1.67.1: METRIC_DESCRIPTION=
1.3.6.1.4.1.111.15.3.1.1.68.1: VALUE=98.343
1.3.6.1.4.1.111.15.3.1.1.69.1: KEY_VALUE=
 1.3.6.1.4.1.111.15.3.1.1.84.1: NUM_KEYS=0
```

Step 2: Assign the notification method to a rule.

You can edit an existing rule (or create a new incident rule), then add an action to the rule that subscribes to the advanced notification method.

3.3 Passing Corrective Action Status Change Information

Passing corrective action status change attributes (such as new status, job name, job type, or rule owner) to PL/SQL procedures or OS commands/scripts allows you to customize automated responses to status changes. For example, you many want to call an OS script to open a trouble ticket for an in-house support trouble ticket system if a critical corrective action fails to run. In this case, you will want to pass status (for example, Problems or Aborted) to the script to open a trouble ticket and escalate the problem.

3.3.1 Passing Corrective Action Execution Status to an OS Command or Script

The notification system passes information to an OS script or executable via system environment variables. Conventions used to access environmental variables vary depending on the operating system:

- UNIX: $ENV_VARIABLE
- MS Windows: %ENV_VARIABLE%

The notification system sets the following environment variables before calling the script. The notification system will set the environment variable $NOTIF_TYPE = NOTIF_CA for Corrective Action Execution. The script can then use any or all of these variables within the logic of the script.

Following table lists the environment variables for corrective action, they are populated when a corrective action is completed for an event.

Table 3–45 Corrective Action Environment Variables

Environment Variable	Description
CA_JOB_STATUS	Corrective action job execution status.
CA_JOB_NAME	Name of the Corrective Action.
CA_JOB_OWNER	Owner of Corrective Action.
CA_JOB_STEP_OUTPUT	The value will be the text output from the Corrective Action execution.
CA_JOB_TYPE	Corrective Action Job type

3.3.2 Passing Corrective Action Execution Status to a PLSQL Procedure

The notification system passes corrective action status change information to PL/SQL procedure - PROCEDURE p(event_msg IN gc$notif_event_msg). The instance gc$notif_corrective_action_job object is defined in event_msg.event_payload. corrective_action if event_msg. msg_info. notification_type is equal to GC$NOTIFICATIONNOTIF_CA. When a corrective action executes, the notification system calls the PL/SQL procedure associated with the incident rule and passes the

populated object to the procedure. The procedure is then able to access the fields of the object that has been passed to it. See Table 3–34, " Corrective Action Job-Specific Attributes" for details.

The following status codes are possible values for the job_status field of the MGMT_NOTIFY_CORRECTIVE_ACTION object.

Table 3–46 Corrective Action Status Codes

Name	Datatype	Value
SCHEDULED_STATUS	NUMBER(2)	1
EXECUTING_STATUS	NUMBER(2)	2
ABORTED_STATUS	NUMBER(2)	3
FAILED_STATUS	NUMBER(2)	4
COMPLETED_STATUS	NUMBER(2)	5
SUSPENDED_STATUS	NUMBER(2)	6
AGENTDOWN_STATUS	NUMBER(2)	7
STOPPED_STATUS	NUMBER(2)	8
SUSPENDED_LOCK_STATUS	NUMBER(2)	9
SUSPENDED_EVENT_STATUS	NUMBER(2)	10
SUSPENDED_BLACKOUT_STATUS	NUMBER(2)	11
STOP_PENDING_STATUS	NUMBER(2)	12
SUSPEND_PENDING_STATUS	NUMBER(2)	13
INACTIVE_STATUS	NUMBER(2)	14
QUEUED_STATUS	NUMBER(2)	15
FAILED_RETRIED_STATUS	NUMBER(2)	16
WAITING_STATUS	NUMBER(2)	17
SKIPPED_STATUS	NUMBER(2)	18
REASSIGNED_STATUS	NUMBER(2)	20

3.4 Passing Job Execution Status Information

Passing job status change attributes (such as new status, job name, job type, or rule owner) to PL/SQL procedures or OS commands/scripts allows you to customize automated responses to status changes. For example, you many want to call an OS script to open a trouble ticket for an in-house support trouble ticket system if a critical job fails to run. In this case you will want to pass status (for example, Problems or Aborted) to the script to open a trouble ticket and escalate the problem. The job execution status information is one of event type - job_status_change event, and its content is in OS command and PL/SQL payload as described in Section 3.2.1.1, "Adding a Notification Method based on an OS Command or Script" and Section 3.2.1.2, "Adding a Notification Method Based on a PL/SQL Procedure".

3.4.1 Passing Job Execution Status to a PL/SQL Procedure

The notification system passes job status change information to a PL/SQL procedure via the event_msg.event_payload object where event_type is equal to job_status_change. An instance of this object is created for every status change. When a job

changes status, the notification system calls the PL/SQL p(event_msg IN `gc$notif_event_msg)` procedure associated with the incident rule and passes the populated object to the procedure. The procedure is then able to access the fields of the event_msg.event_payload object that has been passed to it.

Table 3–47 lists all corrective action status change attributes that can be passed:

Table 3–47 Job Status Attributes

Attribute	Datatype	Additional Information
event_msg.event_payload.source.source_name	VARCHAR2(128)	The job name.
event_msg.event_payload.source.source_owner	VARCHAR2(256)	The owner of the job.
event_msg.event_payload.source.source_sub_type	VARCHAR2(32)	The type of the job.
event_msg.event_payload. event_attrs(i).value where event_attrs(i).name=' execution_status'	NUMBER	The new status of the job.
event_msg.event_payload. event_attrs(i).value where event_attrs(i).name='state_change_guid'	RAW(16)	The GUID of the state change record.
event_msg.event_payload.source.source_guid	RAW(16)	The unique id of the job.
event_msg target.event_payload. event_attrs(i).value where event_attrs(i).name=' execution_id'	RAW(16)	The unique id of the execution.
event_msg.event_payload.target	gc$notif_target	Target Information object..
event_msg.msg_info.rule_owner	VARCHAR2(64)	The name of the notification rule that cause the notification to be sent.
event_msg.msg_info.rule_name	VARCHAR2(132)	The owner of the notification rule that cause the notification to be sent.
event_msg.event_payload. reported_date	DATE	The time and date when the status change happened.

When a job status change occurs for the job, the notification system creates an instance of the event_msg.event_payload. event_attrs(i).value where event_attrs(i).name=' execution_status' object and populates it with values from the status change. The following status codes have been defined as constants in the MGMT_JOBS package and can be used to determine the type of status in the job_status field of the event_msg.event_payload. event_attrs(i).value where event_attrs(i).name=' execution_status' object.

Table 3–48 Job Status Codes

Name	Datatype	Value
SCHEDULED_STATUS	NUMBER(2)	1
EXECUTING_STATUS	NUMBER(2)	2
ABORTED_STATUS	NUMBER(2)	3
FAILED_STATUS	NUMBER(2)	4
COMPLETED_STATUS	NUMBER(2)	5
SUSPENDED_STATUS	NUMBER(2)	6
AGENTDOWN_STATUS	NUMBER(2)	7
STOPPED_STATUS	NUMBER(2)	8
SUSPENDED_LOCK_STATUS	NUMBER(2)	9
SUSPENDED_EVENT_STATUS	NUMBER(2)	10
SUSPENDED_BLACKOUT_STATUS	NUMBER(2)	11
STOP_PENDING_STATUS	NUMBER(2)	12
SUSPEND_PENDING_STATUS	NUMBER(2)	13
INACTIVE_STATUS	NUMBER(2)	14
QUEUED_STATUS	NUMBER(2)	15
FAILED_RETRIED_STATUS	NUMBER(2)	16
WAITING_STATUS	NUMBER(2)	17
SKIPPED_STATUS	NUMBER(2)	18
REASSIGNED_STATUS	NUMBER(2)	20

Example 3–12 PL/SQL Procedure Using a Status Code (Job)

```
CREATE  TABLE job_log (jobid RAW(16), status_code NUMBER(2), occured DATE);

CREATE OR REPLACE PROCEDURE LOG_JOB_STATUS_CHANGE(event_msg IN GC$NOTIF_EVENT_MSG)
IS
  l_attrs gc$notif_event_attr_array;
  exec_status_code NUMBER(2) := NULL;
  occured_date DATE := NULL;
  job_guid RAW(16) := NULL;

BEGIN
  IF event_msg.event_payload.event_type = 'job_status_change'
  THEN
    l_attrs := event_msg.event_payload.event_attrs;
    IF l_attrs IS NOT NULL
    THEN
      FOR i IN 1..l_attrs.COUNT
      LOOP
        IF l_attrs(i).name = 'exec_status_code'
        THEN
          exec_status_code := TO_NUMBER(l_attrs(i).value);
        END IF;
      END LOOP;
    END IF;

    occured_date := event_msg.event_payload.reported_date;
```

```
    job_guid := event_msg.event_payload.source.source_guid;
    -- Log all jobs' status
    BEGIN
      INSERT INTO job_log (jobid, status_code, occured)
      VALUES (job_guid, exec_status_code, occured_date);
    EXCEPTION
    WHEN OTHERS
    THEN
      -- If there are any problems then get the notification retried
      RAISE_APPLICATION_ERROR(-20000, 'Please retry');
    END;
    COMMIT;

  ELSE
    null; -- it is not a job_status_change event, ignore
  END IF;
END LOG_JOB_STATUS_CHANGE;
/
```

3.4.2 Passing Job Execution Status to an OS Command or Script

The notification system passes job execution status information to an OS script or executable via system environment variables. Conventions used to access environmental variables vary depending on the operating system:

- UNIX: $ENV_VARIABLE
- MS Windows: %ENV_VARIABLE%

The notification system sets the following environment variables before calling the script. The script can then use any or all of these variables within the logic of the script.

Table 3–49 Environment Variables

Environment Variable	Description
SOURCE_OBJ_NAME	The name of the job.
SOURCE_OBJ_OWNE	The owner of the job.
SOURCE_OBJ_SUB_TYPE	The type of job.
EXEC_STATUS_CODE	The job status.
EVENT_REPORTED_ TIME	Time when the severity occurred.
TARGET_NAME	The name of the target.
TARGET_TYPE	The type of the target.
RULE_NAME	Name of the notification rule that resulted in the severity.
RULE_OWNER	Name of the Enterprise Manager administrator who owns the notification rule.

3.5 Passing User-Defined Target Properties to Notification Methods

Enterprise Manager allows you to define target properties (accessed from the target home page) that can be used to store environmental or usage context information specific to that target. Target property values are passed to custom notification methods where they can be processed using conditional logic or simply passed as additional alert information to third-party devices, such as ticketing systems. By

default, Enterprise Manager passes all defined target properties to notification methods.

Note: Target properties are not passed to notification methods when short e-mail format is used.

Figure 3–4 Host Target Properties

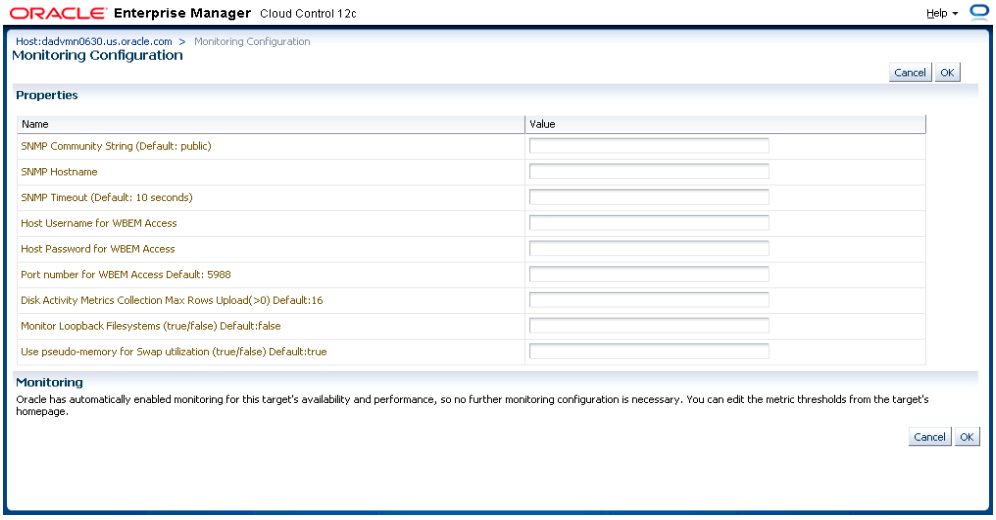

3.6 Management Information Base (MIB)

Enterprise Manager Cloud Control can send SNMP Traps to third-party, SNMP-enabled applications. Details of the trap contents can be obtained from the management information base (MIB) variables. The following sections discuss Enterprise Manager MIB variables in detail.

3.6.1 About MIBs

A MIB is a text file, written in ASN.1 notation, which describes the variables containing the information that SNMP can access. The variables described in a MIB, which are also called MIB objects, are the items that can be monitored using SNMP. There is one MIB for each element being monitored. Each monolithic or subagent consults its respective MIB in order to learn the variables it can retrieve and their characteristics. The encapsulation of this information in the MIB is what enables master agents to register new subagents dynamically — everything the master agent needs to know about the subagent is contained in its MIB. The management framework and management applications also consult these MIBs for the same purpose. MIBs can be either standard (also called public) or proprietary (also called private or vendor).

The actual values of the variables are not part of the MIB, but are retrieved through a platform-dependent process called "instrumentation". The concept of the MIB is very important because all SNMP communications refer to one or more MIB objects. What is transmitted to the framework is, essentially, MIB variables and their current values.

3.6.2 Reading the MIB Variable Descriptions

This section covers the format used to describe MIB variables. Note that the STATUS element of SNMP MIB definition, Version 2, is not included in these MIB variable descriptions. Since Oracle has implemented all MIB variables as CURRENT, this value does not vary.

3.6.2.1 Variable Name

Syntax
Maps to the SYNTAX element of SNMP MIB definition, Version 2.

Max-Access
Maps to the MAX-ACCESS element of SNMP MIB definition, Version 2.

Status
Maps to the STATUS element of SNMP MIB definition, Version 2.

Explanation
Describes the function, use and precise derivation of the variable. (For example, a variable might be derived from a particular configuration file parameter or performance table field.) When appropriate, incorporates the DESCRIPTION part of the MIB definition, Version 2.

Typical Range
Describes the typical, rather than theoretical, range of the variable. For example, while integer values for many MIB variables can theoretically range up to 4294967295, a typical range in an actual installation will vary to a lesser extent. On the other hand, some variable values for a large database can actually exceed this "theoretical" limit (a "wraparound"). Specifying that a variable value typically ranges from 0 to 1,000 or 1,000 to 3 billion will help the third-party developer to develop the most useful graphical display for the variable.

Significance
Describes the significance of the variable when monitoring a typical installation. Alternative ratings are Very Important, Important, Less Important, or Not Normally Used. Clearly, the DBA will want to monitor some variables more closely than others. However, which variables fall into this category can vary from installation to installation, depending on the application, the size of the database, and on the DBA's objectives. Nevertheless, assessing a variable's significance relative to the other variables in the MIB can help third-party developers focus their efforts on those variables of most interest to the most DBAs.

Related Variables
Lists other variables in this MIB, or other MIBs implemented by Oracle, that relate in some way to this variable. For example, the value of this variable might derive from that of another MIB variable. Or perhaps the value of this variable varies inversely to that of another variable. Knowing this information, third-party developers can develop useful graphic displays of related MIB variables.

Suggested Presentation
Suggests how this variable can be presented most usefully to the DBA using the management application: as a simple value, as a gauge, or as an alarm, for example.

3.6.2.2 MIB Definition

You can find the SNMP MIB file at the following location:

$ORACLE_HOME/network/doc/omstrap.v1

The file omstrap.v1 is the OMS MIB.

3.7 Troubleshooting Notifications

To function properly, the notification system relies on various components of Enterprise Manager and your IT infrastructure. For this reason, there can be many causes of notification failure. The following guidelines and suggestions can help you isolate potential problems with the notification system.

3.7.1 General Setup

The first step in diagnosing notification issues is to ensure that you have properly configured and defined your notification environment.

OS Command, PL/SQL and SNMP Trap Notifications

Make sure all OS Command, PLSQL and SNMP Trap Notification Methods are valid by clicking the Test button. This will send a test notification and show any problems the OMS has in contacting the method. Make sure that your method was called, for example, if the OS Command notification is supposed to write information to a log file, check that it has written information to its log file.

E-mail Notifications

- Make sure an e-mail gateway is set up under the Notification Methods page of Setup. The Sender's e-mail address should be valid. Clicking the Test button will send an e-mail to the Sender's e-mail address. Make sure this e-mail is received. Note that the Test button ignores any Notification Schedule.
- Make sure an e-mail address is set up. Clicking the Test button will send an e-mail to specified address and you should make sure this e-mail is received. Note that the Test button ignores any Notification Schedule.
- Make sure an e-mail schedule is defined. No e-mails will be sent unless a Notification Schedule has been defined.
- Make sure a incident rule is defined that matches the states you are interested and make sure e-mail and notification methods are assigned to the rule.

3.7.2 Notification System Errors

For any alerts involving problems with notifications, check the following for notification errors.

- Any serious errors in the Notification System are logged as system errors in the MGMT_SYSTEM_ERROR_LOG table. From the **Setup** menu, choose **Management Services and Repository** to view these errors.
- Check for any delivery errors. You can view them from Incident Manager. From the Enterprise menu, choose Monitoring and then Incident Manager. To view the alert history, from the Enterprise menu, choose Monitoring and then Alert History. Click on the Details icon for more information about the alert. The details will give the reason why the notification was not delivered. Delivery errors are stored in MGMT_NOTIFICATION_LOG with the DELIVERED column set to 'N'.
- Severities will not be displayed in the Enterprise Manager console if no metric values have been loaded for the metric associated with the severity.

3.7.3 Notification System Trace Messages

The Notification System can produce trace messages in sysman/log/emoms.trc file.

Tracing is configured by setting the *log4j.em.notification* property flag using the `emctl set property` command. You can set the trace level to INFO, WARN, DEBUG. For example,

```
./emctl set property -sysman_pwd your_sysman_password -name log4j.em.notification
-value DEBUG
```

Trace messages contain the string "em.notification". If you are working in a UNIX environment, you can search for messages in the emoms.trc and emoms_pbs.trc files using the `grep` command. For example,

```
grep em.notification emoms.trc emoms_pbs.trc
```

What to look for in the trace file.

The following entries in the emoms.trc file are relevant to notifications.

Normal Startup Messages

When the OMS starts, you should see these types of messages.

```
2011-08-17 13:50:29,458 [EventInitializer] INFO  em.notification init.167 - Short
format maximum length is 155
2011-08-17 13:50:29,460 [EventInitializer] INFO  em.notification init.185 - Short
format is set to both subject and body
2011-08-17 13:50:29,460 [EventInitializer] INFO  em.notification init.194 -
Content-Transfer-Encoding is 8-bit
2011-08-17 13:50:29,460 [EventInitializer] DEBUG em.notification
registerAdminMsgCallBack.272 - Registering notification system message call back
2011-08-17 13:50:29,461 [EventInitializer] DEBUG em.notification
registerAdminMsgCallBack.276 - Notification system message callback is registered
successfully
2011-08-17 13:50:29,713 [EventInitializer] DEBUG em.notification
upgradeEmailTemplates.2629 - Enter upgradeEmailTemplates
2011-08-17 13:50:29,735 [EventInitializer] INFO  em.notification
upgradeEmailTemplates.2687 - Email template upgrade is not required since no
customized templates exist.
2011-08-17 13:49:28,739 [EventCoordinator] INFO  events.EventCoordinator logp.251
- Creating event worker thread pool: min = 4 max = 15
2011-08-17 13:49:28,791 [[STANDBY] ExecuteThread: '2' for queue:
'weblogic.kernel.Default (self-tuning)'] INFO emdrep.pingHBRecorder
initReversePingThreadPool.937 - Creating thread pool for reverse ping : min = 10
max = 50
2011-08-17 13:49:28,797 [[STANDBY] ExecuteThread: '2' for queue:
'weblogic.kernel.Default (self-tuning)'] DEBUG emdrep.HostPingCoordinator logp.251
- Creating thread pool of worker thread  for host ping: min = 1 max = 10
2011-08-17 13:49:28,799 [[STANDBY] ExecuteThread: '2' for queue:
'weblogic.kernel.Default (self-tuning)'] DEBUG emdrep.HostPingCoordinator logp.251
- Creating thread pool for output of worker's  output for host ping: min = 2 max =
20
2011-08-17 13:49:30,327 [ConnectorCoordinator] INFO
connector.ConnectorPoolManager logp.251 - Creating Event thread pool: min = 3 max
= 10
2011-08-17 13:51:48,152 [NotificationMgrThread] INFO  notification.pbs logp.251 -
Creating thread pool: min = 6 max = 24
2011-08-17 13:51:48,152 [NotificationMgrThread] INFO  em.rca logp.251 - Creating
RCA thread pool: min = 3 max = 20
```

Notification Delivery Messages

```
2006-11-08 03:18:45,387 [NotificationMgrThread] INFO  em.notification run.682 -
Notification ready on EMAIL1

2006-11-08 03:18:46,006 [DeliveryThread-EMAIL1] INFO  em.notification run.114 -
Deliver to SYSMAN/admin@oracle.com

2006-11-08 03:18:47,006 [DeliveryThread-EMAIL1] INFO  em.notification run.227 -
Notification handled for SYSMAN/admin@oracle.com
```

Notification System Error Messages

```
2011-08-17 14:02:23,905 [NotificationMgrThread] DEBUG notification.pbs logp.251 -
Notification ready on EMAIL1
2011-08-17 14:02:23,911 [NotificationMgrThread] DEBUG notification.pbs logp.251 -
Notification ready on PLSQL4
2011-08-17 14:02:23,915 [NotificationMgrThread] DEBUG notification.pbs logp.251 -
Notification ready on OSCMD14
2011-08-17 14:02:19,057 [DeliveryThread-EMAIL1] INFO  notification.pbs logp.251 -
Deliver to To: my.admin@oracle.com; issue type: 1; notification type: 1
2011-08-17 14:02:19,120 [DeliveryThread-OSCMD14] INFO  notification.pbs logp.251 -
Deliver to SYSMAN, OSCMD, 8; issue type: 1; notification type: 1
2011-08-17 14:02:19,346 [DeliveryThread-PLSQL4] INFO  notification.pbs logp.251 -
Deliver to SYSMAN, LOG_JOB_STATUS_CHANGE, 9; issue type: 1; notification type: 1
2011-08-17 14:02:19,977 [DeliveryThread-PLSQL4] DEBUG notification.pbs logp.251 -
Notification handled for SYSMAN, LOG_JOB_STATUS_CHANGE, 9
2011-08-17 14:02:20,464 [DeliveryThread-EMAIL1] DEBUG notification.pbs logp.251 -
Notification handled for To: my.admin@oracle.com
2011-08-17 14:02:20,921 [DeliveryThread-OSCMD14] DEBUG notification.pbs logp.251 -
Notification handled for SYSMAN, OSCMD, 8
```

3.7.4 E-mail Errors

The SMTP gateway is not set up correctly:

```
Failed to send e-mail to my.admin@oracle.com: For e-mail notifications to be sent,
your Super Administrator must configure an Outgoing Mail (SMTP) Server within
Enterprise Manager. (SYSMAN, myrule)
```

Invalid host name:

```
Failed to connect to gateway: badhost.us.oracle.com: Sending failed;
nested exception is:
javax.mail.MessagingException: Unknown SMTP host: badhost.us.oracle.com;
```

Invalid e-mail address:

```
Failed to connect to gateway: rgmemeasmtp.oraclecorp.com: Sending failed;
nested exception is:
javax.mail.MessagingException: 550 5.7.1 <smpemailtest_ie@oracle.com>... Access
denied
```

Always use the Test button to make sure the e-mail gateway configuration is valid. Check that an e-mail is received at the sender's e-mail address

3.7.5 OS Command Errors

When attempting to execute an OS command or script, the following errors may occur. Use the Test button to make sure OS Command configuration is valid. If there are any errors, they will appear in the console.

Invalid path or no read permissions on file:

```
Could not find /bin/myscript (stacb10.us.oracle.com_Management_Service) (SYSMAN,
myrule )
```

No execute permission on executable:

```
Error calling /bin/myscript: java.io.IOException: /bin/myscript: cannot execute
(stacb10.us.oracle.com_Management_Service) (SYSMAN, myrule )
```

Timeout because OS Command ran too long:

```
Timeout occurred running /bin/myscript (stacb10.us.oracle.com_Management_Service)
(SYSMAN, myrule )
```

Any errors such as out of memory or too many processes running on OMS machine will be logged as appropriate.

Always use the Test button to make sure OS Command configuration is valid.

3.7.6 SNMP Trap Errors

Use the Test button to make sure SNMP Trap configuration is valid.

Other possible SNMP trap problems include: invalid host name, port, or community for a machine running an SNMP Console.

3.7.7 PL/SQL Errors

When attempting to execute an PL/SQL procedure, the following errors may occur. Use the Test button to make sure the procedure is valid. If there are any errors, they will appear in the console.

Procedure name is invalid or is not fully qualified. Example: SCOTT.PKG.PROC

```
Error calling PL/SQL procedure plsql_proc: ORA-06576: not a valid function or
procedure name (SYSMAN, myrule)
```

Procedure is not the correct signature. Example: PROCEDURE event_proc(s IN GC$NOTIF_EVENT_MSG)

```
Error calling PL/SQL procedure plsql_proc: ORA-06553: PLS-306: wrong number or
types of arguments in call to 'PLSQL_PROC' (SYSMAN, myrule)
```

Procedure has bug and is raising an exception.

```
Error calling PL/SQL procedure plsql_proc: ORA-06531: Reference to uninitialized
collection (SYSMAN, myrule)
```

Care should be taken to avoid leaking cursors in your PL/SQL. Any exception due to this condition will result in delivery failure with the message being displayed in the Details section of the alert in the Cloud Control console.

Always use the Test button to make sure the PL/SQL configuration is valid.

4

Metric Extensions

Metric extensions provides you with the ability to extend Oracle's monitoring capabilities to monitor conditions specific to your IT environment. This provides you with a comprehensive view of your environment. Furthermore, metric extensions allow you to simplify your IT organization's operational processes by leveraging Enterprise Manager as the single central monitoring tool for your entire datacenter instead of relying on other monitoring tools to provide this supplementary monitoring.

This chapter covers the following:

4.1 What are Metric Extensions?

Metric extensions also allow you to create metrics on any target type and customize metric thresholds and collections. Unlike user-defined metrics (used to extend monitoring in previous Enterprise Manager releases), metric extensions allow you to create full-fledged metrics for a multitude of target types, such as:

- Hosts
- Databases
- Fusion Applications
- IBM Websphere
- Oracle Exadata databases and storage servers
- Siebel components
- Oracle Business Intelligence components

You manage metric extensions from the Metric Extensions page. This page lists all metric extensions in addition to allowing you to create, edit, import/export, and deploy metric extensions.

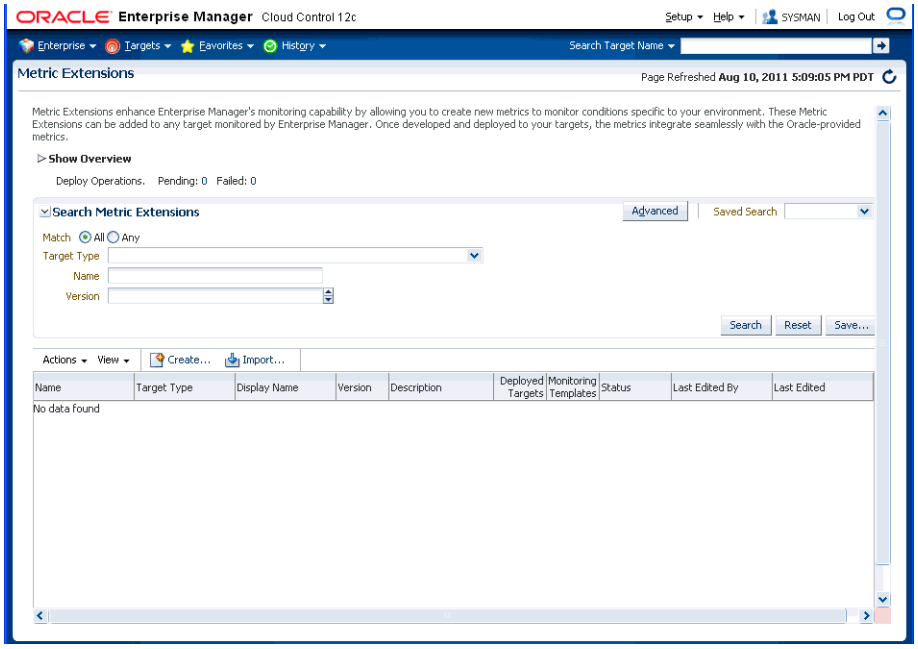

The cornerstone of the metric extension is the Oracle Integration Adapter. Adapters provide a means to gather data about targets using specific protocols. Adapter availability depends on the target type your metric extension monitors.

How Do Metric Extensions Differ from User-defined Metrics?

In previous releases of Enterprise Manager, user-defined metrics were used to extend monitoring capability in a limited fashion: user-defined metrics could be used to collect point values through execution of OS scripts and a somewhat more complex set of values (one per object) through SQL. Unlike metric extensions, user-defined metrics have several limitations:

- **Limited Integration**: If the OS or SQL user-defined metric executed custom scripts, or required addional dependent files, the user needed to manually transfer these files to the target's file system.
- **Limited Application of Query Protocols**: OS user-defined metrics cannot model child objects of servers by returning multiple rows from a metric (this capability only exists for SQL user-defined metrics).
- **Limited Data Collection**: Full-fledged Enterprise Manager metrics can collect multiple pieces of data with a single query and reflect the associated data in alert context. In the case of user-defined metrics, multiple pieces of data can be collected by creating multiple user-defined metrics, however, it is not possible to refer to the related data when alerts are generated because they are collected separately.
- **Limited Query Protocols**: User-defined metrics can only use the "OS" and "SQL" protocols, unlike metric extensions which can use additional protocols such as SNMP and JMX.
- **Limited Target Application**: You can only create OS user-defined metrics against host targets and SQL user-defined metrics agains database targets. No other target types are permitted. User-defined metrics only allow OS user-defined metrics against host targets and SQL user-defined metrics against database targets. If, for example, you want to deploy a user-defined metric against Weblogic instances in

your environment, you will not be able unable to do so, making it impossible to associate suspending of monitoring (blackouts) on these targets when servers are undergoing maintenance periods.

Most importantly, the primary difference between metric extensions and user-defined metrics is that, unlike user-defined metrics, metric exensions are full-fledged metrics similar to Enterprise Manager out-of-box metrics. They are handled and exposed in all Enterprise Manager monitoring features as any Enterprise Manager-provided metric and will automatically apply to any new features introduced.

4.2 Metric Extension Lifecycle

Developing a metric extension involves the same three phases you would expect from any programmatic customization:

- Developing Your Metric Extension
- Testing Your Metric Extension
- Deploying and Publishing Your Metric Extension

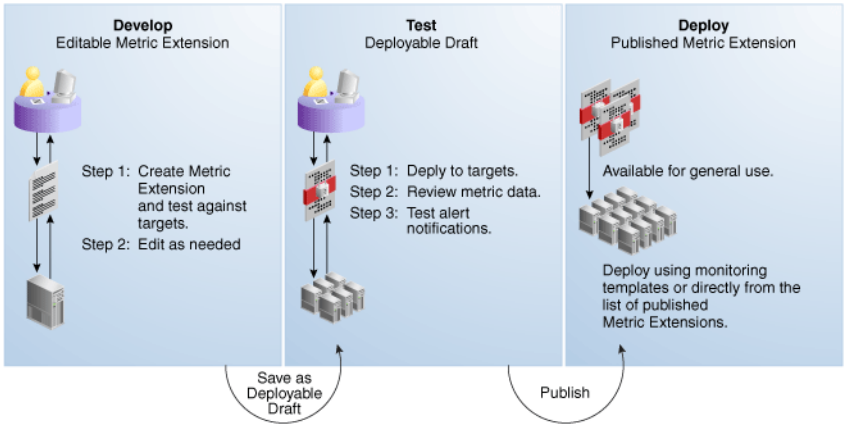

Developing Your Metric Extension

The first step is to define your monitoring requirements. This includes deciding the target type, what data needs to be collected, what mechanism (adapter) can be used to collect that data, and if elevated credentials are required. After making these decisions, you are ready to begin developing your metric extension. Enterprise Manager provides an intuitive user interface to guide you through the creation process.

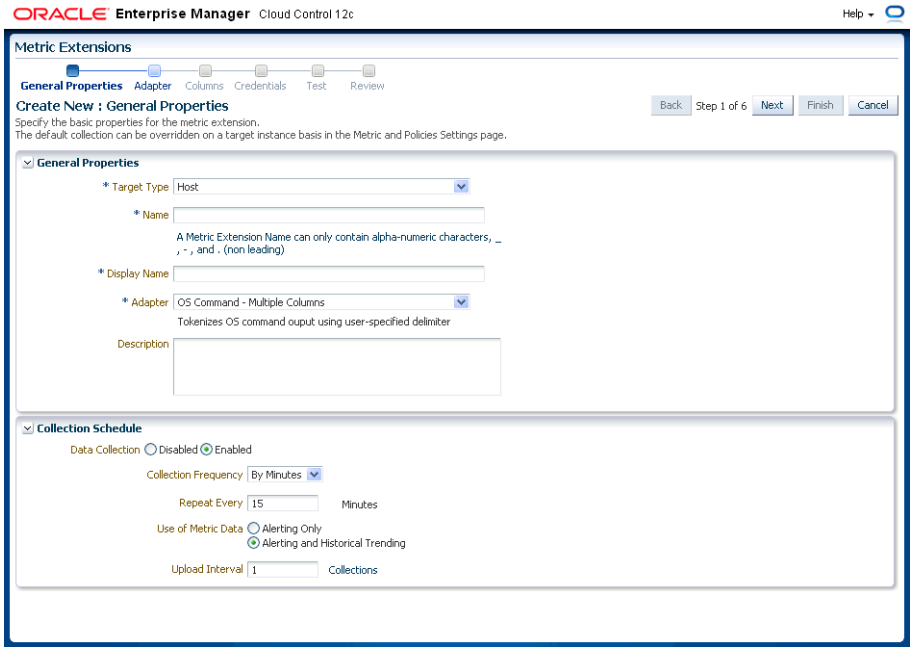

The metric extension wizard allows you to develop and refine your metric extension in a completely editable format. And more importantly, allows you to interactively test your metric extension against selected targets without having first to deploy the extension to a dedicated test environment. The **Test** page allows you to run real-time metric evaluations to ensure there are no syntactical errors in your script or metric extension definition.

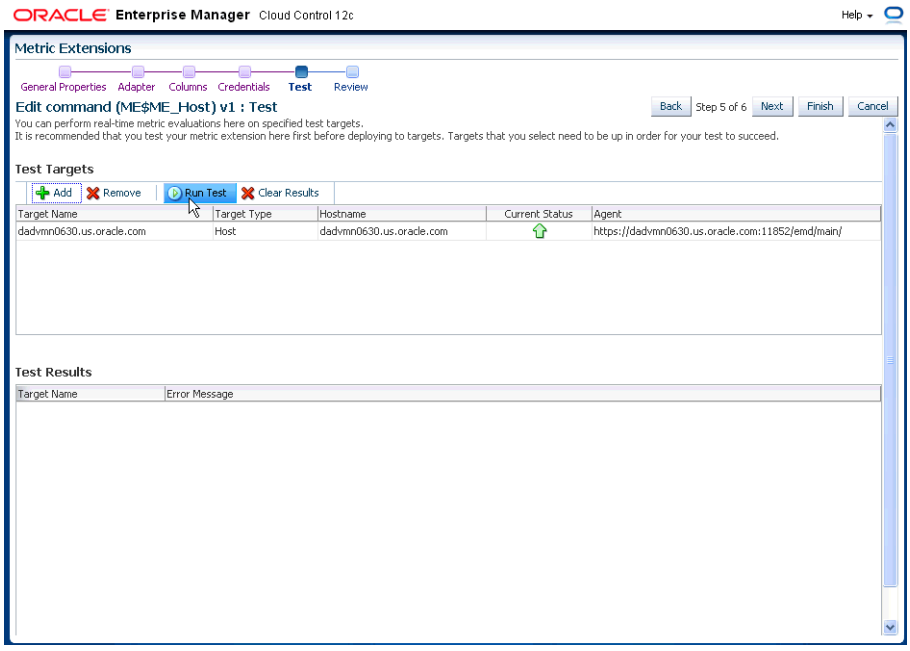

When you have completed working on your metric extension, you can click Finish to exit the wizard. The newly created metric extension appears in the Metric Extension Library where you can edit can be opened for further editing or saved as a deployable draft that can be tested against multiple targets.

Note: You can edit a metric extension only if its status is *editable*. Once it is saved as a deployable draft, you must create a new version to implement further edits.

Testing Your Metric Extension

Once your metric extension returns the expected data during real-time target testing, you are ready to test its robustness and actual behavior in Enterprise Manager by deploying it against targets and start collecting data. At this this point, the metric extension is still private (only the developer can deploy to targets), but is identical to Oracle out-of-box metrics behavior wise.This step involves selecting your editable metric extension in the library and generating a deployable draft.

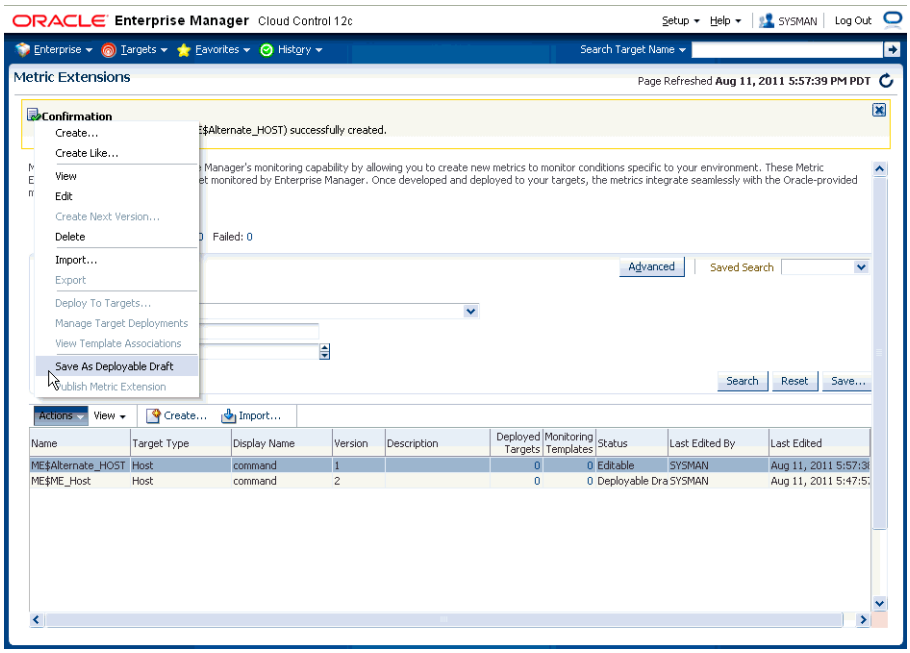

You can now deploy the metric extension to actual targets by going through the "Deploy To Targets…" action. After target deployment, you can review the metric data returned and test alert notifications. As mentioned previously, you will not be able to edit the metric extension once a deployable draft is created: You must create a new version of the metric extension.

Deploying Your Metric Extension

After rigorous testing through mutiple metric extension versions and target deployments, your metric extension is ready for deployment to your production environment. Until this point, your metric extension is only viewable by you, the metric extension creator. To make it accessible to all Enterprise Manager administrators, it must be published.

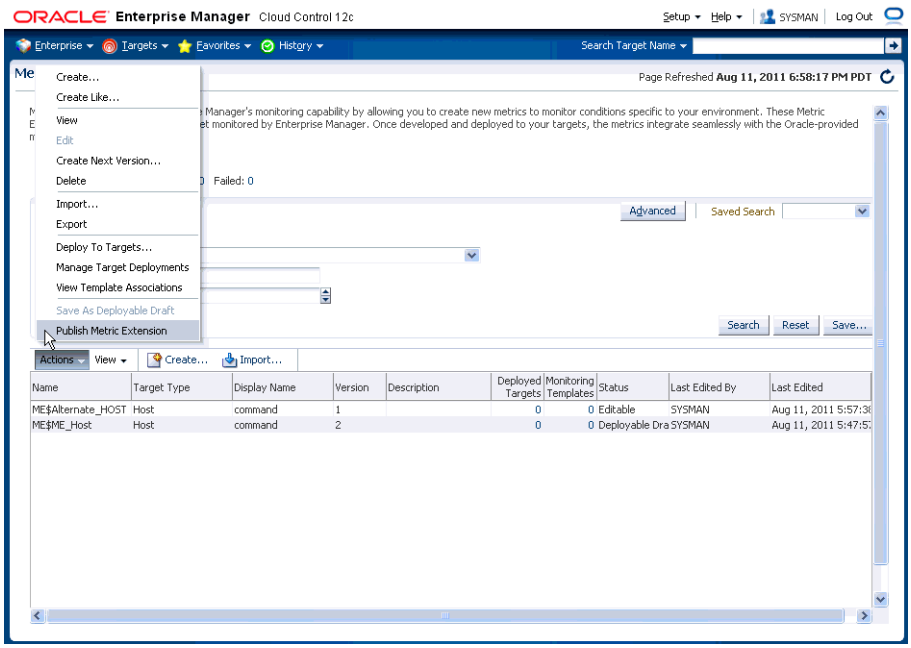

Now that your metric extension has been made public, your metric extension can be deployed to intended production targets. If you are monitoring a small number of targets, you can choose the **Deploy To Targets** menu option and add targets one at a time. For large numbers of targets, you deploy metric extensions to targets using monitoring templates. An extension is added to a monitoring template in the same way a full-fledged metric is added. The monitoring template is then deployed to the targets.

Note: You cannot add metric extensions to monitoring templates before publishing the extension. If you attempt to do so, the monitoring template page will warn you about it, and will not proceed until you remove the metric extension.

4.3 Working with Metric Extensions

Most all metric extension operations can be carried out from the Metric Extension home page. If you need to perform operations on published extensions outside of the UI, Enterprise Manger also provides EMCLI verbs to handle such operations as importing/exporting metric extensions to archive files and migrating legacy user-defined metrics to metric extensions. This section covers metric extension operations carried out from the UI.

4.3.1 Administrator Privilege Requirements

In order to create, edit, view, deploy or undeploy metric extensions, you must have the requisite administrator privileges. Enterprise Manager administrators must have the following privileges:

- **Create Metric Extension**: System level access that:

  Lets administrators view and deploy metric extensions

  Allows administrators to edit and delete extensions.

- **Edit Metric Extension**: Lets users with "Create Metric Extension" prvilege edit and create next versions of a particular metric extensions. The metric extension creator has this privilege by default.

  Note: This privilege must be granted on a per-metric extension basis.

- **Full Metric Extension**: In addition to the Edit Metric Extension privileges, allows deletion of a particular metric extension.

- **Manage Metrics**: Lets users deploy and un-deploy extensions on targets

  Note: The Manage Metrics privilege must be granted on a per-target basis.

4.3.2 Granting Create Metric Extension Privilege

To grant create metric extension prvileges to another administrator:

1. From the **Setup** menu, choose **Security** and then **Administrators**.
2. Choose the Administrator you would like to grant the prvilege to.
3. Click **Edit**.
4. Go to the Resource Privileges tab, and click **Manage Privilege Grants** for the Metric Extension resource type.
5. Under Resource Type Prvileges, click the Create Metric Extension check box.
6. Click **Continue**, review changes, and click **Finish** in the Review tab.

4.3.3 Creating a New Metric Extension

To create a new metric extension:

1. From the **Enterprise** menu, choose **Monitoring** and then **Metric Extensions.**
2. Click **Create New**. Enterprise Manager will determine whether you have the Create Extension privilege and guide you through the creation process.
3. Decide on a metric extension name. Be aware that the name (and Display Name) must be unique across a target type.
4. Enter the general parameters.

   The selected Adapter type defines the properties you must specify in the next step of the metric extension wizard. The following adapter types are available:

   - OS Command Adapter - Single Column

     Executes the specified OS command and returns the command output as a single value. The metric result is a 1 row, 1 column table.

   - OS Command Adapter- Multiple Values

     Executes the specified OS command and returns each command output line as a separate value. The metric result is a multi-row, 1 column table.

   - OS Command Adapter - Multiple Columns

     Executes the specified OS command and parses each command output line (delimited by a user-specified string) into multiple values. The metric result is a mult-row, multi-column table.

   - SQL Adapter

     Executes custom SQL queries or function calls against single instance databases and instances on Real Application Clusters (RAC).

- SNMP (Simple Network Management Protocol) Adapter

  Allow Enterprise Manager Management Agents to query SNMP agents for Management Information Base (MIB) variable information to be used as metric data.

- JMX (Java Management Extensions) Adapter

  Retrieves JMX attributes from JMX-enabled servers and returns these attributes as a metric table.

Refer to the Adapters section for specific information on the selected adapter needed in the Adapter page (step 2) of the wizard.

Note: Be aware that if you change the metric extension Adapter, all your previous adapter properties (in Step 2) will be cleared.

5. From the Columns page, add metric columns defining the data returned from the adapter. Note that the column order should match the order the adapter returns the data in.

   - **Column Type**

     A column is either a Key column, or Data column. A Key column uniquely identifies a row in the table. For example, employee ID is a unique identifier of a table of employees. A Data column is any non-unique data in a row. For example, the first and last names of an employee.

   - **Value Type**

     A value type is Number or String. This determines the alert comparison operators that are available, and how Enterprise Manager renders collection data for this metric column.

   - **Alert Thresholds**

     The Comparison Operation, Warning, and Critical fields define an alert threshold.

   - **Alert Thresholds By Key**

     The Comparison Operation, Warning Thresholds By Key, and Critical Thresholds By Key fields allow you to specify distinct alert thresholds for different rows in a table. This option becomes available if there are any Key columns defined. For example, if your metric is monitoring CPU Usage, you can specify a different alert threshold for each disinct CPU. The syntax is to specify the key column values in a comma seperated list, the "=" symbol, followed by the alert threshold. Multiple thresholds for different rows can be seperated by the semi-colon symbol ";". For example, if the key colums of the CPU Usage metric are cpu_id and core_id, and you want to add a warning threshold of 50% for procecessor1, core1, and a threshold of 60% for processor2, core2, you would specify:
     procecessor1,core1=50;processor2,core2=60

   - **Manually Clearable Alert**

     If this option is set to true, then the alert will not automatically clear when the alert threshold is no longer satisfied. For example, if your metric is counting the number of errors in the system log files, and you set an alert threshold of 50, if an alert is raised once the threshold is met, the alert will not automatically clear once the error count falls back below 50. The alert will

need to be manually cleared in the Alerts UI in the target home page or Incident Manager.

- **Number of Occurances Before Alert**

  The number of consecutive metric collections where the alert threshold is met, before an alert is raised.

- **Alert Message / Clear Message**

  The message that is sent when the alert is raised / cleared. Variables that are available for use are: %columnName%, %keyValue%, %value%, %warning_threshold%, %critical_threshold%

  You can also retrieve the value of another column by surrounding the desired column name with "%". For example, if you are creating a alert for the cpu_usage column, you can get the value of the core_temperature column by using %core_temperature%. Note that the same alert / clear message is used for warning or critical alerts.

  Note: Think carefully and make sure all Key columns are added, because you cannot create additional Key columns in newer versions of the metric extension. Once you **Save As Deployable Draft**, the Key columns are final (edits to column display name, alert thresholds are still allowed). You can still add new Data columns in newer versions. Also be aware that some properties of existing Data columns cannot be changed later, including Column Type, Value Type, Comparison Operator (you can add a new operator, but not change an existing operator), and Manually Clearable Alert.

- **Metric Category**

  The metric category this column belongs to.

6. From the Credentials page, you can override the default monitoring credentials by using custom monitoring credential sets. By default, the metric extension wizard chooses the existing credentials used by Oracle out-of-box metrics for the particular target type. For example, metric extensions will use the dbsnmp user for database targets. You have the option to override the default credentials, by creating a custom monitoring credential set through the "emcli create_credential_set" command. Refer to the *Enterprise Manager Command Line Interface Guide* for additional details. Some adapters may use additional credentials, refer to the Adapters section for specific information.
7. From the Test page, add available test targets.
8. Click **Run Test** to validate the metric extension. The extension is deployed to the test targets specified by the user and a real-time collection is executed. Afterwards, the metric extension is automatically undeployed. The results and any errors are added to the **Test Results** region.
9. Repeat the edit /test cycle until the metric extension returns data as expected.
10. Click **Finish**.

4.3.4 Creating a New Metric Extension (Create Like)

To create a new metric extension based on an existing metric extension:

1. From the **Enterprise** menu, choose **Monitoring**-->**Metric Extensions.**

2. From the **Metric Extensions** page, determine which extensions are accessible. The page displays the list of metric extensions along with target type, owner, production version and deployment information.
3. Select an existing metric extension.
4. From the **Actions** menu, choose **Create Like**. Enterprise Manager will determine whether you have the Create Extension privilege and guide you through the creation process.
5. Make desired modifications.
6. From the **Test** page, add available test targets.
7. Click **Run Test** to validate the metric extension. The extension is deployed to the test targets specified by the user and a real-time collection is executed. Afterwards, the metric extension is automatically undeployed. The results and any errors are added to the **Test Results** region.
8. Repeat the edit /test cycle until the metric extension returns data as expected.
9. Click **Finish**.

4.3.5 Editing a Metric Extension

Before editing an existing metric extension, you must have Edit privileges on the extension you are editing or be the extension creator. Note: Once a metric extension is saved as a deployabled draft, it cannot be edited, you can only create a new version.

To edit an existing metric extension:

1. From the **Enterprise** menu, choose **Monitoring** and then choose **Metric Extensions.**
2. From the **Metric Extensions** page, determine which extensions are accessible. The page displays the list of metric extensions along with target type, owner, production version and deployment information.
3. Select the metric extension to be edited.
4. From the **Actions** menu, choose **Edit**.
5. Update the metric extension as needed.
6. From the **Test** page, add available test targets.
7. Click **Run Test** to validate the metric extension. The extension is deployed to the test targets specified by the user and a real-time collection is executed. Afterwards, the metric extension is automatically undeployed. The results and any errors are added to the **Test Results** region.
8. Repeat the edit /test cycle until the metric extension returns data as expected.
9. Click **Finish**.

4.3.6 Creating the Next Version of an Existing Metric Extension

Before creating the next version of an existing metric extension, you must have Edit privileges on the extension you are versioning or be the extension creator.

To create next version of an existing metric extension:

1. From the **Enterprise** menu, choose **Monitoring** and then choose **Metric Extensions**.

2. From the Metric Extensions page, determine which extensions are accessible. The page displays the list of metric extensions along with target type, owner, production version and deployment information.
3. Select the metric extension to be versioned.
4. From the **Actions** menu, choose **Create Next Version**.
5. Update the metric extension as needed. The target type, and extension name cannot be edited, but all other general properties can be modified. There are also restrictions on metric columns modifications. See Note in Creating a New Metric Extension section for more details.
6. From the Test page, add available test targets.
7. Click **Run Test** to validate the metric extension. The extension is deployed to the test targets specified by the user and a real-time collection is executed. Afterwards, the metric extension is automatically undeployed. The results and any errors are added to the **Test Results** region.
8. Repeat the edit /test cycle until the metric extension returns data as expected.
9. Click **Finish**.

4.3.7 Importing a Metric Extension

Metric extensions can be converted to portable, self-contained packages that allow you to move the metric extension to other Enterprise Manager installations, or for storage/backup. These packages are called Metric Extension Archives (MEA) files.

MEA files are zip files containing all components that make up the metric extension: metric metadata, collections, and associated scripts/jar files. Each MEA file can contain only one metric extension. To add the metric extension back to your Enterprise Manager installation, you must import the metric extension from the MEA.

To import a metric extension from an MEA file:

1. From the **Enterprise** menu, choose **Monitoring** and then **Metric Extensions.**
2. Click **Import**.
3. Browse to file location, and select the MEA file. Enterprise Manager checks if the target type and metric extension name combination is already used in the system. If not, the system will create a new metric extension. If the extension name is already in use, the sytem will attempt to create a new version of the existing extension using the MEA contents. This will require the MEA to contain a superset of all the existing metric extension's metric columns. You also have the option to rename the metric extension.
4. Clicking on OK creates the new metric extension or the new version of an existing metric extension.
5. From the **Actions** menu, choose **Edit** to verify the entries.
6. From the **Test** page, add available test targets.
7. Click **Run Test** to validate the metric extension. The extension is deployed to the test targets specified by the user and a real-time collection is executed. Afterwards, the metric extension is automatically undeployed. The results and any errors are added to the **Test Results** region.
8. Repeat the edit /test cycle until the metric extension returns data as expected.
9. Click **Finish**.

4.3.8 Exporting a Metric Extension

Existing metric extensions can be package as self-contained zip files (exported) for portability and/or backup and storage.

To export an existing metric extension:

1. From the **Enterprise** menu, choose **Monitoring** and then**Metric Extensions.**
2. From the **Metric Extensions** page, determine which extensions are accessible. The page displays the list of metric extensions along with target type, owner, production version and deployment information.
3. Select the metric extension to be exported.
4. From the **Actions** menu, choose **Export**. Enterprise Manager prompts you to enter the name and location of the MEA file that is to be created.
5. Enter the name and location of the package. Enterprise Manager displays the confirmation page after the export is complete. Note: You can only export the production version. Note: You can only export Deployable Draft and Published metric extension versions.
6. Confirm the export file is downloaded.

4.3.9 Deleting a Metric Extension

Initiating the deletion of a metric extension is simple. However, the actual deletion triggers a cascade of activity by Enterprise Manager to completely purge the metric extension from the system. This includes closing open metric alerts, and purging collected metric data (if the latest metric extension version is deleted).

Before a metric extension version can be deleted, it must be undeployed from all targets, and removed from all monitoring templates (including templates in pending apply status).

To delete a metric extension:

1. From the **Enterprise** menu, choose **Monitoring** and then **Metric Extensions.**
2. From the **Metric Extensions** page, determine which extensions are accessible. The page displays the list of metric extensions along with target type, owner, production version and deployment information.
3. Select the metric extension that is to be deleted.
4. From the **Actions** menu, choose **Delete**. Enterprise Manager prompts you to confirm the deletion.
5. Confirm the deletion.

4.3.10 Granting Edit/Full Access to Metric Extensions

Before an Enterprise Manager administrator can be edit, or delete a metric extension created by another administrator, they must have been granted requisite access privileges. Edit privilege allows editing and creating next versions of the extension, and Full privilege allows the above operations and deletion of the extension.

To grant edit/full access to an existing metric extension to another administrator:

1. From the **Setup** menu, choose **Security** and then **Administrators**.
2. Choose the Administrator you would like to grant access to.
3. Click **Edit**.

4. Go to the Resource Privileges tab, and click **Manage Privilege Grants for the Metric Extension** resource type.
5. Under **Resource Prvileges**, you can search for and add existing metric extensions. Add the metric extensions you would like to grant privileges to. This allows the user to edit and create next versions of the metric extension.
6. If you would additionally like to allow delete operations, then click the pencil icon in the **Manage Resource Prvilege Grants** column, and select **Full Metric Extension** privilege in the page that shows up.
7. Click **Continue**, review changes, and click **Finish** in the review tab.

4.3.11 Deploying Metric Extensions to a Group of Targets

A metric extension must be deployed to a target in order for it to begin collecting data.

To deploy a metric extension to one or more targets:

1. From the **Enterprise** menu, choose **Monitoring** and then **Metric Extensions.**
2. From the **Metric Extensions** page, determine which extensions are accessible. The page displays the list of metric extensions along with target type, owner, production version and deployment information.
3. Select the metric extension that is to be deployed.
4. From the **Actions** menu, choose **Manage Target Deployments**. The **Manage Target Deployments** page appears showing you on which target(s) the selected metric extension is already deployed.
5. Return to the **Metric Extensions** page.
6. Select the metric extension.
7. From the Actions menu, choose Deploy to Targets. Enterprise Manager determines whether you have "Manage Target Metrics" privilege, and only those targets where you do show up in the target selector.
8. Add the targets where the metric extension is to be deployed and click Submit. Enterprise Manager submits a job deploying the metric extension to each of the targets. A single job is submitted per deployment request.
9. You are automatically redirected to the Pending Operations page, which shows a list of currently scheduled, executing, or failed metric extension deploy operations. Once the deploy operation completes, the entry is removed from the pending operations table.

4.3.12 Updating Older Versions of Metric Extensions Already deployed to a Group of Targets

When a newer metric extension version is published, you may want to update any older deployed instances of the metric extension.

To update old versions of the metric extension already deployed to targets:

1. From the **Enterprise** menu, choose **Monitoring** and then **Metric Extensions.**
2. From the **Metric Extensions** page, determine which extensions are accessible. The page displays the list of metric extensions along with target type, owner, production version and deployment information.
3. Select the metric extension to be upgraded.

4. From the **Actions** menu, choose **Manage Target Deployments**. The **Manage Target Deployments** page appears showing a list of targets where the metric extension is already deployed.
5. Select the list of targets where the extension is to be upgraded and click **Upgrade**. Enterprise Manager submits a job for the deployment of the newest Published metric extension to the selected targets. A single job is submitted per deployment request.
6. You are automatically redirected to the Pending Operations page, which shows a list of currently scheduled, executing, or failed metric extension deploy operations. Once the deploy operation completes, the entry is removed from the pending operations table.

4.4 Adapters

Oracle Integration Adapters provide comprehensive, easy-to-use monitoring connectivity with a variety of target types. The adapter enables communication with an enterprise application and translates the application data to standards-compliant XML and back.

The metric extension target type determines which adapters are made avaialble from the UI. A complete list of all adapters is shown below.

- OS Command Adapter - Single Column
- OS Command Adapter- Multiple Values
- OS Command Adapter - Multiple Columns
- SQL Adapter
- SNMP (Simple Network Management Protocol) Adapter
- JMX Adapter

4.4.1 OS Command Adapter - Single Column

Executes the specified OS command and returns the command output as a single value. The metric result is a 1 row, 1 column table.

Basic Properties

The complete command line will be constructed as: Command + Script + Arguments.

- **Command** - The command to execute. For example, `%perlBin%/perl`. The complete command line will be constructed as: Command + Script + Arguments.
- **Script** - A script to pass to the command. For example, `%scriptsDir%/myscript.pl`. You can upload custom files to the agent, which will be accessible under the `%scriptsDir%` directory.
- **Arguments** - Additional arguments to be appended to the Command.

Advance Properties

- **Input Properties** - Additional properties can be passed to the command through its standard input stream. This is usually used for secure content, such as username or passwords, that you don't want to be visible to other users. For example, you can add the following Input Property:

`Name=targetName,Value=%NAME%`

which the command can read through it's standard input stream as "STDINtargetName=<target name>".

- **Environment Variables** - Additional properties can be accessible to the command from environment variables. For example, you can add Environment Variable: `Name=targetType,Value="%TYPE%"`, and the command can access the target type from environment variable "ENVtargetType".

Credentials

- **Host Credentials** - The credential used to launch the OS Command.
- **Input Credentials** - Additional credentials passed to the OS Command's standard input stream.

Example 1

Read the contents of a log file, and dump out all lines containing references to the target.

- **Approach 1** - Use the grep command, and specify the target name using %NAME% parameter.

  ```
  Command = /bin/grep %NAME% mytrace.log
  ```

- **Approach 2** - Run a perl script

  ```
  Command = %perlBin%/perl
  ```

  ```
  Script = %scriptsDir%/filterLog.pl
  ```

  Input Properties:

  ```
  targetName = %NAME%
  ```

  ```
  targetType = %TYPE%
  ```

filterLog.pl:

```
require "emd_common.pl";

my %stdinVars = get_stdinvars();
my $targetName = $stdinVars{"targetName"};
my $targetType = $stdinVars{"targetType"};
open (MYTRACE, mytrace.log);
foreach $line (<MYTRACE >)
{
    # Do line-by-line processing
}

close (MYTRACE);
```

Example 2

Connect to a database instance from a perl script and query the HR.JOBS sample schema table.

- Approach 1 - Pass credentials from target type properties into using Input Properties:

  ```
  Command = %perlBin%/perl
  ```

  ```
  Script = %scriptsDir%/connectDB.pl
  ```

Input Properties:

```
EM_DB_USERNAME = %Username%
EM_DB_PASSWORD = %Password%
EM_DB_MACHINE = %MachineName%
EM_DB_PORT = %Port%
EM_DB_SID = %SID%
```

connectDB.pl

```
use DBI;
require "emd_common.pl";

my %stdinVars = get_stdinvars();
my $dbUsername = $stdinVars{"EM_DB_USERNAME"};
my $dbPassword = $stdinVars{"EM_DB_PASSWORD"};
my $dbMachine = $stdinVars{"EM_DB_MACHINE"};
my $dbPort = $stdinVars{"EM_DB_PORT"};
my $dbSID = $stdinVars{"EM_DB_SID"};

my $dbAddress =
"(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=$dbMachine)(Port=$dbPort))(CONNECT_
DATA=(SID=$dbSID)))";

# Establish Target DB Connection
my $db = DBI->connect('dbi:Oracle:', "$dbUsername@".$dbAddress, "$dbPassword",
    {PrintError => 0, RaiseError => 0, AutoCommit => 0})
    or die (filterOraError("em_error=Could not connect to
$dbUsername/$dbAddress: $DBI::errstr\n", $DBI::err));

my $query = "SELECT JOB_TITLE, MIN_SALARY FROM HR.JOBS";
my $st = $db->prepare($query);
$st->execute();

while ( my ($job_title, $min_sal) = $st->fetchrow_array() )
{
    print "$job_title|$min_sal\n";
}

$db->disconnect
    or warn "disconnect $DBI::errstr\n";

exit 0;
```

- Approach 2 - Pass monitoring credential set using Input Credentials

```
Command = %perlBin%/perl
Script = %scriptsDir%/connectDB.pl
```

Input Credentials:

```
dbCreds = MyCustomDBCreds
```

connectDB.pl

```
use DBI;

require "emd_common.pl";
```

```
my %stdinVars = get_stdinvars();
my $credType = getCredType("dbCred", \%stdinVars);
my %credProps = getCredProps("dbCreds", \%stdinVars);
my $dbUsername = $credProps{"DBUserName"};
my $dbPassword = $credProps{"DBPassword"};
```

Example 3

Overiding default monitoring credentials by creating and using a custom monitoring credential set for host target.

We will only show a simple example here. Refer to the Credentials guide of the Admin Guide for more details on creating and configuring custom monitoring credential sets.

Creating host credentials for the host target type:

```
> emcli create_credential_set -set_name=myCustomCreds -target_type=host -auth_
target_type=host -supported_cred_types=HostCreds -monitoring -description='My
Custom Credentials'
```

When you go to the Credentials page of the Metric Extension wizard, and choose to "Specify Credential Set" for Host Credentials, you will see "My Custom Credentials" show up as an option in the drop down list.

Note that this step only creates the Monitoring Credential Set for the host target type, and you need to set the credentials on each target you plan on deploying this metric extension to. You can set credentials from EnterpriseManager by going to Setup, then Security, then Monitoring Credentials. Alternatively, this can be done from the command line.

```
> emcli set_monitoring_credential -target_name=target1 -target_type=host -set_
name=myCustomCreds -cred_type=HostCreds -auth_target_type=host
-attributes='HostUserName:myusername;HostPassword:mypassword'
```

4.4.2 OS Command Adapter- Multiple Values

Executes the specified OS command and returns each command output line as a separate value. The metric result is a multi-row, 1 column table.

For example, if the command output is:

```
em_result=foo
em_result=bar
```

then three columns are populated with values 1,2,3 respectively.

Basic Properties

- **Command** - The command to execute. For example, %perlBin%/perl.
- **Script** - A script to pass to the command. For example, `%scriptsDir%/myscript.pl`. You can upload custom files to the agent, which will be accessible under the `%scriptsDir%` directory.
- **Arguments** - Additional arguments to be appended to the Command.
- **Starts With** - The starting string of metric result lines.

  Example: If the command output is:

  ```
  em_result=4354
  update
  test
  ```

setting *Starts With = em_result* specifies that only lines starting with *em_result* will be parsed.

Advanced Properties

- **Input Properties** - Additional properties to be passed to the command through its standard input stream. For example, you can add Input Property: Name=targetName, Value=%NAME%, which the command can read through its standard input stream as "STDINtargetName=<target name>". See usage examples in OS Command Adapter - Single Columns.
- **Environment Variables** - Additional properties can be accessible to the command from environment variables. For example, you can add Environment Variable: Name=targetType, Value="%TYPE%", and the command can access the target type from environment variable "ENVtargetType". See usage examples in OS Command Adapter - Single Columns.

Credentials

- **Host Credentials** - The credential used to launch the OS Command. See usage examples in OS Command Adapter - Single Columns.
- **Input Credentials** - Additional credentials passed to the OS Command's standard input stream. See usage examples in OS Command Adapter - Single Columns.

4.4.3 OS Command Adapter - Multiple Columns

Executes the specified OS command and parses each command output line (delimited by a user-specified string) into multiple values. The metric result is a mult-row, multi-column table.

Example: If the command output is

```
em_result=1|2|3
em_result=4|5|6
```

and the Delimiter is set as "|", then there are two rows of three columns each:

Table 4–1 Multi-Column Ouptut

1	2	3
4	5	6

Basic Properties

The complete command line will be constructed as: Command + Script + Arguments

- **Command** - The command to execute. For example, %perlBin%/perl.
- **Script** - A script to pass to the command. For example, %scriptsDir%/myscript.pl. You can upload custom files to the agent, which will be accessible under the %scriptsDir% directory.
- **Arguments** - Additional arguments.
- **Delimiter** - The string used to delimit the command output.
- **Starts With** - The starting string of metric result lines.

  Example: If the command output is

```
em_result=4354
foo
bar
```

setting *Starts With = em_result* specifies that only lines starting with *em_result* will be parsed.

- **Input Properties** - Additional properties can be passed to the command through its standard input stream. For example, you can add Input Property: *Name=targetName*, *Value=%NAME%*, which the command can read through it's standard input stream as *STDINtargetName=<target name>*. To specify multiple Input Properties, enter each property on its own line.
- **Environment Variables** - Additional properties can be accessible to the command from environment variables. For example, you can add Environment Variable: *Name=targetType*, *Value="%TYPE%*, and the command can access the target type from environment variable "ENVtargetType".

Advanced Properties

- **Input Properties** - Additional properties can be passed to the command through its standard input stream. For example, you can add Input Property: Name=targetName, Value=%NAME%, which the command can read through its standard input stream as STDINtargetName=<target name>. See usage examples in OS Command Adapter - Single Columns.
- **Environment Variables** - Additional properties can be accessible to the command from environment variables. For example, you can add Environment Variable: *Name=targetType*, *Value="%TYPE%*, and the command can access the target type from environment variable "ENVtargetType". See usage examples in OS Command Adapter - Single Columns.

Credentials

- **Host Credentials** - The credential used to launch the OS Command. See usage examples in OS Command Adapter - Single Columns
- **Input Credentials** - Additional credentials passed to the OS Command's standard input stream. See usage examples in OS Command Adapter - Single Columns.

4.4.4 SQL Adapter

Executes custom SQL queries or function calls supported against single instance databases and instances on Real Application Clusters (RAC).

Properties

- **SQL Query** - The SQL query to execute. Normal SQL statements should not be semi-colon terminated. For example, SQL Query = "select a.ename, (select count(*) from emp p where p.mgr=a.empno) directs from emp a". PL/SQL statements are also supported, and if used, the "Out Parameter Position" and "Out Parameter Type" properties should be populated.
- **SQL Query File** - A SQL query file. Note that only one of "SQL Query" or "SQL Query File" should be used. For example, %scriptsDir%/myquery.sql. You can upload custom files to the agent, which will be accessible under the %scriptsDir% directory.
- **Transpose Result** - Transpose the SQL query result.

- **Bind Variables** - Declare bind variables used in normal SQL statements here. For example, if the SQL Query = "select a.ename from emp a where a.mgr = :1", then you can declare the bind variable as Name=1, Value=Bob.
- **Out Parameter Position** - The bind variable used for PL/SQL output. Only integers can be spcified.

  Example: If the SQL Query is

```
DECLARE
    l_output1 NUMBER;
    l_output2 NUMBER;
BEGIN
    .....
    OPEN :1 FOR
        SELECT l_output1, l_output2 FROM dual;
END;
```

  you can set Out Parameter Position = 1, and Out Parameter Type = SQL_CURSOR
- **Out Parameter Type** - The SQL type of the PL/SQL output parameter. See comment for Out Parameter Position

Credentials

- **Database Credentials** - The credential used to connect to the database.

Example

Overiding default monitoring credentials by creating and using a custom monitoring credential set for database target.

We will only show a simple example here. Refer to the Credentials guide of the Admin Guide for more details on creating and configuring custom monitoring credential sets.

Creating host credentials for the database target type:

```
> emcli create_credential_set -set_name=myCustomDBCreds -target_type=oracle_
database -auth_target_type=oracle_database -supported_cred_types=DBCreds
-monitoring -description='My Custom DB Credentials'
```

When you go to the Credentials page of the Metric Extension wizard, and choose to "Specify Credential Set" for Database Credentials, you will see "My Custom DB Credentials" show up as an option in the drop down list.

Note that this step only creates the Monitoring Credential Set for the host target type, and you need to set the credentials on each target you plan on deploying this metric extension to. You can set credentials from EnterpriseManager by going to **Setup**, then **Security**, then **Monitoring Credentials**. Alternatively, this can be done from the command line.

```
> emcli set_monitoring_credential -target_name=db1 -target_type=oracle_database
-set_name=myCustomDBCreds -cred_type=DBCreds -auth_target_type=oracle_database
-attributes='DBUserName:myusername;DBPassword:mypassword'
```

4.4.5 SNMP (Simple Network Management Protocol) Adapter

Allow Enterprise Manager Management Agents to query SNMP agents for Management Information Base (MIB) variable information to be used as metric data.

Basic Properties

- **Object Identifiers (OIDs)**: Object Identifiers uniquely identify managed objects in a MIB hierarchy. One or more OIDs can be specified. The SNMP adapter will collect data for the specified OIDs. For example, 1.3.6.1.4.1.111.4.1.7.1.1

Advanced Properties

- **Delimiter** - The delimiter value used when specifying multiple OID values for an OID's attribute. The default value is space or \n or \t
- **Tabular Data** - Indicates whether the excpeted result for a metric will have multiple rows or not. Possible values are TRUE or FALSE. The default value is FALSE
- **Contains V2 Types** - Indicates whether any of the OIDs specified is of SNMPV2 data type. Possible values are TRUE or FALSE. The default value is FALSE. For example, if an OID value specified is of counter64 type, then this attribute will be set to TRUE.

4.4.6 JMX Adapter

Retrieves JMX attributes from JMX-enabled servers and returns these attributes as a metric table.

Properties

- **Metric** -- The MBean ObjectName or ObjectName pattern whose attributes are to be queried. Since this is specified as metric metadata, it needs to be instance-agnostic. Instance-specific key properties (such as *servername*) on the MBean ObjectName may need to be replaced with wildcards.
- **ColumnOrder** -- A semi-colon separated list of JMX attributes in the order they need to be presented in the metric.

Advanced Properties

- **IdentityCol** -- The MBean key property that needs to be surfaced as a column when it is not available as a JMX attribute. For example:

  ```
  com.myCompany:Name=myName,Dept=deptName, prop1=prop1Val, prop2=prop2Val
  ```

  In this example, setting *identityCol* as *Name;Dept* will result in two additional key columns representing Name and Dept besides the columns representing the JMX attributes specified in the *columnOrder* property.

- **AutoRowPrefix** -- Prefix used for an automatically generated row. Rows are automatically generated in situations where the MBean *ObjectName* pattern specified in metric property matches multipld MBeans and none of the JMX attributes specified in the *columnOrder* are unique for each. The *autoRowId* value specified here will be used as a prefix for the additional key column created. For example, if the metric is defined as:

  ```
  com.myCompany:Type=CustomerOrder,* columnOrder
  ```

  is

  ```
  CustomerName;OrderNumber;DateShipped
  ```

  and assuming *CustomerName;OrderNumber;Amount* may not be unique if an order is shipped in two parts, setting *autoRowId* as "ShipItem-" will populate an

additional key column for the metric for each row with ShipItem-0, ShipItem-1, ShipItem-2...ShipItem-n.

- **Metric Service** -- True/False. Indicate whether *MetricService* is enabled on a target Weblogic domain. This property would be false (unchecked) in most cases for Metric Extensions except when metrics that are exposed via the Oracle DMS MBean needs to be collected. If *MetricService* is set to true, then the basic property *metric* becomes the *MetricService* table name and the basic property *columnOrder* becomes a semicolon-separated list of column names in the *MetricService* table.

> **Note:** Refer to Chapter 20, "Monitoring Using Web Services and JMX" for an in-depth example of creating a JMX based Metric Extension.

4.5 Converting User-defined Metrics to Metric Extensions

For targets monitored by Enterprise Manager 12c Agents, both older user-defined metrics and metric extensions will be supported. After release 12c, only metric extensions will be supported. If you have existing user-defined metrics, it is recommended that you migrate them to metric extensions as soon as possible to prevent potential monitoring disruptions in your managed environment.

Migration of user-defined metric definitions to metric extensions is not automatic and must be initiated by an administrator. The migration process involves migrating user-defined metric metadata to metric extension metadata.

> **Note:** Migration of collected user-defined metric historic data is not supported.

After the user-defined metric is migrated to the metric extension and the metric extension has been deployed successfully on the target, the user-defined metric should be either disabled or deleted. Disabling the collection of the user-defined metric will retain the metadata(definition of the user-defined metric) but will clear all the open alerts, remove the metric errors and prevent further collections of the user-defined metric. Deleting the user-defined metric will delete the metadata, historic data, clear open alerts and remove metric errors.

4.5.1 Overview

The User Defined Metric (UDM) to Metric Extension (ME) migration replaces an existing UDM with a new or existing ME. The idea behind the migration process is to consolidate UDMs with the same definition that have been created on different targets into a single ME. In addition, MEs support multiple metric columns, allowing the user to combine multiple releated UDMs into a single ME.

This migration process is comprised of the following steps:

1. Identify the UDMs that need to be migrated.
2. Use the provided emcli commands to create or select a compatible metric extension.
3. Test and publish the metric extension.
4. Deploy the metric extension to all targets and templates where the original UDMs are located. Also update the existing notification rules to refer to the ME.

5. Delete the original UDMs. Note that the historical data and alerts from the old UDM is still accessible from the UI, but the new ME will not inherit them.

Note that the credentials being used by the UDM are NOT migrated to the newly created ME. The user interface allows a user to specify the credential set required by the metric extension. If the ME does not use the default monitoring credentials, the user will need to create a new credential set to accommodate the necessary credentials through the relevant emcli commands. This set will then be available in the credentials page of the metric extension wizard.

The migration process is categorized by migration sessions. Each session is responsible for migrating one or more UDMs. The process of migrating an individual session is referred to as a task. Therefore, a session is comprised of one or more tasks. In general terms, the migration involves creating a session and providing the necessary input to complete each tasks within that session. The status of the session and tasks is viewable throughout the workflow.

4.5.2 Commands

A number of emcli commands are responsible for completing the various steps of this process. For a more detailed explanation of the command definition, please use the 'emcli help <command>' option.

- **list_unconverted_udms** - Lists the UDMs that have yet to be migrated and not in a session
- **create_udmmig_session** - Creates a session to migrate one or more UDMs
- **udmmig_summary** - Lists the migration sessions in progress
- **udmmig_session_details** - Provides the details of a specific session
- **udmmig_submit_metricpics** - Provides a mapping between the UDM and the ME in order to create a new ME or use an existing one
- **udmmig_retry_deploys** - Deploys the ME to the targets where the UDM is present. Note that the ME has to be in a deployable draft or published state for this command to succeed
- **udmmig_request_udmdelete** - Deletes the UDM and completing the migration process

Usage Examples

The following exercise outlines a simple use case to showcase the migration

Consider a system with one host (host1) that has one host UDM (hostudm1) on it. The goal is to create a new ME (me1) that represents the UDM. The sequence of commands would be as follows

```
$ emcli list_unconverted_udms

-------------+---------------------+-----------+--------------------
Type         | Name                | Metric    | UDM
-------------+---------------------+-----------+--------------------
host         | host1               |UDM        | hostudm1
```

The command indicates that there is only one UDM that has not been migrated or in the process of migration at this stage. Now proceed with the creation of a session.

```
$ emcli create_udmmig_session -name=migration1 -desc="Convert UDMs for host
target" -udm_choice=hostudm1 -target=host:host1

Migration session created - session id is 1
```

The command creates a migration session with name migration1 and the description "convert UDMs for host target". The udm_choice flag indicates the UDM chosen and the target flag describes the target type and the target on which the UDM resides. Migration sessions are identified by session IDs. The current session has an ID of 1.

```
$ emcli udmmig_summary

------+--------------+------------------+------+------+--------+------+--------
ID    | Name         | Description      |#Tgts |Todo  |#Tmpls  |Todo  |IncRules
------+--------------+------------------+------+------+--------+------+--------
1     |migration1    |Convert UDMS      |      | 1/1  |  0     | -/0  | -/0
------+--------------+------------------+------+------+--------+------+--------
```

The command summarizes all the migrations sessions currently in progress. The name and description fields identify the session. The remaining columns outline the number of targets, templates and incident rules that contain references to the UDM that is being converted to a metric extension. The 'Todo' columns indicate the number of targets, templates and incident rules whose references to the UDM are yet to be updated. Since a migration session can be completed over a protracted period of time, the command provides an overview of the portion of the session that was been completed.

```
$ emcli list_unconverted_udms

There are no unconverted udms
```

Since the UDM is part of a migration session, it no longer shows up in the list of unconverted UDMs.

```
$ emcli udmmig_session_details -session_id=1

Name: migration1
Desc: Convert UDMs for host target
Created: <date> <time>
UDM Pick: [hostudm1]
UDMs being converted:
----------+----------+---------+------+------------+---------+---------+-----
Type      |Name      |UDM      |#MC   |Metric      |Column   |DepS     |DelS
----------+----------+---------+------+------------+---------+---------+-----
host      |host1     |hostudm1 |  0   |            |         |WAIT     |WAIT
----------+----------+---------+------+------------+---------+---------+-----
```

The command provides the status of a single migration session. It lists the name of the UDM and the target type and name of the target on which the UDM resides. In addition, it also outlines the metric extensions currently in the EM instance that match the UDM. The user can elect to use one of the existing choices or create an entirely new metric extension.

The system attempts to find compatible metric extensions by matching the properties of the UDM. For example, in the case of a host UDM, the system tries to find a metric extension that has the same command, script and argument fields. In the case of a database UDM, the system attempts to match the SQL query.

Finally, the DepS column indicates whether the metric extension that was matched to the UDM has been deployed to the target on which the UDM is defined. The DelS

column tells the user whether the UDM has been deleted after the metric extension has been deployed. As the user proceeds with the migration, the above table is updated from left to right. When the delete status column is set to complete, the migration session has ended.

```
$ emcli udmmig_submit_metricpicks -session_id=1 -input_file=metric_picks:filename

Successfully submitted metric picks for migration session
```

The command instructs the EM instance to use an existing metric extension or create a new one to replace the UDM. The various options are presented through a file, which is filename in the above command. The contents of the file are shown below

"host,host1,hostudm1,N,ME$me1,Usage"

Each line in the file represents a mapping from n UDM to an ME. The line provides the target type, the name of the target, the name of the UDM, a flag to indicate whether the metric extension is new (N) or existing (E), the name of the metric extension (note that ME$ must be prefixed) and the column name.

The types of UDMs supported are:

- Host (host)
- Database (oracle_database)
- RAC (rac_database)

A user can only specify the names of the data columns via the collection item portion of the file. A metric extension created through migration will always have two columns to represent the structure of the UDM. The first column is an index column for single column UDMs while the second column uses the column name mentioned in the file. In the case of two column UDMs, the first column of the ME is termed as the 'KEY' column and the collection name is used for the second column.

At this stage, the metric extension has been created and is visible in the metric extensions library.

```
$ emcli udmmig_session_details -session_id=1

Name: migration1
Desc: Convert UDMs for host target
Created: <date> <time>
UDM Pick: [hostudm1]
Udms being converted:
----------+--------+---------+------+----------+------------+---------+-----
Type      |Name    |UDM      |#MC   |Metric    |Column      |DepS     |DelS
----------+--------+---------+------+----------+------------+---------+-----
host      |host1   |hostudm1 |   1  | ME$me1   | Usage      |WAIT     |WAIT
----------+--------+---------+------+----------+------------+---------+-----

        #MC : There are 1 matches for udms in this session.
        Use emcli udmmig_list_matches to list available matches
```

The session details command indicates that there is one matching metric extension for this UDM (the value of the MC column is 1) and that metric extension is named as ME$me1. At this stage, we are ready to test the metric extension through the library page. Once the testing is complete and the user is satisfied with the metric extension that has been created, it is ready to be deployed. In order to deploy, the metric extension has to be minimally saved as a deployable draft.

```
$ emcli udmmig_retry_deploys -session_id=1 -input_file=metric_tasks:filename2

Metric Deployments successfully submitted
```

Note that the system will trigger a job to automatically deploy the metric extension to all targets where the UDM was present once the metric extension is published. If the user is interested in manually controlling the operation, the above command will perform the necessary steps. The command is similar to the submit_metricpicks option in that a file with the UDM to target mapping is provided. It is referred to by filename2 above. The contents of the file are as follows

"host,host1,hostudm1"

Each line in the file is a mapping from the UDM to the targets type and target on which it resides. Once the command is executed, jobs to deploy the metric extensions to various targets have been launched and can be tracked through the user interface.

```
$ emcli udmmig_request_udmdelete -session_id=1 -input_file=metric_tasks:demo_tasks

Udm deletes successfully submitted
```

The final command deletes the UDMs that were migrated to metric extensions. Note that this command might partially finish based on how many of the deployments were completed when the command was run.

```
$ emcli udmmig_session_details -session_id=1

Name: migration1
Desc: Convert UDMs for host target
Created: <date > <time>
Completed: <date > <time>
UDM Pick: [hostudm1]
Udms being converted:
--------+----------+---------+------+------------+------------+---------+-----
Type    |Name      |UDM      |#MC   |Metric      |Column      |DepS     |DelS
--------+----------+---------+------+------  ----+------------+---------+-----
host    |host1     |hostudm1 |  1   | ME$me1     | Usage      |COMP   |COMP
--------+----------+---------+------+------------+------------+---------+-----

        #MC : There are 1 matches for udms in this session.
          Use emcli udmmig_list_matches to list available matches
```

The session details command shows that the migration process is indeed complete.

4.6 Metric Extension Command Line Verbs

Metric extensions can be manipulated outside the UI via the Enterprise Manager Command Line Interface (EMCLI). Two categories of verbs are available:

- Metric Extension Verbs
  - *export_metric_extension*: Export a metric extension to an archive file
  - *get_unused_metric_extensions*: Get a list of unused metric extensions.
  - *import_metric_extension*: Import a metric extension archive file.
  - *publish_metric_extension*: Publish a metric extension for use by all administrators.
  - *save_metric_extension_draft*: Save a deployable draft of a metric extension.
- User-defined Metric Migration Verbs
  - *abort_udmmig_session*: Abort (partially) user-defined metric migration session.

- *analyze_unconverted_udms*: Analyze the unconverted user-defined metrics.
- *create_udmmig_session*: Create a user-defined metric migration session.
- *list_unconverted_udms*: List the user-defined metrics that are not yet in a migration session.
- *udmmig_list_matches*: List the matching metrics per user-defined metric in a specific user-defined metric migration session.
- *udmmig_request_udmdelete*: Request deletion of user-defined metrics from targets.
- *udmmig_retry_deploys*: Retry deployment of metric extensions to targets.
- *udmmig_session_details*: Retrieve the details of a specific user-defined metric migration session.
- *udmmig_submit_metricpicks*: Select the metrics to replace user-defined metrics in a session.
- *udmmig_summary*: Summarize the status of all user-defined metric migration sessions.
- *udmmig_update_incrules*: Update user-defined metric incident rules to include replacement metric references.

Metric Extension Verbs

```
  emcli export_metric_extension
       -file_name=<name of the metric extension archive>
       -target_type=<target type of the metric extension>
       -name=<name of the metric extension
       -version=<version of the metric extension>

  Description:
    Export a metric extension archive file.

  Options:
    -file_name=<file name>
      The name of the metric extension archive file to export into.
    -target_type=<target type>
      Target type of the metric extension.
    -name=<name>
      Name of the metric extension.
    -version=<version>
      Version of the metric extension to be exported.


  emcli get_unused_metric_extensions

  Description:
    Get a  list of metric extensions that are deployed to agents but not attached
to any targets.


  emcli import_metric_extension
        -file_name=<name of the metric extension archive>
        -rename_as=<name of the metric extension to import as>

  Description:
    Import a metric extension archive file.
```

```
  Options:
    -file_name=<file name>
      The name of the metric extension archive file to be imported.
    -rename_as=<metric extension name>
      Import the metric extension using the specified name, replacing the name
given in the archive.


  emcli publish_metric_extension
       -target_type=<target type of the metric extension>
       -name=<name of the metric extension
       -version=<version of the metric extension>

  Description:
    Publish a metric extension for use by all administrators.
  The metric extension must currently be a deployable draft.

  Options:
    -target_type=<target type>
      Target type of the metric extension.
    -name=<name>
      Name of the metric extension.
    -version=<version>
      Version of the metric extension to be published.


  emcli save_metric_extension_draft
       -target_type=<target type of the metric extension>
       -name=<name of the metric extension
       -version=<version of the metric extension>

  Description:
    Save a deployable draft of a metric extension. The metric
  extension must currently be in editable state. Once saved as
  draft, the metric extension will no longer be editable.

  Options:
    -target_type=<target type>
      Target type of the metric extension.
    -name=<name>
      Name of the metric extension.
    -version=<version>
      Version of the metric extension to be saved to draft.
```

User-Defined Metric Verbs

```
  emcli abort_udmmig_session
        -session_id=<sessionId>
        [-input_file=specific_tasks:<complete path to file>]

  Description:
    Abort the migration of user-defined metrics to MEs in a session

  Options:
    -session_id=<id of the session>
      Specify the id that was returned at time of session created,
      or from the output of udmmig_summary
    [-input_file=specific_tasks:<complete file path>]
      This optional parameter points at a file name that contains a
```

```
        target, user-defined metric,
      one per line in the following format:
      <targetType>,<targetName>,<collection name>
      Use targetType=Template to indicate a template
      Use * for collection name to abort all user-defined metrics for a target


  emcli analyze_unconverted_udms [-session_id=<sessionId>]

  Description:
    Analyze user-defined metrics and list unique user-defined metrics, any
possible matches, and
    templates that can apply these matching metric extensions
  Options:
    -session_id=<id of a session to be reanalyzed>
      Not specifying a session id causes the creation of a analysis
      session that contains all unconverted user-defined metrics. You can specify
      this session id in future invocations to get fresh analysis.


  emcli create_udmmig_session
      -name=<name of the session>
      -desc=<description of the session>
      [-udm_choice=<specific udm to convert>]*
      {-target=<type:name of the target to migrate> }*
      | {-input_file=targetList:<complete path to file>};      {-template=<name of
the template to update> }*
      | {-input_file=templateList:<complete path to file>}
      [-allUdms]

  Description:
    Creates a session to migrate user-defined metrics to metric extensions for
targets.

  Options:
    -name=<session name>
      The name of the migration session to be created.
    -desc=<session session description>
      A description of the migration session to be created.
    -udm_choice=<udm name>
      If the session should migrate specific user-defined metrics, specify them
      Otherwise, all user-defined metrics will be migrated
    -target=<type:name of target to migrate>
      The type:name of the target to be updated.
      Multiple values may be specified.
    -input_file=targetList:<complete file path>
      This takes a file name that contains a list of targets,
      one per line in the following format:
      <targetType>:<targetName>
    -template=<name of template to migrate>
      The name of the template to update.Multiple values may be specified
    -input_file=templateList:<complete file path>
      This takes a file name that contains a list of templates,
      one name per line
    -allUdms
      This forces the session to contain all user-defined metrics from targets and
      templates (default behavior just picks those not in a session)

  emcli list_unconverted_udms [-templates_only]
```

```
  Description:
    Get the list of all user-defined metrics that are not yet in a migration
session

  Options:
    -templates_only
      Only lists unconverted user-defined metrics in templates.


  emcli udmmig_list_matches
        -session_id=<sessionId>

  Description:
    Lists the matching metrics per user-defined metric in a migration session
  Options:
    -session_id=<id of the session>
      Specify the id that was returned at time of session created,
      or from the output of udmmig_summary


  emcli udmmig_request_udmdelete
        -session_id=<sessionId>
        -input_file=metric_tasks:<complete path to file>

  Description:
    Delete the user-defined metrics that have been replaced by Metric Extenions

  Options:
    -session_id=<id of the session>
      Specify the id that was returned at time of session created,
      or from the output of udmmig_summary
    -input_file=metric_tasks:<complete file path>
      This takes a file name that contains a target, user-defined metric,
      one per line in the following format:
      <targetType>,<targetName>,<collection name>


  emcli udmmig_retry_deploys
        -session_id=<sessionId>
        -input_file=metric_tasks:<complete path to file>

  Description:
    Retry the deployment of metric extensions to a target

  Options:
    -session_id=<id of the session>
      Specify the id that was returned at time of session created,
      or from the output of udmmig_summary
    -input_file=metric_tasks:<complete file path>
      This takes a file name that contains a target, user-defined metric,
      one per line in the following format:
      <targetType>,<targetName>,<collection name>


  emcli udmmig_submit_metricpicks
        -session_id=<sessionId>
        -input_file=metric_picks:<complete path to file>

  Description:
    Supply the metric picks to use to replace user-defined metrics per target in a
```

```
session

  Options:
    -session_id=<id of the session>
      Specify the id that was returned at time of session created,
      or from the output of udmmig_summary
    -input_file=metric_picks:<complete file path>
      This takes a file name that contains a target, user-defined metric, metric
pick,
      one per line in the following format:
      <targetType>,<targetName>,<collection name>,[N/E],<metric>,<column>
       using N if a new metric should be created or E if an existing
       metric is referenced.


  emcli udmmig_summary
      [-showAll]

  Description:
    Gets the summary details of all migration sessions in progress

  Options:
    -showAll
      This prints out all sessions including those that are complete.
      By default, only in-progress sessions are listed.


  emcli udmmig_update_incrules
        -session_id=<sessionId>
        -input_file=udm_inc_rules:<complete path to file>

  Description:
    Update Incident Rules that reference user-defined metrics with a reference to
    replacing metric extension.

  Options:
    -session_id=<id of the session>
      Specify the id that was returned at time of session created,
      or from the output of udmmig_summary
    -input_file=udm_inc_rules:<complete file path>
      This takes a file name that contains rule, user-defined metric, metric,
      one per line in the following format:
      <ruleset id>,<rule id>,<udm name>,<metric name>
```

5

Administration Groups and Template Collections

Administration groups simiplify adding targets to your managed environment by allowing Enterprise Manger to automatically categorize targets and, upon adding a target to a specific administration group, apply the appropriate monitoring settings.

This chapter covers the following topics:

5.1 What is an Administration Group?

Administration groups are a special type of group used to fully automate application of monitoring settings to targets upon joining the group. When a target is added to the group, Enterprise Manager applies monitoring settings using a template collection consisting of monitoring templates, compliance standards, and cloud policies. This completely eliminates the need for administrator intervention. With regular groups, Enterprise Manager applies the template settings only to those targets that are current members of the group.

Administration groups are mutually exclusive with other administration groups in terms of group membership: A target can only be a member of one administration group at any given point in time. Administration groups can also be used for hierarchically classifying targets in an organization. For example, all production databases located in the Denver data center running financial applications could be part of one administration group. All test databases located at the same data center running identical financial applications could be part of another administration group.

Administration groups are essentially regular (dynamic) Enterprise Manager groups that possess the following additional characteristics.

- Administration group members can be of different target types
- Members of administration groups can be either targets or other administration groups.
- An administration group itself can be a member of at most one administration group

- A system can be added to an administration group. You can specify that only certain members of the system be added to the administration group. By default, all members of the system will not be added to administration group as some system members may belong to other systems.
- Templates can be applied to administration groups and will be automatically applied to new members of the group.
- All administration groups are privilege propagating groups.

Key Advantages of Administration Groups

- Privilege propagation to simplify management of member target privileges.
- Administration groups defined and created based on membership criteria (target properties)
- Enterprise Manager automatically adds targets to an administration group if that target meets membership criteria.
- Users cannot directly add targets to the administration group.

5.2 Before You Begin

As with any management decision, the key to effective implementation is planning and preparation. The same holds true for administration groups.

Step 1: Plan Your Group Hierarchy

A target can directly belong to at most administration group. This prevents any conflicts occurring as a result of joining multiple administration groups with potentially different monitoring settings. To ensure a target belongs to only one administration group, Enterprise Manager only allows you to create a single administration group hierarchy. A target can join only one group in that hierarchy. Each administration group in the hierarchy is defined by membership criteria and a target is added to the group only if it meets the group's membership criteria. Although a target can directly belong to only one administration group, indirectly it can be part of multiple Administration Groups. An administration group will be indirectly part of its parent administration group, in addition to all connected administration groups further up in the hierarchy.

When defining the hierarchy, you should first categorize all managed targets into groups such that targets that are monitored and managed in the same way are part of the same group. The attributes used to define group membership criteria are based on *target properties*. Target properties can include attributes such as:

- Lifecycle Status
- Location
- Line of Business

In the following illustration, two administration groups are created, *Production* and *Test*, because monitoring settings for production targets are different from the monitoring settings for test targets.

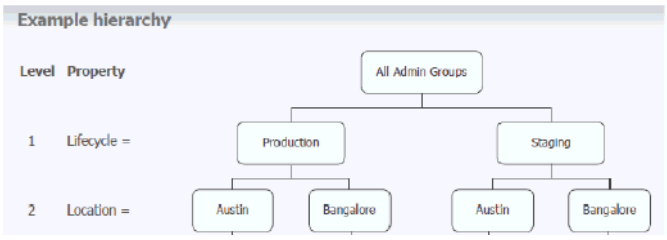

In this example, the group membership criteria are based on the *Lifecycle Status* target property. Targets whose *Lifecycle Status* is 'Production' join the Production group and targets whose *Lifecycle Status* is 'Test' join the Test group. For this reason, *Lifecycle Status* is the target property that determines the first level in the administration group hierarchy, and each value of the Lifecycle Status property determines the membership criteria of each administration group in the first level.

Additional levels in the administration group hierarchy can be added based on other target properties. Typically, additional levels are added if there are additional monitoring (or management) settings that need to be applied and these could be different for different subsets of targets in the administration group. For example, in the *Production* group, there could be additional monitoring settings for targets in *Finance* line of business that are different from targets in *Sales* line of business. In this case, an additional level based on *Line of Business* target property level would be added to the hierarchy as shown in the following illustration.

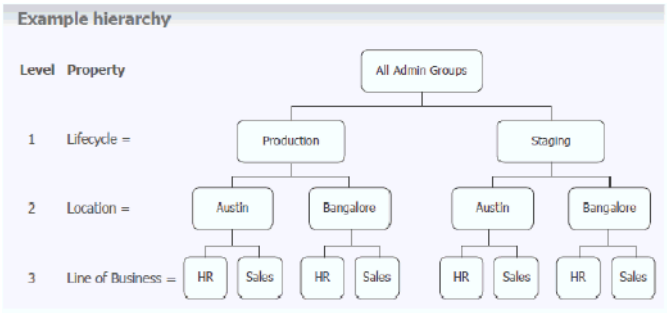

Step 2: Assign Target Properties

After establishing the desired organizational hierarchy, you must make sure properties are set correctly for each target to ensure they join the correct administration group. Using target properties, Enterprise Manager automatically places targets into the appropriate administration group without user intervention.

You can use **All Targets** page to view properties across all targets in your enterprise. To view target properties:

1. From the **Targets** menu, choose **All Targets** to display the All Targets page.
2. From the **View** pull-down menu, choose **Columns** and then choose **Show All**.
3. Alternatively, if you are interested in specific target properties, choose **Colums** and then select **Show More Columns** to display column selector, as shown in following graphic.

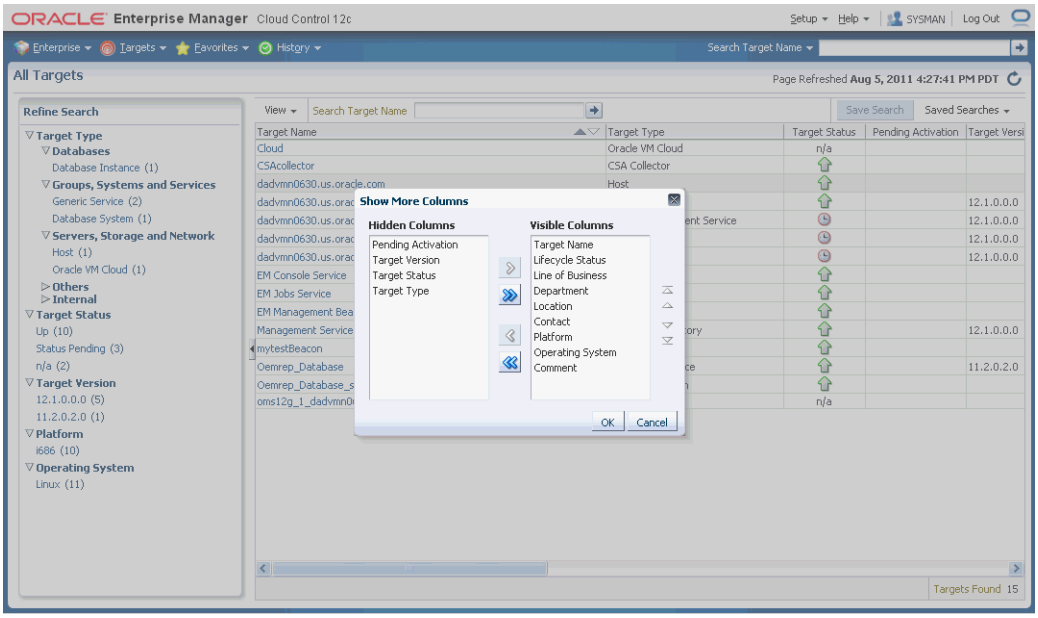

For small numbers of targets, you can change target properties directly from the Enterprise Manager console. You can right-click on any target on the **All Target**s page and choose **Target Setup** and then **Properties** from the context-sensitive menu to display that target's properties.

For large numbers of targets, it is best to use the Enterprise Manager Command Line Interface's (EMCLI) set_target_property_value verb to perform a mass update. For more information about this EMCLI verb, see the Enterprise Manager Command Line Interface guide.

COMMENT: Library link to EMCLI

Step 3: Prepare for Creating Template Collections

Template Collections are sets of Monitoring Templates, Compliance Standards and Cloud Policies that are applied to targets. Ensure all of these collection components are correctly defined before adding them to template collections.

Once you have completed all the planning and preparation steps, you are ready to begin creating an administration group.

5.3 Creating Administration Groups and Template Collections

With the preparatory work complete, you are ready to begin the four step process of creating an administration group and Template Collections. The administration group user interface is organized to guide you through the creation process, with each tab containing the requisite operations to peform each step.

Administration group creation involves:

1. Creating the administration groups hierarchy.
2. Creating template collections.
3. Associating template collections to administration groups.
4. Synchronizing the targets with the selected items.

The following graphic shows a completed administration group hierarchy with associated Template Collections. It illustrates how Enterprise Manager uses this to automate the application of target monitoring settings.

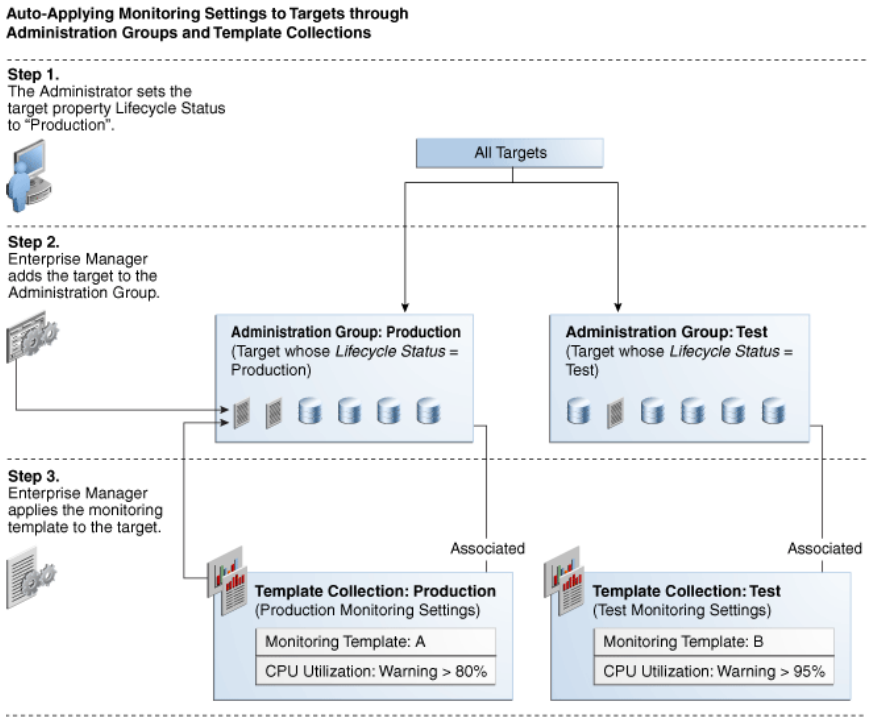

5.3.1 Developing an Administration Group

In order to create an administration group, you must have Enterprise Manager Super Administrator privileges. Super Administrator privileges are also required when editing and deleting administration groups.

To summarize, developing an administration group is performed in two phases:

- **Planning**
  - Plan your administration group hierarchy by creating a group hierarchy based on how you manage your targets.
  - Plan the management settings associated with the administration groups in the hierarchy.
    * Management settings: Monitoring settings, Compliance standard settings, Cloud policy settings
    * For Monitoring settings, you can have additional metric settings or override metric settings lower in your hierarchy
    * For Compliance standards or Cloud policies, additional rules/policies lower in the hierarchy are additive
- **Implementation**
  - Enter the group hierarchy definition and management settings in Enterprise Manager.
    * Create the administration group hierarchy.

* Create the monitoring templates, compliance standards, cloud policies and add these to template collections.
* Associate template collections with administration groups.
* Add targets to the administration group by assigning the appropriate values to the target properties such that Enterprise Manager automatically adds them to the appropriate administration group.

5.3.2 Creating an Administration Group

All administration group operations are performed from the administration group home page.

1. From the **Setup** menu, choose **Add Target-->Administration Groups**. The administration groups homepage displays.

2. Read the relevant information on the **Overview** page. The information contained in this page is also covered in the first part of this chapter.
3. Click **Hierarchy Definition** and define the group hierarchy. See "Defining a Hierarchy" on page 5-6.
4. Click **Template Collections** and define the template collections to be associated with the administration groups. See "Defining Template Collections" on page 5-8.
5. Click **Template/Group Associations** and assign the template collections with the appropriate administration groups. See "Associating Template Collections with Administration Groups" on page 5-10.

5.3.2.1 Defining a Hierarchy

On this page you define the administration group hierarchy that reflects your orgainzational infrastructure and which target properties are associated with a particular hierarchy level.

Adding a Hierarchy Level

1. From the Administration Group page, click the **Hierarchy Definition** tab.

2. From the **Hierarchies Level** menu, choose one of the available level properties used to build a hierarchy.
3. Choose the desired hierarchy level property. The chosen level is added to the hierarchy list.
4. Click on the hierarchy level type. Membership values associated with the property are displayed.
5. Click **Add**. The associated property value add dialog displays. Add the requisite value(s).
6. Continue adding hierarchy level until the group hierarchy is complete. The Preview window dynamically displays any changes you make to your administration group hierarchy.

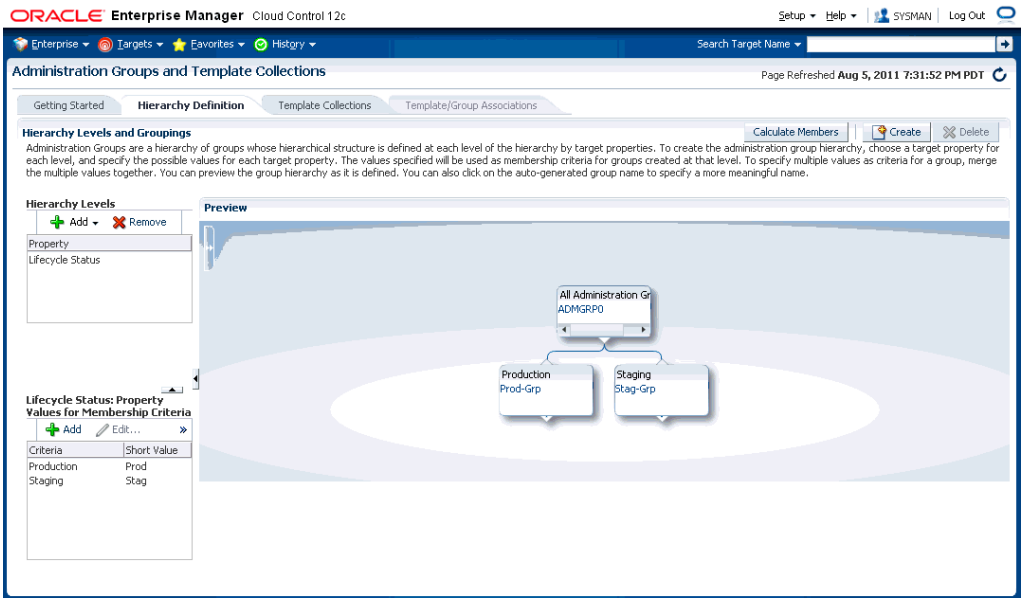

7. Click **Create** to define the hierarchy.

Adding/Removing Levels AFTER the Hierarchy is Created

By default, the hierarchy levels are locked upon creation. If you need to add/remove levels, you must first unlock the levels.

1. From the **Hierarchy Levels** region, click **Modify Levels**. Enterprise Manager displays a warning message stating All administration groups will be deleted and Template Collection associations will be lost.

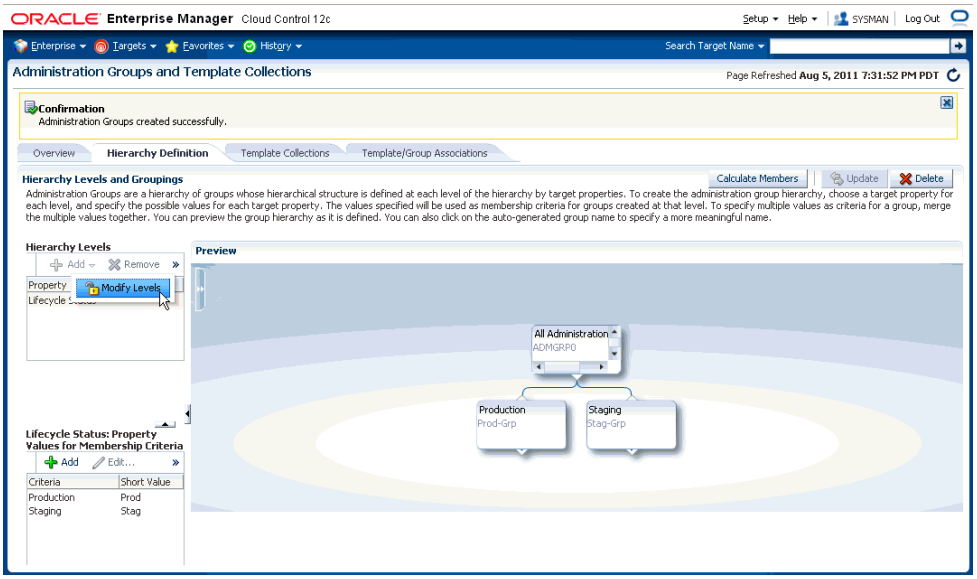

2. Click **Continue** from the confirmation message dialog. Once the levels are unlocked you can add or remove hierarchy levels.

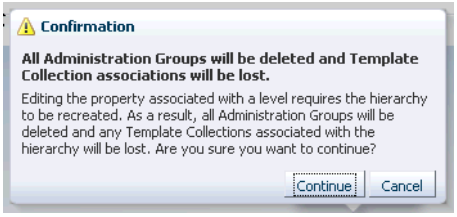

3. Add hierarchy levels as necessary.
4. Click **Recreate**. Once the hierarchy has been recreated, the levels will be locked.

> **Note:** You do not need to unlock the hierarchy levels to change property values.

Merging Hierarchy Property Values

Sometimes it is useful to treat multiple property values as one. This can be accomplished by merging the property values.

1. Select a hierarchy level from the list. The associated property values are displayed.
2. Select two or more property values by holding down the Shift key and clicking on the desired values.
3. Click **Merge**.

5.3.2.2 Defining Template Collections

A template collection is a assemblage of settings used to monitor/manage targets in Enterprise Manager. Multiple monitoring templates can be associated with a single (template collection) administration group. When members targets are added to an

administration group, they automatically inherit monitoring settings defined in the template collections.

> **Important:** Within a template collection, there can only be one template per target type. For example, you can have a template collection containing a template for database and a template for listener, but you cannot have a template collection containing 2 templates for databases.

You create template collections when you define administration groups. Template collections may consist of three types of monitoring/management setting categories:

- Monitoring Templates (monitoring settings)
- Compliance Standards (compliance policy rules)
- Cloud Templates (cloud policies such as determining when to start virtual machines or scale out clusters).

When creating a template collection, you can use the default monitoring templates, compliance standards, or cloud templates supplied with Enterprise Manager or you can create your own.

To create a template collection:

1. Click the **Template Collections** tab. The Template Collection page displays.

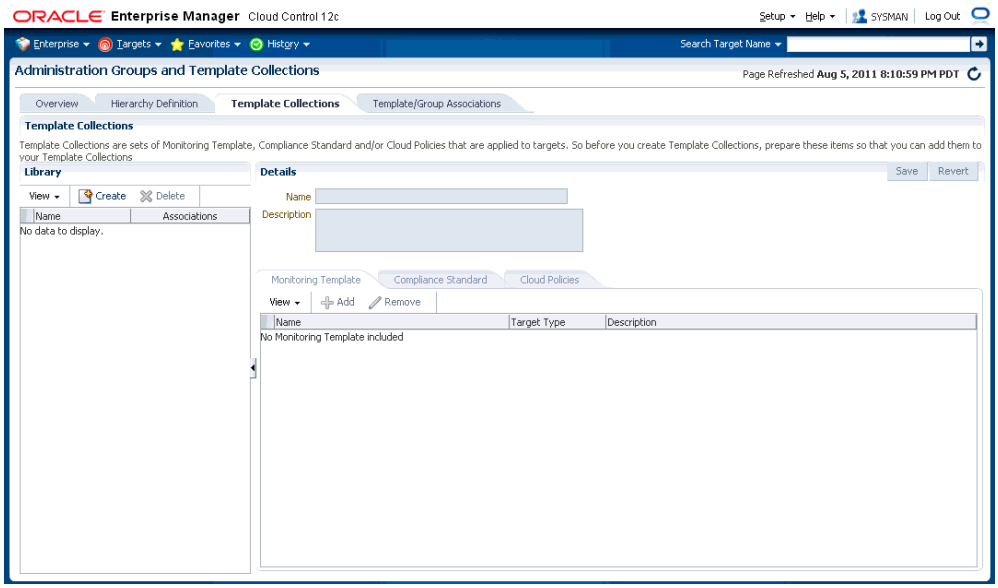

2. In the **Name** field, specify the template collection name.
3. Click the template collection member type you want to add (Monitoring Template, Compliance Standard, Cloud Polices). The requisite definition page appears.
4. Click **Add**. A list of available template entities appears.
5. Select the desired template entities you want added to the template collection.
6. Click **OK**.

7. Continue adding template entities (Monitoring Template, Compliance Standard, Cloud Policies) and as required.
8. Click **Save**. The newly defined collection appears in the **Template Collections Library**.
9. To create another template collection, click **Create** from the **Library** region and create and repeat steps two through eight. Repeat this process until you have created all required template collections.

> **Note:** When editing existing template collections, you can back out of any changes made during the editing session by clicking **Revert**. This restores the template collection to its state when it was last saved.

5.3.2.3 Associating Template Collections with Administration Groups

Once you have defined one or more template collections, you need to assign them to groups in the hierarchy. You can associate a template collection with one or more administration groups. The Template/Group Associations page displays the current administraton group hierarchy diagram. Each administration group in the hierarchy can only be associated with one template collection.

Associating a Template Collection with an Administration Group

1. Click the **Template/Group Associations** tab. The Template/Group Associations page displays.

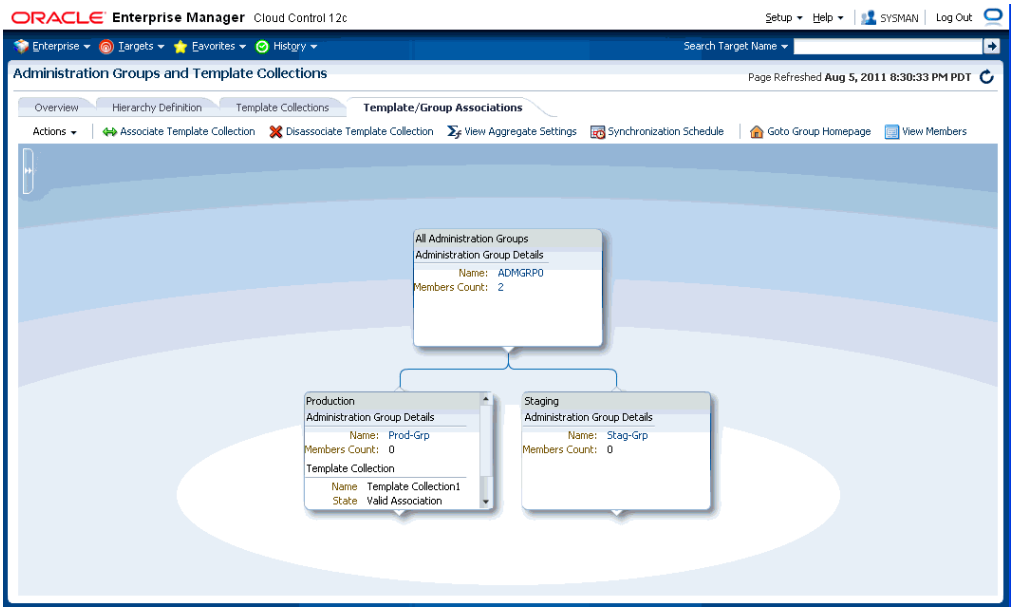

2. Select the desired administration group in the hierarchy.
3. Click **Associate Template Collection**. The **Choose a Template Collection** dialog displays.
4. Choose the desired template collection and click **Select**. The template collection icon appears in the selected group.

Note: All sub-nodes in the hierarchy will inherit the selected template collection.

5. Repeat steps 1-3 until template collections have been associated with the desired groups.

Note: The target privileges of the administrator who performs the association will be used when Enterprise Manager applies the template to the group.

Note: Settings applied at lower levels in the hierarchy override settings inherited from higher levels.

Synchronizing the Template Collections with the Administration Group

After you have associated the template collections with the requisite groups, you need to apply the template settings to member targets of the groups. This is called synchronization.

To synchronize template collections:

1. Click **Synchronization Schedule**. The Synchronization Schedule dialog displays.

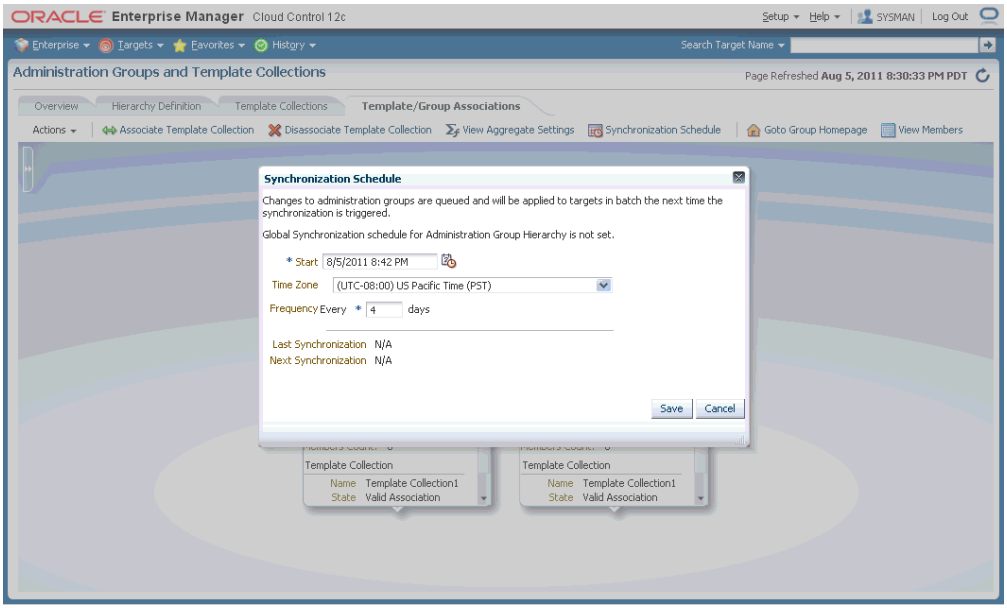

2. Choose a date and time you want the Administration Group-Template synchronization to take place. By default, the current date and time is shown.
3. Click **Save**.

Viewing Synchronization Status

You can check the current synchronization status for a specific administration group directly from the group's homepage.

1. Select an administration group in the hierarchy.

2. Click **Goto Group Homepage**.
3. From the **Synchronization Status** pane, you can view the status of the monitoring template, compliance standard, and/or cloud policies synchronization (In Sync, Pending, or Failed).

   You can initiate an immediate synchronization by clicking **Sync Now**.

Disassociating a Template Collection from a Group

To remove a template collection from an administration group.

1. From the hierarchy diagram, select the administration group with the template collection you wish to remove.
2. Click **Disassociate Template Collection**.

The template collection is immediately removed.

Viewing Aggregate Settings

For any administration group, you can easily view what template collection components (monitoring templates, compliance standards, and/or cloud policies) are associated with individual group members.

1. From the hierarchy diagram, select the desired administration group.
2. Click **View Aggregate Settings**.

   The **Aggregate Settings** page appears. This page displays all Monitoring Templates, Compliance Standards and Cloud Policies associated with the selected administration group (listed by target type).

Viewing the Administration Group Homepage

Like regular groups, each administration group has an associated group homepage providing a comprehensive overview of group member status and/or activity such as synchronization status, job activity, or critical patch advisories. To view administration group homepages:

1. From the hierarchy diagram, select an administration group.
2. Click **Goto Group Homepage.** The homepage for that particular administration group displays.

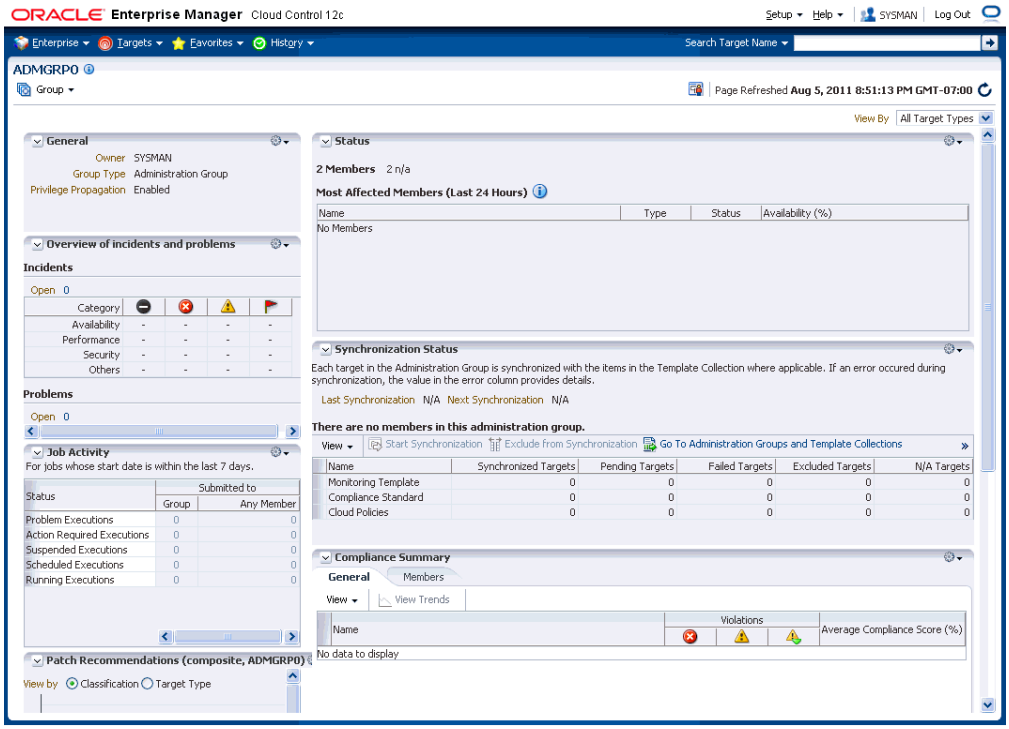

Identifying Targets Not Part of Any Administration Group

You can determine which targets do not belong to any administration group by generating an *Unassigned Targets Report*.

1. From the **Actions** menu, choose **Unassigned Targets Report**. The report lists all the targets that are not part of any administration group. The values for the target properties defining the administration groups hierarchy are shown.
2. From the **View** menu, choose the customization options to display only the desired information.
3. Click your browser 'back' button to return to the **Administration Groups and Template Collections** homepage.

5.4 Modifying Administration Groups

Modifying an administration group occurs at two levels: Editing the group members from the Administration Groups and Template Collections area and editing/configuring the Administration Group from the group homepage.

Editing Administration Group Members

Editing administration group members is accomplished through the target properties that define the group hierarchy.

1. From the **Setup** menu, choose **Add Target-->Administration Groups**.
2. From the **Hierarchy Definition** tab, modify the Hierarchy Levels and/or associated properties as required.
3. Click **Update**. Upon successful update, you are taken to the **Template/Group Association** page.

4. Click **Synchronization Schedule**.
5. From the **Synchronization Schedule** dialog, click **Sync Now** or set a date and time for synchronization.

Editing/Configuring the Administration Group

You edit and configure the administration group itself just as you would a regular group from the group homepage.

1. From the **Setup** menu, choose **Add Target-->Administration Groups**.
2. Click on the **Template/Group Associations** tab.
3. Select the desired group from the hierarchy diagram.
4. Click **Goto Group Homepage**.
5. Edit the group as appropriate.

5.5 Deleting Administration Groups

When you delete an administration group, any stored membership criteria is removed.

1. From the **Setup** menu, choose **Add Target-->Administration Groups**.
2. Click on the **Hierarchy Definition** tab.
3. Select the administration group(s) you want to remove.
4. Click **Discard**.

Note: You cannot delete a compound administration group, since it will break the hierarchy. If you must delete a compound administration group then a leaf-up approach should be followed for deletion.

---

6 Group Management

This chapter introduces the concept of group management and contains the following sections:

6.1 Introduction to Groups

Today's IT operations can be responsible for managing a great number of components, such as databases, application servers, hosts, or other components, which can be time consuming and impossible to manage individually. The Enterprise Manager Cloud Control group management system lets you combine components (called targets in Enterprise Manager) into logical sets, called groups. This enables you to organize, manage, and effectively monitor the potentially large number of targets in your enterprise.

Enterprise Manager Groups can include:

- Targets of the same type, such as:
  - All hosts in your data center
  - All of your production databases
- Targets of different types, such as:
  - The database, application server, listener, and host that are used in your application environment
  - Targets operating within a particular data center region

> **Note:** An Enterprise Manager "System," used specifically to group the components on which a service runs, is a special kind of Enterprise Manager group. Many of the functions and capabilities for groups and systems are similar.

Typically you can gather together targets that you want to manage as a group. If you use the target properties (for example, *Line of Business* or *Deployment Type*) to put

operational information about your targets in Enterprise Manager, you can use these properties when creating groups to locate targets. For example, you could search for all databases of *Deployment Type = Production* and belonging to *Line of Business 'HCM'*. You can also create a group hierarchy and use nested groups.

6.2 Managing Groups

By combining targets in a group, Enterprise Manager offers a wealth of management features that enable you to efficiently manage these targets as one group. Using the Group pages, you can:

- View a summary status of the targets within the group.
- Monitor outstanding alerts and incidents for the group collectively, rather than individually.
- Monitor the overall performance of the group.
- Perform administrative tasks, such as scheduling jobs for the entire group, or blacking out the group for maintenance periods.

You can also customize the console to provide direct access to group management pages.

6.2.1 Using the Groups Page

When you choose Groups from the Targets menu on the Enterprise Manager menu bar, the Groups page appears as seen in Figure 6–1. From the page you can view the currently available groups and perform the following tasks:

- View a list of all the defined groups.
- Search for existing groups and save search criteria for future searches.
- View a roll-up of the outstanding alerts and incidents for members in a group.
- Create administration groups, associate template collections, and disassociate template collections
- Add groups or privilege propagating groups, remove groups, and change the configuration of currently defined groups.
- Drill down from a specific group to collectively monitor and manage its member targets.

Redundancy groups and special high availability groups are not accessed from this Groups page. You can access them from the All Targets page.

Figure 6–1 Groups Home Page

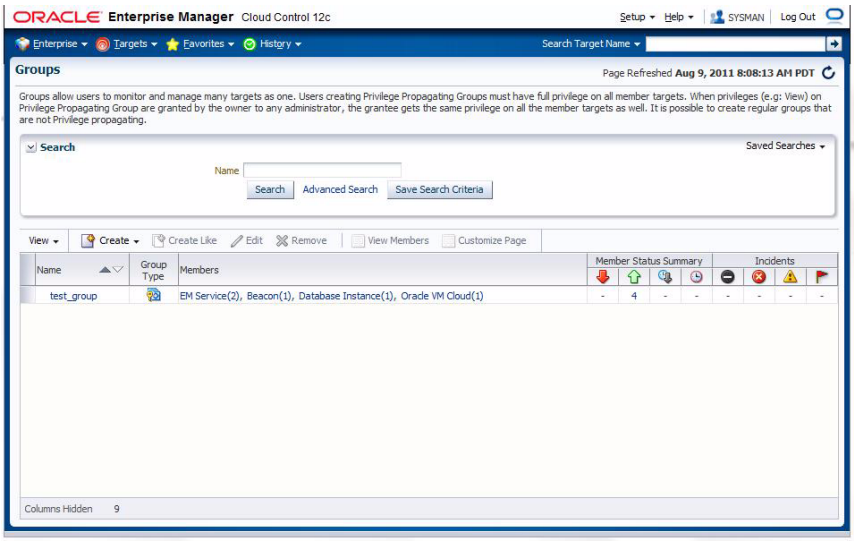

Groups Home Page

****

6.2.2 Group Home Page

The Group Home page, shown in Figure 6–2, enables you to quickly view key information about members of a group, eliminating the need to navigate to individual member targets to check on availability and performance. You can view the entire group on a single screen and drill down to obtain further details. The rolled up numbers include alerts and incidents for all members including those in nested groups. The Group Home page provides the following sections:

- A General section that shows the Owner, Group Type, and Privilege Propagation status.
- A Status section that shows how many member targets are in up, down, and unknown states. For nested groups, this segment shows how many targets are in up, down, and unknown states across all its sub-groups. The status roll up count is based on the unique member targets across all sub-groups. Consequently, even if a target appears more than once in sub-groups, it is counted only once in status roll ups. Click on member names to go to the member Status page.
- An Overview of Incidents and Problems section that displays the number of outstanding critical, warning, and error alerts associated with the current group. For nested groups, this segment shows how many targets are in an alert state across all its sub-groups.

  The rolled up information is shown for all the member targets regardless of their status. The status roll up count is based on the unique member targets across all sub-groups. Consequently, even if a target appears more than once in sub-groups, its alerts are counted only once in alert roll ups.

  Click on the number in the Problems column to go to the Incident Manager page to search, view, and manage exceptions and issues in your environment. By using Incident Manager, you can track outstanding incidents and problems.

- A Compliance Summary section that shows how many of your group members do not comply with Enterprise Manager policy rules. Non-compliant members are indicated with the number of critical, warning, and informational incidents along the Average Compliance Score (as a percentage) for each compliance rule. You can click on the Members tab to see the Member Targets and their types along with any violations and an average score for each member target.

  The numbers include the group-level incidents (if any group-level policy is defined), as well as group members violations. The rolled-up information for all policy categories, including security, is shown for all the member targets regardless of their status.

- A Job Activity section that displays the status for jobs that have started within the previous 7 days. The embedded table shows you the number of executions submitted to the group or any group members listed by status type, such as Problem Executions, Suspended Executions, and so on.

- A Blackouts section that allows you to create blackout periods and view the status of existing blackouts. The table displays the Scheduled and Active blackouts for each group or group member. Click on Create to define primary blackout identification information and assign targets to be blacked out.

- A Patch Recommendations section that shows the total number of Oracle critical patch advisories (including one or more critical patches) that are applicable to your enterprise, and the number of Oracle homes in your enterprise to which those patches should be applied. You can view the information by Classification or by Target Type.

  Click on the Current number link to go to the Group Critical Patch Advisories page. If your Oracle MetaLink Credentials are not configured, click Not Configured to go to the Patching Setup page. After you configure this page, Enterprise Manager collects information about Oracle critical patch advisories that are relevant to your enterprise.

- An Inventory and Usage section where you can view inventory summaries for deployments such as hosts, database installations, and fusion middleware installations on an enterprise basis or for specific targets. You can select an option such as Platform or Version to roll up inventory. Optionally, you can click See Details to navigate to the Inventory and Usage Details page where you can perform more detailed tasks such as viewing trends in inventory counts charted across a time line, revising selections to refresh chart and details based on new selections, or export deployment and details tables to CSV files.

- A Configuration Changes for Last 7 Days section that displays the number for configuration changes and Relationship changes incurred over the previous 7 days. Configuration history is a log of changes to a target, such as a group or beacon, recorded over a period of time. The recorded history includes changes both to configurations and to relationships. Relationships are the associations that exist among managed entities. You can click on the number of changes in either column to view more detailed information about the change.

Figure 6–2 Group Home Page

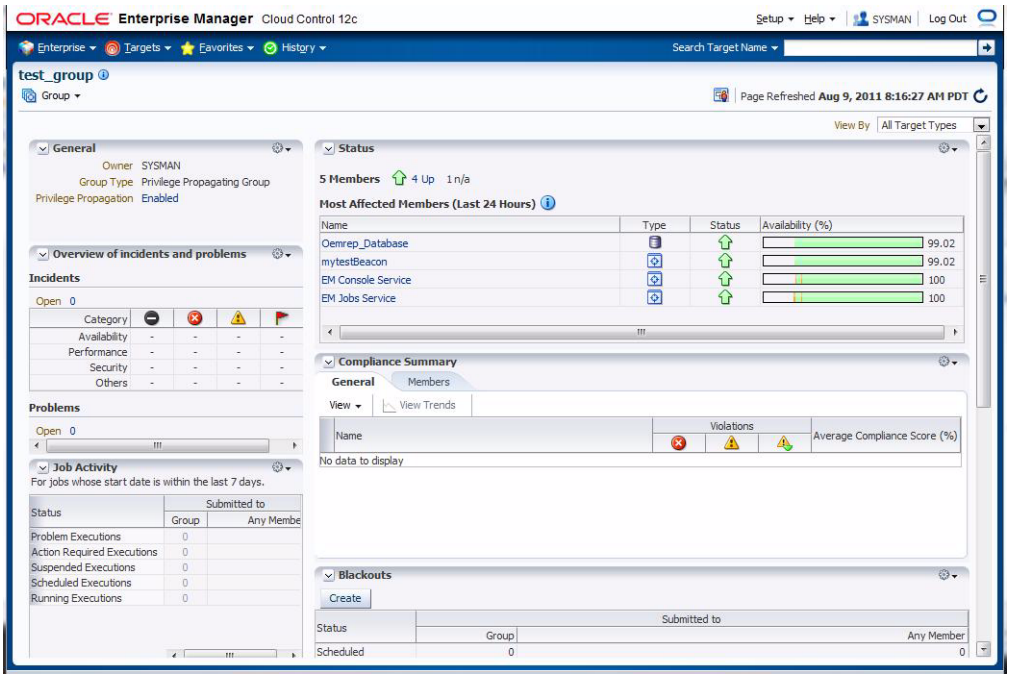

Group Home page

****

6.2.3 Group Charts Page

The Group Charts page, shown in Figure 6–3, enables you to monitor the collective performance of the group. Out-of-box performance charts are provided based on the type of members in the group. For example, when databases are part of the group, a Wait Time (%) chart is provided that shows the top databases with the highest wait time percentage values. You can view this performance information over the last 24 hours, last 7 days, or last 31 days. You can also add your own custom charts to the page.

You can access the Charts page by choosing Charts from the Monitoring sub-menu of the Group menu.

Figure 6–3 Group Charts Page

Group Charts Page

****

6.2.4 Group Members Page

The Group Members page, shown in Figure 6–4, summarizes information about the member targets in the group. It includes information on their current availability status, roll-up of open alerts and incidents, and key performance metrics based on the type of targets in the group.

You can visually assess availability and relative performance across all member targets. You can sort on any of the columns to rank members by a certain criterion (for example, database targets in order of decreasing wait time percentage). Default key performance metrics are displayed based on the targets you select, but you can customize these to include additional metrics that are important for managing your group.

Figure 6–4 Members for Group Page

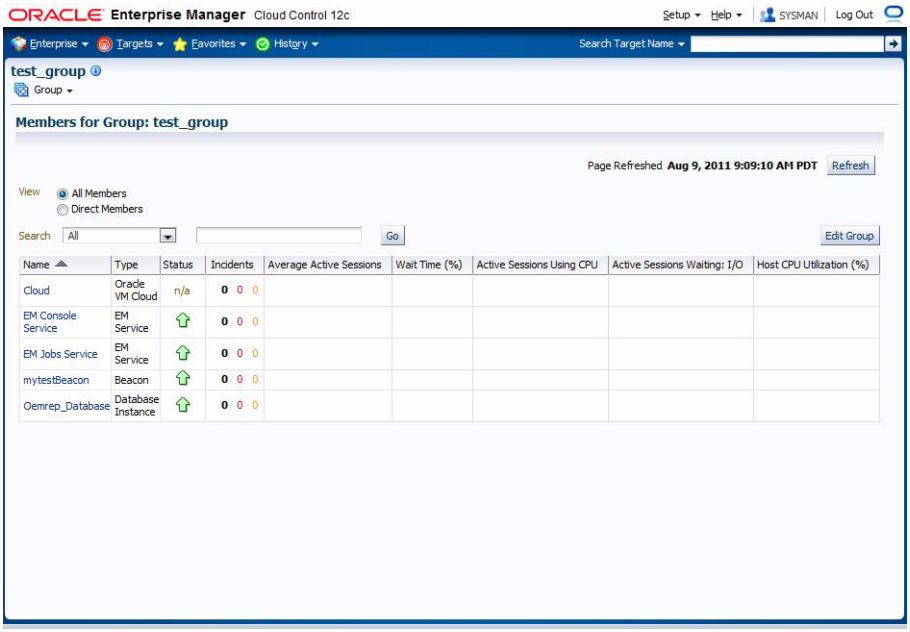

Group Members page

****

6.2.5 Viewing Group Status History

You can use the Group Status History page to view the historical availability of a member during a specified time period, view the current status of all group members, or access the home pages for members.

Bar graphs provide a historical presentation of the availability of group members during a time period you select from the View Data drop-down list. The color-coded graphs can show statuses of Up, Down, Under Blackout, Agent Down, Metric Collection Error, and Status Pending. You can select time periods of 24 hours, 7 days, or 31 days.

To view the current status or a member, you can click on a Status icon to go to the Availability page, which shows the member's current and past availability status within the last 24 hours, 7 days, or 31 days. Click a member Name to go to the member's Home page. You can use this page as a starting point when evaluating the performance of the selected member.

You can access the Group Status History page by choosing Status History from the Monitoring section of the Group menu.

6.2.6 System Dashboard

The System Dashboard, shown in Figure 6–5, enables you to pro-actively monitor the status and alerts in the group as they occur. The color-coded interface is designed to highlight problem areas using the universal colors of alarm—targets that are down are highlighted in red, metrics in critical severity are shown as red dots, metrics in warning severity are shown as yellow dots, and metrics operating within normal boundary conditions are shown as green dots.

Using these colors, you can easily spot the problem areas for any target and drill down for details as needed. An alert table is also included to provide a summary for all open alerts in the group. The alerts in the table are presented in reverse chronological order to show the most recent alerts first, but you can also click on any column in the table to change the sort order.

The Dashboard allows you to drill down for more detailed information. You can click on the following items in the Dashboard for more information:

- A target name to access the target home page
- A group or system name to access the System Dashboard
- Status icon corresponding to specific metric columns to access the metric detail page
- Alerts icon for a group to access the Alerts page for that aggregate object
- Status icon for a metric with key values to access the metric page with a list of all key values
- Status icon for a metric with a specific key value to access the metric detail page with the specified key
- Dashboard header to access the group home page
- Status icon for down, critical or warning alerts to access the Alerts page
- Alerts messages to access the metric detail page containing the alert history for the target

Click **Customize** to access the Edit Group pages. By default, Enterprise Manager takes you to the Edit Group Dashboard page where you can change the target display and data refresh frequency. However, you can also modify any other group properties that affect the content of the System Dashboard. Columns that appear in the Dashboard target area mirror the columns that appear in the Edit Group Columns page. To display additional columns, click Modify on the Edit Group Columns page and add the desired metric columns.

In the "Group by Target Type" mode, the Dashboard displays information of the targets based on the specific target types present in the group or system. The statuses and alerts displayed are rolled up for the targets in that specific target type.

Columns that appear in the Dashboard target area mirror the columns that appear in the Edit Group Columns page. To display additional columns, click Modify on the Edit Group Columns page and add the desired metric columns.

If you minimize the dashboard window, pertinent alert information associated with the group or system is still displayed in the Microsoft Windows toolbar. For example, (#1 X3 !5) denotes there is 1 Target Down Alert, 3 Critical Alerts and 5 Warning Alerts associated with this group or system.

Figure 6–5 System Dashboard Page

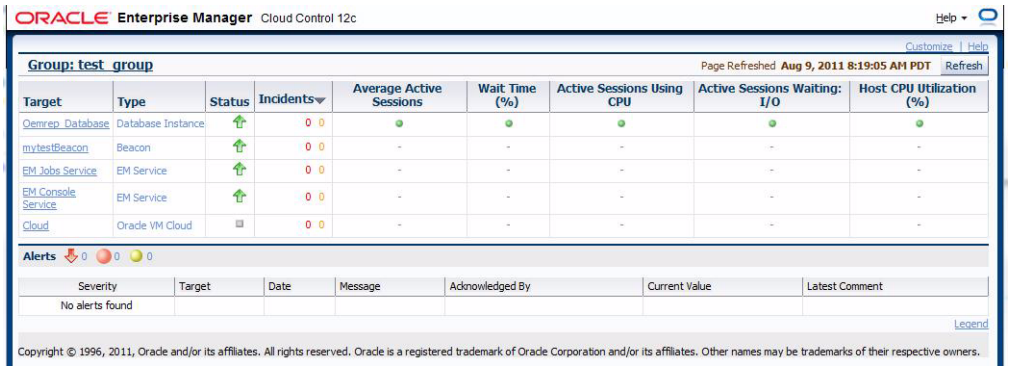

6.3 Out-of-Box Reports

Enterprise Manager provides several out-of-box reports for groups as part of the reporting framework, called Information Publisher. These reports display important administrative information, such as hardware and operating system summaries across all hosts within a group, and monitoring information, such as outstanding alerts and incidents for a group.

You can access these reports from the **Information Publisher Reports** menu item on the Groups menu.

See Also: Chapter 10, "Information Publisher"

6.4 Redundancy Groups

A redundancy group is a group that contains members of the same type that function collectively as a unit. A type of redundancy group functions like a single logical target that supports a status (availability) metric. A redundancy group is considered up (available) if at least one of the member targets is up.

You can create and administer a redundancy group from the All Targets page. Redundancy groups support all group management features previously discussed.

When you define the Redundancy Group, you must choose the member type for the members in the Redundancy Group.

You can define the options for how availability of the redundancy group is calculated by selecting either Number or Percentage:

- Number - When you choose Number, you can specify either the number of member targets that should be up in order for the group to be considered up, or the number of member targets that should be down in order for the group to be considered down.
- Percentage - When you choose Percentage, you can specify either the minimum percentage of member targets that must be up in order for the group to be considered up, or the minimum percentage of member targets that are down in order to consider the group to be down. If you choose Percentage, the required number of member targets will be rounded off to the next integer. For example if you define the Percentage as 50% and the total number of member targets is 5, then the value used for calculating the availability will be 3.

Figure 6–6 shows the Create Redundancy Group page while defining the group using Percentage availability.

Figure 6–6 Redundancy Group Example

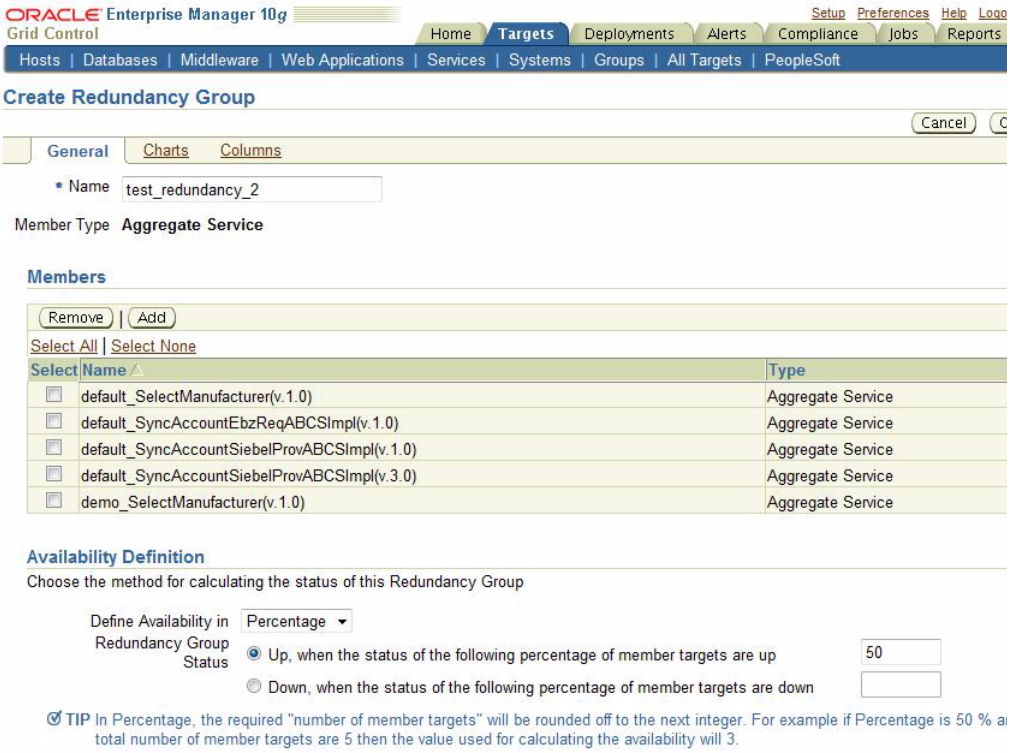

Do not use redundancy groups if the group you want to model is an Oracle Real Application Clusters database, host cluster, HTTP server high availability group, or OC4J high availability group. Instead, you can use the following specialized target types for this purpose:

- Cluster
- Cluster Database
- HTTP HA Group
- OC4J HA Group

6.5 Privilege Propagating Groups

Privilege propagating groups enable administrators to grant privileges to other administrators in a manner in which new administrators get the same privileges as its member targets. For example, granting *operator* privilege on a group to an Administrator grants him the *operator* privilege on its member targets and also to any members that will be added in the future. Privilege propagating groups can contain individual targets or other privilege propagating groups.

Privileges on the group can be granted to an Enterprise Manager user or a role. Use a role if the privileges you want to grant are to be granted to a group of EM users. See Figure 6–7, "Granting Privileges On a Group To a Role".

Figure 6–7 Granting Privileges On a Group To a Role

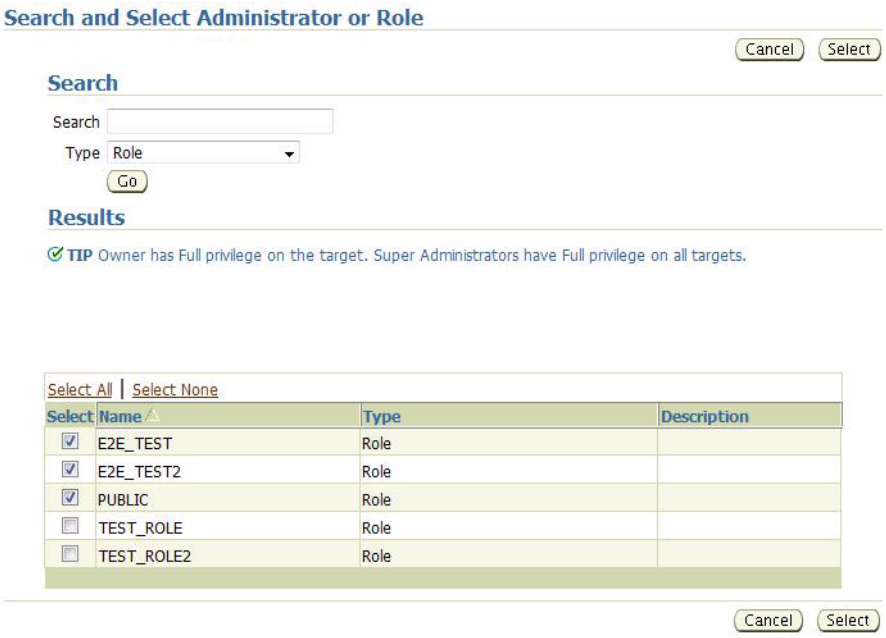

For example, suppose you create a large privilege propagating group and grant a privilege to a role which is then granted to administrators. If new targets are later added to the privilege propagating group, then the administrators receive the privileges on the target automatically. Additionally, when a new administrator is hired, you only need to grant the role to the administrator for the administrator to receive all the privileges on the targets automatically.

6.5.1 Creating Privilege Propagating Groups

The privilege propagating group creation function is a privileged activity. The privilege propagating group feature contains two new privileges:

- Create Privilege Propagating Group

  This privileged activity allows the administrators to create the privilege propagating groups. Administrators with this privilege can create propagating groups and delegate the group administration activity to other users.

- Group Administration

  This privilege can be granted to administrators on specific group targets and is used to delegate the group administration activities to other administrators. It is granted to both conventional and privilege propagating groups.

6.5.2 Using the Group Administration Privilege

The Group Administration Privilege is available for both Privilege Propagating Groups and conventional groups. If you are granted this privilege, you can grant access to the group to other Enterprise Manager users without having to be the SuperAdministrator to grant the privilege.

6.5.3 Adding Members to Privilege Propagating Groups

The target privileges granted on a propagating group are propagated to member targets. The administrator grants target objects scoped to another administrator, and the grantee maintains the same privileges on member targets. The propagating groups maintain the following features:

- The administrator with a Create Privilege Propagating Group privilege will be able to create a propagating group
- To add a target as a member of a propagating group, the administrator must have *Full* target privileges on the target

You can add any non-aggregated target as the member of a privilege propagating group. For aggregated targets in Cloud Control version 12*g*, cluster and RAC databases and other propagating groups can be added as members (cluster and RAC databases must be added via the *emcli* verb). There is no support for this through the Enterprise Manager interface in version 10.2.0.5. Grid Control version 11g, however, supports more aggregated target types, such as redundancy groups, systems and services. These, along with cluster and RAC databases, can be added in version 12*g* via the Cloud Control Console.

If you are not the group creator, you must have at least the *Full* target privilege on the group to add a target to the group.

6.5.4 Converting Conventional Groups to Privilege Propagating Groups

In Enterprise Manager version 12*g* you can convert conventional groups to privilege propagating groups (and vice-versa) through the use of the specified EMCLI verb. Two new parameters have been added in the *modify_group* EMCLI verb:

- *privilege_propagation*

  This parameter is used to modify the privilege propagation behavior of the group. The possible value of this parameter is either true or false.

- *drop_existing_grants*

  This parameter indicates whether existing privilege grants on that group are to be revoked at the time of converting a group from privilege propagation to normal (or vice versa). The possible values of this parameter are yes or no. The default value of this parameter is yes.

These same enhancements have been implemented on the following EMCLI verbs: *modify_system*, *modify_redundancy_group*, and *modify_aggregrate_service*.

The EMCLI verb is listed below:

```
emcli modify_group
   -name="name"
   [-type=<group>]
   [-add_targets="name1:type1;name2:type2;..."]...
   [-delete_targets="name1:type1;name2:type2;..."]...
   [-privilege_propagation = true/false]
   [-drop_existing_grants = Yes/No]
```

For more information about this verb and other EMCLI verbs, see the *EMCLI Reference Manual*.

6.6 Administration Groups

Administration Groups are a special type of group used to fully automate application of monitoring settings to targets upon joining the group. When a target is added to the group, Enterprise Manager applies monitoring settings using a template collection consisting of monitoring templates, compliance standards, and cloud policies. This completely eliminates the need for administrator intervention. With regular groups, Enterprise Manager applies the template settings only to those targets that are current members of the group.

Administration groups are mutually exclusive with other administration groups in terms of group membership: A target can only be a member of one administration group at any given point in time. Administration groups can also be used for hierarchically classifying targets in an organization. For example, all production databases located in Denver datacenter running financial applications could be part of one administration group. All test databases located at the same datacenter running identical financial applications could be part of another administration group.

The following are key attributes of Administration Groups:

- Privilege propagation to simplify management of member target privileges
- Administration groups defined and created based on membership criteria (target properties)
- Enterprise Manager automatically adds targets to an administration group if that target meets membership criteria
- Users cannot directly add targets to the administration group

6.6.1 Working with Administration Groups

Administration groups are essentially regular (dynamic) Enterprise Manager groups that possess the following additional characteristics:

- Administration group members can be of different target types.
- Members of administration groups can be either targets or other administration groups.
- An administration group itself can be a member of at most one administration group.
- You can add a system to an administration group. You can specify that only certain members of the system be added to the administration group by setting target properties for the system members which need to become members of the same administration group. By default, all members of the system will not be added to the administration group as some system members may belong to other systems.
- Templates can be applied to administration groups and will be automatically applied to new members of the group.
- All administration groups are privilege propagating groups. Groups, Generic Systems, Generic Services and any other non-privilege propagating aggregates cannot become members of an administration group.

Use these general steps to create Administration Groups.

Step 1: Design Your Grid's Hierarchy

First, design a way to organize your targets so they make a logical hierarchy of your organization. The hierarchy shown at the bottom is one example. Properties you can

use to manage your hierarchy are global target properties like Contact, Lifecycle Status, Location, Line of Business, Department, and so on.

The order of the properties that make up your hierarchy matters. It determines the order that template collections are applied to groups in the hierarchy. Settings from template collections at the lowest level of the hierarchy override settings from the template collections at higher levels.

Step 2: Assign Properties To Targets

Be sure all of the targets in your organization have the correct property values associated. You can use All Targets page to view the state of each property across all targets in your grid. As for associating properties to targets, while you could use Enterprise Manager to do this, EMCLI will be more efficient, especially if you have a large number of changes to make.

Step 3: Prepare for Creating Template Collections

Template Collections are sets of Monitoring Templates, Compliance Standards and/or Cloud Policies that are applied to targets. So before you create Template Collections, prepare these items so that you can add them to your Template Collections.

Step 4: Manage Administration Groups and Template Collections

Setup the Administration Groups hierarchy. Administration Groups are a hierarchy of groups whose hierarchical structure is defined at each level of the hierarchy by target properties. To create the administration group hierarchy, choose a target property for each level, and specify the possible values for each target property. The values specified will be used as membership criteria for groups created at that level. To specify multiple values as criteria for a group, merge the multiple values together. You can preview the group hierarchy as it is defined. You can also click on the auto-generated group name to specify a more meaningful name.

Create template collections. Template Collections are sets of Monitoring Template, Compliance Standard and/or Cloud Policies that are applied to targets. Therefore before you create Template Collections, prepare these items so that you can add them to your Template Collections. For more information, see Working With Template Collections.

Next, associate Template Collections to Administration Groups.

Finally, synchronize the targets with the selected items.

6.6.2 Working With Template Collections

A template collection is a assemblage of settings used to monitor/manage targets in Enterprise Manager. Template collections are assigned to administration groups. When members targets are added to an administration group, they automatically inherit monitoring settings specified in the template collections.

Within a template collection, there can only be one template per target type. For example, you can have a template collection containing a template for database and a template for listener, but you cannot have a template collection containing 2 templates for databases.

You create template collections when you define administration groups. Template collections may consist of three types of monitoring/management setting categories:

- Monitoring Templates (monitoring settings)
- Compliance Standards (compliance policy rules)

- Cloud Templates (cloud policies such as determining when to start virtual machines or scale out clusters)

When creating a template collection, you can use the default monitoring templates, compliance standards, or cloud templates supplied with Enterprise Manager or you can create your own.

6.6.3 Developing an Administration Group

In order to create an administration group, you must have the requisite privileges. You must be logged in as an Enterprise Manager Super Administrator,. The same privilege requirements apply to editing and deleting administration groups.

Developing an administration group involves the following process:

1. Plan your administration group hierarchy by creating a group hierarchy based on how you manage your targets.

   Example: If production targets are monitored differently from non-production targets, then 'Lifecycle Status' property (Production versus Non-Production) can be used as criteria in the administration group hierarchy.

2. Plan the management settings associated with the administration groups in the hierarchy.
   - Management settings: Monitoring settings, Compliance standard settings, Cloud policy settings
   - For Monitoring settings, you can have additional metric settings or override metric settings lower in your hierarchy
   - For Compliance standards or Cloud policies, additional rules/policies lower in the hierarchy are additive
3. Enter the group hierarchy definition and management settings in Enterprise Manager.
   - Create the administration group hierarchy
   - Create the monitoring templates, compliance standards, cloud policies and add these to template collections
   - Associate template collections with administration groups
4. Add targets to the administration group by assigning the appropriate values to the target properties such that Enterprise Manager automatically adds them to the appropriate administration group.

6.6.4 Creating an Administration Group

Follow these steps to create an Administration Group. Steps one through three are performed on the Administration Group homepage. To access this page:

1. From the **Setup** menu, choose **Administration Groups** from the **Add Target** sub menu. The Administration Groups homepage displays.
2. Read the relevant information on the Overview tab.
3. Click **Hierarchy Definition** and define the group hierarchy. See the Hierarchy Definition page online help for specific task information.
4. Click **Template Collections** and define the template collections to be associated with the administration groups. See the Template Collections page online help for specific task information.

5. Click **Template/Group Associations** and assign the template collections with the appropriate administration groups. See the Template/Group Associations page online help for specific task information.

6.6.5 Modifying Administration Groups

Modifying an administration group occurs at two levels: Editing the group members from the Administration Groups and Template Collections area and editing/configuring the Administration Group from the Group homepage.

6.6.5.1 Editing Administration Group Members

Editing administration group members is accomplished through the target properties that define the group hierarchy.

1. From the **Setup** menu, choose **Administration Groups** from the **Add Target** sub menu.
2. From the Hierarchy Definition tab, modify the Hierarchy Levels and/or associated properties as required.
3. Click **Update**. Upon successful update, you are taken to the Template/Group Association page.
4. Click **Synchronization Schedule**.
5. From the Synchronization Schedule dialog, click **Sync Now** or set a date and time for synchronization.

6.6.5.2 Editing/Configuring the Administration Group

You edit and configure the administration group itself just as you would a regular group from the group homepage.

1. From the **Setup** menu, choose **Administration Groups** from the **Add Target** sub menu.
2. Click on the **Template/Group Associations** tab.
3. Select the desired group from the hierarchy diagram.
4. Click **Goto Group Homepage**.
5. Edit the group as appropriate.

6.6.6 Deleting an Administration Group

When you delete an administration group, any stored membership criteria is removed. Note: You cannot delete a compound administration group, since it will break the hierarchy. If you must delete a compound administration group then a leaf-up approach should be followed for deletion.

1. From the **Setup** menu, choose **Administration Groups** from the **Add Target** sub menu.
2. Click on the **Hierarchy Definition** tab.
3. Select the administration group(s) you want to remove.
4. Click **Discard**.

7

Job System and Corrective Actions

Today's IT environments have many sets of components, so it is beneficial to minimize the time needed to support these IT components and eliminate the human error associated with component maintenance. The Enterprise Manager Cloud Control Job System can automate routine administrative tasks and synchronize components in your environment so you can manage them more efficiently.

This chapter facilitates your usage of the Job System by presenting instructional information in the following sections:

7.1 Job System Purpose and Overview

The Enterprise Manager Job System serves these purposes:

- Automates many administrative tasks; for example: backup, cloning, and patching
- Enables you to create your own jobs using your own custom OS and SQL scripts
- Enables you to create your own multi-task jobs comprised of multiple tasks

A job is a unit of work that you define to automate commonly-run tasks. Scheduling flexibility is one of the advantages of jobs. You can schedule a job to start immediately or start at a later date and time. You can also run the job once or at a specific interval, such as three times every month.

The Job Activity page (Figure 7–1) is the hub of the Job System. From this page, you can:

- Search for existing job runs and job executions filtered by name, owner, status, scheduled start, job type, target type, and target name
- Create a job
- View or edit the job definition
- Create like, copy to library, suspend, resume, stop, and delete a job

- View results, edit, create like, suspend, resume, retry, stop, and delete a job run or execution

Figure 7–1 Job Activity Page

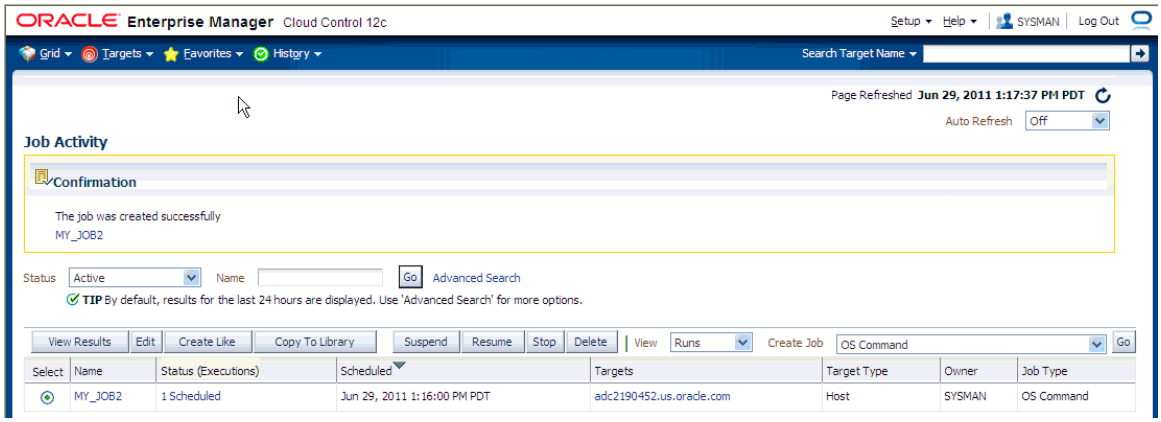

Besides accessing the Job Activity page from the Enterprise menu, you can also access this page from any Database or Cluster Database property page (Home, Performance, Availability, and so forth) by selecting Job Activity from the Oracle Database menu. When you access this page from these alternate locations, rather than showing the entire list of jobs, the Job Activity page shows a subset of the jobs associated with the particular target.

7.1.1 What Are Job Executions and Job Runs?

Job executions are usually associated with one target, such as a patch job on a particular database. When a job is run against multiple targets, each execution may execute on one target.

Job executions are not always a one-to-one mapping to a target. Some executions have multiple targets, such as comparing hosts. A few jobs have no target.

When you submit a job to many targets, it would be tedious to examine the status of each execution of the job against each target. For example, suppose you run a backup job against several databases. A typical question would be: Were all the backup jobs successful, and if not, which jobs failed? If this backup job runs every week, you would want to know which backups were successful and those that failed each week.

With the Job System, you can easily get these answers by viewing the *job run*. A job run is the summary of all job executions of a job that ran on a particular scheduled date. For example, if you have a job scheduled for March 5th, you will have a March 5 job run. The job table that shows the job run provides a roll-up of the status of the executions, such as Succeeded, Failed, or Error.

7.1.2 Operations on Job Executions and Job Runs

Besides supporting the standard job operations of create, edit, create like, and delete, the Job System enables you to:

- **Suspend jobs —**

  You can suspend individual executions or entire jobs. For example, you may need to suspend a job if a needed resource was unavailable, or the job needs to be postponed.

If a job is scheduled to repeat but is suspended past the scheduled repeat time, the execution of this job would be marked "Skipped."

- **Resume jobs** —

  After you suspend a job, any scheduled executions do not occur until you decide to resume the job.

- **Retry failed executions** —

  When analyzing individual executions or entire jobs, it is useful to retry a failed execution after you determine the cause of the problem. This alleviates the need to create a new job for that failed execution. When you use the Retry operation in the Job System, Enterprise Manager provides links from the failed execution to the retried execution and vice versa, should it become useful to retroactively examine the causes of the failed executions. Only the most recent retry is shown in the Job Run page.

With regard to job runs, the Job System enables you to:

- **Delete old job runs**
- **Stop job runs**
- **Retry job runs**

> **See Also:** For more information on job executions and runs, refer to Enterprise Manager Cloud Control online help.

7.2 Preliminary Considerations

Before proceeding to the procedural information presented in Section 7.3, "Creating Jobs" on page 7-4, it is advisable to read these topics presented in the sections below:

- Creating Scripts
- Sharing Job Responsibilities
- Jobs and Groups

7.2.1 Creating Scripts

Besides predefined job tasks, you can define your own job tasks by writing code to be included in OS and SQL scripts. The advantages of using these scripts include:

- When defining these jobs, you can use target properties.
- When defining these jobs, you can use the job library, which enables you to share the job and make updates as issues arise. However, you need to resubmit modified library jobs for them to take effect.
- You can submit the jobs against multiple targets.
- You can submit the jobs against a group. The job automatically keeps up with changes to group membership.
- For host command jobs, you can submit to a cluster.
- For SQL jobs, you can submit to a Real Application Cluster.

7.2.2 Sharing Job Responsibilities

To allow you to share job responsibilities, the Job System provides job privileges. These job privileges allow you to share the job with other administrators. Using privileges, you can:

- Grant access to the administrators who need to see the results of the job.
- Grant Full access to the administrators who may need to edit the job definition or control the job execution (suspend, resume, stop).

You can grant these privileges on an as-needed basis.

7.2.3 Jobs and Groups

Besides submitting jobs to individual targets, you can submit jobs against a group of targets. Any job that you submit to a group is automatically extended to all of its member targets that match the target type of the job, and accounts for the membership of the group as it changes.

For example, if a Human Resources job is submitted to the Payroll group, then a new host is added to this group, the host automatically becomes part of future Human Resources job runs. For instance, for a daily repeating job scheduled for 10:00 a.m. today, if you add a target before that time, the new target would be part of the job run. However, if you add a target after that time today, the target would not be part of today's run, but would be part of the next run. Additionally, if the Payroll group is comprised of diverse targets (for example: databases, hosts, and application servers), the job only runs against applicable targets in the group.

By accessing the Group Home page, you can analyze the job activity for that group.

See Also: Chapter 6, "Group Management"

7.3 Creating Jobs

Your first task in creating a job from the Job Activity page is to choose a job type, which the next section, Selecting a Job Type, explains. The most typical job types are OS command jobs, script jobs, and multi-task jobs, which are explained in these subsequent sections:

- Creating an OS Command Job
- Creating a SQL Script Job
- Creating a Multi-task Job

7.3.1 Selecting a Job Type

Using the Job System, you can create a job by selecting one of the job types from the Create Job drop-down in the Job Activity page. The most commonly used types are as follows:

- **OS Command** — Runs an operating system command or script.
- **SQL Script** — Runs a user-defined SQL or PL/SQL script.
- **Multi-Task** — Use to specify primary characteristics for multi-task jobs or corrective actions. Multi-task jobs enable you to create composite jobs by defining tasks, with each task functioning as an independent job. You edit and define tasks similarly to a regular job.

Blocked Agents Job Type

A blocked Agent is a condition where the Oracle Management Server (OMS) rejects all heartbeat or upload requests from the blocked Agent. Therefore, a blocked Agent cannot upload any alerts or metric data to the OMS. However, blocked Agents continue to collect monitoring data.

On a blocked Agent, the OMS “ignores” requests from the blocked Agent, thereby reducing the workload on the OMS. For example, by using this feature for an Agent that fails to upload properly, you can block the Agent until you can resolve the upload issue.

An Agent can become blocked under the following circumstances:

- The system detects that the Agent is no longer sending the correct state. This can occur after a failed recovery, or when users have corrupted state files. The OMS can detect some of the corruptions, and when it finds one, it blocks the Agent until the problem has been resolved.
- A superuser has blocked an Agent to prevent a "rogue" Agent from flooding the system with errors and bad data.

When an Agent is blocked for a long period of time and the Agent is kept running, it eventually must stop monitoring, because it will run out of local disc space to store all of the results. However, this is not an issue, because the "state" of the Agent was corrupt anyway. Therefore, unless corrective actions were taken, the Agent should remain blocked so that no data then penetrates the system.

7.3.2 Creating an OS Command Job

Use this type of job to run an operating system command or script. Tasks and their dependent steps for creating an OS command are discussed below.

Task 1 Initiate Job Creation

1. From the Enterprise menu, select **Jobs**, then Job Activity.
2. Select **OS Command** from the Create Job drop-down, then click **Go**. The **General property page** of the Create OS Command Job page appears.

Task 2 Specify General Job Information

Perform these steps on the General property page:

1. Provide a required Name for the job, then select a Target Type from the drop-down.

   After you have selected a target of a particular type for the job, only targets of that same type can be added to the job. If you change target types, the targets you have populated in the Targets table disappear, as well as parameters and credentials for the job.

   If you specify a composite as the target for this job, the job executes only against targets in the composite that are of the selected target type. For example, if you specify a target type of host and a group as the target, the job only executes against the hosts in the group, even if there are other non-host targets in the group.
2. Click **Add**, then select one or more targets from the Search and Select: Targets pop-up window. The targets now appear in the Targets table.
3. Click the **Parameters** property page link.

Task 3 Specify Parameters

Perform these steps on the Parameters property page:

1. Select either **Single Operation** or **Script** from the Command Type drop-down.

   The command or script you specify executes against each target specified in the target list for the job. The Management Agent executes it for each of these targets.

   Depending on your objectives, you can choose one of the following options:

   - Single Operation to run a specific command
   - Script to run an OS script and optionally provide an interpreter, which processes the script; for example, `%perlbin%/perl or /bin/sh`.

   Sometimes, a single command line is insufficient to specify the commands to run, and you may not want to install and update a script on all hosts. In this case, you can use the Script option to specify the script text as part of the job.

2. Based on your objectives, follow the instructions in Section 7.3.2.1, "Specifying a Single Operation" or Section 7.3.2.2, "Specifying a Script".

3. Click the **Credentials** property page link.

Task 4 Specify Credentials - (optional)

You do not need to provide input on this page if you want to use the system default of using preferred credentials.

On the Credentials property page, you can specify the credentials that you want the Oracle Management Service to use when it runs the OS Command job against target hosts. The job can use either the job submitter's preferred credentials for hosts, or you can specify other credentials to override the preferred credentials.

You do not need to provide input on this page if you have already set preferred credentials.

- **To use preferred credentials:**

  1. Select the **Preferred Credential** radio button, which is the default selection.

     If the target for the OS Command job is a host or host group, the preferred host credentials are used. You specify these on the Database Preferred Credentials page, and they are different from the host credentials for the host on which the database resides.

  2. Select either **Normal Host Credentials** or **Privileged Host Credentials** from the Host Credentials drop-down.

     You specify these separately on the Preferred Credentials page, which you can access by selecting **Security** from the **Setup** menu, then **Preferred Credentials**. The Preferred Credentials page appears, where you can click the Manage Preferred Credentials button to set credentials.

- **To use named credentials:**

  1. Select the **Named Credential** radio button to override database or host preferred credentials.

     The drop-down list is a pre-populated credential set with values saved with names. These are not linked to targets, and you can use them to provide credential and authentication information to tasks.

- **To use other credentials:**

1. Select the **New Credential** radio button to override previously defined preferred credentials.

   Note that override credentials apply to all targets.

2. Optionally select Sudo or PowerBroker as the run privilege.

   Sudo enables you to authorize certain users (or groups of users) to run some (or all) commands as root while logging all commands and arguments. PowerBroker provides access control, manageability, and auditing of all types of privileged accounts.

   If you provide Sudo or PowerBroker details, they must be applicable to all targets. It is assumed that Sudo or PowerBroker settings are already applied on all the hosts on which this job is to run.

   See your Super Administrator about setting up these features if they are not currently enabled.

   Tip: For information on using Sudo, see the Sudo Manual at:

   `http://www.sudo.ws/sudo/man/1.7.4p6/sudo.man.html`

   For information on using PowerBroker, see the PowerBroker Desktops User Guide at:

   `http://www.ubm-global.com/docs/powerbroker/PBWD_User_Guide_V5%200.pdf`

Task 5 Schedule the Job - (optional)

You do not need to provide input on this page if you want to proceed with the system default of running the job immediately after you submit it.

1. Select the type of schedule:

   - **One Time (Immediately)**

     If you do not set a schedule before submitting a job, Enterprise Manager executes the job immediately with an indefinite grace period. You may want to run the job immediately, but specify a definite grace period in case the job is unable to start for various reasons, such as a blackout, for instance.

     A grace period is a period of time that defines the maximum permissible delay when attempting to start a scheduled job. If the job system cannot start the execution within a time period equal to the scheduled time plus grace period, it sets the job status to Skipped.

   - **One Time (Later)**

     – Setting up a custom schedule:

       You can set up a custom schedule to execute the job at a designated time in the future. When you set the Time Zone for your schedule, the job runs simultaneously on all targets when this time zone reaches the start time you specify. If you select each target's time zone, the job runs at the scheduled time using the time zone of the managed targets. The time zone you select is used consistently when displaying date and time information about the job, such as on the Job Activity page, Job Run page, and Job Execution page.

       For example, if you have targets in the Western United States (US Pacific Time) and Eastern United States (US Eastern Time), and you specify a schedule where Time Zone = US Pacific Time and Start Time = 5:00 p.m., the job runs simultaneously at 5:00 p.m. against the targets in the Western

United States and at 8:00 p.m. against the targets in the Eastern United States. If you specify 5:00 p.m. in the Agent time zone, the executions do not run concurrently. The EST target would run 3 hours earlier.

– Specifying the Grace Period:

The grace period controls the latest start time for the job in case the job is delayed. A job might not start for many reasons, but the most common reasons are that the Agent was down or there was a blackout. By default, jobs are scheduled with indefinite grace periods.

A job can start any time before the grace period expires. For example, a job scheduled for 1 p.m. with a grace period of 1 hour can start any time before 2 p.m., but if it has not started by 2 p.m., it is designated as skipped.

- **Repeating**
  - Defining the repeat interval:

    Specify the Frequency Type (time unit) and Repeat Every (repeat interval) parameters to define your job's repeat interval. The Repeat Until options are as follows:

    Note that both the end date and time determine the last execution. For example, for a job that runs daily at 6 p.m., where...

    Start Time is June 1, 2010 at 6 p.m.
    End Time is June 30, 2010 at 4 a.m.

    ... the last execution runs on June 29, not June 30, since the June 30 end time occurs before the daily time of the job.

  – Specifying the Grace Period:

    The grace period controls the latest start time for the job in case the job is delayed. A job might not start for many reasons, but the most common reasons are that the Agent was down or there was a blackout. By default, jobs are scheduled with indefinite grace periods.

    If the job starts on time, the grace period is ignored. For example, a job scheduled for 1 p.m. with a grace period of 1 hour can start any time before 2 p.m., but if it has not started by 2 p.m., it is designated as skipped.

2. Click the **Access** property page link.

Task 6 Specify Who Can Access the Job - (optional)

You do not need to provide input on this page if you want to proceed with the system default of not sharing the job. The table shows the access that administrators and roles have to the job. Only the job owner (or Super Administrator) can make changes on the Job Access page.

1. Change access levels for administrators and roles, or remove administrators and roles. Your ability to make changes depends on your function.

   If you are a job owner, you can:

   - Change the access of an administrator or role by choosing the Full or View access privilege in the Access Level column in the table.
   - Remove all access to the job for an administrator or role by clicking the icon in the Remove column for the administrator or role. All administrators with Super Administrator privileges have the View access privilege to a job. If you

choose to provide access privileges to a role, you can only provide the View access privilege to the role, not the Full access privilege.

If you are a Super Administrator, you can:

- Grant View access to other Enterprise Manager administrators or roles.
- Revoke all administrator access privileges.

Note: Neither the owner nor a super user can revoke View access from a super user. All super users have View access.

For more information on access levels, see Section 7.3.2.3, "Access Level Rules".

2. Click **Add** to add administrators and roles. The Create Job Add Administrators and Roles page appears.

   a. Specify a **Name** and **Type** in the Search section and click **Go**. If you just click Go without specifying a Name or Type, all administrators and roles in the Management Repository appear in the table.

      The value you specify in the Name field is not case-sensitive. You can specify either * or % as a wildcard character at any location in a string (the wildcard character is implicitly added to the end of any string). For example, if you specify `%na` in the Name field, names such as ANA, ANA2, and CHRISTINA may be returned as search results in the Results section.

   b. Select one or more administrators or roles in the Results section, then click **Select** to grant them access to the job. Enterprise Manager returns to the Create Job Access page or the Edit Job Access page, where you can modify the access of administrators and roles.

3. Define a notification rule.

   You can use the Notification system (rule creation) to easily associate specific jobs with a notification rule. The Cloud Control Notification system enables you to define a notification rule that sends e-mail to the job owner when a job enters one of these chosen states:

   - Scheduled
   - Running
   - Suspended
   - Succeeded
   - Problems
   - Action Required

   Note: Before you can specify notifications, you need to set up your email account and notification preferences. See Chapter 3, "Notifications" for this information.

Task 7 Conclude Job Creation

At this point, you can either submit the job for execution or save it to the job library.

- **Submitting the job —**

Click **Submit** to send the active job to the job system for execution, and then view the job's execution status on the main Job Activity page. If you are creating a library job, Submit saves the job to the library and returns you to the main Job Library page where you can edit or create other library jobs.

If you submit a job that has problems, such as missing parameters or credentials, an error appears and you will need to correct these issues before submitting an active job. For library jobs, incomplete specifications are allowed, so no error occurs.

Note: If you click Submit without changing the access, only Super Administrators can view your job.

- **Saving the job to the library** —

  Click **Save to Library** to the job to the Job Library as a repository for frequently used jobs. Other administrators can then share and reuse your library job if you provide them with access privileges. Analogous to active jobs, you can grant View or Full access to specific administrators. Additionally, you can use the job library to store:

  – Basic definitions of jobs, then add targets and other custom settings before submitting the job.

  – Jobs for your own reuse or to share with others. You can share jobs using views or giving Full access to the jobs.

  – Critical jobs for resubmitting later, or revised versions of these jobs as issues arise.

7.3.2.1 Specifying a Single Operation

Note: The following information applies to step 2 in Task 3, "Specify Parameters" on page 7-6.

Enter the full command in the **Command** field. For example:

```
/bin/df -k /private
```

Note the following points about specifying a single operation:

- You can use shell commands as part of your command. The default shell for the platform is used, which is `/bin/sh` for Linux and `cmd/c` for Windows.

  ```
  ls -la /tmp > /tmp/foobar.out
  ```

- If you need to execute two consecutive shell commands, you must invoke the shell in the Command field and the commands themselves in the OS Script field. You would specify this as follows in the Command field:

  ```
  sleep 3; ls
  ```

7.3.2.2 Specifying a Script

Note: The following information applies to step 2 in Task 3, "Specify Parameters" on page 7-6.

The value you specify in the OS Script field is used as stdin for the command interpreter, which defaults to `/bin/sh` on Linux and `cmd/c` on Windows. You can override this with another interpreter; for example: `%perlbin%/perl`. The shell scripts size is limited to 2 GB.

To control the maximum output size, set the `mgmt_job_output_size_limit` parameter in `MGMT_PARAMETERS` to the required limit. Values less than 10 KB and greater than 2 GB are ignored. The default output size is 10 MB.

You can run a script in several ways:

- **OS Scripts** — Specify the path name to the script in the OS Script field. For example:

  OS Script field: `/path/to/mycommand`
  Interpreter field:

- **List of OS Commands** — You do not need to enter anything in the Interpreter field for the following example of standard shell commands for Linux or Unix systems. The OS's default shell of `/bin/sh` or `cmd/c` will be used.

  ```
  /usr/local/bin/myProg arg1 arg2
  mkdir /home/$USER/mydir
  cp /dir/to/cp/from/file.txt /home/$USER/mydir
  /usr/local/bin/myProg2 /home/$USER/mydir/file.txt
  ```

  When submitting shell-based jobs, be aware of the syntax you use and the targets you choose. This script does not succeed on NT hosts, for example.

- **Scripts Requiring an Interpreter** — Although the OS shell is invoked by default, you can bypass the shell by specifying an alternate interpreter. For example, you can run a Perl script by specifying the Perl script in the OS Script field and the location of the Perl executable in the Interpreter field:

  OS Script field: `<Enter-Perl-script-commands-here>`
  Interpreter field: `%perlbin%/perl`

  The following example shows how to run a list of commands that rely on a certain shell syntax:

  ```
  setenv VAR1 value1
  setenv VAR2 value2
  /user/local/bin/myProg $VAR1 $VAR2
  ```

  You would need to specify csh as the interpreter. Depending on your system configuration, you may need to specify the following string in the Interpreter field:

  ```
  /bin/csh
  ```

  When submitting shell-based jobs, be aware of the syntax you use and the targets you choose. This script would not succeed on NT hosts, for example. However, you do have the option of running a script for a list of Windows shell commands, as shown in the following example. The default shell of `cmd/c` is used for Windows systems.

  ```
  C:\programs\MyApp arg1 arg2
  md C:\MyDir
  copy C:\dir1x\copy\from\file.txt \home\$USER\mydir
  ```

7.3.2.3 Access Level Rules

> **Note:** The following rules apply to Task 6, "Specify Who Can Access the Job - (optional)" on page 7-8.

- Super Administrators always have View access on any job.
- The Enterprise Manager administrator who owns the job can make any access changes to the job, except revoking View from Super Administrators.
- Super Administrators with a View or Full access level on a job can grant View (but not Full) to any new user. Super Administrators can also revoke Full and View from normal users, and Full from Super Administrators.
- Normal Enterprise Manager administrators with Full access levels cannot make any access changes on the job.
- If the job owner performs a Create Like operation on a job, all access privileges for the new job are identical to the original job. If the job owner grants other administrators View or Full job access to other administrators, and any of these administrators perform a Create Like operation on the job, ALL administrators will, by default, have View access on the newly created job.

7.3.3 Creating a SQL Script Job

The basic process for creating a SQL script job is the same as described in Section 7.3.2, "Creating an OS Command Job." The following sections provide supplemental information specific to script jobs.

7.3.3.1 Specifying Targets

You can run a SQL Script job against database and cluster database target types. You select the targets to run the job against by doing the following:

1. Click **Add** in the Targets section.
2. Select the database target(s) from the pop-up.

Your selection(s) now appears in the Target table.

> **Note:** For a cluster host or RAC database, a job runs only once for the target, regardless of the number of database instances. Consequently, a job cannot run on all nodes of a RAC cluster.

7.3.3.2 Options for the Parameters Page

In a SQL Script job, you can specify any of the following in the SQL Script field of the Parameters property page:

- Any directives supported by SQL*Plus
- Contents of the SQL script itself
- Fully-qualified SQL script file; for example:

  ```
  @/private/oracle/scripts/myscript.sql
  ```

  Make sure that the script file is installed in the appropriate location on all targets.

- PL/SQL script using syntax supported by SQL*Plus; for example, one of the following:

```
EXEC plsql_block;
```

or

```
DECLARE
   local_date DATE;
BEGIN
   SELECT SYSDATE INTO local_date FROM dual;
END;
/
```

You can use target properties in the SQL Script field, a list of which appears in the Target Properties table. Target properties are case-sensitive. You can enter optional parameters to SQL*Plus in the Parameters field.

7.3.3.3 Specifying Host and Database Credentials

In the Credentials property page, you specify the host credentials and database credentials. The Management Agent uses the host credentials to launch the SQL*Plus executable, and uses database credentials to connect to the target database and run the SQL script. The job can use either the preferred credentials for hosts and databases, or you can specify other credentials that override the preferred credentials.

- **Use Preferred Credentials —**

  Select this choice if you want to use the preferred credentials for the targets for your SQL Script job. The credentials used for both host and database are those you specify in the drop-down. If you choose Normal Database Credentials, your normal database preferred credentials are used. If you choose SYSDBA Database Credentials, the SYSDBA preferred credentials are used. For both cases, the host credentials associated with the database target are used. Each time the job executes, it picks up the current values of your preferred credentials.

- **Named Credentials —**

  Select this choice if you want to override the preferred credentials for all targets, then enter the named credentials you want the job to use on all targets.

  Many IT organizations require that passwords be changed on regular intervals. You can change the password of any preferred credentials using this option. Jobs and corrective actions that use preferred credentials automatically pick up these new changes, because during execution, Enterprise Manager uses the current value of the credentials (both user name and password). Named credentials are also centrally managed. A change to a named credential is propagated to all jobs or corrective actions that use it.

  For corrective actions, if you specify preferred credentials, Enterprise Manager uses the preferred credentials of the last Enterprise Manager user who edited the corrective action. For this reason, if a user attempts to edit the corrective action that a first user initially specified, Enterprise Manager requires this second user to specify the credentials to be used for that corrective action.

7.3.3.4 Returning Error Codes from SQL Script Jobs

The SQL Script job internally uses SQL*Plus to run a user's SQL or PL/SQL script. If SQL*Plus returns 0, the job returns a status of Succeeded. If it returns any other value, it returns a job status of Failed. By default, if a SQL script runs and encounters an error, it may still result in a job status of Succeeded, because SQL*Plus still returned a

value of 0. To make such jobs return a Failed status, you can use SQL*Plus EXIT to return a non-zero value.

The following examples show how you can return values from your PL/SQL or SQL scripts. These, in turn, will be used as the return value of SQL*Plus, thereby providing a way to return the appropriate job status (Succeeded or Failed). Refer to the *SQL*Plus User's Guide and Reference* for more information about returning EXIT codes.

Example 1

```
WHENEVER SQLERROR EXIT SQL.SQLCODE
select column_does_not_exist from dual;
```

Example 2

```
-- SQL*Plus will NOT return an error for the next SELECT statement
SELECT COLUMN_DOES_NOT_EXIST FROM DUAL;

WHENEVER SQLERROR EXIT SQL.SQLCODE;
BEGIN
  -- SQL*Plus will return an error at this point
  SELECT COLUMN_DOES_NOT_EXIST FROM DUAL;
END;
/
WHENEVER SQLERROR CONTINUE;
```

Example 3

```
variable exit_code number;

BEGIN
 DECLARE
 local_empno number(5);
 BEGIN
  -- do some work which will raise exception: no_data_found
  SELECT 123 INTO local_empno FROM sys.dual WHERE 1=2;
 EXCEPTION
  WHEN no_data_found THEN
    :exit_code := 10;
  WHEN others THEN
    :exit_code := 2;
  END;
 END;
/
exit :exit_code;
```

7.3.4 Creating a Multi-task Job

The basic process for creating a multi-task job is the same as described in Section 7.3.2, "Creating an OS Command Job." The following sections provide supplemental information specific to multi-task jobs.

7.3.4.1 Job Capabilities

Multi-task jobs enable you to create complex jobs consisting of one or more distinct tasks. Because multi-task jobs can run against targets of the same or different type, they can perform ad hoc operations on one or more targets of the same or different type.

The Job System's multi-task functionality makes it easy to create extremely complex operations. You can create multi-task jobs in which all tasks run on a single target. You

can also create a multi-task job consisting of several tasks, each of which has a different job type, and with each task operating on separate (and different) target types. For example:

- Task 1 (OS Command job type) performs an operation on Host 1.
- If Task 1 is successful, run Task2 (SQL Script job type) against Database 1 and Database 2.

7.3.4.2 Specifying Targets for a Multi-task Job

You can run a multi-task job against any targets for which jobs are defined that can be used as tasks. Not all job types can be used as tasks.

The Target drop-down in the General page enables you to choose between running the job against the same targets for all tasks, or different targets for different tasks. Because each task of a multi-task job can be considered a complete job, when choosing the **Same targets for all tasks** option, you add all targets against which the job is to run from the General page. If you choose the **Different targets for different tasks** option, you specify the targets (and required credentials) the tasks will run against as you define each task.

After making your choice from the Target drop-down, you then select the targets to run the job against by clicking Add in the Targets section.

7.3.4.3 Adding Tasks to the Job

You can use the Tasks page to:

- Add, delete, or edit tasks of various job types
- Set task condition and dependency logic
- Add task error handling

You must define at least two tasks in order to set Condition and Depends On options. Task conditions define states in which the task will be executed. Condition options include:

- **Always** — Task is executed each time the job is run.
- **On Success** — Task execution **Depends On** the successful execution of another task.
- **On Failure** — Task execution **Depends On** the execution failure of another task.

The Error Handler Task is often a "clean-up" step that can undo the partial state of the job. The Error Handler Task executes if any task of the multi-task job has an error. Errors are a more severe form of failure, usually meaning that the job system could not run the task. Failures normally indicate that the task ran, but failed. The Error Handler Task does not affect the job execution status. Use the Select Task Type page to specify the job type of the task to be used for error handling.

7.4 Analyzing Job Activity

After you submit jobs, the status of all job executions across all targets is automatically rolled up and available for review on the Enterprise Summary page. Figure 7–2 shows the Jobs section at the bottom of the Enterprise Summary page.

Figure 7–2 Summary of Target Jobs on the Enterprise Summary Page

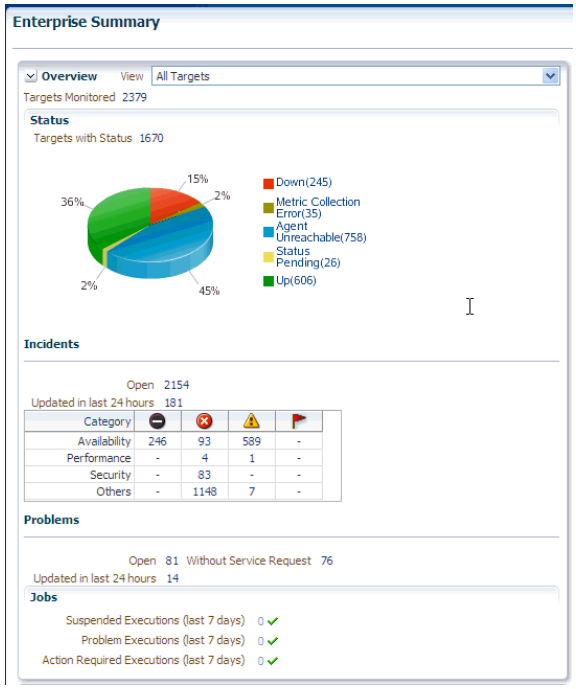

This information is particularly important when you are examining jobs that execute against hundreds or thousands of systems. You can determine the job executions that have failed. By clicking the number associated with a particular execution, you can drill down to study the details of the failed jobs.

7.5 Generating Job Event Criteria

The job system publishes status change events when a job changes its execution status, and these events have different severities based on the execution status.

Use the Job Event Generation Criteria page (Figure 7–3) to set up targets for job event notifications. This page enables you to decide about the jobs or targets or statuses for which you want to raise events or notifications. This ensures that users raise only useful events. Any settings you make on this page do not change the job behavior whatsoever. You can set up notifications on job events through incident rule sets.

Figure 7–3 Job Event Generation Criteria Page

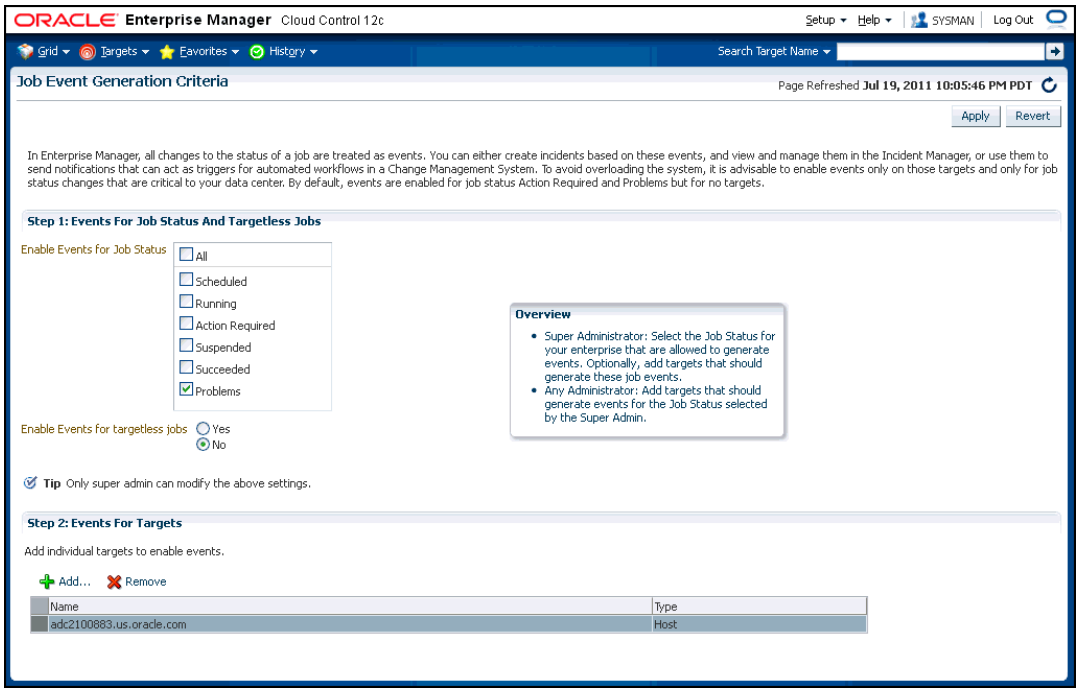

You need to add targets in the target filter if you want to set automatic job event generation for these targets. The job event generation settings you make on this page apply to all users. If you do not add any targets in the target filter, no targets will be set up for automatic job event generation.

On this page, you can do the following:

- Enable events for job status and targetless jobs
- Add targets to generate events for job status

7.5.1 Enabling Events For Job Status and Targetless Jobs

To enable events for job status and targetless jobs, do the following:

1. Ensure that you have Super Administrator privileges to select the job status for which you want to generate events.
2. Ensure that you are an administrator with View Target privileges to add targets for which you want to generate events for the job status set by the Super Administrator.
3. Log into Cloud Control as a Super Administrator.
4. From the **Setup** menu, select **Incidents** and then select **Job Events**. The Job Event Generation Criteria Page is displayed.
5. In the Job Event Generation Criteria page, do the following:
   a. In the Events For Job Status And Targetless Jobs section, from the **Enable Events for Job Status** check boxes, select the status for which you want to publish events. In **Enable Events for targetless jobs**, select **Yes** to create events for jobs that are not associated with any target.

b. In the Events For Targets section, click **Add** to add targets for which you want the job events to be enabled.

6. Click **Apply**.

7.5.2 Adding Targets To Generate Events For Job Status

After a Super Administrator selects events for which job status will be published, administrators can add targets to generate events. To add targets to generate events for job status, do the following:

1. Ensure that you are an administrator with View Target privileges to add targets for which you want to generate events for the job status set by a Super Administrator.
2. Log into Cloud Control as an administrator.
3. From the **Setup** menu, select **Incidents** and then select **Job Events**. The Job Event Generation Criteria Page is displayed.
4. In the Job Event Generation Criteria page, do the following:
   a. In the Events For Job Status And Targetless Jobs section, you can view the status for which events can be published. You can also see if events have been enabled for targetless job filters.
   b. In the Events For Targets section, click **Add** to add targets for which you want the job events to be enabled. You can also remove targets for which you do not want the job events to be enabled by clicking **Remove**.

   Note: Your selected settings in the Events for Targets section are global. Adding or removing targets for events also affect other Enterprise Manager users.

5. Click **Apply**.

7.6 Creating Event Rules For Job Status Change

Enterprise Manager enables you to create and apply rules to events, incidents, and problems. A rule is applied when a newly created or updated event, incident, or problem matches the conditions defined in the rule. The following sections explain how to create event rules for job status change events:

- Creating Job Status Change Event Rules For Jobs
- Creating Job Status Change Event Rules For Targets

7.6.1 Creating Job Status Change Event Rules For Jobs

To create job status change event rules for jobs, do the following:

1. Ensure that the relevant job status is enabled and required targets have been added to job event generation criteria.
2. Ensure that you have administrator privileges to create event rules for job status change events.
3. Log into Cloud Control as an administrator.
4. From the **Setup** menu, select **Incidents** and then **Incident Rules**. The Incident Rules Page is displayed.

5. In the Incident Rules page, click **Create Rule Set** to create rule sets for incidents.
6. Specify the **Name**, **Description**, and select **Enabled** to enable the rule set. Select Type as **Enterprise** if you want to set the rule for all Enterprise Manager users or Private if you want to set the rule for a specific user only. Select **Applies to Job**.

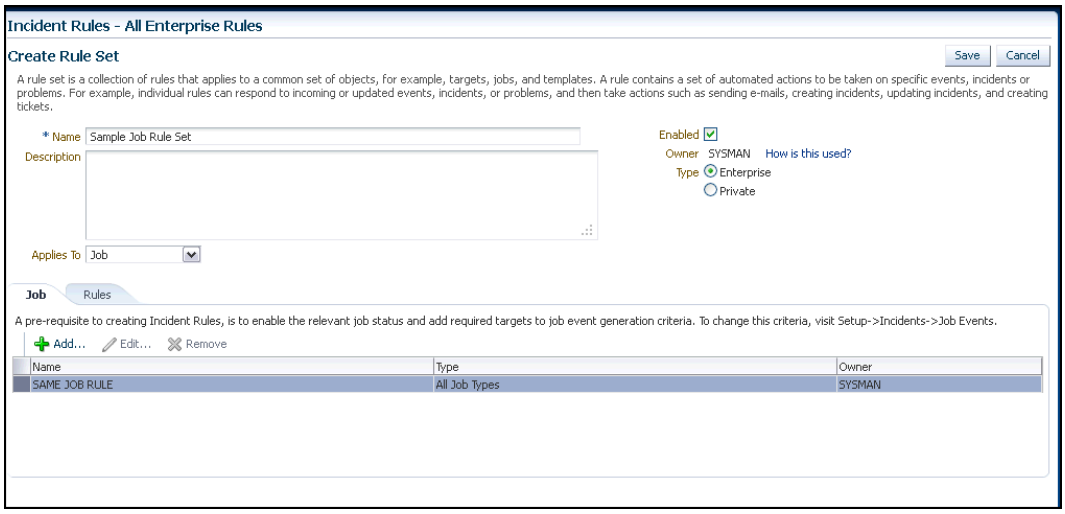

In the Job section, click **Add** to add jobs for which you want to create event rules.

7. In the Add Jobs dialog box, if you select Job **By Pattern**, provide **Job name like** and select the **Job Type**. Specify **Job owner like**. For **Specific jobs**, select the job. Click **OK**.
8. In the **Rules** tab, click **Create**.
9. In the Select Type of Rule to Create dialog box, select from the following choices according to the rule set you want to create:
   - **Incoming events and updates to events** to receive notification or create incidents for job rules. If you are operating on events (for example, if you want to create incidents for incoming events, such as job failed, or notify someone), choose this option.
   - **Newly created incidents or updates to incidents** receive notifications or create rules for incidents even though the events for which incidents are generated do not have associated rules. If you are operating on incidents already created or newly created (for example, you want to direct all incidents related to a group, say foo, to a particular user or escalate all incidents open for more than 3 days), choose this option.
   - **Newly created problems or updates to problems** to receive notifications or create rules for problems even though the incidents for which problems are generated do not have associated rules. This option does not apply for jobs.
10. Select **Incoming events and updates to events**, and in the Create New Rule: Select Events page, do the following:
    - **Select By Type** to **Job Status Change**. Select **All events of type Job Status Change** if you want to take an action for all job state change events for the selected jobs. Select **Specific events of type Job Status Change** if you only want to act on specific job states. If you have selected Specific events of type Job Status Change, select Job Status for events for which you want to create the rule.

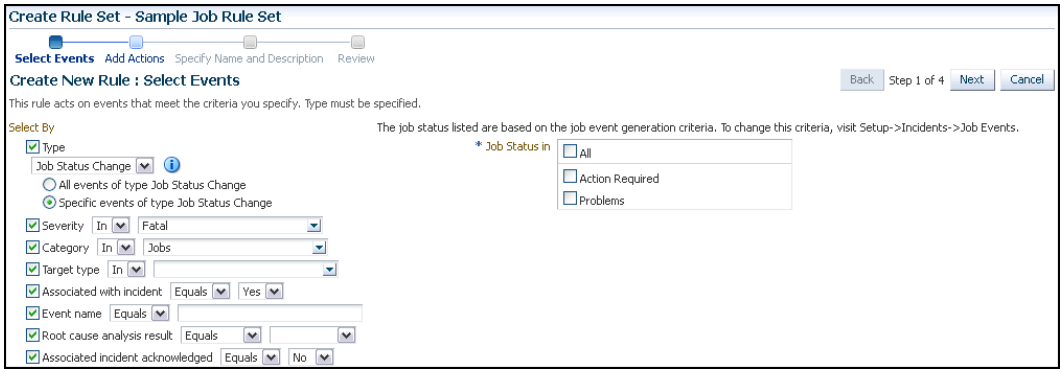

- Set the other criteria for which you want to set the rule as displayed in the above graphic.

11. Select **Newly created incidents or updates to incidents** if you want to create rules for an incident, though the event associated with the incident does not have notification rules. In the Create New Rule: Select Incidents page, select any of the following:

    - **All new incidents and updated incidents** to apply the rule to all new and updated incidents
    - **All new incidents** to apply the rule to all new incidents
    - **Specific incidents** and then select the criteria for the incidents

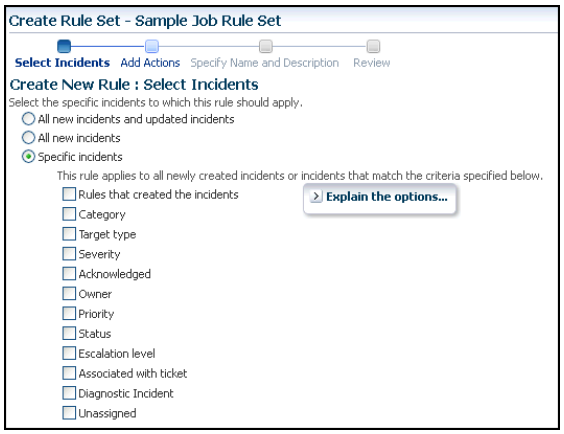

12. In the Create New Rule: Add Actions page, click **Add** to add actions to the rule.
13. In the Add Conditional Actions page, specify actions to be performed when the event matches the rule.

    In the Conditions for actions section, select:

    - **Always execute the actions** to execute actions regardless of event.
    - **Only execute the actions if specified conditions match** to execute actions to match specific criteria.

    When adding actions to events, specify the following:

    - Select **Create Incident** to create an incident for the event to manage and track its resolution.

- In the Notifications section, specify recipients for notifications in the **E-mail To**, **E-mail Cc**, and **Page** fields who will receive e-mail when the event for which a condition is set occurs. If Advanced and Repeat Notifications options have been set, specify them.
- In the Clear events section, select **Clear permanently** if you want to clear an event after the issue that generated the event is resolved.
- If you have configured event connections, in the Forward to Event Connectors section, you can send the events to third-party event management systems.

When adding actions to incidents, specify the following:

- In the Notifications section, specify recipients for notifications in the **E-mail To**, **E-mail Cc**, and **Page** fields who will receive e-mail when the event for which a condition is set occurs. If Advanced and Repeat Notifications options have been set, specify them.
- In the Update Incident section, specify the details to triage incidents when they occur. Specify **Assign to**, **Set priority to**, **Set status to**, and **Escalate to** details.
- In the Create Ticket section, if a ticket device has been configured, specify details to create the ticket.

Click **Continue**.

14. In the Specify Name and Description page, specify a **Name** and **Description** for the event rule. Click **Next**.
15. In the Review page, verify the details you have selected for the event rule and click **Continue** to add this rule in the rule set.
16. On the Create Rule Set page, click **Save** to save the rule set.

7.6.2 Creating Job Status Change Event Rules For Targets

To create job status change event rules for targets, do the following:

1. Ensure that the relevant job status is enabled and required targets have been added to job event generation criteria.
2. Ensure that you have administrator privileges to create event rules for job status change events.
3. Log into Cloud Control as an administrator.
4. From the **Setup** menu, select **Incidents**, then **Incident Rules**. The Incident Rules Page is displayed.
5. In the Incident Rules page, click **Create Rule Set** to create rule sets for incidents.
6. Specify the **Name**, **Description**, and select **Enabled** to enable the rule set. Select Type as **Enterprise** if you want to set the rule for all Enterprise Manager users, or Private if you want to set the rule for a only specific user. Select **Applies to Targets**.

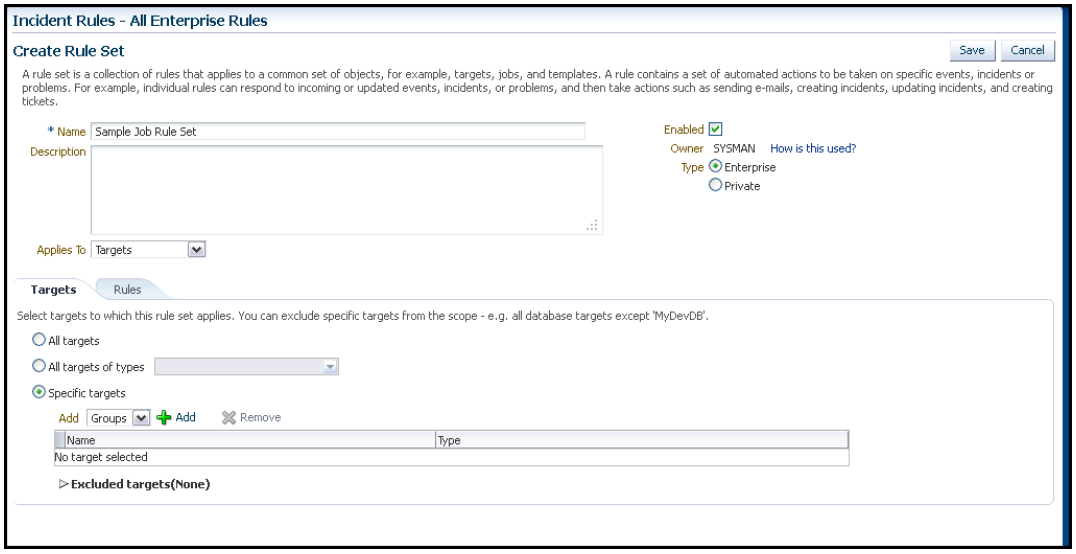

In the **Targets** tab, select one of the following:

- **All targets** to apply to all targets. In the Excluded Targets section, click **Add** to search and select the target that you want to exclude from the rule set. Click **Select**.
- **All targets of types** to select the types of targets to which you want to apply the rule set.
- **Specific targets** to individually specify the targets. Select to Add **Groups** or **Targets** to add groups or targets and click **Add** to search and select the targets to which you want to apply the rule set. Click **Select**. In the Excluded Targets section, click **Add** to search and select the target that you want to exclude from the rule set. Click **Select**.

7. In the **Rules** tab, click **Create**.
8. In the Select Type of Rule to Create dialog box, select from the following choices according to the rule set you want to create:
   - **Incoming events and updates to events** to receive notifications or create incidents for job rules. If you are operating on events (for example, if you want to create incidents for incoming events, such as job failed, or notify someone), choose this option.
   - **Newly created incidents or updates to incidents** receive notifications or create rules for incidents even though the events for which incidents are generated do not have associated rules. If you are operating on incidents already created or newly created (for example, you want to direct all incidents related to a group, say foo, to a particular user or escalate all incidents open for more than 3 days), choose this option.
   - **Newly created problems or updates to problems** to receive notifications or create rules for problems even though the incidents for which problems are generated do not have associated rules. This option does not apply for jobs.
9. Select **Incoming events and updates to events**, and in the Create New Rule: Select Events page, do the following:
   - **Select By Type** to **Job Status Change**. Select **All events of type Job Status Change** if you want to take an action for all job state change events for the

selected jobs. Select **Specific events of type Job Status Change** if you only want to act on specific job states. If you have selected Specific events of type Job Status Change, select Job Status for events for which you want to create the rule.

- Set the other criteria for which you want to set the rule as displayed in the above graphic.

10. Select **Newly created incidents or updates to incidents** if you want to create rules for an incident, though the event associated with the incident does not have notification rules. In the Create New Rule: Select Incidents page, select any of the following:
    - **All new incidents and updated incidents** to apply the rule to all new and updated incidents.
    - **All new incidents** to apply the rule to all new incidents.
    - **Specific incidents** and then select the criteria for the incidents.

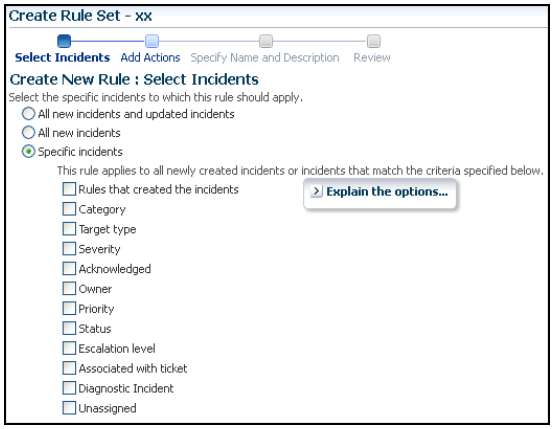

11. In the Create New Rule: Add Actions page, click **Add** to add actions to the rule.
12. In the Add Conditional Actions page, specify actions to be performed when the event matches the rule.

    In the Conditions for actions section, select:
    - **Always execute the actions** to execute actions regardless of event.
    - **Only execute the actions if specified conditions match** to execute actions to match specific criteria.

When adding actions to events, specify the following:

- Select **Create Incident** to create an incident for the event to manage and track its resolution.
- In the Notifications section, specify recipients for notifications in the **E-mail To**, **E-mail Cc**, and **Page** fields who will receive e-mail when the event for which a condition is set occurs. If Advanced and Repeat Notifications options have been set, specify them.
- In the Clear events section, select **Clear permanently** if you want to clear an event after the issue that generated the event is resolved.
- If you have configured event connections, in the Forward to Event Connectors section, you can send the events to third-party event management systems.

When adding actions to incidents, specify the following:

- In the Notifications section, specify recipients for notifications in the **E-mail To**, **E-mail Cc**, and **Page** fields who will receive e-mail when the event for which a condition is set occurs. If Advanced and Repeat Notifications options have been set, specify them.
- In the Update Incident section, specify the details to triage incidents when they occur. Specify **Assign to**, **Set priority to**, **Set status to**, and **Escalate to** details.
- In the Create Ticket section, if a ticket device has been configured, specify the details to create the ticket.

Click **Continue**.

13. In the Specify Name and Description page, specify a **Name** and **Description** for the event rule. Click **Next**.
14. In the Review page, verify the details you have selected for the event rule and click **Continue** to add this rule in the rule set.
15. On the Create Rule Set page, click **Save** to save the rule set.

7.7 Creating Corrective Actions

Corrective Actions enable you to specify automated responses to metric alerts. Corrective Actions ensure that routine responses to metric alerts are automatically executed, thereby saving you time and ensuring problems are dealt with before they noticeably impact end users.

Corrective actions share many features in common with the Job System. By default, a corrective action runs on the target on which the metric alert has triggered. Alternatively, you can specify a corrective action to contain multiple tasks, with each task running on a different target. You can also receive notifications for the success or failure of corrective actions.

You define corrective actions for individual metrics for monitored targets. See Section 7.7.1, "Creating Corrective Actions for Metrics," for instructions on setting up corrective actions.

Credential Requirements

Since corrective actions are associated with a target's metric thresholds, you can define corrective actions if you have been granted OPERATOR or greater privilege on the target. You can define separate corrective actions for both Warning and Critical thresholds. Corrective actions must run using the credentials of a specific user. For this

reason, whenever a corrective action is created or modified, you must specify the credentials that the modified action runs with.

7.7.1 Creating Corrective Actions for Metrics

For any target, the Metric and Collection Settings page shows whether corrective actions have been set for various metrics. For each metric, the Corrective Actions column shows whether Critical and/or Warning severities of corrective actions have been set.

1. From any target's home page menu, select **Monitoring**, then **Metric and Collection Settings**. The Metric and Collection Settings page appears.

   Tip: For instance, on the home page for a host named `dadvmn0630.us.oracle.com`, you would select the Host menu, then Monitoring, then Metric and Collection Settings.

2. Click the pencil icon for a specific metric to access the Edit Advanced Settings page for the metric.
3. In the Corrective Actions section, click **Add** for the metric severity (Warning and/or Critical) for which you want to associate a corrective action.
4. Select the task type on the Add Corrective Actions page, then click **Continue**.
   - If you want to use a corrective action from the library, select **From Library** as the task type. Using a library corrective action copies the description, parameters, and credentials from the library corrective action. You must still define a name for the new corrective action. You can provide corrective action parameters if necessary.
   - If you want to create a corrective action to store in the library, see Section 7.7.2, "Creating a Library Corrective Action."
   - If you want to provide an Agent-side response action, select Agent Response Action as the task type. See Section 7.7.5, "Providing Agent-side Response Actions" for more information.
5. On the Corrective Action page, provide input for General, Parameters, and Credentials as you would similarly do when creating a job.
6. Click **Continue** to save the corrective action and return to the Edit Advanced Settings page, where your corrective action now appears.
7. *Optional*: To prevent multiple instances of a corrective action from operating simultaneously, enable the **Allow only one corrective action for this metric to run at any given time** checkbox.

   This option specifies that both Critical and Warning corrective actions will not run if a severity is reported to the Oracle Management Services when an execution of either corrective action is currently running. This can occur if a corrective action runs longer than the collection interval of the metric it corrects; the value of the metric may be oscillating back and forth across one of the thresholds (leading to multiple executions of the same corrective action), or may be rising or falling quickly past both thresholds (in which case an execution of the Warning corrective action may overlap an execution of the Critical corrective action).

   If you do not select this option, multiple corrective action executions are launched under the aforementioned circumstances. It is the administrator's responsibility to ensure that the simultaneous corrective action executions do not conflict.

8. Click **Continue** when you have finished adding corrective actions to return to the Metric and Collection Settings page.

   The page shows the corrective action value you have provided for the metric in the Corrective Actions column. Possible values are:

   - **None** — No corrective actions have been set for this metric.
   - **Warning** — A corrective action has been set for Warning, but not Critical, alerts for this metric.
   - **Critical** — A corrective action has been set for Critical, but not Warning, alerts for this metric.
   - **Warning and Critical** — Corrective actions have been set for both Warning and Critical alerts for this metric. If an Agent-side response action is associated with the metric, the value is also Warning and Critical, since Agent-side response actions are always triggered on either Critical or Warning alert severities.

9. Continue the process from step 2 forward, then click **OK** on the Metric and Collection Settings page to save your corrective actions and return to the target page you started from in step 1.

7.7.2 Creating a Library Corrective Action

For corrective actions that you use repeatedly, you can define a library corrective action. After a corrective action is in the library, you can re-use the corrective action definition whenever you define a corrective action for a target metric or policy rule.

1. From the Enterprise menu, select **Monitoring**, then **Corrective Actions**. The Corrective Action Library page appears.
2. Select a job type from the **Create Library Corrective Action** drop-down, then click **Go**.
3. Define the corrective action as you would for creating a job in Section 7.3, "Creating Jobs" for General, Parameters, and Credentials. For Access, go to the following optional step.
4. *Optional*: Select **Access** to define or modify the access you want other users to have for this corrective action.

   For more information, see Section 7.7.3, "Specifying Access to Corrective Actions."
5. Click **Save to Library** when you have finished. The Corrective Action Library page reappears, and your corrective action appears in the list.

   You can now create another corrective action based on this one (Create Like button), edit, or delete this corrective action.

You can access this library entry whenever you define a corrective action for a metric severity by selecting From Library as the task type in the Add Corrective Actions page. See step 4 in Section 7.7.1, "Creating Corrective Actions for Metrics," for more information.

7.7.3 Specifying Access to Corrective Actions

As mentioned in the procedure above, you can determine the access to corrective actions by other users. You do not need to provide input for this page if you do not want to share the corrective action.

7.7.3.1 Defining or Modifying Access

The table on the Access page shows the access that administrators and roles have to the corrective action. Only the corrective action owner (or Super Administrator) can make changes on this page.

As the corrective action owner, you can do the following:

- Add other administrators and roles to the table by clicking **Add**, then selecting the appropriate type in the subsequent page that appears.
- Change the access of an administrator or role by choosing the **Full** or **View** access right in the Access Level column in the table.
- Remove all access to the corrective action for an administrator or role by clicking the icon in the **Remove** columns for this administrator or role. All administrators with Super Administrator privileges have the View access right to a corrective action.

If you choose to provide access rights to a role, you can only provide the View access right to the role, not the Full access right.

If you are a Super Administrator, you can:

- Grant View access to other Enterprise Manager administrators or roles.
- Revoke all administrator access privileges.

> **Note:** If a new user is being created, the user should have the CREATE_JOB privilege to create Corrective Actions.

7.7.3.2 Access Level Rules

Access level rules are as follows:

- Super Administrators always have View access for any corrective action.
- The Enterprise Manager administrator who owns the corrective action can make any access changes to the corrective action (except revoking View from Super Administrators).
- Super Administrators with a View or Full access level for a corrective action can grant View (but not Full) access to any new user. Super Administrators can also revoke Full and View access from normal users, and Full access from Super Administrators.
- Normal Enterprise Manager administrators with Full access levels cannot make any access changes on the corrective action.
- If the corrective action owner performs a Create Like operation on a corrective action, all access privileges for the new corrective action become identical to the original corrective action. If the corrective action owner grants other administrators View or Full access to other administrators, and any of these administrators perform a Create Like operation on this corrective action, all administrators will, by default, have View access on the newly created corrective action.

7.7.4 Setting Up Notifications for Corrective Actions

Corrective actions are associated with metrics whose alerts trigger them. Any Enterprise Manager administrator with View or higher privileges on a target can receive notifications following the success or failure of a corrective action.

A single incident rule can contain any combination of alert and corrective action states. All metrics and targets selected by the incident rule are notified for the same alert and corrective action states. Therefore, if you want to be notified of corrective action success or failure for one metric, but only on failure for another, you need to use two incident rules. An incident rule can include corrective action states for metrics with which no corrective actions have been associated. In this case, no notifications are sent.

> **Note:** Notifications cannot be sent for Agent-side response actions, regardless of the state of any incident rules applied to the target.

To create incident rules for notifications:

1. From the Setup menu, select **Incidents**, then **Incident Rules**.
2. Click **Create Rule Set**. The Create Rule Set wizard appears.
3. Provide the requisite information at the top of the Create Rule Set page, then select one of the target choices in the Targets sub-tab, supplying additional information as needed for the "All targets of types" and "Specific targets" choices.
4. Select the **Rules** sub-tab, then click **Create**.
5. In the pop-up that appears, select the default **Incoming events and update to events** choice, the click **Continue**.
6. On the Select Events page, enable the **Type** checkbox, then select **Metric Alert**.
7. Click the **Specific events of type Metric alert** radio button, then click **Add** in the table that appears.
8. In the pop-up that appears, select the Target Type, filter and select the metric, select a severity, then enable the desired corrective action status. Click **OK**.
9. From the Add Actions page, click **Add**.
10. Specify recipients in the Basic Notifications section of the Add Conditional Actions page.
11. Proceed through the final two pages of the wizard, then click **Continue**. Your new rule appears in the Create Rule Set page.
12. Click **Save** to save this rule.

After you have created one or more rule sets, you need to set up notification methods as follows:

1. From the Setup menu, select **Notifications**, then **Notification Methods**.
2. From the Notification Methods page, select **Help**, then **Enterprise Manager Help** for assistance on providing input for this page.

7.7.5 Providing Agent-side Response Actions

Agent-side response actions perform simple commands in response to an alert. When the metric triggers a warning or critical alert, the Management Agent automatically runs the specified command or script without requiring coordination with the Oracle Management Service (OMS). The Agent runs this command or script as the OS user who owns the Agent executable. Specific target properties can be used in the Agent response action script.

Note: Use the Agent-side Response Action page to specify a single command-line action to be executed when a Warning or Critical severity is reached for a metric. For tasks that require alert context, contain more complex logic, or require that notifications be sent on success or failure, corrective actions should be used instead of an Agent-side response action.

To access this page, follow steps 1 through 4 in Section 7.7.1, "Creating Corrective Actions for Metrics."

7.7.5.1 Specifying Commands and Scripts

You can specify a single command or execute a script. You cannot specify special shell command characters (such as > and <) as part of the response action command. If you must include these types of special characters in your response action commands, you should use them in a script, then specify the script as the response action command.

If using a script, make sure the script is installed on the host machine that has the Agent. If using shell scripts, make sure the shell is specified either in the Response Action command line:

Script/Command: `/bin/csh myScript`

... or within the body of the script itself:

Script/Command: `myScript`

... where `myScript` contains the following:

```
!#/bin/csh<
<rest of script>
```

7.7.5.2 Using Target Properties in Commands

You can use target properties in a command. Click **Show Available Target Properties** to display target properties you can use in the Script/Command field. The list of available target properties changes according to the type of target the response action is to run against.

Use Target Properties as command-line arguments to the script or command, then have the script reference these command-line arguments. For example, to use the `%OracleHome%` and `%SID%` target properties, your command might appear as follows:

```
/bin/csh MyScript %OracleHome% %SID%
```

.... and your script, `MyScript`, can reference these properties as command-line arguments. For example:

```
IF $1 = 'u1/bin/OracleHome' THEN...
```

Target properties are case-sensitive. For example, if you want to access the Management Agent's Perl interpreter, you can specify `%perlBin%/perl <my_perl_script>` in the Script/Command field.

7.7.5.3 Advanced Usage

You can get other target properties from the target's XML file in the OracleHome/sysman/admin/metadata directory, where OracleHome is the Oracle home of the Management Agent that is monitoring the target. In the XML file, look for

the PROP_LIST attribute of the DynamicProperties element to get a list of properties that are not listed in the targets.xml entry for the target.

The following example is an excerpt from the hosts.xml file:

```
<InstanceProperties>
 <DynamicProperties NAME="Config" FORMAT="ROW"
 PROP_LIST="OS;Version;OS_patchlevel;Platform;Boottime;IP_address">
   <ExecutionDescriptor>
   <GetTable NAME="_OSConfig"/>
   <GetView NAME="Config" FROM_TABLE="_OSConfig">
   <ComputeColumn NAME="osName" EXPR="Linux" IS_VALUE="TRUE"/>
   <Column NAME="osVersion"/>
   <Column NAME="osPatchLevel"/>
   <Column NAME="Platform"/>
   <Column NAME="Boottime"/>
   <Column NAME="IPAddress"/>
   </GetView>
   </ExecutionDescriptor>
   </DynamicProperties>
   <InstanceProperty NAME="Username" OPTIONAL="TRUE" CREDENTIAL="TRUE">
   <ValidIf>
   <CategoryProp NAME="OS" CHOICES="Linux"/>
   </ValidIf>
   <Display>
   <Label NLSID="host_username_iprop">Username</Label>
   </Display>
   </InstanceProperty>
   <InstanceProperty NAME="Password" OPTIONAL="TRUE" CREDENTIAL="TRUE">
   <ValidIf>
   <CategoryProp NAME="OS" CHOICES="Linux"/>
   </ValidIf>
   <Display>
   <Label NLSID="host_password_iprop">Password</Label>
   </Display>
   </InstanceProperty>
</InstanceProperties>
```

8 Compliance

Compliance is the conformance to standards, or requirements, or both.

Enterprise Manager Compliance Management (EMCM) provides the ability to evaluate the compliance of targets and systems as they relate to business best practices for configuration, security, and storage. This is accomplished by defining, customizing, and managing compliance frameworks, compliance standards, and compliance standard rules. In addition, EMCM provides advice of how to change configuration to bring your targets and systems into compliance.

This chapter explains how EMCM verifies that applications in your enterprise comply with preestablished standards and how to manage the compliance structure. This chapter includes:

- Section 8.1, "Compliance Overview"

  Explains at a high level the compliance features in Enterprise Manager.

- Section 8.2, "Evaluating Compliance"

  Explains how to best use the compliance features to evaluate compliance.

- Section 8.3, "Managing Compliance"

  Explains how to perform the operations on compliance frameworks, compliance standards, and compliance standard rules.

8.1 Compliance Overview

The Oracle Enterprise Manager Compliance Management (EMCM) solution provides the tools to evaluate targets and systems for compliance with business best practices in terms of configuration, security, storage, and so on. In addition, EMCM provides the capability to define, customize, and manage the entities used to evaluate compliance.

The compliance solution:

- Automatically determines if targets and systems have valid configuration settings and whether they are exposed to configuration-related vulnerabilities.
- Advises how to change configurations to bring targets and systems into compliance with respect to best practices.
- Provides real-time monitoring of a target's files, processes, and users to let Enterprise Manager users know where configuration change is taking place in their environment.
- Provides out-of-box compliance frameworks (PCI, for example) and compliance standards to map to compliance standard rules. This mapping makes it possible to

visualize how out-of-compliance settings and actions will affect any compliance framework an organization follows.

- Provides a compliance-focused view if IT configuration and change that is suitable for Line of Business Owners, IT Managers, and Compliance Managers to refer to regularly to check on their organization's compliance coverage.

Before you start using the compliance features, there are a few basics you need to know. See the following for details:

- Section 8.1.1, "Terminology Used in Compliance"
- Section 8.1.2, "Accessing the Compliance Features"
- Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features"

8.1.1 Terminology Used in Compliance

Compliance frameworks, compliance standards, and compliance standard rules are some of the terms used when describing and managing compliance. The following terms are used throughout this chapter when discussing the compliance feature:

- Compliance Framework

  A compliance framework is an industry-specified best practices guideline that deals with the underlying IT infrastructure, applications, business services and processes, and how they are organized, managed and monitored. Compliance frameworks are hierarchical to allow for direct representation of these industry frameworks. A Compliance Framework can be used to represent a framework such as PCI.

  A compliance framework is a way for you to map your standards to a structure similar to the regulatory or standard compliance structure you use in your company.

- Compliance Standard

  A compliance standard is a collection of checks or rules. It is the Enterprise Manager representation of a compliance control that must be tested against some set of IT infrastructure to determine if the control is being followed.

  A compliance standard performs a collection of checks that follow broadly accepted best practices. This ensures that IT infrastructure, applications, business services and processes are organized, configured, managed, and monitored properly. A compliance standard evaluation can provide information related to platform compatibility, known issues affecting other customers with similar configurations, security vulnerabilities, patch recommendations, and more.

- Compliance Standard Rule

  A compliance standard rule is a test to determine if a configuration data change affects compliance. A compliance standard rule is mapped to one or more compliance standards.

  Enterprise Manager 12*c* has the following types of rules.

  – Repository Rule

    Used to perform a check against any metric collection data in the Management Repository

  – WLS Signature Rule

Used to check a WebLogic target for support best practice configurations. This type of rule is not relevant for external/partner plugins.

– Real-time Monitoring Rule

Used to monitor actions to files, processes, and more. Also captures user login/logout activities.

- Compliance Standard Rule Folder

  Compliance standard rule folders are hierarchical structures that contain compliance standard rules.

- Importance

  For compliance frameworks, importance indicates the relative importance of a compliance standard to all other compliance standards in the compliance framework.

  For compliance standards, importance indicates the relative importance of a compliance standard rule to all other compliance standard rules in the compliance standard. The values represent a way of weighting a compliance standard.

- Score

  A target's compliance score for a compliance standard is used to reflect the degree of the target's conformance with respect to a compliance standard. The compliance score is in the range of 0% to 100% inclusive. A compliance score of 100% indicates a target fully complies with the compliance standard.

- Real-time Facets

  The real-time monitoring rule definition includes facets that are used to determine what is important to monitor for a given target type, target properties, and entity type. A facet is a collection of patterns that make up one attribute of a target type.

- Real-Time Observations

  Observations are the actions that were seen on a host or target that were configured to be monitored through real-time monitoring rules. Each distinct user action results in one observation.

In Summary

Compliance standard rules perform single health and real-time monitoring checks. These checks are grouped into compliance standards which together constitute one test of compliance. These compliance standards are then grouped into respective compliance frameworks so that the results of the test can be associated with the areas of the your framework being affected.

8.1.2 Accessing the Compliance Features

To access the compliance features, navigate to the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select one of the following:

1. Results

   Compliance results include evaluation results and errors for compliance frameworks and compliance standards, as well as target compliance.

2. Library

The Compliance Library page contains the entities used for defining standards. From the Compliance Library page you can manipulate compliance frameworks, compliance standards, compliance standard rules, and real-time monitoring facets.

Note that the real-time monitoring facets are only for real-time monitoring rules.

3. Real-Time Observations

   Examination of observations made through Real-time Monitoring.

8.1.3 Privileges and Roles Needed to Use the Compliance Features

To use the compliance standard features, you need to have access to the following privileges and roles.

Privilege	Description
CREATE_COMPLIANCE_ENTITY	Allows you to create compliance standards, compliance standard rules, and Real-time Monitoring facets
FULL_ANY_COMPLIANCE_ENTITY	Allows you to edit and delete compliance standards and compliance standard rules
VIEW_ANY_COMPLIANCE_FWK	Allows you to view compliance framework definition and results
MANAGE TARGET COMPLIANCE	Allows you to associate a compliance standard to a target
VIEW	Allows you to view a single target

Role	Description
EM_COMPLIANCE_DESIGNER	Using this role you can create, modify, and delete compliance standards, compliance standard rules, and Real-time Monitoring facets.
EM_COMPLIANCE_OFFICER	Using this role you can view compliance framework definitions and results.

The following table lists the compliance tasks with the privileges and roles required.

Task	Privileges and Roles Required
Create compliance framework	CREATE_COMPLIANCE_ENTITY privilege VIEW_ANY_COMPLIANCE_FWK privilege
Edit and delete compliance framework	FULL_ANY_COMPLIANCE_ENTITY privilege VIEW_ANY_COMPLIANCE_FWK privilege
Create, edit, and delete compliance framework	EM_COMPLIANCE_DESIGNER role EM_COMPLIANCE_OFFICER role
Associate a compliance standard to a target	MANAGE_TARGET_COMPLIANCE privilege
Import or export a compliance framework	EM_COMPLIANCE_DESIGNER role EM_COMPLIANCE_OFFICER role
Create a real-time monitoring rule	EM_COMPLIANCE_DESIGNER role
Create a real-time monitoring facet	EM_COMPLIANCE_DESIGNER role

Additional restrictions regarding the use of the compliance features include:

- You have privileges to access the compliance standards you will be associating with the target. In particular, you need the Manage Target Compliance privilege on the target.
- When you are working with a compliance framework, ensure that the compliance framework to be edited is defined in the Management Repository
- Ensure you have the privilege to manipulate real-time monitoring facets.

8.2 Evaluating Compliance

Compliance evaluation is the process of testing the compliance standard rules within a compliance standard against a target and recording any violations in the Management Repository.

By evaluating a target against a compliance standard, you are determining whether a target complies with the guidelines of the standard. In the case when a target does not meet the desired state, the test may suggest what changes are required to make that target compliant.

Compliance evaluation generates a score for a target as in how much the target is compliant with the standard. A 100% compliance score means that the target follows all requirements and regulations imposed by the compliance standard.

Because target compliance is required to be monitored regularly, you need to associate a compliance standard with a single target. Evaluation is automatically performed for any associated targets, when their state refreshes, that is when new data has been collected from the target. For Repository Rules, when new data for the target gets loaded into the Management Repository, evaluation happens again. For Real-time Monitoring, evaluation happens when an observation occurs.

What Compliance Evaluation Includes

Compliance evaluation includes:

- Viewing the results of an evaluation

  Only results from the targets for which you have View privilege will be available. The compliance standard rule evaluation results are rolled up in order to produce a compliance standard evaluation state as well as a compliance summary.
- Studying out-of-box reports
- Studying the trend overview as a result of the evaluation

  Use the graphs in the Trend Overview pages to visually determine whether the targets are adhering to or distancing themselves from the compliance best practices.

  To access the Trend Overview pages for compliance standards:

  1. From the **Enterprise** menu select **Compliance**, then select **Results**.
  2. From the **Compliance Standards** tab, choose **Evaluation Results**.
  3. On the Evaluation Results page, choose the compliance standard you want to investigate and click **Show Details**.
  4. On the resulting details page, click the **Trend Overview** tab.

  Note that you can also review Trend Overview pages for compliance frameworks.

What You Can Do To Ensure Compliance

Consider performing the following:

- Ensure your environments match baselines (or each other) by creating rules on top of configuration compare capabilities. Then monitor for configuration drift using real-time monitoring.
- Study the results of the evaluations and make the needed changes to the targets
- Evaluate compliance against best practices
- Evaluate validity of configuration settings
- Evaluate exposure to configuration-related vulnerabilities, storage, and security
- Modify targets and systems to be compliant
- Verify authorization of configuration changes
- Continually test your systems, services, and targets, ensuring the best possible protection and performance your system can have
- Use out-of-box compliance standards and compliance standard rules to determine compliance. Click here to see a demo of this functionality.

The following sections provide additional details:

- Section 8.2.1, "Accessing Compliance Statistics"
- Section 8.2.2, "Viewing Compliance Summary Information"
- Section 8.2.3, "Viewing Target Compliance Evaluation Results"
- Section 8.2.4, "Viewing Compliance Framework Evaluation Results"
- Section 8.2.5, "Investigating Compliance Violations and Evaluation Results"
- Section 8.2.6, "Investigating Evaluation Errors"
- Section 8.2.7, "Compliance Audit by a Compliance Auditor"
- Section 8.2.8, "Compliance Reports"
- Section 8.2.9, "Compliance Score and Importance"

8.2.1 Accessing Compliance Statistics

Compliance statistics are available throughout the Enterprise Manager interface in Compliance Summary regions located on pages such as the Enterprise Summary page and a target's home page.

These regions report the violations and compliance scores for the particular targets. However, the region only reports that there is a violation; it does not give the details. For example, a violation can be against the Secure Port compliance standard rule that is part of the Secure Configuration for Host compliance standard. But you will not know the details just by looking at the Compliance Summary region.

8.2.1.1 How to Determine the Compliance Standard Rule Being Violated and the Target Causing the Violation

Say that you are looking at the Enterprise Summary page and you notice that there are critical violations against the Secure Configuration for Host compliance standard. You need to find what targets are causing the violations. Follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.

2. In the **Evaluations Results** tab for Compliance Standards, highlight the Secure Configuration for Host compliance standard. Click **Show Details**.
3. In the **Summary** tab on the Compliance Standard Result Detail page, you can look at the results either by target or compliance standard rule. For this example, we will use Result by Compliance Standard Rule.
4. In the navigational list, click the **Secure Ports** compliance standard rule. In the resulting **Secure Ports Summary** tab, you will get a list of all the targets that are violating the Secure Ports rule. This is a security issue that needs to be addressed.

8.2.1.2 How to View All the Violations Reported for Your Enterprise

Say that you want to see all the targets that are not compliant with the compliance standards. To access this information:

- From the **Enterprise** menu, select **Compliance**, then select **Results**. You have the option of viewing violations associated with compliance standards and compliance frameworks.
  - Click the **Target Compliance** tab for a roll-up view of all violations across all targets, that is, all those targets that are out of compliance.
  - Click the **Compliance Standards** tab to view the list of compliance standards against which there are violations. From this tab, you can also access the **Errors** tab to view the errors against the compliance standard.
- Navigate to the Home page for a particular target. The **Compliance Standard Summary** region lists the compliance violations according to severity level. Click the name of the compliance standard of interest to view the details of the violations.

8.2.2 Viewing Compliance Summary Information

Compliance summary information is available from the Cloud Control home page and individual target home pages.

To view compliance summary information from the Cloud Control home page, follow these steps:

1. Navigate to the Cloud Control home page.
2. From the **Enterprise** menu, select **Compliance**, then select **Results.**

To view compliance summary information from a target's home page, follow these steps:

1. Navigate to the Cloud Control home page.
2. From the **Targets** menu, select the target type, and click the target.
3. On the target's home page, scroll down to the **Compliance Standards Summary** region.

To view compliance summary information from the target menu on a target's home page, follow these steps:

1. Navigate to the Cloud Control home page.
2. From the **Targets** menu, select the target type, and click the target.
3. On the target's home page, click the target menu located at the top-left of the page.
4. Select **Compliance**, then select **Results**. On the Results page, click **Target Compliance**.

8.2.3 Viewing Target Compliance Evaluation Results

Target compliance evaluation results are available by way of the Cloud Control home page and individual target home pages. When testing a target, the possible evaluation results are:

Evaluation Results	Description
Compliant	Target meets the desired state
Non-Compliant	Target does not meet the desired state. At least one test in the compliance standard detected a deviation from the desired state.
Error	No results returned due to an error. The error may be an unexpected internal error or an error in the test. Examples of errors in the test include attempts to: ■ Divide by zero ■ Invoke a function with incorrect parameter values

To view results using Cloud Control home page, follow these steps:

1. Navigate to the Cloud Control home page.
2. From the **Enterprise** menu, select **Compliance**, then select **Results**.
3. Click the **Target Compliance** tab. The Target Results page displays the targets with their Average Compliance Score.

To view compliance evaluation results from a target's home page, follow these steps:

1. Navigate to the Cloud Control home page.
2. From the **Enterprise** menu, select **Targets**, then select the target type.
3. Click the name of the target in which you are interested.
4. On the target's home page, scroll to the **Compliance Standard Summary** region.

Use the page or region to get a comprehensive view about a target in regards to compliance over a period of time. Using the tables and graphs, you can easily watch for trends in progress and changes.

Note: Trend overview data might take up to six hours, after target discovery, to display in the time series charts.

8.2.4 Viewing Compliance Framework Evaluation Results

To effectively use a compliance framework, organize the framework to reflect the compliance framework you use in your organization.

Oracle provides an out-of-box framework for Payment Card Industry (PCI), as well as one for the Oracle Generic Compliance. These out-of-box frameworks can be used as a starting point for you to create your own frameworks to match your needs.

To view the results of a compliance framework evaluation, use the Evaluations Results page accessed through the Compliance Frameworks tab.

1. From the Enterprise Manager Cloud Control Home page, select **Enterprise**, select **Compliance**, then select **Results**.
2. On the Compliance Results page, click the **Compliance Frameworks** tab and highlight the compliance framework of interest.

Tips on Using Compliance Frameworks

Here are a few tips on how to best use compliance frameworks:

- Manage your compliance framework to match your company's framework
- Specify or manage compliance standards to define all your compliance tests
- Manage compliance standard rules
- Use the Results page, accessed from the **Compliance** menu, to:
  - Browse and Search Compliance Framework Evaluation Results
  - Browse and Search Compliance Framework Errors

Benefits of Using Compliance Frameworks

Compliance standards are defined to perform tests, for example, test if a configuration is set properly, test to see if real-time changes are occurring, and so on. In turn, a compliance framework is a way to map how different areas of your compliance initiative are going to be affected by the results of those tests.

For example, an organization may choose to define a compliance framework that extends an out-of-box compliance framework. This is accomplished by creating a new compliance framework like the out-of-box compliance framework and include new or existing compliance standards.

8.2.5 Investigating Compliance Violations and Evaluation Results

Here are a few suggestions for investigating compliance violations. Attend to the most critical violations or those that have the biggest impact on your enterprise.

- Study the statistics on the Enterprise Summary Home page. In particular, look at the statistics in the Compliance Summary region. The compliance violations with "Critical" severity should be dealt with first.
- Address targets that have the lowest compliance scores.
- For the compliance violations of a particular target, examine the home page for that target. The Compliance Standard Summary region provides overview information, but it also gives you access to the Trend for that target.
- To deal with compliance standards regardless of the target, from the **Enterprise** menu, select **Compliance.** Using this option, you have access to all the compliance violations events for the enterprise (Result option), the compliance associations, the compliance standard library (Library option), and compliance evaluation errors.

  Note: Only results from those targets for which you have View privilege will be available for viewing.

  - Navigate to the Results page for a particular compliance standard. In the navigation tree, click the name of the compliance standard and a summary page lists all the targets along with the number of violations.
  - Navigate to the Trend Overview page to see charts relating to the number of targets evaluated, the average violation count per target, number of targets by compliance score, and the average compliance score.

Using the Compliance Standards Evaluation Results Page

Use the Compliance Standards Evaluation Results page to:

- View a summary of how well targets, that are expected to comply with a compliance standard, are actually adhering to the standard.
- View the detailed evaluation results of the compliance standard.
- Study the details of how well the compliance standard within the targets complied with the compliance standard. The results reflect the hierarchy within the compliance standard as defined by its folders.
- View how the targets are complying with this compliance standard. By studying the graphs on the Trend Overview page, you can watch for trends and changes in the compliance of the targets to the compliance standard.
- Study the compliance scores, violation count, and targets evaluated for all the elements of a compliance standard.
- Study the impact of violations and recommendations. The impact explains why the compliance standard is important. The recommendation explains how to bring a system back into compliance with the compliance standard.

Note: When viewing compliance evaluation results, the most recent results are provided. The results of a compliance evaluation overwrite the previous evaluation's results

8.2.6 Investigating Evaluation Errors

The Evaluation Errors page reports the deviations from the norm, that is, statistics about the problems encountered during the evaluation. On initial display, the Evaluation Errors page shows all the evaluation errors.

- Use the Evaluation Errors page to view the errors that occurred as a result of metric collection, as well as those that occurred during the last evaluation.
- Use the search filter to view only those evaluation errors that meet a set of search criteria that you specify.
- Click the message in the Message column to decide what your course of action should be to resolve the error.
- Normally the results of an evaluation overwrite the previous evaluation's results. However, in the case of evaluation failure or data provider collection failure, the previous results are left untouched.

Once the underlying problem is fixed, the error is no longer reported.

Example of Search Filter

By default, all the evaluation errors in your enterprise configuration appear in the results table. However, you can specify a set of search criteria and then perform a search that will display only the evaluation errors that meet those criteria in the results table.

For example, if you choose Host in the Target Type list, contains in the Target Name list, and "-sun" in the adjacent Target Name text field, and then click **Go**, Enterprise Manager displays, in the results table, only the compliance standard rule evaluation errors for the hosts that contain "-sun" in their names.

8.2.7 Compliance Audit by a Compliance Auditor

Before you perform an audit, ensure that the compliance manager or line of business manager has associated the necessary compliance standards with compliance frameworks that are being followed. The IT administrator then ensures that the compliance standard is associated to the appropriate targets in the environment. Also ensure that you have Enterprise Manager login and view target privileges. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To verify that targets are compliant, follow these steps:

1. Determine how compliant the targets are with respect to various compliance frameworks.

   From the **Enterprise** menu on the Cloud Control home page, select **Compliance,** then select **Results**. Click **Compliance Frameworks**, then select **Evaluation Results**.

   Analyze the evaluation errors and violations raised.

2. Determine how compliant the target is with respect to various associated and evaluated compliance standards for that target.

   Click **Target Compliance.**

   Analyze the evaluation errors and violations raised.

8.2.8 Compliance Reports

Enterprise Manager provides reports specific to compliance: Descriptions and Results. These compliance reports are available by selecting **Reports** on the **Enterprise** menu, selecting **Information Publisher Reports**, and then scrolling to the Compliance section.

- Descriptions reports

  The Descriptions reports list all the available compliance standards, compliance frameworks, and compliance standard rules available in the Compliance Library. These reports enable you to decide whether additional compliance standards and compliance frameworks need to be defined for your enterprise to attain and maintain its compliance to the standards.

- Results reports

  The Results reports provide details of the various evaluations against compliance standards and compliance frameworks. Using the Results reports you can view, in one place, all the statistics regarding the compliance of your enterprise against the defined standards. To view the target that is most likely in need of your immediate attention, view the Target with Lowest AVG COMPLIANCE SCORE report. The following are examples of the reports provided:

  – Compliance Standard Results Details

    Displays the compliance summary for all the compliance standards evaluated against a target. Data includes compliance score, compliant and non-compliant rules, violations, and last evaluation date.

  – Compliance Standard Result Summary

    Displays the compliance summary of a particular compliance standard. For example, if there are three targets each reporting on Security Recommendations for Oracle Products compliance, the Result Summary rolls up the information into one report. Data includes average compliance score,

the number of targets that need immediate attention, and the number of rules that are non-compliant.

8.2.9 Compliance Score and Importance

A target's compliance score for a compliance standard is used to reflect the degree of the target's conformance with respect to compliance standard. The compliance score is in the range of 0% to 100% inclusive. A compliance score of 100% indicates a target fully complies with the compliance standard.

During an evaluation, a target is found to be compliant or non-compliant with that compliance standard.

Types of Importance

For compliance frameworks, importance indicates the relative importance of a compliance standard to all other compliance standards in the compliance framework.

For compliance standards, importance indicates the relative importance of a compliance standard rule to all other compliance standard rules in the compliance standard. The values represent a way of weighting a compliance standard.

However, just because a compliance standard rule has an importance of 'low' does not mean that it can safely be ignored.

Importance is used to roll up results bottom up in a compliance standard hierarchy.

The following sections provide examples of how the compliance score is calculated.

8.2.9.1 How Compliance Score of a Compliance Standard Rule-Target Is Calculated

Note: This calculation is used for WebLogic Server Signature rules and Repository rules.

Compliance score of a compliance standard rule-target is calculated by taking the severity and importance of the compliance standard rule and multiplying the result by the total number of violations divided by the total number of rows evaluated for that target.

The formula is:

```
hirange - (hirange - lorange) * (number of violations / number of rows evaluated)
```

The following table provides the combination of the severity and importance values used to calculate a compliance score.

Table 8–1 Importance and Severity Ranges

Importance	Critical Severity (*1*)	Warning Severity (*1*)	Minor Warning Severity (*1*)
High	0-25 (2)	66-75	95-96
Normal	26-50	76-85	97-98
Low	51-75	86-95	99-99

(1) low range and high range of the severity

(2) 0 is the lorange; 25 is the hirange

8.2.9.2 How Compliance Score of a Real-time Monitoring Rule is Calculated

The compliance score of a real-time monitoring rule is a rule-based score that is the number of observation bundles that have violations rather than all observation bundles that have happened over time. (**Note:** There can only be one violation per bundle.)

When calculating the count of past observation bundles, the most recent bundles are rated and they have a different rating as they get older. For example, if there had been 1,000,000 observation bundles (all of which have no violations) over the history of the Enterprise Manager installation and then one day a new bundle comes in that has a violation, then the score would have been 999,999/1,000,000, or 100% when rounding.

This one violation, though in the context of other bundles that came in just in the last few days, may be really important. To continue the example, say in the last week there has only been 10 bundles. Then this one comes in, 9/10 of the observations are good, or 90% score. To keep track of the older observations, observation bundles are weighted by how old they are.

The score is calculated using the formula:

```
1 - V/T
  where T is the sum of all the weighted counts
    and V is the sum of the current violations (which is the same as the number
    of bundles in violation at that time)
```

The result of the calculation of 1-V/T will be a number around 1 as V is 0 or will be a number near 0 when V is close to the value of T.

8.2.9.3 How Compliance Score of a Compliance Standard for a Target Is Calculated

The compliance score of a compliance standard for each target is calculated by taking the individual compliance score of each rule - target and multiplying it by its importance. This multiplication is repeated for each rule then the resulting products are added. The sum of the products is then divided by the sum of the importance of each rule. See Figure 8–1.

Figure 8–1 How Compliance Score of a Compliance Standard-Target Is Calculated

Key:

CS: compliance standard
Rule: compliance standard rule. There are 3 rules: Rule1, Rule2, and Rule3.
i: importance
  i1: importance for Rule1
  i2: importance for Rule2
  i3: importance for Rule3
S: compliance score of the rule
  S1: compliance score for rule1-target
  S2: compliance score for rule2-target
  S3: compliance score for rule3-target

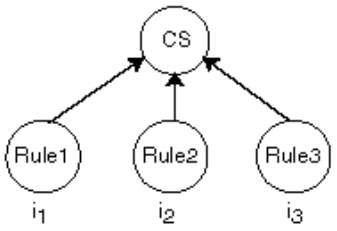

$$\text{Compliance Score of Compliance Standard-Target} = \frac{(S_1 \times i_1) + (S_2 \times i_2) + (S_3 \times i_3)}{(i_1 + i_2 + i_3)}$$

8.2.9.4 How Compliance Score of a Compliance Framework Is Calculated

The compliance framework score is a rolled up weighted average of all compliance standard-target scores across all compliance standards within the compliance framework hierarchy. The weight is based on the importance of a compliance standard. In Figure 8–2, compliance framework CF has 2 standards CS1 and CS2. CS1 is associated and evaluated on targets t1 and t2 and CS2 is associated and evaluated on targets t3 and t4.

Figure 8–2 How Compliance Score of a Compliance Framework Is Calculated

Key:

CF: compliance framework
CS: compliance standard
  CS_1: compliance standard 1
  CS_2: compliance standard 2
t: target
i: importance
  i_{cs1}: importance of CS1
  i_{cs2}: importance of CS2
ST: compliance score of a compliance standard for a target
  ST1: compliance standard score for CS1-t1
  ST2: compliance standard score for CS1-t2
  ST3: compliance standard score for CS2-t3
  ST4: compliance standard score for CS2-t4

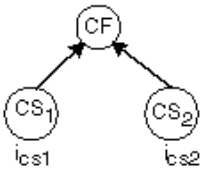

$$\text{Compliance Score of Compliance Framework} = \frac{(ST_1 \times i_{cs1}) + (ST_2 \times i_{cs1}) + (ST_3 \times i_{cs2}) + (ST_4 \times i_{cs2})}{(i_{cs1} + i_{cs1} + i_{cs2} + i_{cs2})}$$

8.2.9.5 How Compliance Score of a Parent Node Is Calculated

The compliance score of a hierarchy node/parent node is calculated as shown in Figure 8–3. Compliance standards are hierarchical, thus the top node in the tree is known as the parent node.

Figure 8–3 Compliance Score of Parent Node

$$\text{Compliance Score of Parent} = \frac{\sum_{\forall i} S_i \times I_i}{\sum_{\forall i} I_i}$$

In Figure 8–3:

- i represents the number of children
- S is the score of the child node
- I is the importance of the child node

8.3 Managing Compliance

Before you can use the compliance features, compliance frameworks, compliance standards, and compliance standard rules must be defined for your enterprise.

The following sections describe how to define and maintain these compliance entities.

- Section 8.3.1, "About Compliance Frameworks"
- Section 8.3.2, "Operations on Compliance Frameworks"
- Section 8.3.3, "About Compliance Standards"

8.3.1 About Compliance Frameworks

A compliance framework is a hierarchical structure comprised of one or more compliance standards, compliance standard rule folders, and compliance standard rules. It is a way for you to map your standards to a structure similar to the regulatory or standard compliance structure you use in your company.

Accessing Compliance Frameworks

To access compliance frameworks, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click **Compliance Frameworks** tab.
3. Highlight the compliance framework you want to manage and choose the action you want to perform.

Frameworks Provided by Oracle and User-Defined Compliance Frameworks

There are compliance frameworks provided by Oracle and user-defined compliance frameworks.

- Compliance frameworks provided by Oracle
  - PCI DSS 2.0 (Payment Card Industry Data Security Standard) is a standard which you can use to evaluate your managed targets compliance with security and best practices standards.
  - Oracle Generic Compliance Framework is a standard set of compliance standards and associated controls for tracking changes and events taking place across your IT infrastructure for determining how well your organization is in compliance with your IT policies.
- User-defined compliance frameworks

  You can define a compliance framework to satisfy the needs of your organization. You can also create a user-defined framework by performing a create-like on an out-of-box framework.

Compliance frameworks provided by Oracle cannot be deleted or edited. However, if you want to extend these frameworks, use the Create Like functionality to create your own user-defined frameworks based on the out-of-box frameworks and then edit the new frameworks.

Recommendation: It is highly recommended that you create a top level compliance framework like the ones provided for PCI and Oracle Generic compliance.

Example of Using the PCI Standard

If you follow the Payment Card Industry (PCI) standards framework, you may have a multiple level structure that mirrors the structure of PCI as follows:

- PCI DSS 2.0 - Payment Card Industry Data Security Standards compliance framework which contains:

  - Build and Maintain a Secure Network (PCI 2.0) compliance standard which contains:
    * Encrypt all administrative access using SSH, VPN, or SSL/TLS (PCI 2.3) compliance standard rule

The compliance standard (PCI 2.0) contains a number of compliance standard rules that are specific to a target type. A single compliance score will be calculated for that compliance standard and then can be rolled up to all the compliance frameworks as well. The top level compliance framework (PCI 2.0) will always be treated as the actual compliance framework that is used.

Benefits of Using Compliance Frameworks

Compliance standards are defined to perform tests, for example, test if a configuration is set properly, test to see if real-time changes are occurring, and so on. In turn, a compliance framework is a way to map how different areas of your compliance initiative are going to be affected by the results of those tests.

For example, an organization may choose to define a compliance framework that extends an out-of-box compliance framework. This is accomplished by creating a new compliance framework like the out-of-box compliance framework and include new or existing compliance standard rules.

Reasons for Using Compliance Frameworks

There are a number of reasons for creating compliance frameworks including:

- Mapping underlying IT violations to the regulatory and standard compliance controls used by your company
- Compliance auditing at compliance specification level (for example, Payment Card Industry (PCI))
- Auditing, security evaluation, and trend analysis

What Compliance Frameworks Can Do

A compliance framework can:

- Represent industry-wide standards or can be created to match your internal frameworks in use.

  Many companies may start by using an industry-wide framework, but modify it according to their own needs and auditing requirements. For example, an organization may choose to define a compliance framework that extends an out-of-box compliance framework. This is accomplished by creating a new compliance framework like the out-of-box compliance framework and include new or existing compliance standard rules.
- Be used as a reference compliance framework or a certified compliance framework
- Be a collection of compliance standards describing best practices in an enterprise. Compliance standards are defined to perform tests, for example, test if a configuration is set properly, test to see if real-time changes are occurring, and so on. In turn, a compliance framework is a way to map how different areas of your compliance initiative are going to be affected by the results of those tests.

Compliance Frameworks and Compliance Scores

The compliance framework is the entry point when looking at compliance scores from a high level view such as the Compliance Results dashboard. Each entity of a compliance framework should have a user-defined importance that is assigned for

reporting/compliance score roll up. The importance can be set for all internal nodes in a compliance framework or compliance standard hierarchy.

This importance at the top compliance framework is the default, but you may decide that more importance should be placed on one compliance sub group over another.

Compliance frameworks can include subgroups and nested subgroups.

Usage Notes

- Evaluation Results for a repository rule may become invalidated if a compliance standard rule within a compliance framework is modified or deleted. Evaluation of a compliance standard always references the current compliance standard rule definition for each compliance standard rule within the standard.
- Compliance frameworks can include subgroups and nested subgroups.
- The compliance framework is the entry point when looking at compliance scores from a high level view such as the Compliance Results dashboard. Each entity of a compliance framework should have a user-defined importance that is assigned for reporting/compliance score roll up. The importance can be set for all internal nodes in a compliance framework or compliance standard hierarchy.
- The importance at the top compliance framework is the default, but you may decide that more importance should be placed on one compliance sub group over another.
- Compliance frameworks can include compliance standards of different target types.

8.3.2 Operations on Compliance Frameworks

You can perform the following operations on a compliance framework:

- Section 8.3.2.1, "Creating a Compliance Framework"
- Section 8.3.2.2, "Creating Like a Compliance Framework"
- Section 8.3.2.3, "Editing a Compliance Framework"
- Section 8.3.2.4, "Deleting a Compliance Framework"
- Section 8.3.2.5, "Exporting a Compliance Framework"
- Section 8.3.2.6, "Importing a Compliance Framework"
- Section 8.3.2.7, "Browsing Compliance Frameworks"
- Section 8.3.2.8, "Searching Compliance Frameworks"

The following sections explain these operations.

8.3.2.1 Creating a Compliance Framework

Before you create a compliance framework, ensure you have privileges to access the compliance standards you will be including during the definition of the framework. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To make the creation for the compliance framework easier, ensure that the compliance standards, which will be referred to by the compliance framework, are already defined in the Enterprise Manager. The compliance standards you add to a compliance framework may be system-defined and user-defined standards as displayed on the Compliance Standard Library page. If you do not define the compliance standards before hand, you must add them later.

To create a compliance framework, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Frameworks** tab.
3. Click **Create** button.
4. Provide the Name and Author and click **OK**.
5. Once you have provided the information on the definition page, look at the options available when you right-click the name of the compliance framework (located at the top-left of the page). From this list you can create subgroups, include compliance standards, and so on.
6. Click **Save**.

Usage Notes

- Lifecycle status can be either Development or Production.
  - Development

    Indicates a compliance framework is under development and that work on its definition is still in progress. While in development mode, all management capabilities of compliance frameworks are supported including editing of the compliance framework and deleting the compliance framework. Results of development compliance standards will NOT be viewable in target and console home pages, and the compliance dashboard.

    Lifecycle status default is Development. It can be promoted to Production only once. It cannot be changed from Production to Development.
  - Production

    Indicates a compliance framework has been approved and is of production quality. When a compliance framework is in production mode, its results are rolled up into a compliance dashboard, target and console home page.

    Production compliance frameworks can only refer to Production compliance standards. A production compliance framework can be edited to add/delete references to production compliance standards ONLY!

    Lifecycle status cannot be changed from Production to Development.
- All compliance frameworks with the same keyword will be grouped together when sorted by the Keyword column.
- If you modify a compliance standard that has been added to a compliance framework, either by editing the compliance standard directly, or by using Import to overwrite the compliance standard with new settings, the existing evaluations become invalid. That is, if this modified compliance standard was included in a compliance framework that was previously evaluated, and has evaluation results, these results are no longer viewable.

Adding a Compliance Standard to a Compliance Framework

Use the Include Compliance Standard Reference page to select one or more compliance standards to be added to the compliance framework.

Use the search criteria to minimize the number of compliance standards that display in the Select list.

Once you make your selections, click **Continue**. The Include Compliance Standard Reference page appears with the compliance standards you chose on the Include Compliance Standard Reference page.

Editing Importance

After you add the compliance standards that are to be part of the compliance framework, the Create Compliance Framework page appears listing the compliance standards you chose to add to the compliance framework. At this time you can edit the importance of each compliance standard.

The importance impacts the overall compliance score. Those compliance standards with higher importance elevate the importance of the compliance framework, whereas those compliance standards with less importance lower the importance of the compliance framework.

See Section 8.2.9, "Compliance Score and Importance" for details on how this score is computed.

8.3.2.2 Creating Like a Compliance Framework

To create a compliance framework like another compliance framework, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Frameworks** tab.
3. On the Compliance Framework Library page, highlight the compliance framework you want to use as the base and click the **Create Like** button.
4. Customize the fields as needed.

   Ensure that the Compliance Framework name is different from the original compliance framework and any other existing compliance frameworks.
5. Click **Save**.

8.3.2.3 Editing a Compliance Framework

Use the edit compliance framework feature to add new compliance standard rules to a compliance framework, or edit details of existing compliance frameworks, or delete compliance standard rules from the compliance framework.

Before you edit a compliance framework, ensure that you have privileges to access the compliance framework to be edited. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To edit a compliance framework, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Frameworks** tab.
3. Highlight the compliance framework you want to edit and click the **Edit** button.
4. Update the properties as needed.

   To add standards and subgroups, right-click the name of the framework located at the top left of the page.
5. Click **Save**.

Usage Notes

- Changing a compliance framework definition may impact trend analysis.
- The compliance standards you add to a compliance framework may be system-defined and user-defined compliance standards as displayed on the Compliance Standard Library page.
- If you modify a compliance standard that has been added to a compliance framework, either by editing the compliance standard directly, or by using Import to overwrite the compliance standard with new settings, the existing evaluations become invalid. That is, if this modified compliance standard was included in a compliance framework that was previously evaluated, and has evaluation results, these results are no longer viewable. The compliance framework evaluation results will again become visible after the next evaluation happens. The new evaluation includes the changes to the compliance standard within the compliance framework.
- The importance impacts the overall compliance score. Those compliance standards with higher importance elevate the importance of the compliance framework, whereas those compliance standards with less importance lower the importance of the compliance framework.
- A compliance standard can be added to more than one compliance framework, and can have a different importance when added to a different compliance framework. For example, you could have a compliance standard called Check Password Expired which flags user accounts with expired passwords. This compliance standard may be a member of two compliance frameworks: All System Passwords Secure and 30-day Password Validation. The All System Passwords compliance framework verifies a password's security, whereas the 30-day Password Validation compliance framework checks the date that this password was last set.
  - The Check Password Expired compliance standard could have Extremely High importance for the 30-day Password Validation compliance framework, since this check is warning users that their passwords are about to expire.
  - In the All System Passwords Secure compliance framework, the Check Password Expired compliance standard could have a Normal importance, and other added compliance standards that do security checks could have a higher importance within the compliance framework.

8.3.2.4 Deleting a Compliance Framework

Before you delete a compliance framework, ensure that you have privileges to access the compliance framework to be deleted. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To delete a compliance framework, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Frameworks** tab.
3. Highlight the compliance framework you want to delete, click **Delete** button.
4. Confirm that you want to delete the compliance framework by clicking **OK**.

Usage Notes

- You can delete a single compliance framework or a list of compliance frameworks. When you delete a compliance framework, the associated metadata and evaluation results are also deleted.
- **YOU CANNOT DELETE COMPLIANCE FRAMEWORKS DEFINED BY ORACLE.**

8.3.2.5 Exporting a Compliance Framework

Exporting allows you to re-use a compliance framework that you already have, that is, minimize duplication of effort.

Before you export a compliance framework, ensure that you have privileges to access the compliance framework to be exported. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To export a compliance framework, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Frameworks** tab.
3. Highlight the compliance framework you want to export.
4. From the **Actions** menu, select **Export**.
5. Provide the file name to which the compliance framework definition is to be exported. Determine whether is to be a shallow or deep export. In a shallow export, no leaf level rules or compliance standards are to be exported. In a deep export, all leaf level rules and compliance standards are exported.

The system generates an XML representation of the compliance framework in the directory and file you specify.

8.3.2.6 Importing a Compliance Framework

Importing allows you to re-use a compliance framework that you already have, that is, minimize duplication of effort.

Before you import a compliance framework, ensure the compliance framework to be imported is defined in a file. The location of the file is independent of Enterprise Manager. Also ensure that you have privileges to access the compliance framework definition XML file to be imported. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To import a compliance framework, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Frameworks** tab.
3. From **Actions** menu, select **Import**.
4. Provide the file name from which the compliance framework definition (as per Compliance Framework XSD) will be imported. Specify whether to override an existing definition if one already exists. Specify whether to import referring content as well, that is, shallow or deep import. In a shallow import, no leaf level rules or compliance standards are to be imported. In a deep import, all leaf level rules and compliance standards are imported. In a deep import, real-time monitoring facets are also imported for real-time monitoring type of rules.

5. Click **OK**.

8.3.2.7 Browsing Compliance Frameworks

Before you browse compliance frameworks, ensure you have privileges to access the compliance framework definitions you will be browsing. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To browse a compliance framework, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Frameworks** tab.
3. To view the details of a particular compliance framework, highlight the compliance framework and click **Show Details**.

8.3.2.8 Searching Compliance Frameworks

Before you search compliance frameworks, ensure you have privileges to access the compliance framework definitions you will be searching. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To search a compliance framework, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Frameworks** tab.
3. In the Search portion of the page, provide criteria to use to narrow the search.
4. Click **Search**.

8.3.2.9 Browsing Compliance Framework Evaluation Results

Before you browse compliance framework evaluation results, ensure you have privileges to access the compliance framework definitions you will be browsing. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To browse compliance framework evaluation results, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Results.**
2. Click the **Compliance Frameworks** tab and then the **Evaluation Results** tab.
3. Highlight the compliance framework and click **Show Details** to view the details of a particular compliance framework.

Results include the following:

- Average compliance score for different targets evaluated for compliance standards referred to by the compliance framework
- Count of target evaluations (critical, warning, compliant) for different compliance standards referred to by the compliance framework
- Count of violations (critical, warning, minor warning) related to compliance standards referred to by the compliance framework

8.3.2.10 Searching Compliance Framework Evaluation Results

Before you search compliance framework evaluation results, ensure you have privileges to access the compliance framework evaluation results you will be

searching. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To search compliance framework evaluation results, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Results.**
2. Click the **Compliance Frameworks** tab and then the **Evaluation Results** tab.
3. In the Search portion of the page, provide criteria to use to narrow the search.
4. Click **Search**.

8.3.2.11 Browsing Compliance Framework Errors

Before you browse compliance frameworks, ensure you have privileges to access the compliance framework evaluation errors you will be browsing. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To browse compliance framework errors, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Results.**
2. Click the **Compliance Frameworks** tab and then the **Errors** tab.

Usage Notes

The error may be an unexpected internal error or an error in the test.

Evaluation errors can often be due to configuration and installation issues. See the following manuals for information:

- *Oracle Enterprise Manager Cloud Control Basic Installation Guide*
- *Oracle Enterprise Manager Cloud Control Advanced Installation and Configuration Guide*

If the installation and configuration are correct and the errors persist, call Oracle for assistance.

8.3.2.12 Searching Compliance Framework Errors

Before you search compliance framework errors, ensure you have privileges to access the compliance framework evaluation errors you will be searching. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To search for compliance framework errors, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Results.**
2. Click the **Compliance Frameworks** tab and then the **Errors** tab.
3. In the Search portion of the page, provide criteria to use to narrow the search.
4. Click **Search**.

Usage Notes

The error may be an unexpected internal error or an error in the test.

Evaluation errors can often be due to configuration and installation issues. See the following manuals for information:

- *Oracle Enterprise Manager Cloud Control Basic Installation Guide*

- *Oracle Enterprise Manager Cloud Control Advanced Installation and Configuration Guide*

If the installation and configuration are correct and the errors persist, call Oracle for assistance.

8.3.2.13 Verifying Database Targets Are Compliant with Compliance Frameworks

For auditors to verify that database targets are in compliance with the compliance frameworks, the Enterprise Manager structure needs to be defined. The steps to provide this structure include the following:

1. Super Administrator creates three Enterprise Manager users: Compliance Author, IT Administrator, and Compliance Auditor.
2. Super Administrator assigns the appropriate roles and privileges to the Compliance Author and IT Administrator.
3. Super Administrator assigns the same target privileges to IT Administrator and Compliance Auditor.
4. Compliance Author logs in to Enterprise Manager and views out-of-box compliance frameworks, compliance standards, and compliance standard rules.

   He then enables and disables the appropriate compliance standard rules and creates new compliance standard rules.
5. IT Administrator logs in to Enterprise Manager and associates the targets for which he has target privileges with the appropriate compliance standards.
6. IT Administrator sets up the correct configuration parameters and settings for the compliance frameworks, compliance standards, and compliance standard rules for a particular target.

   He then creates a monitoring template from this target and applies it to the other targets, to which he has privileges, that require compliance standards.
7. Compliance Auditor logs in to Enterprise Manager to view the violations and errors at the Enterprise level, for which he has view privileges, and at each target level.

   He would then take the necessary actions to rectify the errors and violations.

8.3.3 About Compliance Standards

A compliance standard is a collection of checks or rules. It is the Enterprise Manager representation of a compliance control that must be tested against some set of IT infrastructure to determine if the control is being followed.

Compliance standards are made up of the following in a hierarchical structure (see Figure 8–4):

- Compliance standard rules
- Rule folders that can include nested rule folders and individual compliance standard rules.

  Rule Folders are hierarchical structures that contain compliance standard rules. A rule folder has an importance attribute that denotes the importance of the rule folder relative to its siblings at the same level. This importance is considered when determining compliance scores being rolled up from other sibling rule folders. A certain rule folder may have multiple tests that occur, in this way a certain test can be given more weight than other tests.

- Included compliance standards. A compliance standard can include other compliance standards.

Figure 8–4 Compliance Standard Definition

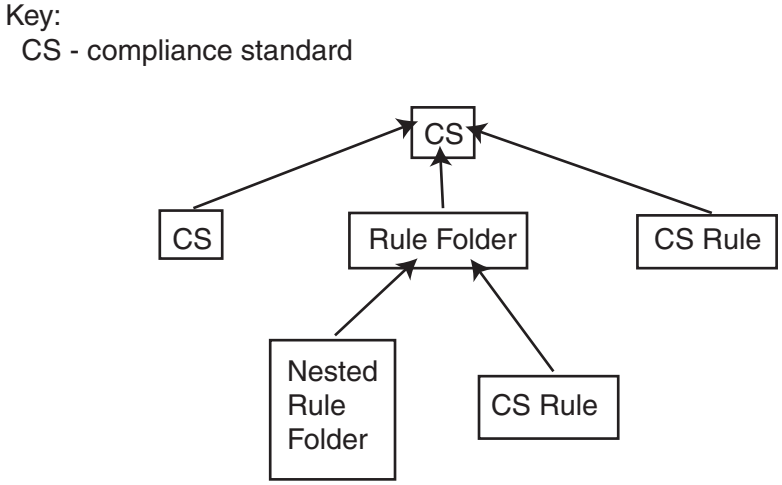

What Compliance Standards Can Do

- Can represent industry-wide standards. A compliance standard is applicable to a single target type.
- Be used as a reference configuration or a certified configuration
- Be a collection of compliance standard rules describing best practices in an enterprise

For example, when a target fails to adhere to a compliance standard, the target is not in compliance with the compliance standard.

Accessing Compliance Standards

The compliance standards, including those provided by Oracle, are available on the Compliance Standard Library page. To access this page, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library**.
2. Click the **Compliance Standards** tab.

To view the compliance standard rules associated with the compliance standard, click the name of the compliance standard and click **Show Details**. Once the Compliance Standard Detail page appears, right click the name of the standard located at the top left of the page, and select either **Expand** or **Expand All Below**.

Note: The compliance standards defined by Oracle cannot be changed. However, you can create a standard similar to the one provided by Oracle by using the Create Like feature.

General Usage Notes for Compliance Standards

You can override an existing compliance standard by checking the Overwrite existing compliance standards check box. As a result:

- If you override a compliance standard, the override deletes all target and template associations, as well as evaluation results for that compliance standard.

- If the overwritten compliance standard is part of a compliance framework, the compliance standard is updated in the compliance framework. However, the evaluation results for that compliance standard within the compliance framework are invalidated.
- Evaluations of compliance standards happen after the compliance standards are associated to a target.

  For repository compliance standards, the evaluation happens after the standard is associated with a target. For WebLogic Server compliance standards, evaluation happens when the Agent-side evaluation metric is refreshed. The refresh occurs once every 24 hours for Oracle WebLogic Domain, Oracle WebLogic Java EE Server, and Oracle WebLogic Cluster targets.

  For Real-time Monitoring Rules, an evaluation is when a compliance standard is associated to a target. A violation occurs when an observation bundle contains at least one observation that is unauthorized

Usage Note Specific to Repository Rules

If you manually type a WHERE clause in the compliance standard rule XML definition, then the < (less than) symbol must be expressed as <, to create a valid XML document. For example:

```
<WhereClause>:status &lt; 100</WhereClause>
```

Example of How to Set Up Compliance Standards for Auditing Use

For auditors to verify that database targets are in compliance with the compliance frameworks, the Enterprise Manager structure needs to be defined. The steps to provide this structure includes the following:

1. Super Administrator creates three Enterprise Manager users: Compliance Author, IT Administrator, and Compliance Auditor.
2. Super Administrator assigns the appropriate roles and privileges to the Compliance Author and IT Administrator.
3. Super Administrator assigns the same target privileges to IT Administrator and Compliance Auditor.
4. Compliance Author logs in to Enterprise Manager and views out-of-box compliance frameworks, compliance standards, and compliance standard rules.

   He then enables and disables the appropriate compliance standard rules and creates new compliance standard rules.
5. IT Administrator logs in to Enterprise Manager and associates the targets for which he has target privileges with the appropriate compliance standards.
6. IT Administrator sets up the correct configuration parameters and settings for the compliance frameworks, compliance standards, and compliance standard rules for a particular target.

   He then creates a monitoring template from this target and applies it to the other targets, to which he has privileges, that require compliance standards.
7. Compliance Auditor logs in to Enterprise Manager to view the violations and errors at the Enterprise level, for which he has view privileges, and at each target level.

   He would then take the necessary actions to rectify the errors and violations.

8.3.4 Operations on Compliance Standards

You can perform the following operations on a compliance standard:

- Section 8.3.4.1, "Creating a Compliance Standard"
- Section 8.3.4.2, "Creating Like a Compliance Standard"
- Section 8.3.4.3, "Editing a Compliance Standard"
- Section 8.3.4.4, "Deleting a Compliance Standard"
- Section 8.3.4.5, "Exporting a Compliance Standard"
- Section 8.3.4.6, "Importing a Compliance Standard"
- Section 8.3.4.7, "Browsing Compliance Standards"
- Section 8.3.4.8, "Searching Compliance Standards"

The following sections explain these operations.

8.3.4.1 Creating a Compliance Standard

You can use the compliance standards provided by Oracle, for example, Security Configuration for Oracle Database, or create your own standard.

Before creating a compliance standard, ensure the compliance standards and compliance standard rules, which will be referred to by the compliance standard, are defined in the Management Repository. Also ensure the you have privileges to access the compliance standards and compliance standard rules you will be including in the compliance standard. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To create a compliance standard, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standards** tab.
3. Click the **Create** button. You will prompted for the Name, Author, target type to which the standard is applicable., and the type of compliance standard (Repository, WebLogic Server Signature, Real-time Monitoring). Click **OK**.
4. On the resulting Compliance Standard Detail page, provide the property values.

   Click **Add** to either add a keyword by which this standard is identified or use an existing keyword.
5. To further define the compliance standard, right-click the name of the compliance standard located at the top left of the page. From this menu, you can create rule folders, add rules, and compliance standards.

   By using rule folders, you can view the summary of results, categorized by the targets that were evaluated against the selected rule folder and the Compliance Standard Rules evaluated for the selected rule folder.
6. Click **Save**.

Once you define the compliance standard, associate the standard with a target and define the target type settings.

1. While on the Compliance Standards Library page, ensure the correct compliance standard is highlighted.
2. Click the **Associate Target** button.

3. On the **Target Association for Compliance Standard** page, click **Add** to choose the target to be evaluated against the standard.
4. In the **Search and Select: Targets** popup, choose the appropriate target.
5. Click **Select**.

After you associate the target with the compliance standard, you can edit the parameters associated with the target.

1. While on the **Target Association for Compliance Standard** page, click **Edit**.
2. On the **Customize Compliance Standard Parameters** page, change the parameters as needed.

Note: You can also associate a compliance standard with a target. At the top left of the target's home page, right click the name of the target. On the resulting menu, select **Compliance**, then select **Standard Associations**.

In addition you can, edit and remove existing associations. See Section 8.2.9, "Compliance Score and Importance" for additional information.

Adding a Compliance Standard to Another Compliance Standard

Use the Include Compliance Standard page to select one or more compliance standards to be added to the compliance standard. This list is prefiltered by the target type of the compliance standard.

To add a compliance to another compliance standard:

1. From the Compliance Standard Library page, highlight the compliance standard to which you want to add another compliance standard.
2. Click the **Edit** button.
3. On the Properties page, right-click the node, located at the top left of the page.
4. On the resulting menu, select **Add Standards**.
5. Select the compliance standard to include. Click **OK**.

   When you include a compliance standard within another top level compliance standard, the included standard must be of the same target type as the top level compliance standard. For composite target types, one of the member target types of the composite target type of the top level standard is a member target type within the top level composite target type.

   Note that a root compliance standard is associated to a root target (of composite target type). Compliance standards are associated to member targets of the same applicable target type and target filter criteria.
6. On the **Properties** page, choose the **Importance** for the compliance standard you just included. Click **Save**.
7. After the compliance standard is included, highlight the root compliance standard. The Properties page displays a set of parameters.

   A parameter is a variable that can be used by one or more compliance standard rules contained in that compliance standard. When a compliance standard rule references a parameter, the parameter's actual value is substituted at compliance

standard rule evaluation time. It is through the use of parameters that customizations of compliance standards is supported.

Usage Notes

- Because compliance standards are hierarchical, the top node in the tree is known as the parent node.
- When you create a compliance standard, the version is 1.
- Lifecycle status default is Development. It can be promoted to Production only once. It cannot be changed from Production to Development.
  - Development

    Indicates a compliance standard is under development and that work on its definition is still in progress. While in Development mode, all management capabilities of compliance standards are supported including complete editing of the compliance standard, deleting the compliance standard, and so on. However, while the compliance standard is in Development mode, its results are not viewable in Compliance Results nor on the target or Cloud Control home page.
  - Production

    Indicates a compliance standard has been approved and is of production quality. When a compliance standard is in production mode, you have limited editing capabilities, that is, you can add references to production rules, and you can delete references to rules ONLY from a compliance standard. All other management capabilities such as viewing the compliance standard and deleting the compliance standard will be supported. Results of production compliance standards are viewable in target and console home pages, and the compliance dashboard. Production compliance standards can only refer to production compliance standards and production compliance standard rules.

    Once the mode is changed to Production, then its results are rolled up into compliance dashboard, target home page, and Cloud Control home page. Production compliance standards can only refer to other production compliance standards and production compliance standard rules. A production compliance standard can be edited to add and delete references to production compliance standards and production compliance standard rules ONLY.

8.3.4.2 Creating Like a Compliance Standard

Before creating a compliance standard like another compliance standard, ensure that you have privileges to access the compliance standard you will be copying from. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To create a compliance standard like another compliance standard, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standards** tab.
3. Click the **Create Like** button.
4. Customize the fields as needed.

   The name has to be different than an existing Compliance Standard.
5. Click **Save**.

8.3.4.3 Editing a Compliance Standard

You can customize compliance standards by editing the existing compliance standard rule settings. You can change the importance for the compliance score calculation, prevent template override, override default parameter values (when possible), and exclude objects from a compliance standard rule's evaluation (when possible).

Before editing a compliance standard, ensure that you have privileges to access the compliance standard to be edited. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

Note: You cannot edit an out-of-box compliance standard, that is, a compliance standard defined by Oracle.

To edit a compliance standard, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standards** tab.
3. Highlight the standard you want to edit and click the **Edit** button.
4. Update the parameters as needed.
5. Click **Save**.

8.3.4.4 Deleting a Compliance Standard

Before you delete a compliance standard, ensure you have privileges to access the compliance standard to be deleted. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".) Also ensure the compliance standard is not in use by a compliance framework. You must remove any references to the compliance standard in all compliance frameworks.

Note: You cannot delete an out-of-box compliance standard, that is, a compliance standard provided by Oracle.

To delete a compliance standard, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standards** tab.
3. Highlight the compliance standard you want to delete, click **Delete** button.
4. Confirm that you want to delete the standard by clicking **OK**.

8.3.4.5 Exporting a Compliance Standard

The Export feature provides a mechanism for transporting user-defined compliance standard definitions across Management Repositories and Cloud Control instances. The export stores the definitions in an operating system file. Because the exported compliance standard definitions are in XML format, they conform to the Oracle Compliance Standard Definition (XSD) format. You can then change the definition of the compliance standard and re-import the generated compliance standard definitions into another Management Repository.

Before you export a compliance standard, ensure the compliance standard to be exported is defined in the Management Repository. Also ensure that you have privileges to access the compliance standard to be exported. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To export a compliance standard, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standards** tab.
3. Highlight the standard you want to export.
4. From the **Actions** menu, select **Export**.
5. Provide the file name to which the standard definition is to be exported. Determine whether is to be a shallow or deep export. In a shallow export, no leaf level rules or compliance standards are to be exported. In a deep export, all leaf level rules and compliance standards are exported.
6. The XML representation of the compliance standard is generated. The file is located in the directory you specify.

8.3.4.6 Importing a Compliance Standard

The Import feature uploads an XML-based compliance standard definition file containing definitions of a single user-defined compliance standard or a list of user-defined compliance standards. This upload creates a new user-defined compliance standard or a list of user-defined compliance standards. This compliance standard must have been previously exported.

The compliance standard xml definition must comply to the compliance standard XML Schema Definition (XSD) as defined in User-Defined Compliance Standard XML Schema Definition.

After importing a user-defined compliance standard, you can edit the standard.

Before importing a compliance standard, ensure the compliance standard to be imported is defined in a file. Also ensure that you have privileges to access the compliance standard definition XML file to be imported. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To import a compliance standard, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standards** tab.
3. From the **Actions** menu, select **Import**.
4. Provide the file name from which the compliance framework definition (as per Compliance Framework XSD) will be imported. Specify whether to override an existing definition if one already exists. Specify whether to import referring content as well, that is, shallow or deep import.
5. Click **OK**.

8.3.4.7 Browsing Compliance Standards

Before browsing compliance standards, ensure you have privileges to access the compliance standard definitions you will be browsing. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To browse a compliance standard, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standards** tab.

3. To view the details of a particular standard, highlight the standard and click **Show Details**.

8.3.4.8 Searching Compliance Standards

Before you search the compliance standards, ensure you have privileges to access the compliance standard definitions you will be searching. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To search for compliance standards, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standards** tab.
3. In the Search portion of the page, provide criteria to use to narrow the search.
4. Click **Search**.

8.3.4.9 Browsing Compliance Standard Evaluation Results

Before you browse the compliance standard evaluation results, ensure you have privileges to access the compliance standard evaluation results you will be browsing. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To browse compliance standard evaluation results, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Results.**
2. Click the **Compliance Standards** tab and then the **Evaluation Results** tab.
3. Highlight the compliance standard and click **Show Details** to view the details of a particular standard.

   Results include the following:

   - Average compliance score for different targets
   - Count of target evaluations (critical, warning, compliant)
   - Count of violations (critical, warning, minor warning)

8.3.4.10 Searching Compliance Standard Evaluation Results

Before you search the compliance standard evaluation results, ensure you have privileges to access the compliance standard evaluation results you will be searching. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To search for compliance standard evaluation results, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Results.**
2. Click the **Compliance Standards** tab and then the **Evaluation Results** tab.
3. In the Search portion of the page, provide criteria to use to narrow the search.
4. Click **Search**.

8.3.4.11 Browsing Compliance Standard Errors

Before you browse compliance standard evaluation errors, ensure you have privileges to access the compliance standard evaluation errors you will be browsing. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To browse compliance standard evaluation errors, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Results.**
2. Click the **Compliance Standards** tab and then the **Errors** tab.

8.3.4.12 Searching Compliance Standard Errors

Before you search compliance standard evaluation errors, ensure you have privileges to access the compliance standard evaluation errors you will be searching. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To search for compliance standard errors, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Results.**
2. Click the **Compliance Standards** tab and then the **Errors** tab.
3. In the Search portion of the page, provide criteria to use to narrow the search.
4. Click **Search**.

Usage Notes

- Use the Evaluation Errors page to view the errors that occurred as a result of metric collection, as well as those that occurred during the last evaluation.
- Use the search filter to view only those evaluation errors that meet a set of search criteria that you specify.
- Click the message in the Message column to decide what your course of action should be to resolve the error.
- On initial display, the Evaluation Errors page shows all the evaluation errors.
- Normally the results of an evaluation overwrite the previous evaluation's results. However, in the case of evaluation failure or data provider collection failure, the previous results are left untouched.

Once the underlying problem is fixed, the error is no longer reported.

Example of Search Filter

By default, all the evaluation errors in your enterprise configuration appear in the results table. However, you can specify a set of search criteria and then perform a search that will display only the evaluation errors that meet those criteria in the results table.

For example, if you choose Host in the Target Type list, contains in the Target Name list, and "-sun" in the adjacent Target Name text field, and then click **Go**, Enterprise Manager displays, in the results table, only the compliance standard rule evaluation errors for the hosts that contain "-sun" in their names.

8.3.4.13 Associating a Compliance Standard with Targets

After you create a compliance standard, you can associate the standard with a target. As part of the association, you can customize parameters, that is, the importance of the standard in relation to the target, status of the compliance standard evaluation, reason for changing the evaluation status, and the thresholds.

Before you associate a compliance standard with a target, ensure you have privileges to access the targets you want to associate compliance standards to. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To associate a compliance standard with a target, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standards** tab.
3. Highlight the compliance standard you want to associate with various targets. Click the **Associate Target** button.
4. Select the targets you want to associate with this compliance standard. Click **OK**.
5. With the compliance standard still highlighted, click the **Override Target Type Settings** button.
6. Customize the critical and warning thresholds and importance as needed.

   By changing critical and warning thresholds, you signify how the Compliance standard score event is generated. For example, if the actual score is less than the critical threshold, then a critical score event is raised.

   Changing the importance can change the compliance score. The importance denotes how important the compliance standard is in the hierarchy.
7. Click **OK**.

To further customize the evaluation of a compliance standard against a target, you can alter compliance standard parameters: importance, critical threshold, and warning threshold. Customizations can also be made on the compliance standard rules used within the compliance standards. For example, for the Secure Ports compliance standard rule, DFLT_PORT is an override parameter. You can change the default value of the port. You can also exclude objects from the evaluation, for example a particular port from the evaluation.

Note: For real-time monitoring, you can change parameters that are used in facet patterns. You can also change Automatic Change Management reconciliation settings.

By changing critical and warning thresholds, you signify how the Compliance standard score event is generated. For example, if the actual score is less than the critical threshold, then a critical score event is raised.

8.3.4.14 Compliance Standards Provided by Oracle

Compliance standards serve as standards by which targets are measured. Compliance standards report deviations and enable closed loop remediation by optionally taking action to bring systems back into compliance. Oracle provides a number of compliance standards including:

- Section 8.3.4.14.1, "Basic Security Configuration for Oracle Database"
- Section 8.3.4.14.2, "High Security Configuration for Oracle Listener"
- Section 8.3.4.14.3, "Storage Best Practices for ASM"
- Section 8.3.4.14.4, "Secure Configuration for Host"

These standards represent best practices and allow you to maintain consistency across enterprise systems and configurations. The trend analysis feature allows fine grained tracking of compliance progress over time.

The following sections provide the highlights of each compliance standard.

8.3.4.14.1 Basic Security Configuration for Oracle Database The Basic Security Configuration For Oracle Database compliance standard provides a benchmark by which to test the targets in your enterprise for compliance to Oracle database security standards.

The compliance standard rules associated with this compliance standard comply with the Oracle recommended security checklist. This standard includes:

- Ensuring well-known accounts are locked and expired
- Ensuring that all profiles have been set to a reasonable of days
- Data dictionary protection has been enabled
- Principle of least privileges is being practiced
- Access controls are effective
- Clients are properly authenticated

8.3.4.14.2 High Security Configuration for Oracle Listener The High Security Configuration for Oracle Listener compliance standard tests Oracle Listeners against the Oracle recommended security checklist.

This compliance standard adheres to the security standards available for the Oracle Listener. This standard ensures that:

- Access to the Listener is restricted, making it more difficult for an operating system user to attack the database
- Network configuration parameter settings are secure.
- Access to the listener configuration tasks are secure. For example, access to the listener is password protected and that no runtime modifications to the listener configuration are allowed.

8.3.4.14.3 Storage Best Practices for ASM This compliance standard checks the Automatic Storage Management (ASM) settings to ensure that customers are correctly setting up the disk groups and therefore avoiding potential space and performance problems. This compliance standard ensures that the disk group:

- With NORMAL or HIGH Redundancy has mirrored or parity protected disks
- Contains disks of significantly different sizes
- Contains disks with different redundancy attributes
- Is checked for disks that are not mirrored or parity protected

8.3.4.14.4 Secure Configuration for Host The Secure Configuration for Host compliance standard tests hosts against operating system threats and attacks.

This compliance standard ensures adherence with best-practice security configuration settings that help protect against operating-system-related threats and attacks, providing a more secure operating environment. This compliance standard ensures that:

- OS configuration parameter, which enables execution of code on the user stack, is not enabled.
- No unintended ports are left open
- There are no insecure services (for example, telnet and ftp) running on the server

- File system on a Windows operating system uses the New Technology File System (NTFS).

8.3.5 About Compliance Standard Rule Folders

Rule Folders are hierarchical structures used to group similar compliance standard rules within a compliance standard. The same compliance standard rules can be added to different Rule Folders within a compliance standard. Rule Folders can be nested within a compliance standard.

A rule folder has an importance attribute that denotes the importance of the rule folder relative to its siblings at the same level. This importance is considered when determining compliance scores being rolled up from other sibling rule folders. A certain rule folder may have multiple tests that occur, in this way a certain test can be given more weight than other tests.

The following topics address compliance standard rule folders:

- Section 8.3.5.1, "Creating Rule Folders"
- Section 8.3.5.2, "Managing Rule Folders in a Compliance Standard"

8.3.5.1 Creating Rule Folders

To create a rule folder, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library**.
2. Click the **Compliance Standards** tab.
3. Either create a compliance standard or edit an existing compliance standard.
4. On the Compliance Standard Library page, highlight the compliance standard and click **Edit**.
5. On the **Properties** page, right-click the name of the compliance standard. The name of the standard is located in the top-left corner of the page.
6. Select **Create Rule Folder**.
7. Type the name of the folder and click **OK**.
8. On the **Properties** page, provide a description, ReferenceUrl, and importance. See Section 8.2.9, "Compliance Score and Importance" for additional information regarding importance.

8.3.5.2 Managing Rule Folders in a Compliance Standard

After you create a rule folder and populate it with compliance standard rules, you can perform the following actions on the folder:

- Edit the tree structure by re-ordering the Rule Folder, Rule Reference, and Compliance Standard Reference nodes in the tree or by deleting any of these nodes.
- Select any node (except the top-level Compliance Standard node) object and then click **Remove** menu item from context menu. The Remove option is disabled on the root node. You can also select multiple objects and click **Remove** to delete multiple nodes.

8.3.6 About Compliance Standard Rules

A compliance standard rule is a test to determine if a configuration data change affects compliance. It describes how to take something that is observed in a target and associate it with a rule group or folder structure so that a compliance score can be built for each level of the rule folder.

This score can then be rolled up and reported by compliance framework. These rule compliance scores are rolled up to compute the compliance standard score and then this score can be rolled up and reported along with the compliance framework scores.

Types of Compliance Standard Rules

The types of compliance standard rules are: Repository rules, WebLogic Server Signature rules, and Real-time Monitoring rules.

- Repository Rules

  Used to perform a check against any metric collection data in the Management Repository.

  Used for checking the configuration state of one or multiple targets. A rule is said to be compliant if it is determined that the configuration items do in fact meet the desired state; that is, the rule test failed to identify any violations. Otherwise, a rule is said to be non-compliant if it has one or more violations. The data source that is evaluated by a compliance standard rules test condition can be based on a repository query. A compliance standard rules test condition can be implemented using a threshold condition based on the underlying metrics (or queries) column value or SQL expression or a PLSQL function

  A repository-check based rule checks the configuration state of one or multiple targets. A rule is said to be compliant if the test fails to identify a violation. In other words, the test determines that the configuration item is in the desired state or has the prescribed value. A rule that uncovers any violation is said to be noncompliant.

  The data source that is evaluated by a rules test condition can be based on a repository query. A rules test condition can be implemented using a threshold condition based on the underlying metrics/queries column value or SQL expression or a PL/SQL function.

  Integration points in this area include:

  – Define Compliance Standard Rules, Compliance Standards, and Compliance Frameworks

  – Replace out-of-box policy groups (10.2.x/11.1) with Compliance Standards you create that can refer to Compliance Standard Rules

  – Map your compliance standards to the appropriate Compliance Frameworks

  – Define BI Publisher reports for compliance

- WebLogic Server Signature Rules

  Used to check a WebLogic target for supporting best practice configurations. This type of rule is not relevant for external/partner plugins.

  WebLogic Server signature rules describe potential problems based on information about WebLogic Servers and the environment in which they are deployed, including Java Virtual Machines (JVMs), operating systems, and databases. Signature rules contain executable logic that can identify specific versions of these products, as well as their configuration settings.

See Section 8.4, "WebLogic Server Signature Rules" for additional information.

- Real-time Monitoring Rules

  Real-time monitoring rules monitor actions to files, processes, and more that users perform on targets. These actions may lead to configuration changes. The actions are detected in real-time as observations. Also captures user login/logout activities.

  These rules monitor Process, OS User, Database tables, views, index, user, Windows Registry key, Active Directory Group, and so on. These rules contain configuration parameters specifying what entities they will be monitoring, for instance, what files to monitor, how to monitor (for example, operations (read/write)), when to do the monitoring (time-period), who to monitor (user name).

  See Section 8.5.3, "Operations on Real-time Monitoring Rules" for additional information.

Importance

A rule has an importance attribute that denotes the importance of the rule, which is considered when determining a compliance score.

Importance is used in compliance score rollup function for both rule/target, and compliance standard/target score. Importance is per node in the hierarchy. Weighted average of the child nodes is used to compute the score of the parent node.

The rule can also have a severity level, which could be Critical (serious issue if this rule is violated), Warning (not a serious issue if violated), or Minor Warning (a minor issue if violated). Severity impacts the compliance score.

Considerations When Creating Compliance Standards

A compliance standard will refer to one or more Compliance Standard Rules. When creating a compliance standard, the standard should be granular enough that it can be appropriately mapping to one or more related Compliance Frameworks. For example, consider this Compliance Framework structure that exists in Enterprise Manager based on PCI:

- PCI - Payment Card Industry Compliance Framework
  - PCI Requirement 10 - Regularly monitor and test networks
    - PCI 10.5 - Secure audit trails

Many compliance standards will exist that should mapped to this part of the Compliance Framework structure, each with their own rules to address this specific requirement. One may check that audit settings are set properly. Another may be used to check in real-time if anyone changes an auditing configuration. Another standard may check that regular users are not trying to read from an audit trail.

In this example, the "audit trail" referenced in the Compliance Framework can relate to many different types of targets. Oracle Database, WebLogic, Enterprise Manager, EBS, and Peoplesoft all have their own types of audit trails that all need to be secured. Any Standards created to monitor these target-specific audit trails would map to the same Compliance Framework named "PCI 10.5 - Secure Audit Trails."

If compliance standards are structured in a granular way so that they can map to existing and future compliance frameworks, then violations in a rule can be rolled up to impact the score of the compliance framework properly.

Usage Notes

Compliance standard rules are mapped to one or more compliance standards.

8.3.7 Operations on Compliance Standards Rules

The following sections explain the operations you can perform on compliance standard rules.

8.3.7.1 Creating a Compliance Standard Rule

Before you create a compliance standard rule, ensure that you have privileges to create compliance standard rules. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To create a compliance standard rule, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standard Rules** tab.
3. Click the **Create** button.
4. In the Create Rule popup, select the type of compliance standard rule you want to create:
   - Repository rule

     Checks if the target has the desired configuration state based on configuration data collected in the Management Repository.
   - WebLogic Server Signature rule

     Preemptively identifies WebLogic Server configuration problems. The purpose of the WebLogic Server Signature rules is to evaluate at the WebLogic Server if certain configuration data satisfies some conditions (or checks) and the evaluation results are sent as violation information to the Oracle Management Service.

     Detailed information about how to identify problems is specified in the WebLogic Server Signature rule definition. The WebLogic Server Signature rule definition includes Dataspec and XQuery logic that are used to determine what is important to collect and evaluate for a given target type and target properties. A Dataspec is a group of MBeans used to collect from a WebLogic Server. The XQuery logic contains the check on the collected data (by the MBeans). The WebLogic Server Signature rule can be associated with one or more specific Web Logic targets: Web Logic Domain, Web Logic Java EE Server, and Web Logic Cluster.

Version-specific details include:

- To enable data collection for the WebLogic Server signature-based rules on WebLogic Server targets earlier than v10.3.3, you need a copy of bea-guardian-agent.war. You can find a copy of this war file in your OMS installation's work directory:

  ```
  $T_WORK/middleware/wlserver_10.3/server/lib/bea-guardian-agent.war
  ```

- For WebLogic Server v9 and v10.0

  Install and deploy bea-guardian-agent.war to all servers in the domain. Do not change the context root. See `http://<host>:<port>/console-help/doc/en-us/com/bea/wlserver/core/index.html` for more information on installing a web application.

- For WebLogic Server v10.3 up to and including v10.3.2

  Copy the war file from your OMS installation into each target's $WL_HOME/server/lib directory. Restart all the servers in the target domain.

- For WebLogic Server v.10.3.3 and higher

  No action is required.

- Real-time Monitoring rule

  Monitors operating system and database level entities that store configuration data. Real-time monitoring rules define the entities to monitor, user actions to watch for, and any types of filters to apply to the monitoring. Monitoring can be filtered by: when changes occurred, who made the changes, and what process made the changes.

  The real-time monitoring rule definition includes facets that are used to determine what is important to monitor for a given target type, target properties, and entity type. A facet is a collection of patterns that make up one attribute of a target type. For example, you may choose to define a facet that lists all of the critical configuration files for the Host target type. These configuration files would be the ones that, if changed, would most likely result in instability of the host. You may also create a facet that lists all users which are DBA users.

  The real-time monitoring rule can be part of a compliance standard that is associated with one or more targets. The monitoring can occur on any operating system level entity, for example, file, process, user, registry, and so on. Real-time monitoring rules can additionally specify whether observations captured by the rule are automatically reconciled. This reconciliation determines whether the actions observed were authorized or not.

  Change Request Management reconciliation compares open change requests to actions performed on targets. If there is a match of expected actions to actual actions, then those actions are authorized, otherwise they are unauthorized. Authorizations can also be done manually. All observations are captured and bundled by rule, target and user. Attributes can be set on the frequency of observation data collection. For additional information, refer to the *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

5. Click **OK**.

Usage Notes

- Rules are visible in the global rule library.

- All rules are visible to all users.
- Users can create compliance standards based on these rules. **Note:** Rules cannot be evaluated directly. These rules are evaluated in context of a compliance standard, and their violations are viewed in context of a compliance standard they are referred from.
- One rule can be referred to by multiple compliance standards.
- The association of rules with targets can be customized per compliance standard rule, in context of the compliance standard from which the rule is included.
- Because the user-defined compliance standard rule is defined by a privileged user, only privileged users can modify the compliance standard rule. Violation results are available to all users.
- To share this user-defined compliance standard rule with other privileged users, provide the XML schema definition (using the Export feature) so they can import the compliance standard rule to their Management Repository.
- To minimize scrolling when reading the Description, Impact, and Recommendation information, restrict the text to 50 characters per line. If more than 50 characters are needed, start a new line to continue the text.
- Once the compliance standard rule is created, it is not automatically evaluated. Consider adding the rule to a compliance standard.
- Look at the context-sensitive help for information for each page in the Compliance Standard Rule wizard.
- A compliance standard rule can be added to more than one compliance standard, and can have a different importance when added to a different standard. For example, you could have a compliance standard rule called Check Password Expired which flags user accounts with expired passwords. This compliance standard rule may be a member of two compliance standards: All System Passwords Secure and 30-day Password Validation. The All System Passwords compliance standard verifies a password's security, whereas the 30-day Password Validation compliance standard checks the date that this password was last set.
  - The Check Password Expired compliance standard rule could have Extremely High importance for the 30-day Password Validation compliance standard, since this check is warning users that their passwords are about to expire.
  - In the All System Passwords Secure compliance standard, the Check Password Expired compliance standard rule could have a Normal importance, and other added compliance standard rules that do security checks could have a higher importance within the compliance standard.

8.3.7.2 Creating Like a Compliance Standard Rule

Before you create a compliance standard rule like another compliance standard rule, ensure that you have privileges to access the compliance standard rule you will be copying from. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To create a compliance standard rule like another compliance standard rule, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standard Rules** tab.

3. Highlight the rule you want to replicate.
4. Click **Create Like** button.
5. Customize the fields as needed.
6. Click **Save**.

8.3.7.3 Editing a Compliance Standard Rule

Before you edit a compliance standard rule, ensure that you have privileges to access the compliance standard rule to be edited. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To edit a compliance standard rule, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standard Rules** tab.
3. Highlight the rule you want to edit and click the **Edit** button.
4. Update the parameters as needed.
5. Click **Save**.

Usage Notes

- You can change all the rule properties except the rule name and target type. Additionally for real-time monitoring rules, you cannot change entity type or target properties.
- If you change the critical rule properties for a repository rule, for example, rule query, violation condition, parameters, or severity, then editing the rule invalidates the results for compliance standards which refer to the rule. The compliance standards compliance score will be reevaluated at the next rule evaluation.
- For rules in production mode, you have a choice to either create and save a draft of the rule or to overwrite the existing production rule. If you create a draft, you can edit the draft rule, at a later point in time, test it, and then overwrite and merge it back to the original production rule the draft was made from. **Note:** You cannot include a draft rule into any compliance standard. After you successfully test a draft rule, you can overwrite the original production rule from which the draft was created.
- For WebLogic Server Signature rule or Real-time Monitoring rule, if the rule being edited is referred to by a compliance standard which is associated with a target, then the rule definition will be deployed to the Management Agent monitoring the target, so that the Management Agent can evaluate the latest definition of the rule. In the case where the Management Agent is down or unreachable, the rule definition changes will be propagated to the Management Agent as soon as the Management Agent is available.

8.3.7.4 Deleting a Compliance Standard Rule

Before you delete a compliance standard rule, ensure that you have privileges to access the compliance standard rule to be deleted. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

Also, ensure that compliance standard rule references have been removed from compliance standards before deleting the compliance standard rule.

To delete a compliance standard rule, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standard Rules** tab.
3. Highlight the rule you want to delete, click **Delete** button.
4. Confirm that you want to delete the rule by clicking **OK**.

Note: You can only delete rules that are not referred to, or used by, any compliance standard.

8.3.7.5 Exporting a Compliance Standard Rule

Before you export a compliance standard rule, ensure the compliance standard rule to be exported is defined in the Management Repository. Also ensure that you have privileges to access the compliance standard rule to be exported. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To export a compliance standard rule, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standard Rules** tab.
3. Highlight the rule you want to export.
4. From the **Actions** menu, select **Export**.
5. Provide the file name to which the standard rule is to be exported.
6. The XML representation of the compliance standard rule is generated and placed in the directory and file you specified.

8.3.7.6 Importing a Compliance Standard Rule

Before you import a compliance standard rule, ensure the compliance standard rule to be imported is defined in a file. Also ensure that you have privileges to access the compliance standard rule definition XML file to be imported. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To import a compliance standard rule, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standard Rules** tab.
3. From **Actions** menu, select **Import**.
4. Provide the file name from which the rule definition (as per Compliance Standard Rule XSD) will be imported. Specify whether to override an existing definition if one already exists.
5. Click **OK**.

8.3.7.7 Browsing Compliance Standard Rules

Before you browse compliance standard rules, ensure you have privileges to access the compliance standard rule definitions you will be browsing. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To browse compliance standard rules, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standard Rules** tab.
3. To view the details of a particular standard rule, highlight the rule and click **Show Details**.

8.3.7.8 Searching Compliance Standard Rules

Before you search compliance standard rules, ensure you have privileges to access the compliance standard rule definitions you will be searching. (See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features".)

To search for compliance standard rules, follow these steps:

1. From the **Enterprise** menu on the Cloud Control home page, select **Compliance**, then select **Library.**
2. Click the **Compliance Standard Rules** tab.
3. In the Search portion of the page, provide criteria to use to narrow the search.

   By default, all the compliance standard rules in the compliance standard rule library appear in the results table. However, you can specify a set of search criteria and then perform a search that will display only the compliance standard rules that meet those criteria in the results table.

   For example, if you choose Security in the Category list, contains in the Compliance Standard Rule list, "port" in the adjacent Compliance Standard Rule text field, Host in the Target Type list, and then click **Go**, Enterprise Manager displays, in the results table, only the compliance standard rules for the host security category that contain "port" in their names.
4. Click **Search**.

8.3.7.9 Compliance Standard Rules Provided by Oracle

Oracle provides over 1600 compliance standard rules.

8.4 WebLogic Server Signature Rules

WebLogic Server (WLS) signature rules deliver pre-emptive support to WebLogic customers by scanning WebLogic installations for vulnerabilities and violations, based primarily on in-depth knowledge of common pitfalls and best practices. This is the Enterprise Manager compliance solution for WebLogic Server. The following sections explain WLS Signature rules in detail.

8.4.1 About WLS Signature Rules

A signature describes a potential problem in a WebLogic installation. It consists of categorization metadata, a user-readable description of the problem, and an XQuery expression for evaluating whether the problem exists at the target.

A WLS Signature rule is an agent-side rule that checks a signature definition against an associated target for the existence of the problem the signature defines. WebLogic Server targets include: WLS Domain; WLS Cluster; WebLogic Managed Server. The first two are composite target types: logical groupings of instances of simple WebLogic Server targets. Rules must be evaluated against the whole domain or cluster to render meaningful violation results.

WLS Signature rules, like other compliance rules, are grouped into Compliance Standards, which are logical groupings based on signature metadata such as severity and remedy.

The general workflow is as follows:

- Rule creation takes place in the wizard, where the rule is grouped into its logical Compliance Standard and is associated with a target.
- The rule, in the context of its Compliance Standard, gets deployed to the agent monitoring the associated target.
- The standard/rule combination gets evaluated agent-side against a metric collected specifically for the Compliance Standard and target type to determine compliance.
- The evaluation generates violations (if any).
- Violations are uploaded to OMS, from where they are subsequently uploaded into violations repository tables.
- Violations are then viewable in compliance results pages and the dashboard.

8.4.2 Creating a WLS Signature Rule

There are several hundred out-of-box WLS signature rules designed to uncover compliance violations known to occur in WebLogic installations. You can supplement this rules repertoire to expose additional, lesser-known violations by creating your own WLS signature rules.

Use the WebLogic Signature Rules wizard to create custom WLS signature rules. To access the wizard:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Select **Compliance Standard Rules**.
3. Click **Create** and select **WebLogic Server Signature Rule**. Click **Continue**.
4. The wizard opens. Complete the wizard process as described in the Cloud Control online help.

Alternatively, you can select a WLS signature rule on the **Compliance Standard Rules** tab and click the **Create Like** button to base rule creation on an existing rule. After you name the new rule, the wizard opens, displaying the contents of the existing rule. Edit the rule as described in the Cloud Control online help.

8.4.3 WLS Signature Rule Example

Instructions you provide to the wizard shape the makeup of a WLS signature rule, but a rule's nerve center is the code you provide as the signature definition. A signature definition consists of a list of managed beans (MBeans) and an XQuery expression. Managed beans represent the configuration data to collect. They define a type and the attributes within the type to collect. They also declare which attributes to consider in determining whether there are violations. The XQuery expression defines the logic to use in evaluating the collected data for compliance. An XML example signature definition follows.

```
<SignatureDefinition>
     <MBeanList>
         <MBean scoreBase="true" mBeanType="ServerRuntime">
             <AttributeName>Name</AttributeName>
             <AttributeName>WeblogicVersion</AttributeName>
```

```
        </MBean>
     </MBeanList>
     <XQueryLogic>declare function
local:getServerRuntimesEqualToVersionWithPatch($targetData, $major as xs:integer,
$minor as xs:integer, $servicePack as xs:integer, $crNumber as xs:string) {
 for $ServerRuntime in $targetData/DataCollection/ServerRuntime
 let $weblogicVersion := fn:replace($ServerRuntime/@WeblogicVersion,
&amp;quot;WebLogic Server Temporary Patch&amp;quot;, &amp;quot;&amp;quot;)
 let $majorVersion :=
     let $spaceParts :=  fn:tokenize(fn:substring-after($weblogicVersion,
&amp;quot;WebLogic Server &amp;quot;), &amp;quot; &amp;quot;)
     let $majorVersionParts := fn:tokenize($spaceParts[1], &amp;quot;\.&amp;quot;)
     return
       $majorVersionParts[1] cast as xs:integer
 let $SP_MP :=
         if ($majorVersion = 8) then
             &amp;quot;SP&amp;quot;
         else
             if ($majorVersion &amp;gt;= 9) then
                 &amp;quot;MP&amp;quot;
             else &amp;quot; &amp;quot;
 let $minorVersion :=
     let $spaceParts := fn:tokenize(fn:substring-after($weblogicVersion,
&amp;quot;WebLogic Server &amp;quot;), &amp;quot; &amp;quot;)
     let $minorVersionParts := fn:tokenize($spaceParts[1], &amp;quot;\.&amp;quot;)
     return
       $minorVersionParts[2] cast as xs:integer
 let $servicePackVersion :=
     let $spaceParts := fn:tokenize(fn:substring-after($weblogicVersion,
&amp;quot;WebLogic Server &amp;quot;), &amp;quot; &amp;quot;)
     let $servicePackParts :=  fn:substring-after($spaceParts[2], $SP_MP)
     return
       if ($servicePackParts = &amp;quot;&amp;quot;) then
        0
       else
        $servicePackParts cast as xs:integer
 where $majorVersion = $major and $minorVersion = $minor and $servicePackVersion =
$servicePack and

fn:contains(fn:upper-case($ServerRuntime/@WeblogicVersion),fn:upper-case($crNumber
))
 return
  $ServerRuntime
};
for $server in
local:getServerRuntimesEqualToVersionWithPatch(/,10,0,1,&amp;quot;CR366527&amp;quot;) |
local:getServerRuntimesEqualToVersionWithPatch(/,10,0,0,&amp;quot;CR366527&amp;quot;)
return &amp;lt;Server Name=&amp;quot;{fn:data($server/@Name)}&amp;quot;/&amp;gt;</XQueryLogic>
 </SignatureDefinition>
```

Effectively, this definition collects the server name and WebLogic version of all runtime servers. Much of the definition iterates over the preciseness of the version—major and minor patch, service pack, CR number, and so forth. A violation occurs if any server has either of the stated patches (10.0.1 CR366527 or 10.0.0 CR 366527), in which case return the name of the server to be reported in violation. Hence, the rule definition must include a column to account for display of the server name. The version is irrelevant in the context of the display. Those alerted are interested only in which servers are in violation.

8.5 Real-time Monitoring Facets

The real-time monitoring rule definition includes facets that are used to determine what is important to monitor for a given target type, target properties, and entity type. A facet is a collection of patterns that make up one attribute of a target type.

The following sections explain real-time monitoring facets in detail:

- Section 8.5.1, "About Real-time Monitoring Facets"
- Section 8.5.2, "Operations on Facets"
- Section 8.5.3, "Operations on Real-time Monitoring Rules"

8.5.1 About Real-time Monitoring Facets

A facet makes up a target type. A target type has several facets to it. A target type will have a facet of which files are critical configuration files, which files are log files, which files are executables, which database tables have sensitive configuration data, and so on. The sum of all of these facets for a given target type makes up everything that is important to monitor for the given target type in terms of compliance.

For a given target type, you can create any number of facets. A facet is not only for a specific target type, but for a specific target type plus a combination of some number of target type properties. For instance, creating a facet for a Host Target Type on Windows is different than creating a facet for a Host Target type on Linux. A facet can have several target type properties. If no target type criteria are set, it is assumed that a facet applies to all criteria (any target of this type).

Real-time Monitoring facets based on target types are used to specify the entities to monitor. Facets are reusable. They can be created on their own, or created inline with a Real-time Monitoring rule. No matter how they are created, they can be used again at a later time in any number of rules.

Facets are used for monitoring and also for filtering.

As an example, if monitoring a host for file changes, a facet can be a list of distinct single files, patterns with wildcards that would include many files, or simply an entire directory. These patterns can also include parameters that have a default, but can be overridden as needed for each target. There are also built-in parameters, such as ORACLE_HOME that will be dynamically filled in for each target. For instance, if you wanted to specify monitoring the database configuration file *tnsnames.ora*, your pattern may be *{ORACLE_HOME}/network/admin/tnsnames.ora*.

When performing continuous real-time monitoring, it is important to scope your monitoring only to critical entities. Monitoring more activity than is important to the organization will result in higher CPU loads on the agent as well as a very large amount of data to be processed/stored by the Oracle Enterprise Manager servers.

8.5.1.1 Facet Entity Types

Each facet has an *entity type* which defines what kind of entities the facet describes. For example, for OS level monitoring, there is File, Process, User, Windows Registry, and several Active Directory elements. For database monitoring, the entity types include Table, View, Index, Procedure among others. The possible entity types are fixed by the continuous real-time configuration change monitoring capabilities available from the Management Agent.

Creation of facets is possible through the Facet Library screen. In this screen, you can add/edit patterns for facets, and see which facets are being consumed by rules.

When you specify a real-time monitoring rule, you must first decide what entity type on a host will be monitored. You can use Enterprise Manager to monitor the following entity types with Real-time Monitoring Rules:

Table 8–2 Monitored Entity Types

Entity Types		
OS File	Oracle Database Table	Oracle Database Package
OS Process	Oracle Database View	Oracle Database Library
OS User	Oracle Database Procedure	Oracle Database Trigger
Microsoft Windows Registry	Oracle Database User	Oracle Database Tablespace
Microsoft Active Directory User	Oracle Database Index	Oracle Database Materialized View
Microsoft Active Directory Computer	Oracle Database Sequence	Oracle Database Cluster
Microsoft Active Directory Group	Oracle Database Function	Oracle Database Link
Oracle Database Dimension	Oracle Database Profile	Oracle Database Public DB Link
Oracle Database Synonym	Oracle Database Public Synonym	Oracle Database Segment
Oracle Database Type	Oracle Database Role	Oracle Database SQL Query Statement

8.5.1.2 Facet Patterns

A facet contains one or more patterns. These patterns can express inclusion or exclusion filters. For instance, you may define a facet for critical configuration files that looks like the following:

Include c:\myapp1\config

Exclude c:\myapp1\config\dummy.cfg

In this case, everything under c:\myapp1\config will be considered to be a member of this facet except for the individual file c:\myapp1\config\dummy.cfg. In general there are some rules to how patterns work given the most common use cases listed below. Each entity type might have special cases or special formats of patterns.

- Patterns of the same specificity with one being include and one being exclude, the include will win.
- Patterns that are more specific override (like in the above example, exclude dummy.cfg overrides the inherited include c:\dummy.cfg from the first pattern.)
- If there are no patterns at all, exclude * is assumed (for example, no entities in the facet)

For each pattern that you add to a facet, an optional description field is available to let you document their patterns.

8.5.2 Operations on Facets

The following sections explain the operations you can perform on facets:

- Section 8.5.2.1, "Creating and Editing Facets"

- Section 8.5.2.2, "Deleting a Facet"
- Section 8.5.2.3, "Using Create Like to Create a New Facet"
- Section 8.5.2.4, "Importing and Exporting Facets"
- Section 8.5.2.5, "Changing Base Facet Attributes Not Yet Used In a Rule"
- Section 8.5.2.6, "Viewing the Facet Library"

8.5.2.1 Creating and Editing Facets

When a you create a facet and subsequently use a facet in a Real-time Monitoring Compliance Standard Rule, the compliance rule only references the facet. If the content changes, then the rule will use the new content automatically. Because a rule references a facet, the facet can change and the rule always uses the current facet definition.

The content of the facet being used becomes important when the compliance standard is assigned to a target.

Each facet is assigned a description that allows you to document the facet. Each pattern also has an optional description field.

Ensure you have the privileges to create, delete, and modify facets as these configurations relate to the compliance monitoring. See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features" for information.

To create or edit a facet, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**. Choose the **Real-time Monitoring Facets Library** tab.

   Enterprise Manager displays the Facet Library page that lists all exiting facets along with their target type, entity type, and other details about the facet. From this page you can perform administrative tasks such as create, create like, view, delete, import and export.

2. Click **Create** to create a new facet.

   The Create Facet page displays.

3. Enter the name you want to assign to the facet in the **Facet Name** field, then choose the target type for the facet you are creating from the drop-down list in the **Target Type** field. Once you choose the Target Type, you can enter values in the Target Property Filter fields.

   The target properties you add here limit which targets to which this facet can ultimately be assigned. For instance, you could define a facet to work only for Linux version 5 on 64-bit servers.

4. Choose the **Entity Type** from the drop-down.

5. Enter a description for the facet in the **Description** field.

6. The Create Facet page contains three tabs you can use to enter more information and parameters for the facet you create. Use the Patterns tab to add patterns to be either **Included** or **Excluded** when this facet is used by a Real-time Monitoring Compliance Standard Rule. Use the **Add** or **Delete** buttons to add additional patterns or to remove a selected pattern from the facet definition.

7. If you are defining a facet for the OS File entity type, there is an optional ability to browse a host to find the files you want to monitor. The right side of the page has an area where you can choose the host to use as the basis for looking for files. In the pattern area, you can click the Browse button to interactively browse the files

on the selected host and select the files to include in the pattern. After selecting patterns from a host, you can continue to manually add more or edit existing ones.

8. Use the Parameters tab to view parameters that are part of the new facet. Oracle provides a set of predefined parameters based on target parameters (such as ORACLE_HOME) that are defined out of the box. These parameters do not require a default value and are always set according to the target's value. Parameters will appear under this tab when they are used in a pattern. To start using a new parameter, simply add the parameter to the pattern by enclosing it in curly brackets {}. For instance, a pattern of {INSTALL_DIR}\config\main.conf would result in a parameter of INSTALL_DIR being listed under this tab. All parameters must have a default value that will be automatically used for all targets against which this facet is used. This value can be overridden when associating a compliance standard containing a real-time monitoring rule to one or more targets. The Parameters tab displays the Parameter Name, Default Value, Used in Pattern, and Description. Used in Pattern indicates that the parameter is currently in use. This parameter may have been defined at some point in a pattern and then removed. The pattern will still be available for use again at a later time even if the pattern is not currently in use. If the entity for which you are adding a pattern includes a "{" or "}", you can escape these characters by using "{{}" and "{}}" in the pattern respectively. These will not be counted as parameters.
9. The Operation Time Window tab is only available if the facet being created/edited is of entity type Operations Time Window. A facet of this entity type is only usable as a filter in a Real-time monitoring rule. For instance, you can specify in the rule that you only want to monitor an item during a specific time, for example, "Production Hours". In the Duration section, choose either a **24 Hour Interval** or **Limit Hours to**, which allows you to enter a Start time and an Interval in Hours and Minutes. In the Repeating section, you can choose either **All the time** or you can select **Repeat** and then choose which days of the week to repeat the operation.
10. Choose **OK** to create the facet.

8.5.2.2 Deleting a Facet

Deleting a facet is not possible as long as the facet is in use either as a monitoring facet in a rule or as a filter facet in a rule. If this facet is not in use in any rules, then the facet can be deleted. If a facet is in use, the user is alerted to the current use and not allowed to delete the facet until the rules using it are modified to no longer include it.

When deleting a facet, the data will no longer be referenced to the facet and instead it will show "(Deleted Facet)" as the name of the facet to which it is related. The data will only be available through the Search Observations page, not the Browse pages.

Ensure you have the privilege to delete a facet. See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features" for information.

For compliance users, customers would want to keep the unused facet available so the compliance data is not lost. You can also remove the patterns as long as you keep the actual facet to maintain collected observations. Then only after the compliance data related to this old facet is no longer available, you can delete the facet without any data loss.

To delete a facet, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**. Choose the **Real-time Monitoring Facets Library** tab.

Enterprise Manager displays the Facet Library page that lists all exiting facets along with their target type, entity type, and other details about the facet. From this page you can perform administrative tasks such as create, create like, view, delete, import and export.

2. Select the facet from the list of facets in the table on the page.
3. Click **Delete** to delete the facet. You will be prompted to confirm that you want to delete the facet.

8.5.2.3 Using Create Like to Create a New Facet

Facets that ship with the product or with a plug-in cannot be changed. If you want to enhance or modify the out-of-box content, you must use the create-like functionality to make your own copy of the facet which can then subsequently be edited.

Ensure you have the privilege to create a rule and also create and edit a facet. See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features" for information

An important limitation to the Create Like function is that you cannot change the target type or entity type. The patterns contained in the facet may be dependent on target type or entity type. Should you want to use Create Like and change these attributes, you should use Export to export the original facet, edit the name, target type, entity type in the XML, and then import the new facet.

To use create like to create a new facet, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**. Choose the **Real-time Monitoring Facets Library** tab.

   Enterprise Manager displays the Facet Library page that lists all exiting facets along with their target type, entity type, and other details about the facet. From this page you can perform administrative tasks such as create, create like, view, delete, import and export.
2. Choose the facet from the facet table that you want to use as the basis for the new facet you want to create.
3. Click **Create Like**.

   Enterprise Manager displays the Create Facet page. All the values that were applicable to the facet you want to clone are entered. Use the page to edit the values for the new facet and click **OK**.

   It is important to understand that if the original base facet you used in the create like activity is changed, that change will not be reflected in the newly created facet. There is no relationship maintained when using Create Like.
4. For more information about using the Create Facet page, see Section 8.5.2.1, "Creating and Editing Facets".

8.5.2.4 Importing and Exporting Facets

You can select facets and export or import them. All selected facets will be exported into one output file.

On import, if a facet of the same name/target type/entity type combination already exists, the import fails with an error that the facet already exists. The user must change the import file to remove the duplicate name and retry the import.

The combination of name, target type, and entity type define a unique facet. You can have the same name facet across different target types and entity types.

To export a facet, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Facet Library**.

   Enterprise Manager displays the Facet Library page that lists all exiting facets along with their target type, entity type, and other details about the facet. From this page you can perform administrative tasks such as create, create like, view, delete, import and export.

2. Select one or more facets from the list of facets on the Facet Library page that you want to export and then click **Export**.

3. On the Open dialog box, you can choose to open or save the facet xml file using an XML editor of your choice and then either edit or save the file to another location.

To import a facet, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**. Choose the **Real-time Monitoring Facets Library** tab.

   Enterprise Manager displays the Facet Library page that lists all available facets along with their target type, entity type, and other details about the facet. From this page you can perform administrative tasks such as create, create like, view, delete, import and export.

2. Click **Import** and choose the facet XML file you want to import into the Facet Library.

3. Enterprise Manager imports all facets specified in the imported XML file. You can then edit the facet or use any other action on it as you would any other facet in the library.

8.5.2.5 Changing Base Facet Attributes Not Yet Used In a Rule

After a facet is in use in at least one rule (either as a monitoring facet or as a filter facet), you cannot change the facet name, target type, entity type, or target criteria of the facet since the rules that have been created are already bound to these attributes. The only attributes that can be changed are the facet patterns, parameters and description fields. Although the rule is not dependent on the facet name, users have used them in their rules based on the name of the facet. Allowing the name of the facet to change after consumption will only lead to confusion of the rule authors.

If a facet is not currently in use but has been in use in the past, then it is treated the same as an in-use facet.

You cannot make changes to the out-of-box facets that ship with the Enterprise Manager product. If you want to use an out-of-box facet with changes, you can perform a “Create Like” operation and then modify it as needed.

To change base facet attributes, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**. Choose the **Real-time Monitoring Facets Library** tab.

   Enterprise Manager displays the Facet Library page that lists all exiting facets along with their target type, entity type, and other details about the facet. From

this page you can perform administrative tasks such as create, create like, view, delete, import and export.

2. Choose the facet from which you want to create a new facet with modified attributes. Click **Create Like**.
3. Enter a new Facet Name and change whatever attributes to create a new facet based on the previous facet.

8.5.2.6 Viewing the Facet Library

Any user that can view observation data is able to also view the facet library and see the facet history for any facet.

To view the facet library, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**. Choose the **Real-time Monitoring Facets Library** tab.

   Enterprise Manager displays the Facet Library page that lists all exiting facets along with their target type, entity type, and other details about the facet. From this page you can perform administrative tasks such as create, create like, view, delete, import and export if you have the audit author role.
2. The Facet Library page displays the Facet Name, Author, Target Type, Entity Type, Rules Using the facet, Description, and the Last Updated time of the facet. You can see the details of any facet by selecting it from the table and clicking **Show Details**.
3. You can choose which columns to display in the table by clicking **View** and then choosing **Columns**. You can either choose to **Show All** columns or you can select individually the columns you want to appear in the table. You can reorder the columns by clicking **Reorder** after you click View and then changing the order in which the columns appear by moving them up or down using the arrow keys.
4. You can view a history of a selected facet by choosing it from the table and then clicking **History**. The View History page displays.

8.5.3 Operations on Real-time Monitoring Rules

The following sections explain the operations you can perform on real-time monitoring rules:

- Section 8.5.3.1, "Creating a Real-time Monitoring Rule"
- Section 8.5.3.2, "Editing a Rule"
- Section 8.5.3.3, "Viewing Real-time Monitoring Rules"
- Section 8.5.3.4, "Deleting a Rule"
- Section 8.5.3.5, "Saving a Development Copy of a Rule Prior to Production"
- Section 8.5.3.6, "Importing and Exporting Rules"
- Section 8.5.3.7, "Setting Severity Levels for Rules"
- Section 8.5.3.8, "Setting Target Criteria For a Rule"
- Section 8.5.3.9, "Selecting the Types of Actions You Want to Monitor"
- Section 8.5.3.10, "Using Facets As Filters In Rules"
- Section 8.5.3.11, "Controlling Observation Bundle Lifetimes"
- Section 8.5.3.12, "Creating a Facet Inline While Creating a Rule"

8.5.3.1 Creating a Real-time Monitoring Rule

Use the following steps to create a Real-time monitoring rule:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.
3. At the top of the Rules table, choose **Create** from the **Actions** menu.
4. From the Create Rule panel that displays, choose **Real-time Monitoring Rule** from the list of selections then click **OK**.

   Cloud Control displays the first page of the Create Rule Real-time Monitoring wizard.
5. Use the Details page to set basic properties for the Real-time monitoring rule.
6. Use the Entitites to Monitor page to choose existing facets or create new facets. Facets available for monitoring in this rule are based on the entity type and target type chosen in the previous screen. Facets contain the patterns of entities that will be monitored.
7. Use the Actions to Monitor page to choose one or more possible user actions to monitor. When associating a Compliance Standard using this rule to target(s), various audit settings on the target may need to be enabled. Refer to the documentation for any operating system specific audit settings required.
8. Use the Filters page to add filters if you want to only perform monitoring under specific conditions. Based on the entity type you chose for the rule, there maybe various filters that can be applied. You can use existing facets as filters or create new facets.
9. Use the Settings page to configure Change Request Management reconciliation as well as advanced settings related to how observation bundles are created.
10. Use the Review page to review all settings for the rule before promoting it to production or saving it as a draft.

8.5.3.2 Editing a Rule

For any rule that is in use in a compliance standard, the rule name, target type, target properties, and entity type cannot be changed. You must first either do a Create Like on the rule and replace the compliance standard association with this new rule, or unassociate the rule from all compliance standards before modifying the rule.

For out of the box rules that are included in the Enterprise Manager product, you can change the following pieces of the rule.

- Facets patterns and parameter defaults – but cannot change the facets monitored
- Actions to Monitor
- Filter
- Settings

All other parts of the rule cannot be changed. To use this facet with modifications, you can use Create Like and then modify it.

To edit a rule, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.

2. Click on the **Compliance Standard Rule** tab.

   Enterprise Manager displays the Compliance Standard Rule Library page.

3. Select the rule you want to edit from the Rules table and click **Edit**.

Usage Notes

You can change all the rule properties except the rule name and entity type.

If you change the critical rule properties for a repository real-time monitoring rule, for example, rule query, violation condition, parameters, or severity, then editing the rule invalidates the results for compliance standards which refer to the rule. The compliance standards compliance score will be reevaluated at the next rule evaluation.

For rules in production mode, you have a choice to either create and save a draft of the rule or to overwrite the existing production rule. If you create a draft, you can edit the draft rule, at a later point in time, test it, and then overwrite and merge it back to the original production rule the draft was made from. Note: You cannot include a draft rule into any compliance standard. After you successfully test a draft rule, you can overwrite the original production rule from which the draft was created.

For a WebLogic Server Signature rule or Real-time monitoring rule, if the rule being edited is referred to by a compliance standard which is associated with a target, then the rule definition will be deployed to the Management Agent monitoring the target, so that the Management Agent can evaluate the latest definition of the rule. In the case where the Management Agent is down or unreachable, the rule definition changes will be propagated to the Management Agent as soon as the Management Agent is available.

8.5.3.3 Viewing Real-time Monitoring Rules

Use the following steps to view an existing Real-time monitoring rule:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.
3. In the Rules table, choose the real-time monitoring rule you want to view.
4. From the Action menu, choose **Show Details**.

   Cloud Control displays the Compliance Rule Details page that shows you all the detailed information about the rule.

8.5.3.4 Deleting a Rule

A compliance standard rule can only be deleted if it is not actively part of any compliance standard. When deleting a rule, any facets that were associated with the rule are unassociated. The facet itself is not changed. The facet remains and any other rules using the facets are not affected.

Any data that references this rule will remains. When you view it from the Observation Search, it simply displays "(Deleted Rule)". You cannot access the data through the two Browse-by screens

For compliance users, a very common use case would be that customers would want to keep the unused rule around so the compliance data is not lost. The user can simply remove the reference to the rule from compliance standards to avoid future evaluations of the rule. Then only after the compliance data related to this old rule is no longer available, they can delete the rule without any data loss.

To delete a rule, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.

   Enterprise Manager displays the Compliance Standard Rule Library page.
3. Select the rule you want to delete from the Rules table and click **Delete**.

Usage Notes

You can only delete rules that are not referred to, or used by, any compliance standard.

8.5.3.5 Saving a Development Copy of a Rule Prior to Production

When first creating a rule, the rule may be in development mode if you select this as a production rule in the first page of the rule creation wizard. When you are finished with the rule creation process, on the last screen of the wizard you can decide to promote the rule to production. This means the rule can then be used in a compliance standard.

Any time during the rule creation step, you can save the rule. Some key fields for the rule, for example the rule name, are required.

This makes the Configuration Change CS Rule consistent with the other two types of Compliance Standard rules.

To save a development copy of a rule prior to production, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.

   Enterprise Manager displays the Compliance Standard Rule Library page.
3. Click **Create**.

   The Create Rule dialog box opens where you can select the type of Compliance Standard Rule you want to create.
4. Select **Real-time Monitoring Rule** and click **OK**.

   The Create Rule: Real-time Monitoring wizard opens to the Details page. Fill in all appropriate information for the rule you want to create.
5. Proceed through the rule creation wizard. On the Review page, click **Finish** while leaving the Lifecycle State as "Development".

8.5.3.6 Importing and Exporting Rules

You can select rules to export or import.

To import rules, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.

   Enterprise Manager displays the Compliance Standard Rule Library page.
3. From the Actions menu, click **Import**.

   Enterprise Manager displays the Import Rule dialog box that allows you to select a file for import.

4. Enter the location and name of the file you want to import, or use the Browse function to navigate to the file on your system or network. Click **OK** to import the rule.
5. Optionally you can use the **Overwrite Option** to overwrite a file with the same name upon import.

To export a rule, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.

   Enterprise Manager displays the Compliance Standard Rule Library page.
3. Select the rule from the Rule table that you want to export.
4. From the Actions menu, click **Export**.

   You can choose to open the file using the default or a chosen XML editor, or you can save the file to a location on your system or network.

8.5.3.7 Setting Severity Levels for Rules

Like all compliance standard rules, you can set a severity level for the rule so that any violations can be weighted for compliance score calculation.

To set the severity level for a rule, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.
3. If you are setting the severity level for an existing rule, choose the Real-time Monitoring Rule for which you want to set a severity level and click **Edit**. Enterprise Manager opens the Edit Rule: Real Time Monitoring wizard and displays the Details page of the wizard. In the **Severity** field, choose the severity level from the drop-down list. You can select Critical, Minor Warning, or Warning.

   If you are creating a rule, you can similarly set the Severity Level on the Details page of the Create Rule: Real-time Monitoring wizard.

8.5.3.8 Setting Target Criteria For a Rule

When creating a rule, you must choose a target type for the rule. Since the Real-time monitoring capabilities on the agent have some dependencies on operating system and versions of operating systems, when creating a rule for real-time configuration change monitoring, you must be allowed to set the criteria for a rule. The target may be different on a target type, so patterns in the facets may be different. For instance, Oracle Database on Microsoft Windows is not the same as it is on the UNIX operating system.

Why would I set target criteria?

If target criteria is not set, all rule options are available then at target-cs association time, if a target's settings do not match, then that rule/facet is ignored. If you only set, for example, the platform name, but not version, then only the options that are common across all versions of the platform are available.

The list of facets that are selectable when creating a rule are filtered by the target criteria that a facet is created to support. For instance if you have a facet, FACET1, that works on Linux or HPUX and you create a rule for Windows, this facet for Linux and

HPUX will not be available to select for your rule. This applies both when selecting the monitoring facet or using a facet as a filter. However if you create a rule for either Linux or HPUX, FACET1 will be available because the criteria for the rule at least overlapped with that of the facet.

To set the target criteria for a rule, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.
3. Click **Create** to create a new rule.

   Enterprise Manager opens the Create Rule: Real-time Monitoring wizard and displays the Details page.
4. After defining the Rule name, Lifecycle State, Severity level and the target type in the Applicable to field, if necessary, expand the Target Property Filter section.

   Use this section to define the criteria of which targets could be associated If the filter is not defined, all targets will be evaluated against this rule. However, without a Compliance Standard and Compliance Standard associated to a target, no targets will be evaluated. This just limits which targets could be associated, but it does not simply enable all monitoring at this time.
5. Enter the **Lifecycle State**, **Version**, and **Platform** information for the filter.

8.5.3.9 Selecting the Types of Actions You Want to Monitor

When creating a rule, you can decide which types of observations or user actions are important to be monitored and reported back to Enterprise Manager. The Management Agent has a specific set of observations that are possible for each entity type. Some options may be specific to certain operating system platforms or versions. You can select one or more of these options.

The observation types that you may be able to select can also be limited by the target properties/criteria selected for the rule. For instance, some operating systems may not have every monitoring capability for files. When building the list of available observation types available, the target type, entity type, and target properties are all taken into consideration to come up with the resulting available observation types.

To select the type of observations you want to monitor in a rule, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.
3. If you want to select observations for a currently existing rule, click on the Real-time Monitoring rule in the Rules table and then click **Edit**.

   Enterprise Manager opens the Edit Rule: Real-time Monitoring wizard and displays the Details page. Move to the Observations page.

   If you want to select observations while creating a new rule, click **Create** to create a new rule. Enterprise Manager opens the Create Rule: Real-time Monitoring wizard and displays the Details page. After entering relevant information on the Details and Facets pages of the wizard, move to the Observations page.
4. On the Observations page, select one or more activities to be observed from the list that appears. During target association for this rule, auditing must be enabled to capture selected details. It is important to note that different operating systems and different capabilities have specific auditing requirements.

5. In the Parameters section, if there are additional observation parameters, you can review and update the parameters.

8.5.3.10 Using Facets As Filters In Rules

When creating a rule, facets can be used in two ways. The first is to use the facet to specify what entities to monitor in the rule. The second is to use the facet as a filter to apply on top of activities detected by the agent.

You can use the same facet as a monitoring facet in one rule and a filtering facet in another rule. The benefit is once you define a collection of patterns, for example to define your administrative users, you can use that collection in many ways without having to redefine the collection again.

Filters in rules are set up to reduce the observations that are captured and reported to the OMS. If there are no filters defined, then all observations related to the monitoring facet(s) selected in the rule are captured. When selecting a facet as a filter, the default is to only include observations that have attributes that match. The following example IT compliance control demonstrates an example for the filtering:

IT Control: Monitor all changes to critical OS configuration files by administrators during production hours.

To implement this IT control, you can create a compliance standard rule with the following:

1. Create a rule and select the file facet "Critical OS configuration files" for the monitoring facet that has patterns covering all critical OS configuration files.
2. Select "content change" as the observation types to capture
3. Add an OS Users filter selecting facet "Administrators" that lists patterns describing all of the OS user accounts that are considered administrators.
4. Add a Time Window filter selecting facet "Production Hours" that lists patterns describing the times of the week that are considered to be production hours. For example, Every day 4am-2pm PST.

When the agent sees any content change to the patterns in Critical OS configuration files, it will only report these changes back to the OMS if the change happened during production hours and if any user described in the Administrator's facet is the one making the change. Filters can also be inverted to monitor anyone not in the administrators group or for changes outside of production hours.

To use facets as filters in rules, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.
3. If you want to use facets as filters in a currently existing rule, click on the Real-time Monitoring rule in the Rules table and then click **Edit**.

   Enterprise Manager opens the Edit Rule: Real-time Monitoring wizard and displays the Details page. Move to the Observation Filters page.

   If you want to use facets as filters while creating a new rule, click **Create** to create a new rule. Enterprise Manager opens the Create Rule: Real-time Monitoring wizard and displays the Details page. After entering relevant information on the Details, Facets, and Observations pages of the wizard, move to the Observation Filters page.
4. To add a facet as a filter, click **Add**.

Enterprise Manager opens the Add Observation Filter dialog box. From here you can enter a Filter Type for the target type you are using, and then can choose to either **Apply the filter as specified** or **Invert the filter to monitor everything except what is specified in the filter.**

At the bottom of the filter step in the wizard is a readable description of what will be monitored to specify what affect adding filters will have on monitoring.

From the list of available facets in the table, select (or multi-select) the facets you want to use as filters. Click **OK**.

5. You can delete any of the facets from the rule by highlighting them in the Facets table and clicking **Remove from Rule**.
6. You can create an inline facet to use as filter for a rule by clicking **Create**. Enterprise Manager displays the Create (Inline Rule) Filter page where you can create the facet.

8.5.3.11 Controlling Observation Bundle Lifetimes

Observation bundles are logical groupings of observations that occur over a relatively short period of time against the same rule on the same target and by the same user. The last three factors cannot be configured by the user because they will be how the agent groups observations before sending them back to the Enterprise Manager server.

The user creating the rule however does have three variables that they need to be able to configure:

1. Idle timeout: The amount of time after the user has no more activity from their last activity on an entity in a given rule on a given target. The use case for this is that a user logs into a server, starts making a few file changes and then no more file changes are made after 15 minutes. This 15 minute waiting period is the idle timeout. After this idle timeout period is reached, the current observation bundle is closed and sent to the Enterprise Manager server. The next time a new observation is detected, a new group will be started and the process starts over.
2. Maximum lifespan of a group: If a user were to set the idle timeout to 15 minutes and a user on a host was making one file change every 10 minutes for an indefinite period of time (say through a script or even manual), the observation bundle will never close and therefore never get sent to the Enterprise Manager server for reporting/processing. Setting the maximum lifespan of a group tells the agent to only allow a group to accumulate for a maximum specific time. For example, this maximum lifespan may be 30 minutes or an hour.
3. Maximum number of observations in a group: If a rule is being triggered because of an activity that is causing a lot of observations to be detected, it may be desirable for the user to not group every observation together if there are too many. Groups have a management lifecycle to them where observations can be set to authorized/unauthorized, and so on. Having observation bundles with tens of thousands of observations in it could become hard to manage.

These three fields must have a value.

The user creating a rule cannot choose to turn off grouping, but if they desired to reduce delays in observation reporting to Enterprise Manager server, they could set the idle timeout and maximum lifespan of a group to be lower.

The event/incident subsystem will track only the observation bundles, not each individual observation.

Observation bundles are built at the agent and will only be sent to the Enterprise Manager server when the group is complete. In most compliance use cases, this is

acceptable because you will not need to view the results immediately. Capturing and grouping results together is more important.

When an observation becomes part of two or more groups on the agent because the same facet is used in multiple rules or multiple targets on the same host monitor the same facet with shared entities, then whenever the first group either hits its ending criteria (idle timeout, group maximum life, or maximum group entries), then all of the groups containing these shared observations are closed at the same time.

To control observation bundle lifetimes, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.
3. If you want to control observation bundle lifetimes in a currently existing rule, click on the Real-time Monitoring rule in the Rules table and then click **Edit**.

   Enterprise Manager opens the Edit Rule: Real-time Monitoring wizard and displays the Details page. Move to the Settings page. You must expand the section to view these settings.

   If you want to control observation bundle lifetimes while creating a new rule, click **Create** to create a new rule. Enterprise Manager opens the Create Rule: Real-time Monitoring wizard and displays the Details page. After entering relevant information on the Details, Facets, Observations, and Observation filters pages of the wizard, move to the Settings page.
4. Under the Collection Settings section, enter the values as discussed above into the appropriate fields.

8.5.3.12 Creating a Facet Inline While Creating a Rule

You can create a fact inline while creating a rule. This flow is useful when the person creating a rule also knows the facets that need to be created. In some cases, the person that creates the rule may be a different person than whomever must populate the pattern content for the facet. This is because the knowledge of a specific facet definition is owned by a team other than the compliance team that creates rules. In this flow, one person can create a rule and specify none or some of the facet patterns and allow someone else to come later to populate the facet content.

When you create a new facet inline with a rule creation, this facet becomes part of the global facet library and can be used by other users in other rules as well. The facet can be created/edited in line with selecting a facet for monitoring or when selecting facets to be used as filters.

When you edit a facet either that was created inline or one that was already in the global library, the changes made are automatically applied to all consumers of the facet automatically.

Follow the steps below to create a facet inline while creating a rule:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.

   Enterprise Manager displays the Compliance Standard Rule Library page.
3. Click **Create**.

   The Create Rule dialog box opens where you can select the type of Compliance Standard Rule you want to create.

4. Select **Real-time Monitoring Rule** and click **OK**.

   The Create Rule: Real-time Monitoring wizard opens to the Details page. Fill in all appropriate information for the rule you want to create.

5. Move to the Facets page of the wizard.

   The Facets table on this page lists all facets associated with the selected or created rule. You can choose to associate existing facets or you can create a new facet inline.

6. Click **Create** to create a new facet inline.

   Enterprise Manager displays the Create (Inline Rule) Facet page where you can define a new facet to be used in your definition of the rule you are creating. For more information about creating a facet, see Section 8.5.2.1, "Creating and Editing Facets".

8.6 Real-Time Observations

Observations are the actions that were seen on a host or target that were configured to be monitored through real-time monitoring rules. Each distinct user action results in one observation. Observations are additionally bundled if there are multiple observations done in a short period of time by the same user on the same target and against the same real-time monitoring rule.

8.6.1 Viewing Observations

The following sections explain the various methods by which to view observations.

- Section 8.6.1.1, "Viewing Observations By Systems"
- Section 8.6.1.2, "Viewing Observations By Compliance Framework"
- Section 8.6.1.3, "Viewing Details of an Incident"
- Section 8.6.1.4, "Viewing a Summary of Observation Bundle Reconciliation Results From Incident Pages"
- Section 8.6.1.5, "Viewing Observations With Compliance Violations From the Target Search Page"

8.6.1.1 Viewing Observations By Systems

When observations occur, they can be marked as authorized or unauthorized automatically. This provides one way you to find observations that are important for you to look into. However, if a rule is not configured to reconcile observations with a change management server or if there is a large number of observations, it is difficult to find the observations that are important to you through only an attribute search. Being able to view observations by business application and drilling down into observation details allows you to discover where there may be issues that should be investigated.

Typically, IT managers and line of business owners must identify when unwanted configuration drift occurs in their business applications. By browsing observations by systems, you can easily see which changes affect specific business applications. Observations can be filtered by whether they are authorized, unauthorized, unaudited or both. They can also be filtered by time.

This begins with you choosing one or more business applications and being able to see the relative counts of observations taking into account some number of filters that

have occurred over periods of time. You will most likely not know what a target is which is why the view starts at business applications. A business application is modeled in Enterprise Manager as a composite target/target group with type of "generic system."

If you are more technical, this instance may be about being able to start from the business application view and drilling down to the actual observations themselves to see the details of what is changing.

To view observations by systems, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Real-Time Observations**.

   Enterprise Manager displays the Real-Time Observations page that lists three options you can choose:

   - Browse Observations by Compliance Framework

     Allows you to select your compliance framework to see where observations are occurring in each area. You can drill down through various levels to narrow the view to highlight observations that need investigation.

   - Browse Observations by System Targets

     Allows you to select your business applications modeled as System Targets to see where observations are occurring in each area. You can drill down through various levels to narrow the view to highlight observations that need investigation.

   - Search Observations

     Allows you to search across any of the attributes of an observation for actions that occurred.

2. Click **Browse Observations by System Targets**

   Enterprise Manager displays the Select Root Target(s) page that lists the Target Name for each group.

3. You can filter the Targets using the search criteria at the top of the page. Alternately, you can also either use a **Saved Search** by choosing it from the drop-down list, or create a Saved Search by conducting a search and then clicking **Save** to save the criteria. Clicking Save opens the Create Saved Search dialog box where you can add a name for the Saved Search and then choose to Set as Default, Run Automatically, or Save Results Layout. Click **OK** to save the search.

4. You can submit a target name to a journal page by selecting the target and clicking **Submit to Journal Page**.

   You will see counts for each system target selected. Click on the system target name to drill down and show the counts by each target that comprises the system target. Continuing to drill down provides more specific information.

   Clicking on the count displays a screen that shows the actual observations that occurred during that time period.

8.6.1.2 Viewing Observations By Compliance Framework

The ability to view observations as they relate to a compliance standard structure is something that is typically done by a non-technical role such as an IT Manager, Line of Business Owner, Compliance Manager, or Executive.

You can identify some set of Compliance Frameworks that reflect the IT compliance framework that the organization follows. Observations can be filtered by whether they are authorized, unauthorized, unaudited or both. They can also be filtered by time.

This instance begins with you choosing one or more top level Compliance Frameworks and being able to see the relative counts of observations that have occurred over periods of time.

To view observations by Compliance Framework, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Real-Time Observations**.

   Enterprise Manager displays the Real-Time Observations page that lists three options you can choose:

   - Browse Observations by Compliance Frameworks

     Allows you to select your compliance framework modeled as Compliance Frameworks to see where observations are occurring in each area. You can drill down through various levels to narrow the view to highlight observations that need investigation.

   - Browse Observations by System Targets

     Allows you to select your business applications modeled as System Targets to see where observations are occurring in each area. You can drill down through various levels to narrow the view to highlight observations that need investigation.

   - Search Observations

     Allows you to search across any of the attributes of an observation for actions that occurred.

2. Click **Browse Observations by Compliance Frameworks**

   Enterprise Manager displays the Select Compliance Framework(s) page that lists the Compliance Framework Name for each group.

3. You can filter the Compliance Frameworks using the search criteria at the top of the page. Alternately, you can also either use a **Saved Search** by choosing it from the drop-down list, or create a Saved Search by conducting a search and then clicking **Save** to save the criteria. Clicking Save opens the Create Saved Search dialog box where you can add a name for the Saved Search and then choose to Set as Default, Run Automatically, or Save Results Layout. Click **OK** to save the search.

4. Click on any Compliance Framework and choose Submit to Journal Page.

Usage Notes

There are filters available to allow you to narrow the observation counts and observations you can view. Some of these filters are:

Time range: A time range the report covers

User: A user that would have had any activity

Authorized Status: Shows all observations, only authorized, only unauthorized, only unaudited, etc.

1. The first level shows the top level Compliance Frameworks, children Compliance Frameworks, and compliance standards in a hierarchical structure. Along with

each cs group, a count of all observations made during the time range will be shown. Selecting some number of cs groups takes the user to the next level

2. The second level shows counts of observations based on filter settings for all selected compliance standards along with counts of observations by compliance standard and time bucket. Selecting a compliance standard allows the user to drill into the next level
3. The third level shows the targets that as associated with the selected compliance standard along with counts of observations for these targets and compliance standard by time bucket. Selecting a target allows the user to drill into the next level.
4. The fourth level shows the entity type. The fifth level shows the facets that are monitored on the target along with counts of observations by face and time bucket. Selecting a facet allows the user to drill into the next level
5. The sixth level shows the entities in a hierarchy fashion with counts of observations by entity. As an example, if the entity type is file, then the hierarchy shows the disk structure to the files involved. Clicking on an entity takes the user to the next level
6. The last level shows the actual observation details in a table with each column showing the various attributes of the observation. The user can also get to this screen by clicking on a count from any of the previous levels.

This drill-down capability provided by these screens makes it easy for you to easily find where observations are occurring. When you have an environment with tens of thousands of targets across hundreds of business applications, it is impossible to view observations simply using a table and search unless you know exactly the search conditions they are looking for. In a matter of an hour, with this large of an environment even with little activity, there can be tens of thousands of observations.

8.6.1.3 Viewing Details of an Incident

Observations are logically bundled together based on the compliance standard rule, target and user that performed the action.

The Observations page shows the list of observations in an observation bundle. You can filter on various attributes for each observation, including but not limited to the authorized/unauthorized status, user, time, entity, entity type, observation type, and so on.

You can use this page to show the user the observations in a bundle that are unauthorized. You can navigate to this page by clicking on some link from another page showing violations or notation of unauthorized activity.

When the Incident Manager Console opens a change request in a Change Management server for an observation bundle that has an unauthorized observation, the details in that incident may have a link that will take you to the bundle detail page. The URL allows you to display a page that shows each unauthorized observation in the observation bundle.

When looking at a compliance standard violation, you can see the detail of the observation bundle involved in the violation. Any observation bundle that has at least one unauthorized observation will be considered a violation. From the Event Details page under the Incident Management console, you can jump to this page to list all observations and filter on the unauthorized observations.

In an observation bundle, there can be a mix of observations; some number that are authorized, some number that are unauthorized and some that are unaudited,

meaning they were not reconciled with a change management server and you have not manually set an authorized/unauthorized status.

To view the details of an incident, follow these steps:

1. From the Cloud Control Home page, navigate to the Incident page by clicking **Enterprise**, then clicking **Monitoring**, then **Incident Manager**.

   Enterprise Manager displays the Incident Manager page.

2. In the All Open Events table, highlight the event for which you want to see incident details. If no event rules have been created to raise incidents, the real-time observation violations will show up under "Events without Incidents".

3. In the details section below, on the General tab you can click on the **Details** arrow to view the details of the incident.

8.6.1.4 Viewing a Summary of Observation Bundle Reconciliation Results From Incident Pages

When an Enterprise Manager incident is created because there was a violation, it is because an observation bundle had at least one observation that was unauthorized. This unauthorized observation could have been automatically reconciled to become unauthorized or you could have manually set the status to unauthorized. When looking at the incident history in the incident pages for an incident related to a real-time configuration change compliance standard rule, the incident will relate to a single observation bundle. The incident has a region that displays a summary of the statuses of the observations in the group. For instance, X authorized, Y unauthorized, Z unaudited out of N total observations. A link is also provided that takes you to the page that displays all of the observations in the observation bundle along with the details of the observations.

8.6.1.5 Viewing Observations With Compliance Violations From the Target Search Page

The Target Search page can display targets with compliance violations. These violations may be of any of the three types of compliance standard rules (config standards, guardian, or real-time configuration change monitoring. For real-time monitoring only, violations can either be for an entire group or for a single observation. You can view how many violations there are on the target related to real-time configuration change monitoring as well as navigate to a page that displays all the details of the observations that lead to the violation.

To view observations with compliance violations from the Target Search page, follow these steps:

1. Navigate to the Target Search page.

2. Click on the number that represents the compliance violation. To get to real-time compliance violations, you must continue to drill down.

8.6.2 Authorized and Unauthorized Real-Time Observations

The following sections describe authorized and unauthorized real-time observations. For additional information, refer to the *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

- Section 8.6.2.1, "Automatically Specifying Whether Real-time Observation Is Authorized"

8.6.2.1 Automatically Specifying Whether Real-time Observation Is Authorized

You can reconcile an observation against a change management system to determine if that observation is authorized or unauthorized.

Multiple observations can belong to the same Observation Bundle. Even though an observation is part of group, the determination of authorized vs. unauthorized is done for a single observation, not at the group level. If a group has at least one observation that is marked as “unauthorized”, then the group is considered to be a “violation” and an event or incident can be raised for this group violation.

This task provides an automated way to determine whether an observation was authorized.

To automatically specify whether real-time observation is authorized, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.

   Enterprise Manager displays the Compliance Standard Rule Library page.
3. You can specify whether Real-time Observation is authorized in an existing rule or when you create a new rule. In an existing rule select the rule you want to edit from the Rules table and click **Edit** and advance through the Edit Rule wizard until you get to the Settings page. Similarly, for a new rule enter the appropriate properties and attributes and advance to the Settings page.
4. On the Settings page of the Edit Rule wizard or Create Rule wizard, select **Authorize Observations Automatically using Change Management System**.

   You must choose a **Change Management Connector** from the drop-down list that has been configured. Optionally you can choose the option that allows you to annotate a ticket with Authorized Details if a change request exists.

8.6.2.2 Annotating Change Requests With Observation Details for Authorized Observations

When an observation is detected on the agent and comes into the OMS, if the rule that caused the observation to be detected had integration with a change request management system set up, and the rule specified automatic reconcile changes with

open change requests on the change request management server, then the observation is automatically determined to be authorized or unauthorized.

An observation can be determined to be authorized if one or more change requests match correlation criteria of the observation.

One user-configured option when selecting Change Request Management integration for a rule is to annotate the change requests with the details of the observations that were authorized by this change request. If an observation is authorized by more than one change request, all of these change requests will be annotated with details of this observation.

To specify that you want to annotate change requests with observation details, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.

   Enterprise Manager displays the Compliance Standard Rule Library page.
3. In an existing rule select the rule you want to edit from the Rules table and click **Edit** and advance through the Edit Rule wizard until you get to the Settings page. Similarly, for a new rule enter the appropriate properties and attributes and advance to the Settings page.
4. On the Settings page of the Edit Rule wizard or Create Rule wizard, select **Authorize Observations Automatically using Change Management System**.

   You must choose a **Change Management Connector** from the drop-down list that has been configured.
5. Choose the option that allows you to annotate a ticket with Authorized Details if a change request exists.

   Whenever you get an authorized observation it updates the authorizing change request with the details.

8.6.2.3 Treating Observation Bundles With Unauthorized Observations As Compliance Violations

When an observation bundle has at least one unauthorized observation (either through automatic reconciliation or manually set by the user), this observation bundle is considered to be a compliance violation.

Compliance violations from either of the three available compliance standard rules provide a single measurable value to affect the compliance score for the compliance standard(s) to which this rule belongs.

The compliance score that is calculated for a real-time monitoring rule is based on the number of observation bundles that have unauthorized observations versus all observation bundles. Older Observation bundles are weighted less in the scoring.

Observation Bundle Score = 1-(# of unauthorized/total observations)

For example, if an observation bundle has 20 observations in it and one is unauthorized, the score for this bundle is 0.95. If this same bundle had 20 unauthorized observations, then the score would be 0.

If there are no unauthorized observations in a bundle, then the score for this bundle would be 1.0 (meaning it is 100% compliant).

8.6.2.4 Overriding the Automatic Determination of Authorized or Unauthorized For an Observation

You can override the audit status of an observation if you investigate the user action and determine that the activity should have resulted in a different audit status.

You can change an automatically authorized observation to unauthorized or vice-versa. You can also change an authorized or unauthorized observation into unaudited which is the same as it would have been if it was not checked against any change management requests (for example, the rule did not enable CM reconciliation).

If you change an unauthorized observation into an authorized observation, then you have the option of entering a change request ID that is known to authorize the change. This change request ID should match a request that already exists in your change request management system. You can also enter a comment. If a change request ID is provided, then the change request is annotated with the change just as if the system had automatically authorized it. If an incident had been created for the observation bundle, then the event/incident is updated with the new number of unauthorized observations.

If there was only one observation that was unauthorized in the group and you manually changed it to authorized, then the incident will automatically close. If after some time of the incident closing, you set one of the authorized observations back to unauthorized, a new event is automatically opened and a new incident may be created based on the event rule definitions. The system does not reopen the previous event/incidents, but creates a new one based on the observation bundle again becoming a violation.

If you change an authorized observation into an unauthorized or unaudited observation, any annotations that were made to any change requests are rolled back. If there was already an incident raised for the observation bundle, then the annotation is changed to update the number of unauthorized observations in the incident. If this is the first unauthorized observation in a group, then an event is created an incident is raised. You can provide a comment for the change.

Although the status will remain as authorized or unauthorized or unaudited, internally in the change tracking, the system notes who made the annotation. This is something that is visible when a customer goes to view the reconciliation history for a given observation. If there is no user name (or some system user name is used) then the annotation was done automatically by the reconciliation engine. Because automatic actions are done as 'sysman', if it has a real Enterprise Manager user name, then it was a manual setting.

When you manually set the observation to be authorized and enter a change request ID and the rule has change management integration enabled, no attributes of the change request are compared with the observation. The change request is simply updated with the observation details.

When rolling back annotations in the change management server, the observation annotations are marked as rolled-back instead of actually removing the annotation. This occurs to avoid user confusion not knowing possibly why the annotations were removed. Also, if the observation later becomes authorized again, the rolled-back marking can simply be removed to bring the annotation back.

8.6.2.5 Manually Setting an Observation As Authorized Or Not Authorized

For users that do not have a change management system integrated for CM reconciliation, all observations are by default given a reconciliation notation of unaudited – meaning they were not reconciled. You can manually set these observations as authorized or unauthorized.

If you are changing the observation from unaudited to authorized, you can provide a change request ID and a comment. Since the customer did not integrate with the CM server, no request updates are made to the change request to annotate the request.

If you changed the observation from unaudited to unauthorized, you can provide a comment. The observation bundle is treated as a compliance violation, resulting in an automatic Enterprise Manager event being raised. You can still decide in the Enterprise Manager event rule that an incident should be raised for the compliance violation. It is possible that even though CM integration was not configured for reconciliation that a change management incident ticket can still be created through the Enterprise Manager incident management system like any other incident.

Although the status remains as authorized or unauthorized or unaudited, internally in the change tracking, the system notes who made the annotation. Because system annotations are done as 'sysman', if there is no user name (or some system user name is used) then the annotation was done automatically by the reconciliation engine. If it has a real Enterprise Manager user name, then it was a manual setting.

To manually set an observation as authorized or not authorized when not integrating with a Change Management System, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Real-Time Observations**.

   Enterprise Manager displays the Real-Time Observations page that lists three options you can choose:

   - Browse Observations by Compliance Framework

     Allows you to select your compliance framework to see where observations are occurring in each area. You can drill down through various levels to narrow the view to highlight observations that need investigation.

   - Browse Observations by System Targets

     Allows you to select your business applications modeled as System Targets to see where observations are occurring in each area. You can drill down through various levels to narrow the view to highlight observations that need investigation.

   - Search Observations

     Allows you to search across any of the attributes of an observation for actions that occurred.

2. Click **Search Observations**

   Enterprise Manager displays the Observations Search page that lists the observations returned by the search criteria you enter.

3. You can modify the search fields used to filter your query by clicking **Advanced**.

   Enterprise Manager displays the Advanced Search Criteria dialog box that allows you to select the fields that appear in the search criteria.

4. Update the audit status of any observation by selecting the observation or multiple observations from the search results and clicking **Update Audit Status**.

   Enterprise Manager displays the Update Audit Status dialog box. Choose the Action as either Authorized or Unauthorized, optionally add a Comment, and click **OK**.

8.6.2.6 Notifying a User When An Observation Occurs Without Change Management Integration

If a compliance standard rule is created and you do not use change management reconciliation with the rule, then there will be no automated authorized/unauthorized check done on the observations. You can specify for this rule that each observation bundle should result in informational event being generated for the observation bundle.

The event will have a notation. From the Incident Management console the user can look at events and incidents. When looking at a single event, there is a link available to see the observations associated with this observation bundle's event. Each observation bundle can only have one event. If at least one observation in the bundle is unauthorized, then the bundle is considered to be in violation which results in the event being generated.

Since this notification does not require user intervention or follow-up action, it is treated as informational. If at a later time, someone changes one of these unaudited observations into an authorized or unauthorized one, a new informational event for the unaudited observations will not be re-delivered. It is delivered only once for the observation bundle. However if one of the observations is manually set to unauthorized, then a violation is raised for the entire observation bundle.

When at least one observation in a bundle is in an unauthorized state, a violation is created. That violation becomes an event in the Incident Manager Console. Use the Incident Manager feature to set up a notification. For more information about this, on the Incident Manager page, click on the online help link, *Setting Up Notifications With Rules* under the Setting Up Notifications section under Getting Started.

8.6.2.7 Notifying a User When An Authorized Observation Occurs

When an authorized observation occurs, it is not a typical for you to receive a notification on these observations since the activity that caused the observation was expected. If you are using change management reconciliation, you have an option to annotate the authorizing change request with the observation details. The updates to the change request is one way customers can learn of authorized activity. You can set filters in their change management system to let them know that a change request has had authorized activity against it.

Enterprise Manager has no other way to get an alert on an authorized observation.

If you are not using automatic reconciliation with a change management server, then to notify a user for every observation bundle with an informational event, follow the steps below:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Library**.
2. Click on the **Compliance Standard Rule** tab.

   Enterprise Manager displays the Compliance Standard Rule Library page.
3. In an existing rule select the rule you want to edit from the Rules table and click **Edit** and advance through the Edit Rule wizard until you get to the Settings page. Similarly, for a new rule enter the appropriate properties and attributes and advance to the Settings page.
4. On the Settings page of the Edit Rule wizard or Create Rule wizard, select **Authorize Observations Manually**.

5. Under Collection Settings, expand the Advanced Settings box and choose the option that allows you to Generate an informational event during manual authorizations.

   An informational event is only created using Manual Observations. When you return to the Rules page and manually change the status, an event is created. However the event is informational only and will not affect the compliance score of the rule now will it display on the Target Home page as an issue. You can then set up a notification on the Incident Manager Console to be sent to specific individuals.

8.6.2.8 Determining Whether Notifications On Unauthorized Observations Have Been Acknowledged, Reassigned, Or Escalated

Since notifications from events/incidents related to unauthorized events are typically indicative of a compliance issue, you should be able to identify notifications that have not been responded to after some period of time. If a notification is not responded to promptly, you should be able to escalate or reassign the notification to allow someone else to react to it.

A notification of an unauthorized observation (violation) typically results in one of two actions happening. Either the observation is investigated and found to be authorized through some change request or through the corporate process or the observation will actually be determined to be unauthorized. In the latter case, the result may be to create a change request to roll back a change or to change the IT corporate policies to ensure the problem does not happen again in the future.

The Incident Manager tracks notifications and the history of events. See the online help for the Incident Manager functionality to learn more about how to determine whether notifications or unauthorized observations have been acknowledged, reassigned, or escalated.

8.6.3 Operations on Observations

The following sections explain the possible operations on observations.

- Section 8.6.3.1, "Selecting Observation Attributes to Display When Viewing Observations"
- Section 8.6.3.2, "Searching Observations By Observation Attributes"
- Section 8.6.3.3, "Changing Status and Annotating Observations"

8.6.3.1 Selecting Observation Attributes to Display When Viewing Observations

An observation presents many attributes or details about an observation. Some observations may have attributes that others do not simply because the entity type is different. For instance, a file change has some attributes that a process start does not have and vice-versa. You can customize the fields that you want to display. You can also designate the order of the fields.

Different types of users will want to see different details. Once you set the fields these settings can be saved per user per screen.

To select observation attributes to display when viewing observations, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Real-Time Observations**.

2. Click **Search Observations**

   Enterprise Manager displays the Observations Search page that lists the observations returned by the search criteria you enter.

3. You can choose which columns to display on the table. Click **View**, then click **Columns**. From here you can either click **Show All** to display all available columns for the entity type, or you can click **Show More Columns** to see a list of columns to display.

4. You can reorder the columns as they appear in the table by clicking View, then **Reorder Columns**. The Reorder Columns dialog box displays where you can highlight columns and use arrow keys to move columns up, down, to the top or bottom of the list.

8.6.3.2 Searching Observations By Observation Attributes

The ability to search observations is something that is typically done by a technical role such as a systems administrator, operator, or DBA.

You can select a set of criteria and search for observations matching these criteria. This is most effective when you know several attributes on which you want to search.

Use a search by attributes when you already know the specific types of observations for which you are looking. You can search based on time and user if you are trying to find the activity of a specific user. You can also search by time and a specific target to perform root-cause-analysis for a system failure.

The result of the search is a table of observations where each row is one observation and each column is one attribute of the observation.

To search observations by observation attributes, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Real-Time Observations**.

   Enterprise Manager displays the Real-Time Observations page that lists three options you can choose:

   - Browse Observations by Compliance Frameworks

     Allows you to select your compliance framework modeled as Compliance Frameworks to see where observations are occurring in each area. You can drill down through various levels to narrow the view to highlight observations that need investigation.

   - Browse Observations by System Targets

     Allows you to select your business applications modeled as System Targets to see where observations are occurring in each area. You can drill down through various levels to narrow the view to highlight observations that need investigation.

   - Search Observations

     Allows you to search across any of the attributes of an observation for actions that occurred.

2. Click **Search Observations**

   Enterprise Manager displays the Observations Search page that lists the observations returned by the search criteria you enter.

3. You can modify the search fields used to filter your query by clicking **Advanced**.

Enterprise Manager displays the Advanced Search Criteria dialog box that allows you to select the fields that appear in the search criteria.

4. You can update the audit status of any observation by selecting the observation from the search results and clicking **Update Audit Status**.

8.6.3.3 Changing Status and Annotating Observations

When viewing observations, you can see the current authorized/unauthorized status, you can change the status and also add a comment. This function is only available on a page that shows individual observations.

You can select one or more rows on the screen and apply the same change to all of the selected rows. When doing more than one row at a time, the same comment is applied to each observation.

Although the status remains as authorized or unauthorized or unaudited, internally in the change tracking, it is noted as a manual setting rather than an automatic reconciliation setting.

Ensure you have the privileges to change the status. See Section 8.1.3, "Privileges and Roles Needed to Use the Compliance Features" for information.

To change the status and annotate observations, follow these steps:

1. From the Enterprise Manager Cloud Control home page, click **Enterprise**, then choose **Compliance** from the drop-down menu, and then select **Real-Time Observations**.

   Enterprise Manager displays the Real-Time Observations page that lists three options you can choose:

   - Browse Observations by Compliance Frameworks

     Allows you to select your compliance framework modeled as Compliance Frameworks to see where observations are occurring in each area. You can drill down through various levels to narrow the view to highlight observations that need investigation.

   - Browse Observations by System Targets

     Allows you to select your business applications modeled as System Targets to see where observations are occurring in each area. You can drill down through various levels to narrow the view to highlight observations that need investigation.

   - Search Observations

     Allows you to search across any of the attributes of an observation for actions that occurred.

2. Click **Search Observations**

   Enterprise Manager displays the Observations Search page that lists the observations returned by the search criteria you enter.

3. Conduct a search displaying the results of the search including the Audit Status column. If the Audit Status column is not part of the column listing, you can add it by clicking **View**, then **Columns**, and choosing **Audit Status** from the list of available columns.

4. Select the row or rows containing the Audit Status you want to update and click **Update Audit Status**. Enterprise Manager displays the Update Audit Status dialog box. In the Action field, choose the audit status you want to assign to the

observation. The dialog box also lists the User and a Comment field where you can add annotated text.

5. You can view the audit history of an observation by clicking the Audit Status link in the table row. Enterprise Manager displays the Audit History dialog box that shows the Time, Status, and other information about each audit event.

9

Enterprise Manager Cloud Control Mobile

This chapter describes how to set up and use an iDevice to remotely connect to Enterprise Manager for the purpose of managing incidents and problems in Cloud Control.

The chapter contains the following sections:

9.1 System Requirements

Enterprise Manager Cloud Control Mobile has the following system requirements:

- iDevice (iPhone, iPod touch, or iPad) running iOS 4.2.x or later
- An installed Enterprise Manager Cloud Control 12c
- A Wi-Fi or 3G connection to a network that has access to Enterprise Manager (Cloud Control Mobile supports connections over VPN)
- An Apple account with which to download the app from the iTunes App Store

9.2 Initial Setup

Initial setup involves the following tasks:

- Connect to a Wi-Fi or 3G network
- Install and configure VPN
- Download the Cloud Control Mobile app and sync with your iDevice
- Add a Cloud Control URL to connect to the installed Enterprise Manager

9.3 Connecting the First Time

When you first install the app, there is no default Enterprise Manager connection, so you must supply a Cloud Control URL. There are two ways to do this:

- Use the iDevice Settings app
- Launch the Cloud Control Mobile app

In either case, first log in to VPN if required before proceeding with the instructions below. Without the VPN connection, the login screen will not appear.

iDevice Settings

Define a default Cloud Control URL as follows:

1. Tap the Settings icon on the Home screen.
2. Tap Cloud Control in the apps list.
3. On the Cloud Control screen, enter a name to identify the site and type the Cloud Control URL to which to connect. The URL should be of the form:

   ```
   https://www.yoursite.com/em
   ```

4. Tap **Settings** to store the information and return to the list of apps.

You can now launch the Cloud Control Mobile app to log in.

Initial App Launch

Define a default Cloud Control URL as follows:

1. Tap the Cloud Control Mobile icon on the Home screen.
2. On the Add Site screen, enter a name to identify the site and type the Cloud Control URL to which to connect. The URL should be of the form:

   ```
   https://www.yoursite.com/em
   ```

   Before you can type in the name field you may first have to clear the field by tapping the X at the right.
3. Tap **Done** to store the information.
4. Tap **Done** on the Sites screen. Note that you also have the option to add additional sites before exiting this screen.
5. Tap **Settings** on the Sites navigation bar to close the Sites list screen.
6. Tap **Save** on the Settings navigation bar to complete the action.

Proceed with the login.

9.4 Login Screen

You encounter the login screen under the following conditions:

- After supplying a default Cloud Control URL upon initial launch
- Anytime you subsequently launch the app
- When you change the default site
- When you log out

Tap the **Settings** icon to see a list of sites or to change the default login site. See "Manage Settings" for more information.

Specify your credentials and tap **Login**; the Incident Manager opens, displaying the my open incidents and problems view.

If your Enterprise Manager installation does not have a site certificate signed by a valid certificate authority, an alert overlays the login screen noting an invalid certificate. You have the option to continue with the login or change the URL to which you are trying to connect.

Note: If your installed Enterprise Manager uses single sign-on, the SSO process supplants site login. Upon completion of single sign-on, the workflow proceeds to the my open incidents and problems view.

9.5 Manage Settings

Cloud Control Mobile has its own settings interface apart from the iDevice Settings app that you use to manage all apps.

In managing Cloud Control Mobile app settings, you perform the following actions:

- Add a site
- Edit a site
- Delete a site
- Change the default site

Each action starts with the same basic steps:

1. Tap the actions icon on the right of the navigation bar.
2. Tap **Settings** in the action sheet.
3. Tap the Edit Sites table row.
4. Tap **Edit**. The Sites management screen appears:

Then proceed as described below for each individual action.

Add a Site

1. Tap the + sign on the left of the Sites navigation bar.
2. Type a name and a URL for the site to be added.
3. Tap **Done** on the Add Site navigation bar to close the screen.
4. Tap **Done** on the Sites navigation bar to exit edit mode.
5. Tap **Settings** on the Sites navigation bar to close the Sites list screen.
6. Tap **Save** on the Settings navigation bar to complete the action.

Edit a Site

1. Tap the blue arrow to the right of the URL to be edited.

2. Change the values as appropriate.
3. Tap **Done** on the Edit Site navigation bar to close the screen.
4. Tap **Done** on the Sites navigation bar to exit edit mode.
5. Tap **Settings** on the Sites navigation bar to close the Sites list screen.
6. Tap **Save** on the Sites navigation bar to complete the action.

Delete a Site

1. Tap the red circle to the left of the URL to be deleted.
2. Tap **Delete** that appears on the right in the table row.
3. Tap **Done** on the Sites navigation bar to close the screen.
4. Tap **Settings** on the Sites navigation bar to close the Sites list screen.
5. Tap **Save** on the Settings navigation bar to complete the action.

Change the Default Site

1. Tap the site to be the new default. The check mark to the right in the table row confirms your selection.
2. Tap **Settings** to close the Sites list screen.
3. Tap **Save** to complete the action.
4. After a brief moment, the Login screen appears. Specify credentials to log in to the new site.

Note that you also can change the default site in iDevice Settings for Cloud Control.

9.6 What You Can Do in Incident Manager Using Cloud Control Mobile

Connecting to Cloud Control remotely, you can do the following in Incident Manager:

- View your open incidents and problems; drill down to incident and problem details, including associated updates and events
- See the list of incidents for a given problem; link to these incidents and their details
- Acknowledge incidents and problems
- Manage incident and problem workflow for better tracking (change status, assign owner, escalate, and so forth)
- See who has been notified about an issue and what comments have been added by administrators

In addition, The FAQ that follows may help in understanding differences, subtleties, and nuances between the mobile app and its browser-based counterpart.

How do I view more issues?

The app displays five rows at a time. Use the next and previous controls at the bottom of the display to scroll the list. Or tap the number range itself (1 - 5) to pick from a list of increments.

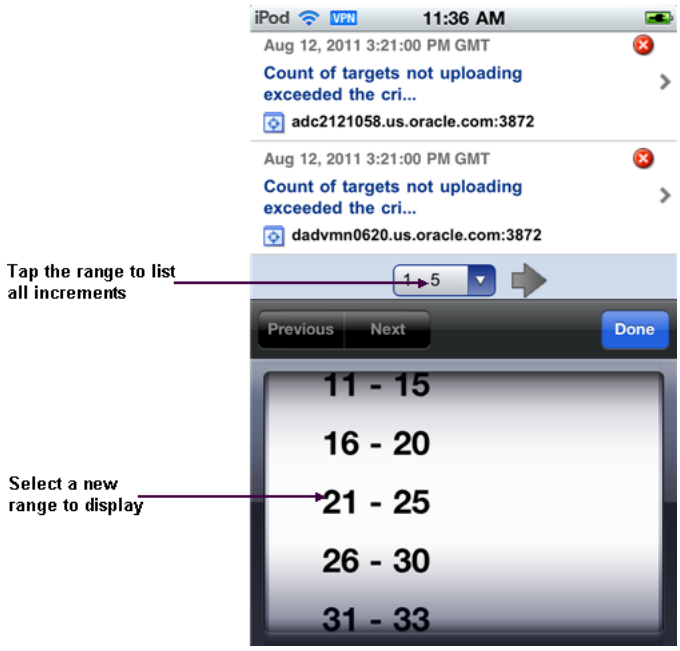

How do I view issue details?

Simply tap the line that identifies the incident or problem. On the Details screen, you can continue to drill down to updates and events as well as expand the summary description.

Are the views the same in the mobile app as in the browser-based product?

Yes, except for event-related views (standard or custom), which are not available in the app. Any custom views created in the browser-based product augment the list of standard views.

Can I refresh the view?

Yes, just tap the refresh icon on the left of the navigation bar. When you do, the timestamp reflects the date and time of the refresh.

Can I view target details?

Yes, first drill down to incident details, then tap the target name to jump to target details. Tap **Incident** to return to the incident details.

Can I set search criteria or create a custom view?

No, you cannot set search criteria or create a custom view in the mobile app, but you can create and manage views in the browser-based product, which are then available in the mobile app.

Can I invoke the incident rules feature?

No, but you can receive notifications generated by incident rules on your mobile device, provided your Enterprise Manager account has the appropriate notification preferences.

Can I connect to My Oracle Support?

If a problem has an assigned SR number, you can click the number to view the SR details in the My Oracle Support (MOS) Mobile app.

Can I access guided resolution information and diagnostics?

No.

Do all iDevices work the same way with the mobile app?

Pretty much. The one difference you will note on an iPad is that if you tap a link to an issue or a target in an external source such as Safari or an e-mail message, Safari launches, pointing to the relevant page in the browser-based product, where you are greeted with the usual Cloud Control login screen. With the other devices, tapping a link in an external source launches the mobile app.

9.7 Working in Cloud Control Mobile

This section covers the following operational tasks in Cloud Control Mobile:

- Viewing incidents and problems
- Changing views
- Performing actions

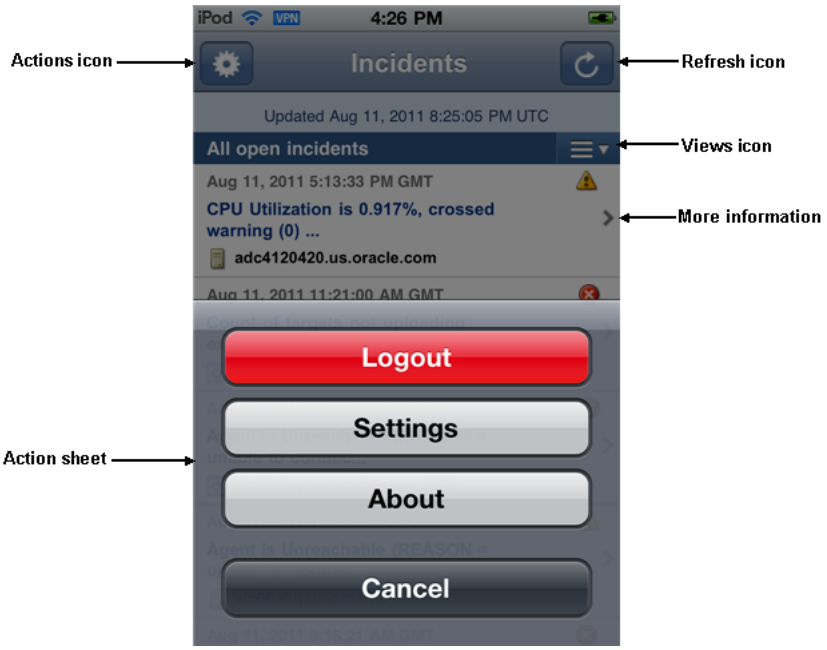

9.7.1 Viewing Incidents and Problems

Although navigation is intuitive, the following sections offer guidance on viewing incidents and problems. As the interactions are slightly different, there is a separate section for each type of issue.

Viewing Incidents

Use the following guidelines as you view incidents in the list:

- An arrow on the right indicates availability of additional information.
- Tap anywhere in the incident row to drill down to incident details.
- The summary appears at the top. As summaries can be lengthy, you may need to tap the opening lines of the summary to view the complete summary. Tap **Incident** to return to incident details.

- Problem ID in incident details is a link to problem details. If you follow the link, tap **Incidents** there to return to the starting point; that is, the original list view where you first opened the incident.
- In the incident details view, target name is a link to target details. Tap **Incident** there to return to incident details.
- Tracking information appears below target name in the incident details view.
- Scroll down in incident details and tap **All Updates** to see the equivalent of the **Updates** tab in the browser-based product. Tap **Incident** there to return to incident details.
- Scroll further and tap **Event List** to see the equivalent of the **Events** tab in the browser-based product. Tap **Incident** there to return to incident details.

Viewing Problems

Use the following guidelines as you view problems in the list:

- An arrow on the right indicates availability of additional information.
- Tap anywhere in the problem row to drill down to problem details.
- The summary appears at the top. As summaries can be lengthy, you may need to tap the opening lines of the summary to view the complete summary. Tap **Problem** to return to problem details.
- If the problem has an SR number assigned, tap the number to log in to the My Oracle Support (MOS) Mobile app using your MOS credentials. You can then view the SR details and take appropriate action. Go to the Home screen and tap the Cloud Control Mobile app icon to return to problem details.
- In the problem details view, target name is a link to target details. Tap **Problem** there to return to problem details.
- Tracking information appears below target name in the problem details view.
- Scroll down in problem details and tap **All Updates** to see the equivalent of the **Updates** tab in the browser-based product. Tap **Problem** there to return to problem details.
- Scroll further to see additional details such as first and last incident and number of incidents in which the problem has occurred.
- Scroll further and tap **Incident List** to see the equivalent of the **Incidents** tab in the browser-based product. Tap **Problem** there to return to problem details.
- Each incident summary in the list is a link to the details of the incident. If you follow the link, tap **Incidents** there to return to the starting point; that is, the original list view where you first opened the issue.

9.7.2 Changing Views

When you first log in, the my open incidents and problems view appears by default. You can change the view as follows:

1. Open the Views menu by tapping the views icon (three horizontal lines to the right of the current views title).
2. Tap the view you want. The check mark to the right confirms your selection.
3. Tap **Incidents** to display the new view.

9.7.3 Performing Actions

You can perform the following actions while viewing incident or problem details:

- Acknowledge the issue
- Manage the workflow of the issue

To acknowledge an incident or problem:

1. While viewing the details, tap **Actions**.
2. Tap **Acknowledge** in the action sheet.

   A message confirms the update on the Details screen.

Note that the Acknowledge action may not appear in the action sheet for a variety of reasons; for example, the issue has already been acknowledged or closed, or you do not have the right permissions to acknowledge the issue.

To manage the workflow of an incident or problem for better tracking:

1. While viewing the details, tap **Actions**.
2. Tap **Manage** in the action sheet.
3. Complete the Manage dialog the same as you would in the browser-based product. The only thing missing is the ability to add styles and formatting to the comment.
4. Tap **Save** to complete the action.

9.8 Tips and Tricks

Use a touch-and-hold gesture at any time within the app to display on the bottom of the screen the current site to which you are logged in or about to log in. If the site name is unavailable, the URL appears. With the site identity displayed, tap the information icon on the right to access the Settings screen, where you can manage your sites and change the default site. Repeat the touch-and-hold gesture to remove the site display from the bottom of the screen.

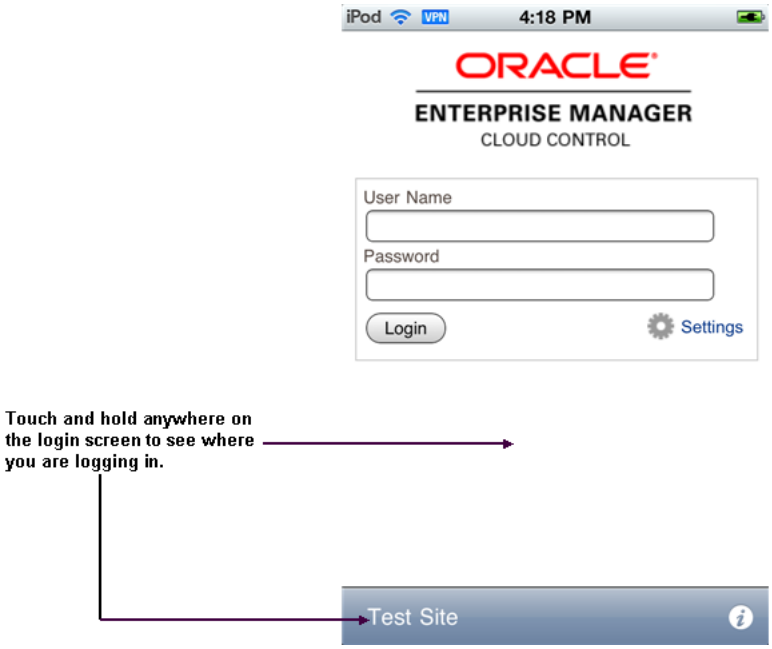

If you have logged out of the app and find that you are stuck on a page trying to return to Cloud Control Mobile, you have a couple of options to resolve the issue:

- Open iDevice Settings and change the Cloud Control default URL.
- Force quit the app and restart Cloud Control Mobile. For example, press and hold the On/Off button on top of the device until the power off slider appears, and then press the Home button until the app closes.

10

Information Publisher

Information Publisher, Enterprise Manager's reporting framework, makes information about your managed environment available to audiences across your enterprise. Strategically, reports are used to present a view of enterprise monitoring information for business intelligence purposes, but can also serve an administrative role by showing activity, resource utilization, and configuration of managed targets. IT managers can use reports to show availability of sets of managed systems. Executives can view reports on availability of applications (such as corporate email) over a period of time.

> **Note:** The Information Publisher (IP) reporting framework is still supported for Enterprise Manager 12c, however, new report development using this framework has been deprecated for Enterprise Manager 12c.

The reporting framework allows you to create and publish customized reports: Intuitive HTML-based reports can be published via the Web, stored, or e-mailed to selected recipients. Information Publisher comes with a comprehensive library of predefined reports that allow you to generate reports out-of-box without additional setup and configuration.

This chapter covers the following topics:

10.1 About Information Publisher

Information Publisher provides powerful reporting and publishing capability. Information Publisher reports present an intuitive interface to critical decision-making information stored in the Management Repository while ensuring the security of this information by taking advantage of Enterprise Manager's security and access control.

Information Publisher's intuitive user-interface allows you to create and publish reports with little effort. The key benefits of using Information Publisher are:

- Provides a framework for creating content-rich, well-formatted HTML reports based on Management Repository data.

- Out-of-box reports let you start generating reports immediately without any system configuration or setup.
- Ability to schedule automatic generation of reports and store scheduled copies and/or e-mail them to intended audiences.
- Ability for Enterprise Manager administrators to share reports with the entire business community: executives, customers, and other Enterprise Manager administrators.

Information Publisher provides you with a feature-rich framework that is your central information source for your enterprise.

10.2 Out-of-Box Report Definitions

The focal point of Information Publisher is the report definition. A report definition tells the reporting framework how to generate a specific report by defining report properties such as report content, user access, and scheduling of report generation.

Information Publisher comes with a comprehensive library of predefined report definitions, allowing you to generate fully formatted HTML reports presenting critical operations and business information without any additional configuration or setup. .

Generating this HTML report involved three simple steps:

Step 1: Click **Availability History** (Group) in the report definition list.

Step 2: Select the group for which you want to run the report.

Step 3: Click **Continue** to generate the fully-formed report.

Supplied report definitions are organized by functional category with each category covering key areas.

To access the Information Publisher home page, from the **Enterprise** menu, choose **Reports** and then **Information Publisher**.

10.3 Custom Reports

Although the predefined report definitions that come with Information Publisher cover the most common reporting needs, you may want to create specialized reports. If a predefined report comes close to meeting your information requirements, but not quite, you can use Information Publisher's Create Like function to create a new report definition based on one of the existing reports definitions.

10.3.1 Creating Custom Reports

To create custom reports:

1. Choose whether to modify an existing report definition or start from scratch. If an existing report definition closely matches your needs, it is easy to customize it by using Create Like function.
2. Specify name, category, and sub-category. Cloud Control provides default categories and sub-categories that are used for out-of-box reports. However, you can categorize custom reports in any way you like.
3. Specify any time-period and/or target parameters. The report viewer will be prompted for these parameters while viewing the report.

4. Add reporting elements. Reporting elements are pre-defined content building blocks, that allow you to add a variety of information to your report. Some examples of reporting elements are charts, tables, and images.
5. Customize the report layout. Once you have assembled the reporting elements, you can customize the layout of the report.

10.3.2 Report Parameters

By declaring report parameters, you allow the user to control what data is shown in the report. There are two types of parameters: target and time-period.

Example: If you are defining a report that will be used to diagnose a problem (such as a memory consumption report), the viewer will be able to see information for their target of interest.

By specifying the time-period parameter, the viewer will be able to analyze historical data for their period of interest.

Analyzing Historical Data

Information Publisher allows you to view reports for a variety of time-periods:

- Last 24 Hours/ 7 Days/ 31 Days
- Previous X Days/ Weeks/ Months/ Years (calendar units)
- This Week/ This Month/ This Year (this week so far)
- Any custom date range.

10.3.3 Report Elements

Report elements are the building blocks of a report definition. In general, report elements take parameters to generate viewable information. For example, the Chart from SQL element takes a SQL query to extract data from the Management Repository and a parameter specifying whether to display the data in the form of a pie, bar, or line chart. Report elements let you "assemble" a custom report definition using the Information Publisher user interface.

Information Publisher provides a variety of reporting elements. Generic reporting elements allow you to display any desired information, in the form of charts, tables or images. For example, you can include your corporate Logo, with a link to your corporate website. Monitoring elements show monitoring information, such as availability and alerts for managed targets. Service Level Reporting elements show availability, performance, usage and achieved service levels, allowing you to track compliance with Service Level Agreements, as well as share information about achieved service levels with your customers and business executives.

10.4 Scheduling Reports

Enterprise manager allows you to view reports interactively and/or schedule generation of reports on a flexible schedule. For example, you might want to generate an "Inventory Snapshot" report of all of the servers in your environment every day at midnight.

10.4.1 Flexible Schedules

Cloud Control provides the following scheduling options:

- One-time report generation either immediately or at any point in the future
- Periodic report generation
  - Frequency: Any number of Minutes/ Hours/ Days/ Weeks/ Months/ Years
  - You can generate copies indefinitely or until a specific date in the future.

10.4.2 Storing and Purging Report Copies

Enterprise manager allows you to store any number of scheduled copies for future reference.

You can delete each stored copy manually or you can set up automated purging based on either the number of stored copies or based on retention time. For example, you can have Enterprise Manager purge all reports that are more than 90 days old.

10.4.3 E-mailing Reports

You can choose for scheduled reports to be e-mailed to any number of recipients. You can specify reply-to address and subject of the e-mail.

10.5 Sharing Reports

Information Publisher facilitates easy report sharing with the entire user community. Enterprise Manager administrators can share reports with other administrators and roles. However, there may be cases when you need to share reports with non-Enterprise Manager administrators, such as customers and/or business executives. To facilitate information sharing with these users, Enterprise Manager renders a separate reporting website that does not require user authentication.

Note: To ensure that no sensitive information is compromised, only Enterprise Manager administrators with a special system privilege are allowed to publish reports to the Enterprise Manager reports website.

Information Publisher honors Enterprise Manager roles and privileges, ensuring that only Enterprise Manager administrators can create reports on the information they are allowed to see.

When sharing reports, administrators have an option of allowing report viewers to see the report with the owner's privileges. For example, as a system administrator you might want to share a host's performance information with a DBA using your server, but you do not want to grant the DBA any privileges on your host target. In this case, you could create a host performance report, and allow the DBA to view it with your privileges. This way, they only see the information you want them to see, without having access to the host homepage.

11 Enterprise Manager Security

This chapter describes how to configure Oracle Enterprise Manager Security. Specifically, this chapter contains the following sections:

11.1 About Oracle Enterprise Manager Security

Oracle Enterprise Manager provides tools and procedures to help you ensure that you are managing your Oracle environment in a secure manner. The goals of Oracle Enterprise Manager security are:

- To be sure that only users with the proper privileges have access to critical monitoring and administrative data.

  This goal is met by requiring username and password credentials before users can access the Enterprise Manager consoles and appropriate privileges for accessing the critical data.

- To be sure that all data transferred between Enterprise Manager components is transferred in a secure manner and that all data gathered by each Oracle Management Agent can be transferred only to the Oracle Management Service for which the Management Agent is configured.

  This goal is met by enabling Enterprise Manager Framework Security. Enterprise Manager Framework Security automates the process of securing the Enterprise Manager components installed and configured on your network.

- To be sure that sensitive data such as credentials used to access target servers are protected.

  This goal is met by Enterprise Manager's encryption support. The sensitive data is encrypted with an **emkey**. By following the best practice, even the repository owner and the `SYSDBA` will not be able to access the sensitive data.

- To be sure that access to managed targets is controlled through user authentication and privilege delegation.

  This goal is met by configuring the Management Agent with PAM and LDAP for user authentication and using privilege delegation tools like Sudo and PowerBroker.

11.2 Enterprise Manager Authentication

Enterprise Manager authentication is the process of determining the validity of the user accessing Enterprise Manager. The authentication feature is available across the different interfaces such as Enterprise Manager console and Enterprise Manager Command Line Interface (EMCLI).

Enterprise Manager's authentication framework consists of pluggable authentication schemes that let you use the type of authentication protocol best suited to your environment.

The following authentication schemes are available:

- **Oracle Access Manager (OAM) SSO** - Oracle Access Manager is the Oracle Fusion Middleware single sign-on solution. The underlying identity stores will be the Enterprise Directory Identity Stores being supported by Oracle Access Manager. For more information about OAM, see *Oracle® Fusion Middleware Administrator's Guide for Oracle Access Manager 12c Release 1 (11.1.1)*.
- **Repository-Based Authentication**: This is the default authentication option. An Enterprise Manager administrator is also a repository (database) user. By using this option, you can take advantage of all the benefits that this authentication method provides like password control via password profile, enforced password complexity, password life time, and number of failed attempts allowed. During the password grace period, the administrator is prompted to change the password but when the password has expired, it must be changed.For more details, refer to Repository-Based Authentication.
- **SSO-Based Authentication**: The single sign-on based authentication provides strengthened and centralized user identity management across the enterprise. After you have configured Enterprise Manager to use the Oracle Application Server Single Sign-On, you can register any single sign-on user as an Enterprise Manager administrator. You can then enter your single sign-on credentials to access the Oracle Enterprise Manager console. For more details, refer to Single Sign-On Based Authentication.
- **Enterprise User Security Based Authentication**: The Enterprise User Security (EUS) option enables you to create and store enterprise users and roles for the Oracle database in an LDAP-compliant directory server. Once the repository is configured with EUS, you can configure Enterprise Manager to use EUS as its authentication mechanism as described in Enterprise User Security Based Authentication. You can register any EUS user as an Enterprise Manager administrator.

  EUS helps centralize the administration of users and roles across multiple databases. If the managed databases are configured with EUS, the process of logging into these databases is simplified. When you drill down to a managed database, Enterprise Manager will attempt to connect to the database using Enterprise Manager credentials. If successful, Enterprise Manager will directly connect you to the database without displaying a login page.

- **Oracle Internet Directory (OID) Based Authentication** - When using an authentication scheme based on Oracle Internet Directory as the identity store, you can plug in the OID-based authentication scheme to have your applications authenticate users against the OID.
- **Microsoft Active Directory Based Authentication** - When using a Microsoft Active Directory as an identity store, you can plug in this scheme to have your applications authenticate users against the Microsoft Active Directory.

11.2.1 Repository-Based Authentication

Enterprise Manager allows you to create and manage new administrator accounts. Each administrator account includes its own login credentials as well as a set of roles and privileges that are assigned to the account. You can also assign a password profile to the administrator. To create, edit, or view an administrator account:

1. From the **Setup** menu, choose **Security** and then **Administrators**.
2. Click the appropriate task button on the Administrators page. The following screen is displayed:

Figure 11–1 Create / Edit Administrator

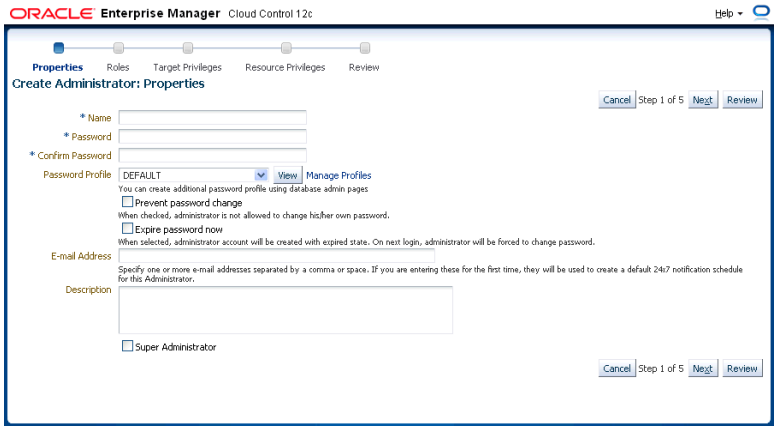

On this page, you can specify the type of administrator account being created and select the password profile. The password cannot be changed by the administrator if the **Prevent Password Change** checkbox is selected.

If you select the **Expire Password Now** checkbox, the password for administrator account will be set to an expired state. If the password has expired, when you login the next time, the following screen is displayed and you are prompted to change the password.

Figure 11–2 Password Expiry Page

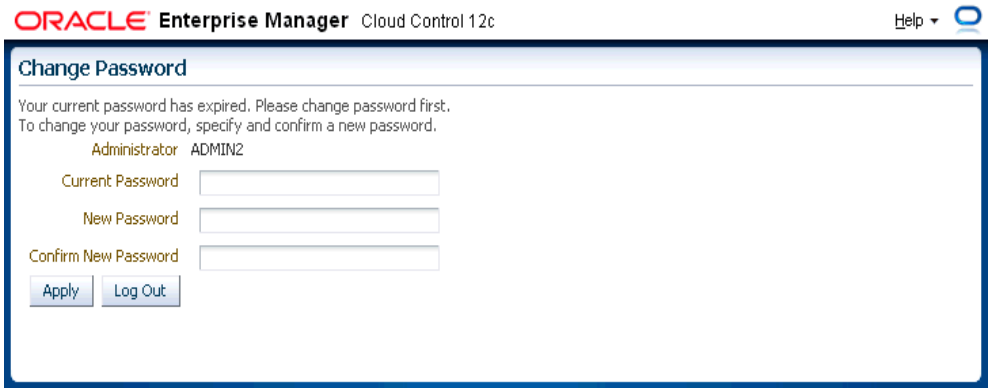

Enter your current password and the new password and click **Apply**. You can now start using Enterprise Manager.

11.2.2 Oracle Access Manager Single Sign-On

When using an Oracle Access Manager Single Sign-On authentication scheme, the underlying identity stores will consist of Enterprise Directory Identity Stores supported by Oracle Access Manager. This section provides instructions on how to configure OAM SSO-based authentication schemes.

Prerequisites

Oracle access manager is installed.

The Oracle Access Manager Single Sign-On server is configured with Oracle HTTP server, Web Gate, and the Oracle Access Manager Identity Store.

1. Run the `emctl config auth` command.

```
emctl config auth oam [-sysman_pwd <pwd>] -oid_host <host> -oid_port <port>
-oid_principal <principal> [-oid_credential <credential>]
-user_base_dn <dn> -group_base_dn <dn>
-oam_host <host< -oam_port <port> [-logout_url <url>] [-is_oam10g] [-user_dn
<dn>] [-group_dn <dn>]
```

   Note: Pass -is_oam10g option only if the OAM version is 10g.

2. Stop each OMS.

```
emctl stop oms -all
```

3. Restart each OMS.

```
emctl start oms
```

11.2.3 Single Sign-On Based Authentication

If you are currently using Oracle Application Server Single Sign-On to control access and authorization for your enterprise, you can extend those capabilities to the Enterprise Manager console.

By default, Enterprise Manager displays the main login page. However, you can configure Enterprise Manager so it uses Oracle Application Server Single Sign-On to authenticate your Enterprise Manager users. Instead of seeing the Enterprise Manager login page, users will see the standard Oracle Application Server Single Sign-On login

page. From the login page, administrators can use their Oracle Application Server Single Sign-On credentials to access the Oracle Enterprise Manager 12c Cloud Control console.

Note:

- You can configure Enterprise Manager to use one of the default Oracle Application Server Single Sign-On or Enterprise User Security features, but not multiple.
- When Enterprise Manager is configured to use Single Sign-On with Server Load Balancer, make sure that the correct monitoring settings have been defined. For details, refer to the chapter on *Cloud Control Common Configurations*.

The following sections describe how to configure Enterprise Manager as an OracleAS Single Sign-On Partner Application:

- Registering Enterprise Manager as a Partner Application
- Removing Single Sign-On Configuration
- Registering Single Sign-On Users as Enterprise Manager Administrators
- Bypassing the Single Sign-On Logon Page

11.2.3.1 Registering Enterprise Manager as a Partner Application

To register Enterprise Manager as a partner application manually, follow these steps:

1. Stop all OMSs by running `emctl stop oms` on each OMS.
2. Enter the following URL to navigate to the SSO Administration page.

   ```
   https://sso_host:sso_port/pls/orasso
   ```

3. Login as `orcladmin` user and click on **SSO Server Administration**.
4. Click **Administer Partner Applications** and then click **Add Partner Application**.
5. Enter the following information on the Add Partner Application page.

   ```
   Name: <EMPartnerName>
   Home URL: protocol://em_host:em_port
   Success URL: protocol://em_host:em_port/osso_login_success
   Logout URL: protocol://em_host:em_port/osso_logout_success
   Administrator Email: user@host.com
   ```

   Note1: host, port, and protocol refer to the Enterprise Manager Host, port and the protocol (http or https) used.

   Note2: The em_host, em_port, email and Enterprise Manager PartnerName need to be replaced appropriately and not typed as shown in this example.

6. Go back to Administer Partner Applications page and click on the Edit icon for <EMPartnerName>.

   Record the values of ID, Token, Encryption Key, Login URL, Single Sign-Off URL, Home URL and write the following in a file osso.txt:

   ```
   sso_server_version: v1.2
   cipher_key=<value of EncryptionKey>
   site_id=<value of ID>
   ```

```
site_token=<value of Token>
login_url=<value of Login URL>
logout_url=<value of Single Sign-Off URL>
cancel_url=<value of Home URL>
sso_timeout_cookie_name=SSO_ID_TIMEOUT
sso_timeout_cookie_key=9E231B3C1A3A808A
```

7. Set the ORACLE_HOME environment variable to WebTier Oracle Home location.

   ```
   setenv ORACLE_HOME /scratch/12c/MWHome/Oracle_WT
   ```

   Then, run the following:

   ```
   $ORACLE_HOME/ohs/bin/iasobf <location of osso.txt> <location
   of osso.conf>
   ```

8. Run the following command on each OMS:

   ```
   emctl config auth sso -ossoconf <osso.conf file loc> -dasurl <DAS URL>
   [-unsecure] [-sysman_pwd <pwd>] [-domain <domain>]-ldap_host <ldap host> -ldap_
   port <ldap port> -ldap_principal <ldap principal> [-ldap_credential <ldap
   credential>] -user_base_dn <user base DN> -group_base_dn <group base DN>
   [-logout_url <sso logout url>]
   ```

   where ldap_host, ldap_port, ldap_principal and ldap_credential are the details of SSO's LDAP.

   The sample output for this command is shown below:

   ```
   Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
   Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
   SSO Configuration done successfully. Please restart Admin & Managed Servers.
   ```

9. Run the following commands on each OMS:

   ```
   emctl stop oms -all
   emctl start oms
   ```

11.2.3.2 Removing Single Sign-On Configuration

To remove the single sign-on configuration, perform the following:

1. Run the following command on each OMS:

   ```
   emctl config auth repos [-sysman_pwd <pwd>]
   ```

   Sample command output:

   ```
   Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
   Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
   Configuring Repos Authentication ... Started
   Configuring Repos Authentication ... Successful
   If you have updated files like httpd.conf (for example, while installing
   WebGate), rollback them.
   If this is a multi-OMS environment, execute this command on remaining servers.
   After that, restart OMS(s) using: 'emctl stop oms -all' and 'emctl start oms'
   ```

2. Bounce all OMSs by issuing the following on each OMS:

   ```
   emctl stop oms -all
   emctl start oms
   ```

11.2.3.3 Registering Single Sign-On Users as Enterprise Manager Administrators

After you have configured Enterprise Manager to use the Single Sign-On logon page, you can register any Single Sign-On user as an Enterprise Manager administrator. You can register single sign-on users using:

- Enterprise Manager Graphical User Interface
- Enterprise Manager Command Line Interface

11.2.3.3.1 Registering Single Sign-On Users Using the Graphical User Interface

You can use the graphical user interface to register single sign-on users by following these steps:

1. Go the Enterprise Manager Console URL.

   The browser is redirected to the standard Single Sign-On Logon page.

2. Enter the credentials for a valid Single Sign-On user. Note: This step requires that an SSO user is already registered with Enterprise Manager.

   If no SSO user is yet registered as Enterprise Manager user, you can create them using the following procedure:

   1. Log in to Enterprise Manager by connecting to Managed Server (MS) directly. eg: https://ms_host:ms_https_port/em.

   2. Log in as a Repository user.

   3) Goto Setup -> Security -> Administrator

   4. Create SSO users.

3. Log in to Enterprise Manager as a Super Administrator.

4. From the **Setup** menu, choose **Security** and then **Administrators** to display the Administrators page.

   Because Enterprise Manager has been configured to use Single Sign-On, the first page in the Create Administrator wizard now offers you the option of creating an administrator either as an External User or as Repository User.

5. Select **External User Identity Store** and advance to the next page in the wizard.

6. Enter the name and e-mail address of the External User Identity Store user, or click the flashlight icon to search for a user name in the Oracle Internet Directory.

7. Use the rest of the wizard pages to define the roles, system privileges, and other characteristics of the Enterprise Manager administrator and then click **Finish**.

   Enterprise Manager displays a summary page that lists the characteristics of the administrator account.

8. Click **Finish** to create the new Enterprise Manager administrator.

   The External User Identity Store user is now included in the list of Enterprise Manager administrators. You can now verify the account by logging out of the Cloud Control console and logging back in using the External User Identity Store user credentials on the Single Sign-On logon page.

11.2.3.3.2 Registering Single Sign-On Users Using EMCLI s

You can use the following EMCLI command to create Single Sign-On users:

```
emcli create_user -name=ssouser -type=EXTERNAL_USER
```

This command creates a user with the name **ssouser** who is authenticated against the single sign-on user.

Argument	Description
-name	Name of the administrator.
-type	The type of user. The default value for this parameter is EM_USER. The other possible values are: ■ EXTERNAL_USER: Used for single-sign-on based authentication. ■ DB_EXTERNAL_USER: Used for enterprise user based security authentication.
-password	The password for the administrator.
-roles	The list of roles that can be granted to this administrator.
-email	The list of email addresses for this administrator.
-privilege	The system privileges that can be granted to the administrator. This option can be specified more than once.
-profile	The name of the database profile. This is an optional parameter. The default profile used is DEFAULT.
-desc	The description of the user being added.
-expired	This parameter is used to set the password to "expired" status. This is an optional parameter and is set to False by default.
-prevent_change_password	When this parameter is set to True, the user cannot change the password. This is an optional parameter and is set to False by default.
-input_file	This parameter allows the administrator to provide the values for any of these arguments in an input file. The format of value is `name_of_argument:file_path_with_file_name`.

Example 1

```
emcli create_user
        -name="new_admin"
        -email="first.last@oracle.com;joe.shmoe@shmoeshop.com"
        -roles="public"
        -privilege="view_job;923470234ABCDFE23018494753091111"
        -privilege="view_target;<host>.com:host"
```

This example creates an Enterprise Manager administrator named new_admin. This administrator has two privileges: the ability to view the job with ID `923470234ABCDFE23018494753091111` and the ability to view the target `<host>.com:host`. The administrator `new_admin` is granted the PUBLIC role.

Example 2

```
emcli create_user
      -name="User1"
      -type="EXTERNAL_USER"
      -input_file="privilege:/home/user1/priv_file"
```

```
Contents of priv_file are:
  view_target;<host>.com:host
```

This example makes `user1` which has been created externally as an Enterprise Manager user. `user1` will have view privileges on `<host>.com:host`.

Example 3

```
emcli create_user
      -name="User1"
      -desc="This is temp hire."
      -prevent_change_password="true"
      -profile="MGMT_ADMIN_USER_PROFILE"
```

This example sets `user1` as an Enterprise Manager user with some description. The `prevent_change_password` is set to true to indicate that the password cannot be changed by `user1` and the `profile` is set to `MGMT_ADMIN_USER_PROFILE`.

Example 4

```
emcli create_user
      -name="User1"
      -desc="This is temp hire."
      -expire="true"
```

This example sets `user1` as an Enterprise Manager with some description. Since the password is set to expire immediately, when the user logs in for the first time, he is prompted to change the password.

11.2.3.4 Bypassing the Single Sign-On Logon Page

If the OMS is configured with SSO or OAM or some other authentication method, you may want to by-pass the Single Sign-On or OAM authentication under certain circumstances.

To bypass the SSO logon page, connect to the following URL:

1. Connect to https://ms_host:ms_https_port/em

   ms_host & ms_https_port are WLS-managed server's hostname & port#. These parameters can be found in the EM_INSTANCE_HOME/emgc.properties file. They are listed as EM_INSTANCE_HOST & MS_HTTPS_PORT in this file.

2. Log in using a repository user's credentials.

11.2.3.5 Restoring the Default Authentication Method

1. Run the following command on each OMS:

   ```
   emctl config auth repos [-sysman_pwd <pwd>]
   ```

   Sample command output:

   ```
   Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
   Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
   Configuring Repos Authentication ... Started
   Configuring Repos Authentication ... Successful
   If you have updated files like httpd.conf (for example, while installing
   WebGate), rollback them.
   If this is a multi-OMS environment, execute this command on remaining servers.
   After that, restart OMS(s) using: 'emctl stop oms -all' and 'emctl start oms'
   ```

2. Run the following commands on each OMS:

```
emctl stop oms -all
emctl start oms
```

11.2.4 Enterprise User Security Based Authentication

Enterprise User Security enables you to create and store Oracle database information as directory objects in an LDAP-compliant directory server. For example, an administrator can create and store enterprise users and roles for the Oracle database in the directory, which helps `centralize t`he administration of users and roles across multiple databases.

> **See Also:** Enterprise User Security Configuration Tasks and Troubleshooting in the *Oracle Database Advanced Security Administrator's Guide*

If you currently use Enterprise User Security for all your Oracle databases, you can extend this feature to Enterprise Manager. Configuring Enterprise Manager for use with Enterprise User Security simplifies the process of logging in to database targets you are managing with the Oracle Enterprise Manager console.

To configure Enterprise Manager for use with Enterprise User Security:

1. Ensure that you have enabled Enterprise User Security for your Oracle Management Repository database, as well as the database targets you will be managing with the Cloud Control console. Refer to *Oracle Database Advanced Security Administrator's Guide* for details.
2. Using the `emctl set property` command, set the following properties:

```
oracle.sysman.emSDK.sec.DirectoryAuthenticationType=EnterpriseUser
oracle.sysman.emSDK.sec.eus.Domain=<ClientDomainName> (For
example:mydomain.com)
oracle.sysman.emSDK.sec.eus.DASHostUrl=<das_url> (For example:
oracle.sysman.emSDK.sec.eus.DASHostUrl=http://my.dashost.com:7777 )
```

   For example:

```
emctl set property -name oracle.sysman.emSDK.sec.DirectoryAuthenticationType
-value EnterpriseUser
```

3. Stop the Oracle Management Service.

> **See Also:** Controlling the Oracle Management Service on page 24-4

4. Start the Management Service.

The next time you use the Oracle Enterprise Manager console to drill down to a managed database, Enterprise Manager will attempt to connect to the database using Enterprise User Security. If successful, Enterprise Manager will connect you to the database without displaying a login page. If the attempt to use Enterprise User Security fails, Enterprise Manager will prompt you for the database credentials.

11.2.4.1 Registering Enterprise Users as Enterprise Manager Users

After you have configured Enterprise Manager to use Enterprise Users, you can register existing enterprise users as Enterprise Manager Users and grant them the necessary privileges so that they can manage Enterprise Manager effectively.

You can register existing enterprise users by using:

- Enterprise Manager Graphic User Interface
- Enterprise Manager Command Line Interface

11.2.4.1.1 Registering Enterprise Users Using the Graphical User Interface

You can use the graphical user interface to register enterprise users by following these steps:

1. Log into Enterprise Manager as a Super Administrator.
2. From the **Setup** menu, choose **Security** and then **Administrators** to display the Administrators page. Since Enterprise Manager has been configured to use Enterprise Users, the first page of the Create Administrator wizard will provide the option to create an administrator based on a registered Oracle Internet Directory user or a normal database user.
3. Select Oracle Internet Directory and click **Continue** to go to the next page in the wizard.
4. Enter the name and e-mail address of the Oracle Internet Directory user or click the flashlight icon to search for a user name in the Oracle Internet Directory.
5. Use the rest of the wizard pages to define the roles, system privileges, and other characteristics of the Enterprise Manager administrator and then click **Finish**. Enterprise Manager displays a summary page that lists the characteristics of the administrator account.
6. Click **Finish** to create the new Enterprise Manager administrator.

   The OID user is now included in the list of Enterprise Manager administrators. You can now verify the account by logging out of the Cloud Control console and logging back in using the OID user credentials on the Single Sign-On logon page.

11.2.4.1.2 Registering Enterprise Users Using the Command Line Interface

To register Enterprise Users as Enterprise Manager users using EMCLI, enter the following command:

```
emcli create_user -name=eususer -type=DB_EXTERNAL_USER
```

This command registers the `eususer` as an Enterprise Manager user where `eususer` is an existing Enterprise User. For more details, refer to Registering Single Sign-On Users Using EMCLI.

11.2.5 Microsoft Active Directory Based Authentication

Enterprise Manager uses the authentication capabilities provided by the Oracle WebLogic Server that is part of the OMS. If you are using Microsoft Active Directory as an identity store, you will need to configure it with the Oracle WebLogic Server which is part of the OMS. The following procedure demonstrates how to set up Enterprise Manager authentication using Microsoft Active Directory.

Prerequisites

- Ensure Enterprise Manager Cloud Control 12c is installed and configured properly and that you can log in as a user with Super Admin privileges.
- Ensure Microsoft Active Directory is installed and configured properly.

- Obtain te following from your Microsoft Active Directory administrator. Below is an example of a simple configuration. More complex configurations can be implemented with additional knowledge of LDAP search filters.
  - Active Directory Port
  - Active Directory Prinicpal (User created to authenticate with Active Directory for the Oracle WebLogic Server.
  - Active Directory Principal Password
  - User Base Distinguised Name (DN)
  - Group Base DN

		Example	Your Value
Host	The Active Directory host	server.oracle.com	
Port	The Active Directory Port	389 (LDAP) or 636 (LDAPS)	
Principal User/Password	The Principal User created in Active Directory that will be used to authenticate WebLogic Server. It must be in the Administrators group and belong to the correct Organizational Unit designated in the User base DN. Ensure the "User must change password at next logon" is not checked during setup.		emgcadmin/Welcome11
User Base DN	The User Base Distinguished Name is the container location of valid users who will be granted access to ENTERPRISE MANAGER. Using the default Users container will allow all Active Directory Users to login to ENTERPRISE MANAGER (though they may not have permissions to see/do anything). Using an Organizational Unit will allow you to further restrict access.		
User Base Filter From Name			
User Name Attribute:		sAMAccountName	

	Example	Your Value
User From Name Filter:	(&(sAMAccountName=%u)(objectclass=user))	

11.2.5.1 Configure WebLogic Server Authentication

1. As the Weblogic/Enterprise Manager Administrator, backup WLS config.xml located at ../gc_inst/user_projects/domains/GCDomain/config/config.xml
2. Login to Weblogic Admin Console as weblogic (you can find the admin console URL in $ORACLE_HOME/install/setupinfo.txt)
3. Under Domain Structure on the left, click on Security Realms
4. Click on myrealm and then click on the tab Providers
5. Click Lock & Edit to enable editing
6. Click New to add a Provider
7. Enter a Name for your Provider (for example, MS Active Directory)
8. Select ActiveDirectoryAuthenticator for Type, Click OK.

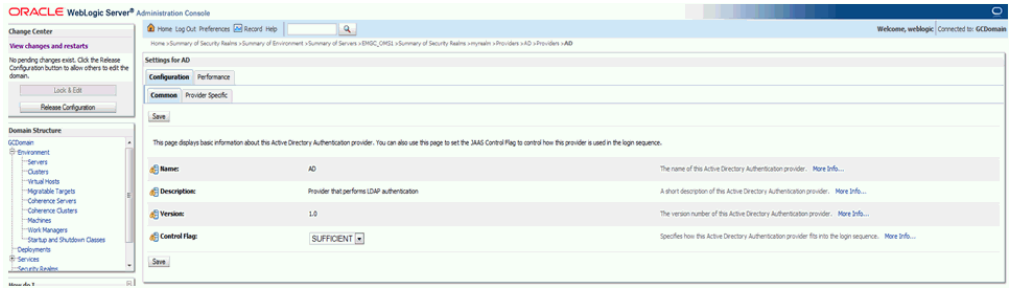

9. Back on the Providers screen, click the new Provider link to Edit
10. Set the Control Flag to Sufficient, Click Save
11. Click on Provider Specific tab
12. In Connection Section

    Host: AD Server Host

    Port: 389 (default for LDAP, or 636 for LDAPS)

    Principal:CN=EMGCADMIN,CN=Users,DC=Cloudcontrol,DC=local

    NOTE: This is the User created in AD steps above and added to Administrators group. The CN/DC string must be confirmed with your AD Administrator.

    Credential: pwd for principal

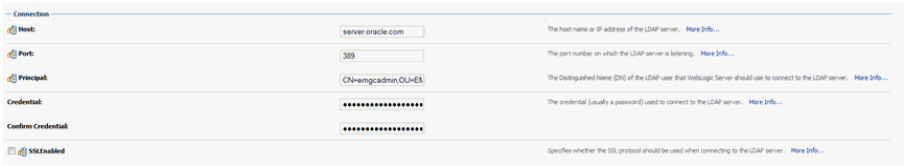

13. In Users section set the User Base DN to the value provided by your Active Directory Administrator. This is the Group or Organization Unit which will have access to Enterprise Manager. To restrict access to a specific set of users, you must use an Organization Unit.

    User Base dn: cn=users,dc=Cloudcontrol,dc=local

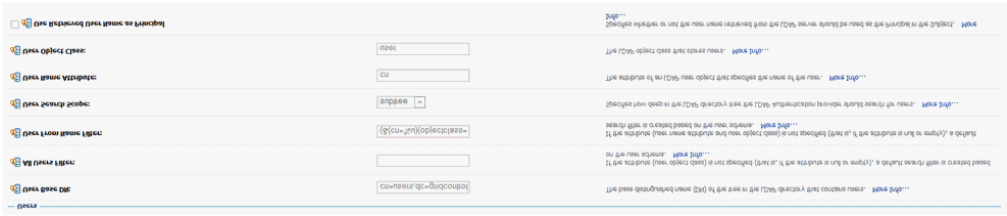

    Note: This information must be obtained from the AD Administrator

14. If you want to use the login name instead of the Account Name (which is typically First Last) then you'll need to set the User From Name Filter and User Name Attribute as follows:

    User Name Attribute: sAMAccountName

    User From Name Filter: (&(sAMAccountName=%u)(objectclass=user))

15. In Groups section

    Group Base dn: cn=Users,dc=Cloudcontrol,dc=local

    Note: This information must be obtained from the AD Administrator

16. In General section:

    Click Propagate Cause For Login Exception

17. Click Save

18. Back on Providers, Click Reorder and move your new provider to the top of the list.

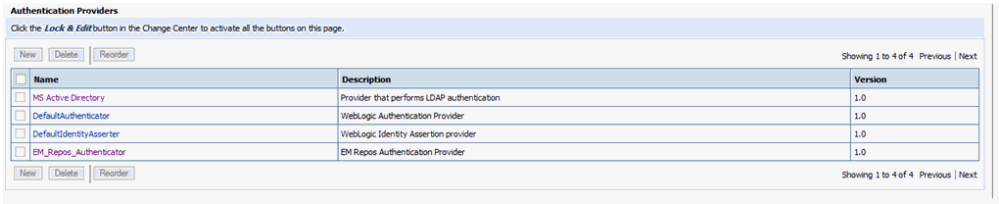

19. Click Apply & Activate Changes

20. There are two options to provision users to Enterprise Manager. You can set a flag to auto-provision all users, or you can manually create them as external users using EMCLI.

    1. To set Auto Provisioning to true

       $ bin/emctl set property -name "em.security.auth.autoprovisioning" -value "true"

       Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.0.0

       Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.

SYSMAN password:

Property em.security.auth.autoprovisioning for oms server.oracle.com:4890_Management_Service has been set to value true

OMS restart is required to reflect the new property value

2. If you don't want all users created automatically, you must manually create using EMCLI (after restart)

   $ bin/emcli create_user -name='TEST' -type='EXTERNAL_USER'

   User "TEST" created successfully

21. Restart OMS

    $ bin/emctl stop oms -all

    $ bin/emctl start oms

22. The users will not show up in Enterprise Manager Administrators until they have logged in once.

11.2.5.2 Manage Active Directory Users with External Roles

To assign a group of privileges to the LDAP users, you can create an external role with the same name as the LDAP group. Once the users are authenticated, they will inherit the permissions and privileges granted to the external role automatically.

1. Create a Group in Active Directory and assign users to the group
2. In Enterprise Manager, click Setup -> Security -> Roles
3. Click Create
4. Enter the name of the Active Directory group and a description
5. Check the box for External, click Next
6. Assign additional Roles, click Next
7. Assign target privileges, click Next
8. Assign resource privileges, click Next
9. Review and click Finish

11.2.5.3 Password Management for Active Directory Users

Password management for Active Directory users must be handled through Active Directory. Password changes are not allowed via Enterprise Manager or WebLogic Server.

11.2.5.4 Remove Active Directory Users

An Active Directory user must be deleted from Enterprise Manager to remove access to Cloud Control. If the user remains in Active Directory, they should be removed from any Groups assigned privileges through External Roles to ensure they cannot login again if auto-provisioning is enabled.

11.2.5.5 Remove Active Directory Authentication

Removing Active Directory authentication will remove all Active Dirctory user accounts from Enterprise Manager.

1. Login to WebLogic Server console.

2. Under Domain Structure on the left, click on Security Realms.
3. Click on myrealm and then click on the tab Providers.
4. Click Lock & Edit to enable editing.
5. Click on the NT Authenticator provider.
6. Click Delete.
7. Save and Activate.
8. Restart OMS.

11.3 Enterprise Manager Authorization

Giving the same level of access to all systems to all administrators is dangerous, but individually granting access to tens, hundreds, or even thousands of targets to every new member of the group is time consuming. With Enterprise Manager's administrator privileges and roles feature, this task can be performed within seconds, instead of hours. Authorization controls the access to the secure resources managed by Enterprise Manager via system, target, and object level privileges and roles.

This section describes Enterprise Manager's Authorization model including user classes, roles, and privileges assigned to each user class. The following topics are described:

- Classes of Users
- Privileges and Roles

11.3.1 Authentication Scheme

An authentication scheme is the type of authentication supported by a target type. For example, a host can support a username/password-based authentication, Public Key authentication or Kerberos authentication. In fact, each target type in an enterprise may support different authentication schemes. To accommodate the many authentication schemes that can exist in a managed environment, Enterprise Manger allows you to configure the credentials for these authentication schemes as well.

11.3.2 Classes of Users

Oracle Enterprise Manager supports different classes of Oracle users, depending upon the environment you are managing and the context in which you are using Oracle Enterprise Manager.

The Enterprise Manager administrators you create and manage in the Could Control console are granted privileges and roles to log in to the Cloud Control console and to manage specific target types and to perform specific management tasks. The default super administrator for the Cloud Control Console is the SYSMAN user, which is a database user associated with the Oracle Management Repository. You define the password for the SYSMAN account during the Enterprise Manager installation procedure.

By restricting access to privileged users and providing tools to secure communications between Oracle Enterprise Manager 12c components, Enterprise Manager protects critical information in the Oracle Management Repository.

The Management Repository contains management data that Enterprise Manager uses to help you monitor the performance and availability of your entire enterprise. This data provides you with information about the types of hardware and software you

have deployed, as well as the historical performance and specific characteristics of the applications, databases, applications servers, and other targets that you manage. The Management Repository also contains information about the Enterprise Manager administrators who have the privileges to access the management data.

You can create and manage Enterprise Manager administrator accounts. Each administrator account includes its own login credentials, as well as a set of roles and privileges that are assigned to the account. There are three administrator access categories:

- **Super Administrator**: Powerful Enterprise Manager administrator with full access privileges to all targets and administrator accounts within the Enterprise Manager environment. The Super Administrator, SYSMAN is created by default when Enterprise Manager is installed. The Super Administrator can create other administrator accounts.
- **Administrator**: Regular Enterprise Manager administrator.
- **Repository Owner**: Database administrator for the Management Repository. This account cannot be modified, duplicated, or deleted.

The types of management tasks that the administrator can perform and targets that he can access depends on the roles, system privileges, and target privileges that he is granted. The Super Administrator can choose to let certain administrators perform only certain management tasks, or access only certain targets, or perform certain management tasks on certain targets. In this way, the Super Administrator can divide the workload among his administrators.

11.3.3 Privileges and Roles

User privileges provide a basic level of security in Enterprise Manager. They are designed to control user access to data and to limit the kinds of SQL statements that users can execute. When creating a user, you grant privileges to enable the user to connect to the database, to run queries and make updates, to create schema objects, and more.

When Enterprise Manager is installed, the SYSMAN user (super administrator) is created by default. The SYSMAN Super Administrator then creates other administrator accounts for daily administration work. The SYSMAN account should only be used to perform infrequent system wide, global configuration tasks.

The Super Administrator divides workload among his administrators by filtering target access, or filtering access to management task, or both through the roles, System Privileges, and Target Privileges he grants them. For example, he can allow some administrators to view any target and to add any target in the enterprise and other administrators to only perform specific operations such as maintaining and cloning on a target for which they are responsible.

A role is a collection of Enterprise Manager resource privileges, or target privileges, or both, which you can grant to administrators or to other roles. These roles can be based upon geographic location (for example, a role for Canadian administrators to manage Canadian systems), line of business (for example, a role for administrators of the human resource systems or the sales systems), or any other model. Administrators do not want to perform the task of individually granting access to tens, hundreds, or even thousands of targets to every new member of their group.

By creating roles, an administrator needs only to assign the role that includes all the appropriate privileges to his team members instead of having to grant many individual privileges. He can divide workload among his administrators by filtering target access, or filtering access to management task, or both.

Out-of-Box Roles: Enterprise Manager Cloud Control 12c comes with predefined roles to manage a wide variety of resource and target types. The following table lists these roles along with their function.

Table 11–1 Out-of-the-Box Roles

Roles	Description
EM_ALL_ ADMINISTRATOR	Role has privileges to perform Enterprise Manager administrative operations. It provides Full privileges on all secure resources (including targets)
EM_ALL_DESIGNER	Role has privileges to design Enterprise Manager operational entities such as Monitoring Templates.
EM_ALL_OPERATOR	Role has privileges to manage Enterprise Manager operations.
EM_ALL_VIEWER	Role has privileges to view Enterprise Manager operations.
EM_CBA_ADMIN	Role to manage Chargeback Objects. It gives the capability to create and view chargeback plans, chargeback consumers, assign chargeback usage, and view any CaT targets.
EM_CLOUD_ ADMINISTRATOR	Enterprise Manager user for setting up and managing the infrastructure cloud. This role could be responsible for deploying the cloud infrastructure (servers, pools, zones) and infrastructure cloud operations for performance and configuration management.
EM_COMPLIANCE_ DESIGNER	Role has privileges for create, modify and delete compliance entities.
EM_COMPLIANCE_ OFFICER	Role has privileges to view compliance framework definition and results.
EM_CPA_ADMIN	Role to manage Consolidation Objects. It gives the capability to create and view consolidation plans, consolidation projects and view any CaT targets.
EM_HOST_DISCOVERY_ OPERATOR	Role has privileges to execute host discovery
EM_INFRASTRUCTURE_ ADMIN	Role has privileges to manage the Enterprise Manager infrastructure such as managing plugin lifecycle or managing self update.
EM_PATCH_ ADMINISTRATOR	Role for creating, editing, deploying, deleting and granting privileges for any patch plan.
EM_PATCH_DESIGNER	Role for creating and viewing for any patch plan
EM_PATCH_OPERATOR	Role for deploying patch plans
EM_PLUGIN_AGENT_ ADMIN	Role to support plug-in lifecycle on Management Agent
EM_PLUGIN_OMS_ ADMIN	Role to support plug-in lifecycle on Management Server
EM_PLUGIN_USER	Role to support view plug-in console
EM_PROVISIONING_ DESIGNER	Role has privileges for provisioning designer
EM_PROVISIONING_ OPERATOR	Role has privileges for provisioning operator
EM_SSA_ ADMINISTRATOR	EM user with privilege to set up the Self Service Portal. This role can define quotas and constraints for self service users and grant them access privileges.

Table 11–1 (Cont.) Out-of-the-Box Roles

Roles	Description
EM_SSA_USER	This role grants EM user the privilege to access the Self Service Portal.
EM_TARGET_ DISCOVERY_OPERATOR	Role has privileges to execute target discovery.
EM_TC_DESIGNER	Role has privileges for creating Template Collections
EM_USER	Role has privilege to access Enterprise Manager Application.
PUBLIC	PUBLIC role is granted to all administrators. This role can be customized at site level to group privileges that need to be granted to all administrators.

Public Role: Enterprise Manager creates one role by default called **Public**. This role is unique in that it is automatically assigned to all new non-super administrators when they are created. By default it has no privileges assigned to it. The Public role should be used to define default privileges you expect to assign to a majority of non-super administrators you create. Privileges need not be assigned to Public initially - they can be added at any time. The role may be deleted if your enterprise does not wish to use it. If deleted, it can be added back in later if you later decide to implement it.

11.3.3.1 Granting Privileges

A privilege is a right to perform management actions within Enterprise Manager. Privileges can be divided into two categories:

- Target Privileges
- Resource Privileges

Target Privileges: These privileges allow an administrator to perform operations on a target. The Target Privileges page shows a list of targets for which privileges can be granted. Select the check box to specify the privileges that are to be granted and click **Next**.

Table 11–2 Target Privileges Applicable to All Targets

Privilege Name	Privilege Display Name	Description
FULL_ANY_TARGET	Full any Target	Ability to do all operations on all the targets, including delete the target
PERFORM_OPERATION_ AS_ANY_AGENT	Execute Command as any Agent	Execute any OS Command as the Agent User at any Agent
PUT_FILE_AS_ANY_ AGENT	Put File as any Agent	Put any File to any Agent's Filesystem as the Agent User
PERFORM_OPERATION_ ANYWHERE	Execute Command Anywhere	Execute any OS Command at any Agent
OPERATOR_ANY_TARGET	Operator any Target	Privilege to grant operator access on all targets
CONNECT_ANY_VIEW_ TARGET	Connect to any viewable target	Ability to connect and manage any of the viewable target
USE_ANY_BEACON	Use any beacon	Ability to register with any Beacon

Table 11–2 (Cont.) Target Privileges Applicable to All Targets

Privilege Name	Privilege Display Name	Description
EM_MONITOR	EM Monitor	Ability to view any EM Repository targets
VIEW_ANY_TARGET	View any Target	Ability to view any target
GRANT_VIEW_ORACLE_VM_MANAGER	Grant View Oracle VM Manager Privilege	Ability to grant View Oracle VM Manager privilege
GRANT_VIEW_ORACLE_VM_ZONE	Grant View Zone Privilege	Ability to grant View Zone privilege
GRANT_VIEW_ORACLE_CLOUD_ZONE	Grant View Database Zone Privilege	Ability to grant view privilege on Database Zone targets
CREATE_PROPAGATING_GROUP	Create Privilege Propagating Group	Ability to create privilege propagating groups.Privileges granted on a privilege propagating group will be automatically granted on the members of the group
CREATE_TARGET	Create Target	Ability to create a target

Table 11–3 Target Privileges Applicable to Specific Targets

Privilege Name	Privilege Display Name	Description
GROUP_ADMINISTRATION	Group Administration	Ability to administor groups
FULL_TARGET	Full Target	Ability to do all operations on the target, including delete the target
FMW_DEPLOY_APP_TARGET	Deploy Fusion Middleware	Ability to deploy Fusion Middleware components
CONNECT_READONLY_TARGET	Connect Target Readonly	Ability to connect to target in readonly mode
CONNECT_TARGET	Connect Target	Ability to connect and manage target
MANAGE_TARGET_COMPLIANCE	Manage Target Compliance	Ability to manage compliance of the target
PERFORM_OPERATION_AS_AGENT	Execute Command as Agent	Execute any OS Command as the Agent User
PUT_FILE_AS_AGENT	Put File as Agent	Put any File to the Agent's Filesystem as the Agent User
MANAGE_TARGET_ALERTS	Manage Target Events	Ability to clear events, re-evaluate metric alert events, create incidents, add events to incidents, and define what actions the administrator can perform on individual incidents, such as acknowledgment or escalation.
PERFORM_OPERATION	Execute Command	Execute any OS Command
CONFIGURE_TARGET	Configure target	Ability to edit target properties and modify monitoring configuration
MANAGE_TARGET_PATCH	Manage Target Patch	Privilege to Analyze, Apply and Rollback patches on the target

Table 11–3 (Cont.) Target Privileges Applicable to Specific Targets

Privilege Name	Privilege Display Name	Description
MANAGE_TC_ OPERATION	Manage Template Collection Operations	Ability to associate a template collection to a adminimistration group and Sync targets with the associated template collections.
MANAGE_TARGET_ METRICS	Manage Target Metrics	Ability to edit threshold for metric and policy setting, apply monitoring templates, and manage User Defined Metrics
BLACKOUT_TARGET	Blackout Target	Ability to create, edit, schedule and stop a blackout on the target
OPERATOR_TARGET	Operator Target	Ability to do normal administrative operations on the target, such as configure a blackout and edit the target properties
FMW_OPERATOR_PRIV	Operator Fusion Middleware	"Ability to perform operations, such as start and shutdown and view logs for Fusion Middleware targets
FMW_PROCESS_ CONTROL_TARGET	Process Control Fusion Middleware	Ability to start or shutdown Fusion Middleware target
FMW_VIEW_LOG_DATA_ TARGET	View Fusion Middleware logs	Ability to view Fusion Middleware diagnostics data
VIEW_ORACLE_CLOUD_ ZONE	View Database Zone	Ability to view Database Zone
VIEW_ORACLE_VM_ MANAGER	View Oracle VM Manager	Ability to view Oracle VM Manager
VIEW_ORACLE_VM_ ZONE	View Oracle VM Zone	Ability to view Oracle VM Zone
VIEW_TARGET	View Target	Ability to view properties, inventory and monitor information about a target

Resource: These privileges allow a user to perform operations against specific types of resources. To set Resource Privileges, from the **Setup** menu, choose **Administrators**. Select an administrator from the list and click **Edit**. The Edit Administrator wizard is displayed. Click **Next** to navigate through the wizard to see the System Privileges page. The following table lists all available resource privileges.

Resource Type	Display Name	Description	Privileges Required to Grant
ACCESS	Access Enterprise Manager	Ability to access Enterprise Manager interfaces	ACCESS
AD4J	JVM Diagnostics User	Gives capability to view the JVM Diagnostic data	SUPER_USER

Resource Type	Display Name	Description	Privileges Required to Grant
AD4J	JVM Diagnostics Administrator	Gives capability to manage all JVM Diagnostic Administrative operations	SUPER_USER
ASREPLAY_ ENTITY_MGMT	Application Replay Operator	View, create, and edit any Application Replay entity.	SUPER_USER
ASREPLAY_ ENTITY_MGMT	Application Replay Viewer	View any Application Replay entity.	SUPER_USER
BTM	Request Monitoring User	Gives capability to view the Request Monitoring Data	SUPER_USER
BTM	Request Monitoring Administrator	Gives capability to manage all Request Monitoring Administrative Operations	SUPER_USER
CA	Full Corrective Action	Internal privilege, not for granting	
CA	View Corrective Action	Internal privilege, not for granting	VIEW
CCS_SECURE_ CLASS	Manage custom configurations owned by any user	Ability to create new and edit/delete Custom Configuration specification owned by any user	
CCS_SECURE_ CLASS	Manage custom configurations owned by the user	Ability to create new and edit/delete Custom Configuration specification owned by the user	
CHANGE_PLAN	Manage change plans	Create and delete Change Manager Change Plans	FULL
CHANGE_PLAN	Edit change plan	Edit a Change Manager Change Plan	EDIT
CHANGE_PLAN	View change plan	View a Change Manager Change Plan	VIEW
CHARGEBACK_ AND_ CONSOLIDATION	Manage Chargeback Plans	Ability to Create and Modify Chargeback Plans.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Manage Any Consolidation Plan	Ability to Manage any Consolidation Plans.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	View Chargeback and Consolidation Target	Ability to View Chargeback and Consolidation Target.	SUPER_USER

Resource Type	Display Name	Description	Privileges Required to Grant
CHARGEBACK_ AND_ CONSOLIDATION	View Any Chargeback and Consolidation Target	Ability to View Any Chargeback and Consolidation Target.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Manage Chargeback and Consolidation Target	Ability to Manage a Chargeback and Consolidation Target.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Manage Any Chargeback and Consolidation Target	Ability to Add/Delete Target to Chargeback and Assign Chargeplan to Target or Add Target to Consolidation Project.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Setup Chargeback and Consolidation	Ability to Setup CAT.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	View Any Chargeback Consumers	Ability to View Any Chargeback Consumers.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Manage Chargeback Consumers	Ability to Create and Modify Chargeback Consumers.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Assign Chargeback Usage	Ability to Assign Chargeback Usage to Consumers.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Assign Chargeback Plan	Ability to Assign Chargeback Plan to CAT Targets.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	View Any Chargeback Plan	Ability to view all the Chargeback Plans.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	View Any Consolidation Plan	Ability to view the Consolidation Plans.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	View Any Consolidation Project	Ability to View any Consolidation Project.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Manage Any Consolidation Project	Ability to Manage any Consolidation Project.	SUPER_USER
CLOUDPOLICY	Full Policy	Privilege required to View, Modify, Delete a Policy	FULL
CLOUDPOLICY	Modify Policy	Ability to Modify a Policy	EDIT
CLOUDPOLICY	View Policy	Ability to View a Policy	VIEW
CLOUDPOLICY	View any Policy	Ability to View any Policy	VIEW
CLOUDPOLICY	Create any Policy	Ability to Create any Policy	CREATE

Resource Type	Display Name	Description	Privileges Required to Grant
CLOUDPOLICYGROUP	Full Policy Group	Privilege required to View, Modify, Delete a Policy Group	FULL
CLOUDPOLICYGROUP	Modify Policy Group	Ability to Modify a Policy Group	EDIT
CLOUDPOLICYGROUP	View Policy Group	Ability to View a Policy Group	VIEW
CLOUDPOLICYGROUP	View any Policy Group	Ability to View any Policy Group	VIEW
CLOUDPOLICYGROUP	Create Policy Group	Ability to Create Policy Group	CREATE
COMPLIANCE_FWK	Create Compliance Entity	Ability to create compliance framework, standard, rules	CREATE
COMPLIANCE_FWK	Full any Compliance Entity	Ability to edit/delete compliance framework, standard, rules	FULL
COMPLIANCE_FWK	View any Compliance Framework	Ability to view compliance framework definition and results	VIEW
DISCOVERY	Can Scan Network	Privilege to create, edit and delete host discovery configuration	
DISCOVERY	View Any Discovered Hosts	Privilege to view any discovered hosts	
DISCOVERY	View Any Discovered Targets On Host	Privilege to view any discovered targets on host	
DP	Grant full privilege	Ability to grant upto full privilege on deployment procedures.	GRANT
DP	Grant launch privilege	Ability to grant launch privilege on deployment procedures.	GRANT
DP	Import	Ability to create deployment procedures and ability to import/export customized deployment procedures.	CREATE

Resource Type	Display Name	Description	Privileges Required to Grant
DP	Full	Ability to perform launch, create like, edit structure and delete operations on a Deployment Procedure.	GRANT_FULL_DP
DP	Create	Ability to create deployment procedures.	CREATE
DP	Launch	Ability to perform launch and create like operations on a Deployment Procedure.	GRANT_LAUNCH_DP
EMHA_SECURE_CLASS	Enterprise Manager High Availability Administration	Gives access to manage Enterprise Manager High Availability	ADMIN
EVENT	Manage Events	Manage events privilege object	MANAGE_EVENT
EVENT	View Events	View events privilege object	VIEW
FMW_DIAG_SEC_CLASS	Create Object	Ability to manage the offline diagnostic object lifecycle	SUPER_USER
FMW_DIAG_SEC_CLASS	View object	Ability to view the offline diagnostics objects	SUPER_USER
ISSUE	Manage Problems	Manage problems privilege object	MANAGE_PROBLEM
ISSUE	Manage Incidents	Manage incidents privilege object	MANAGE_INCIDENT
ISSUE	View Issues - (Incidents and Problems)	View issues - Incidents and Problems privilege object	VIEW
JOB	Full	Ability to perform all the valid operations on job, library job, deployment procedure configuration and on deployment procedure instance.	FULL
JOB	Grant view privilege	Ability to grant view privilege on jobs.	GRANT

Resource Type	Display Name	Description	Privileges Required to Grant
JOB	Manage	Ability to perform various operations except edit and delete on job, library job, deployment procedure configuration and on deployment procedure instance.	EDIT
JOB	View	Ability to view, do create like on a job, launch deployment procedure configuration and view deployment procedure instance.	GRANT_VIEW_JOB
JOB	Create	Ability to submit jobs, create library jobs, create deployment procedure instance and create deployment procedure configuration.	CREATE
MEXT_SECURE_CLASS	Full MEXT	Gives complete access to edit, and delete metric extension object	
MEXT_SECURE_CLASS	Edit MEXT	Can edit or create the next version of a metric extension object, but cannot delete it	
MEXT_SECURE_CLASS	Create New Metric Extension	Create or import new metric extensions	
NAMED_CREDENTIALS	Create new Named Credential	Ability to create new named credentials	
NAMED_CREDENTIALS	View Credential	View Credential	
NAMED_CREDENTIALS	Edit Credential	User can update credential but cannot delete it.	
NAMED_CREDENTIALS	Full Credential	Full Credential	
PATCH	Privileges for Patch Setup	Privilege to grant privileges any Patching plan object	
PATCH	Manage privileges on any Patching Plan	Privilege to grant or revoke privileges on any Patching plan object	MANAGE

Resource Type	Display Name	Description	Privileges Required to Grant
PATCH	Full privileges on any Patching Plan	Privilege to view, modify, execute and delete any Patching plan object	FULL
PATCH	Manage privileges on a Patching Plan	Privilege to grant or revoke privileges on a Patching plan object	MANAGE
PATCH	View any Patching Plan	Privilege to view any Patching plan object	VIEW
PATCH	Full Patch Plan	Privilege to view, modify, execute and delete a Patching plan object	MANAGE_PRIV_ ANY_PATCH_PLAN
PATCH	View any Patching Plan Template	Privilege to view any Patching Plan Template object	VIEW
PATCH	Create Patch Plan	Privilege for creating a Patching Plan object	
PATCH	View Patching Plan	Privilege to View a Patching Plan Object	MANAGE_PRIV_ ANY_PATCH_PLAN
PATCH	Create Patch Plan Template	Privilege for creating a Patching Plan Template object	
PLUGIN	Plug-in view privilege	Gives access to manage Enterprise Manager plug-in life cycle console	USER
PLUGIN	Plug-in Agent Administrator	Gives access to manage Enterprise Manager plug-in on Agent	ADMIN
PLUGIN	Plug-in OMS Administrator	Gives access to manage Enterprise Manager plug-in on Management Server	ADMIN
REPORT_DEF	View Report	Ability to view report definition and stored reports, generate on demand reports and do a create like	VIEW
REPORT_DEF	Publish Report	Ability to publish reports for public viewing	
RULESET_SEC	Edit Business Ruleset	Edit Business Ruleset	EDIT
RULESET_SEC	Create Business Ruleset	Create Business Ruleset	CREATE
SBRM_BACKUP_ CONFIG	Create Backup Configuration	Ability to create a backup configuration.	SUPER_USER

Resource Type	Display Name	Description	Privileges Required to Grant
SBRM_BACKUP_CONFIG	Use Backup Configuration	Ability to use a backup configuration.	SUPER_USER
SBRM_BACKUP_CONFIG	Edit Backup Configuration	Ability to edit a backup configuration.	SUPER_USER
SBRM_BACKUP_CONFIG	Full Access	Full access to a backup configuration.	SUPER_USER
SELFUPDATE_SECURE_CLASS	Self Update Administrator	Gives access to manage Enterprise Manager Update	FULL
SELFUPDATE_SECURE_CLASS	View any Enterprise Manager Update	Gives access to view any Enterprise Manager Update	VIEW
SSA	Access Cloud Self Service Portal	Users with this privilege have access to Cloud Self Service Portal.	SUPER_USER
SSA	Setup Cloud Self Service Portal	Privilege to perform Cloud Self Service Portal setup like defining quotas for roles, publishing assemblies etc.	SUPER_USER
SWLIB_ADMINISTRATION	Software Library Storage Administration	Ability to manage upload and reference file storage locations, import and export entities, and purge deleted entities	FULL
SWLIB_ENTITY_MGMT	View any Assembly Entity	View any Assembly Entity	SWLIB_GRANT_ANY_ENTITY_PRIV
SWLIB_ENTITY_MGMT	View any Template Entity	View any Template Entity	SWLIB_GRANT_ANY_ENTITY_PRIV
SWLIB_ENTITY_MGMT	Grant Any Entity Privilege	Ability to grant view, edit and delete privilege on any Software Library entity. This privilege is required if the user granting the privilege on an entity is not a super administrator or owner of the entity.	GRANT
SWLIB_ENTITY_MGMT	Manage Entity	Ability to view, edit and delete a Software Library entity	SWLIB_GRANT_ANY_ENTITY_PRIV
SWLIB_ENTITY_MGMT	View Software Library Entity	Ability to view a Software Library entity	SWLIB_GRANT_ANY_ENTITY_PRIV

Resource Type	Display Name	Description	Privileges Required to Grant
SWLIB_ENTITY_MGMT	Edit an Software Library Entity	Ability to edit a Software Library entity	SWLIB_GRANT_ANY_ENTITY_PRIV
SWLIB_ENTITY_MGMT	Create Any Software Library Entity	Ability to create any Software Library entity	CREATE
SWLIB_ENTITY_MGMT	View Any Software Library Entity	Ability to view any Software Library entity	VIEW
SWLIB_ENTITY_MGMT	Edit Any Software Library Entity	Ability to edit any Software Library entity	EDIT
SWLIB_ENTITY_MGMT	Manage Any Software Library Entity	Ability to create, view, edit and delete any Software Library entity	FULL
SWLIB_ENTITY_MGMT	Import Any Software Library Entity	Ability to import any Software Library entity from a Provisioning Archive (PAR) file	IMPORT
SWLIB_ENTITY_MGMT	Export Any Software Library Entity	Ability to view and export any Software Library entity to a Provisioning Archive (PAR) file	EXPORT
SYSTEM	Super User	Provides all the privileges to any target in the system	
TEMPLATE	View Template	Ability to access a template and apply it to any target on which you have Manage Target Metrics	
TEMPLATE	View Template	Ability to view a template and apply it to any target on which you have Manage Target Metrics	VIEW
TEMPLATE	View any Monitoring Template	View any Monitoring Template.	VIEW
TEMPLATECOLLECTION	Full Template Collection	Ability to edit and delete Template Collection	FULL
TEMPLATECOLLECTION	View Template Collection	Ability to view Template Collection	VIEW
TEMPLATECOLLECTION	View any Template Collection	Ability to view any Template Collection	VIEW
TEMPLATECOLLECTION	Create any Template Collection	Ability to create any Template Collection	CREATE

Select the check box to select the resource privilege to be granted to the administrator and click **Next**.

11.4 Configuring Secure Communication (SSL) for Cloud Contro

This section contains the following topics:

11.4.1 About Enterprise Manager Framework Security

Enterprise Manager Framework Security provides safe and secure communication channels between the components of Enterprise Manager. For example, Framework Security provides secure connections between your Oracle Management Service and its Management Agents.

> **See Also:** *Oracle Enterprise Manager Concepts* for an overview of Enterprise Manager components

Enterprise Manager Framework Security implements the following types of secure connections between the Enterprise Manager components:

- HTTPS and Public Key Infrastructure (PKI) components, including signed digital certificates, for communications between the Management Service and the Management Agents.

  > **See Also:** *Oracle Security Overview* for an overview of Public Key Infrastructure features, such as digital certificates and public keys

- Oracle Advanced Security for communications between the Management Service and the Management Repository.

  > **See Also:** *Oracle Database Advanced Security Administrator's Guide*

11.4.2 Enabling Security for the Oracle Management Service

To enable Enterprise Manager Framework Security for the Management Service, you use the `emctl secure oms` utility, which is located in the following subdirectory of the Management Service home directory:

```
ORACLE_HOME/bin
```

The `emctl secure oms` utility performs the following actions:

- Generates a Root Key within your Management Repository. The Root Key is used during distribution of Oracle Wallets containing unique digital certificates for your Management Agents.
- Modifies your WebTier to enable an HTTPS channel between your Management Service and Management Agents, independent from any existing HTTPS configuration that may be present in your WebTier.
- Enables your Management Service to accept requests from Management Agents using Enterprise Manager Framework Security.

To run the `emctl secure oms` utility you must first choose an Agent Registration Password. The Agent Registration password is used to validate that future installation of Oracle Management Agents are authorized to load their data into this Enterprise Manager installation.

To enable Enterprise Manager Framework Security for the Oracle Management Service:

1. Stop the Management Service, the WebTier, and the other application server components using the following command:

   ```
   OMS_ORACLE_HOME/bin/emctl stop oms
   ```

2. Enter the following command:

   ```
   OMS_ORACLE_HOME/bin/emctl secure oms
   ```

3. You will be prompted for the Enterprise Manager Root Password. Enter the `SYSMAN` password.
4. You will be prompted for the Agent Registration Password, which is the password required for any Management Agent attempting to secure with the Management Service. Specify an Agent Registration Password for the Management Service.
5. Restart the OMS.
6. After the Management Service restarts, test the secure connection to the Management Service by browsing to the following secure URL using the HTTPS protocol:

   `https://`*`hostname.domain`*`:https_console_port/em`

   Note: The Enterprise Manager console URL can be found by running the "emctl status oms -details" command.

   For example:

   ```
   https://mgmthost1.acme.com:7799/em
   ```

   If the Management Service security has been enabled, your browser displays the Enterprise Manager Login page.

Example 11–1 Sample Output of the emctl secure oms Command

```
emctl secure oms
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
Securing OMS... Started.
Securing OMS... Successful
```

Example 11–2 Usage of the emctl secure oms Command (II)

```
emctl secure oms [-sysman_pwd <sysman password>] [-reg_pwd <registration
password>] [-host <hostname>] [-slb_port <slb port>] [-slb_console_port <slb
console port>] [-reset] [-console] [-lock] [-lock_console] [-secure_port <secure_
port>] [-upload_http_port <upload_http_port>] [-root_dc <root_dc>] [-root_country
<root_country>] [-root_email <root_email>] [-root_state <root_state>] [-root_loc
<root_loc>] [-root_org <root_org>] [-root_unit <root_unit>] [-wallet <wallet_loc>
-trust_certs_loc <certs_loc>] [-key_strength <strength>] [-cert_validity
<validity>] [-protocol <protocol>] [-force_newca] [-ms_hostname <Managed Server
hostname>] [-sign_alg <md5|sha1|sha256|sha384|sha512>]
```

Valid values for <protocol> are the allowed values for Apache's SSLProtocol directive

The parameters are explained below:

- `sysman_pwd` - Oracle Management Repository user password.
- `reg_pwd` - The Management Agent registration password.
- `host` - The host name to be used in the certificate used by the Oracle Management Service. You may need to use the SLB host name if there is an SLB before the Management Service.
- `reset` - A new certificate authority will be created. All the Agents and Oracle Management Services need to be resecured.
- `secure_port` - Specify this to change HTTPS Upload port on WebTier
- `upload_http_port` - Specify this to change HTTP Upload port on WebTier
- `slb_port` - This parameter is required when Server Load Balancer is used. It specifies the secure upload port configured in the Server Load Balancer.
- `slb_console_port` - This parameter is required when Server Load Balancer is used. It specifies the secure console port configured in the Server Load Balancer.
- `root_dc` - The domain component used in the root certificate. The default value is com.
- `root_country` - The country to be used in the root certificate. The default value is **US**.
- `root_state` - The state to be used in the root certificate. The default value is **CA**.
- `root_loc` - The location to be used in the root certificate. The default value is **EnterpriseManager on <hostname>**.
- `root_org` - The organization name to be used in the root certificate. The default value is EnterpriseManager on <hostname>.
- `root_unit` - The organizational unit to be used in the root certificate. The default value is EnterpriseManager on <hostname>.
- `root_email` - The email address to be used in the root certificate. The default value is EnterpriseManager@<hostname>.
- `wallet`: This is the location of the wallet containing the third party certificate. This parameter should be specified while configuring third party certificates.
- `trust_certs_loc` - The location of the `trusted_certs.txt` (required when third party certificates are used).
- `key_strength`: The strength of the key to be used. Valid values are 512, 1024, 2048, and 4096.

- `cert_validity`: The number of days for which the self-signed certificate is valid. The valid range is between 1 to 3650.
- `protocol`: This parameter is used to configure Oracle Management Service in TLSv1-only or SSLv3-only or mixed mode (default). Valid values are the allowed values as per **Apache's SSLProtocol** directive.

> **Note:** The `key_strength` and `cert_validity` parameters are applicable only when the `-wallet` option is not used.

- `force_newca` - If specified, any Agents that are still configured with an older Certificate Authority are ignored.
- `ms_hostname` - Managed Server's hostname.
- `sign_alg` - Signature algorithm.
- `lock`: Locks the Upload
- `lock_console`: Locks the Console
- `console`: If specified, certificate is re-created for HTTPS Console port as well

11.4.2.1 Creating a New Certificate Authority

You may need to create a new Certificate Authority (CA) if the current CA is expiring or if you want to change the key strength. A unique identifier is assigned to each CA. For instance, the CA created during installation may have an identifier as ID 1, subsequent CAs will have the IDs 2,3, and so on. At any given time, the last created CA is active and issues certificates for OMSs and Agents.

Example 11–3 Creating a New Certificate Authority

```
emctl secure createca [-sysman_pwd <pwd>] [-host <hostname>] [-key_
strength<strength>] [-cert_validity <validity>] [-root_dc <root_dc>] [-root_
country <root_country>] [-root_email <root_email>] [-root_state <root_state>]
[-root_loc <root_loc>] [-root_org <root_org>] [-root_unit <root_unit>]
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
Creating CA... Started.
Successfully created CA with ID 2
```

Example 11–4 Viewing Information about a Certificate Authority

```
emcli get_ca_info -ca_id="1;2" -details
Info about CA with ID: 1
CA is not configured
DN: CN=myhost.mydomain.com, C=US
Serial# : 3423643907115516586
Valid From: Tue Mar 16 11:06:20 PDT 2011
Valid Till: Sat Mar 14 11:06:20 PDT 2020
Number of Agents registered with CA ID 1 is 1
myhost.mydomain.com:3872

Info about CA with ID: 2
CA is configured
DN: CN=myhost.mydomain.com, C=US, ST=CA
Serial# : 1182646629511862286
Valid From: Fri Mar 19 05:17:15 PDT 2011
Valid Till: Tue Mar 17 05:17:15 PDT 2020
```

```
There are no Agents registered with CA ID 2
```

The WebLogic Administrator and Node Manager passwords are stored in the Administration Credentials Wallet. This is present in the `EM_INSTANCE_HOME/sysman/config/adminCredsWallet` directory. To recreate Administrator Credentials wallet, run the following command on each machine on which the Management Service is running:

```
emctl secure create_admin_creds_wallet [-admin_pwd <pwd>]
[-nodemgr_pwd <pwd>]
```

11.4.2.2 Viewing the Security Status and OMS Port Information

To view the security status and OMS port information, use the following command

Example 11–5 emctl status oms -details

```
> emctl status oms -details
Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.
Enter Enterprise Manager Root (SYSMAN) Password : ******
Console Server Host : omshost1.example.com
HTTP Console Port : 7802
HTTPS Console Port : 5416
HTTP Upload Port : 7654
HTTPS Upload Port : 4473
OMS is not configured with SLB or virtual hostname
Agent Upload is locked.
OMS Console is locked.
Active CA ID: 1
Console URL: https://omshost1.example.com:5416/em
Upload URL: https://omshost1.example.com:4473/empbs/upload

WLS Domain Information
Domain Name : EMGC_DOMAIN
Admin Server Host: omshost1.example.com

Managed Server Information
Managed Server Instance Name: EMGC_OMS1
Managed Server Instance Host: omshost1.example.com
```

11.4.2.3 Configuring Transparent Layer Security

The Oracle Management Service can be configured in the following modes:

- **TLSv1-only mode**: To configure the OMS to use only TLSv1 connections, do the following:

  1. Stop the OMS by entering the following command:

     ```
     OMS_ORACLE_HOME/bin/emctl stop oms
     ```

  2. Enter the following command:

     ```
     emctl secure oms -protocol TLSv1
     ```

  3. Append `-Dweblogic.security.SSL.protocolVersion=TLS1` to `JAVA_OPTIONS` in `Domain_Home/bin/startEMServer.sh/cmd`. If this property already exists, update the value to TLS1.

4. Restart the OMS with the following command:

   ```
   OMS_ORACLE_HOME/bin/emctl start oms
   ```

- **SSLv3 Only Mode**: To configure the OMS to use SSLv3 connections only, do the following:

  1. Stop the OMS by entering the following command:

     ```
     OMS_ORACLE_HOME/bin/emctl stop oms
     ```

  2. Enter the following command:

     ```
     emctl secure oms -protocol SSLv3
     ```

  3. Append `-Dweblogic.security.SSL.protocolVersion=SSL3` to `JAVA_OPTIONS` in `Domain_Home/bin/startEMServer.sh` or `startEMServer.cmd` on Windows. If this property already exists, update the value to SSL3.

  4. Restart the OMS with the following command:

     ```
     OMS_ORACLE_HOME/bin/emctl start oms
     ```

- **Mixed Mode**: To configure the OMS to use both SSLv3 and TLSv1 connections, do the following:

  1. Stop the OMS by entering the following command:

     ```
     OMS_ORACLE_HOME/bin/emctl stop oms
     ```

  2. Enter the following command:

     ```
     emctl secure oms
     ```

  3. Append `-Dweblogic.security.SSL.protocolVersion=ALL` to `JAVA_OPTIONS` in `Domain_Home/bin/startEMServer.sh`. If this property already exists, update the value to `ALL`.

  4. Restart the OMS with the following command:

     ```
     OMS_ORACLE_HOME/bin/emctl start oms
     ```

Note: By default, the OMS is configured to use the Mixed Mode. To configure the Management Agent in TLSv1 only mode, set `allowTLSOnly=true` in the `emd.properties` file and restart the Agent.

11.4.3 Securing the Oracle Management Agent

When you install the Management Agent on a host, you must identify the Management Service that will be used by the Management Agent. To enable Enterprise Manager Framework Security for the Management Agent, use the `emctl secure agent` utility, which is located in the following directory of the Management Agent home directory:

AGENT_HOME`/bin (UNIX)`
AGENT_HOME`\bin (Windows)`

The `emctl secure agent` utility performs the following actions:

- Obtains an Oracle Wallet from the Management Service that contains a unique digital certificate for the Management Agent. This certificate is required in order for the Management Agent to conduct SSL communication with the secure Management Service.
- Obtains an Agent Key for the Management Agent that is registered with the Management Service.
- Configures the Management Agent so it is available on your network over HTTPS and so it uses the Management Service HTTPS upload URL for all its communication with the Management Service.

To enable Enterprise Manager Framework Security for the Management Agent:

1. Ensure that your Management Service and the Management Repository are up and running.
2. Change directory to the following directory:

   ```
   AGENT_HOME/bin (UNIX)
   AGENT_HOME\bin (Windows)
   ```

3. Stop the Management Agent:

   ```
   emctl stop agent
   ```

4. Enter the following command:

   ```
   emctl secure agent (UNIX)
   emctl secure agent (Windows)
   ```

   The `emctl secure agent` utility prompts you for the Agent Registration Password, authenticates the password against the Management Service, and reconfigures the Management Agent to use Enterprise Manager Framework Security.

   Example 11–6 shows sample output of the `emctl secure agent` utility.

5. Restart the Management Agent:

   ```
   emctl start agent
   ```

6. Confirm that the Management Agent is secure by checking the Management Agent home page.

   Note: You can also check if the Agent Management is secure by running the `emctl status agent -secure` command, or by checking the Agent and Repository URLs in the output of the `emctl status agent` command.

   In the Management Agent home page , the **Secure Upload** field indicates whether or not Enterprise Manager Framework Security has been enabled for the Management Agent.

Example 11–6 Sample Output of the emctl secure agent Utility

```
emctl secure agent
Oracle Enterprise Manager 12c Release 1 Cloud Control.
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
Securing agent...   Started
Securing agent...   Successful.
```

Example 11–7 Sample Output of the emctl status agent secure Command

```
emctl status agent -secure
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
Checking the security status of the Agent at location set in
/private/home/oracle/product/102/em/agent10g/sysman/config/emd.properties...
Done.
Agent is secure at HTTPS Port 3872.
Checking the security status of the OMS at
https://cloudcontrol.oraclecorp.com:4889/em/upload/...  Done.
OMS is secure on HTTPS Port 4888
```

11.4.4 Enabling Security with Multiple Management Service Installations

Because you have already established at least one Agent Registration Password and a Root Key in your Management Repository, they must be used for your new Management Service. Your secure Management Agents can then operate against either Management Service.

All the registration passwords assigned to the current Management Repository are listed on the Registration Passwords page in the Oracle Enterprise Manager 12c Cloud Control Console.

If you install a new Management Service that uses a new Management Repository, the new Management Service is considered to be a distinct enterprise. There is no way for the new Management Service to partake in the same security trust relationship as another Management Service that uses a different Management Repository. Secure Management Agents of one Management Service will not be able to operate against the other Management Service.

11.4.5 Restricting HTTP Access to the Management Service

Note: The Oracle Management Service is locked (both console & upload) by default begining with Enterprise Manager 12c.

It is important that only secure Management Agent installations that use the Management Service HTTPS channel are able to upload data to your Management Repository and Cloud Control console is accessible via HTTPS only.

To restrict access so Management Agents can upload data to the Management Service only over HTTPS:

1. Stop the Management Service, the WebTier, and the other application server components:

   ```
   cd ORACLE_HOME/opmn/bin
   emctl stop oms
   ```

2. Change directory to the following location in the Management Service home:

   ```
   ORACLE_HOME/bin
   ```

3. Enter the following command to prevent Management Agents from uploading data to the Management Service over HTTP:

   ```
   emctl secure lock -upload
   ```

   To lock the console and prevent HTTP access to the console, enter the following command:

```
emctl secure lock -console
```

To lock both, enter either of the following commands:

```
emctl secure lock or
emctl secure lock -upload -console
```

To lock both the console access and uploads from Agents while enabling security on the Management Service, enter the following command:

```
emctl secure oms -lock [other options]
```

4. Restart the Management Service, the WebTier, and the other application server components:

```
cd ORACLE_HOME/bin
emctl start oms
```

5. Verify that you cannot access the OMS upload URL using the HTTP protocol:

   For example, navigate to the following URL:

```
http://hostname.domain:4889/empbs/upload
```

   You should receive an error message similar to the following:

```
Forbidden
You are not authorised to access this resource on the server.
```

6. Verify that you can access the Management Agent Upload URL using the HTTPS protocol:

   For example, navigate to the following URL:

```
https://hostname.domain:4888/empbs/upload
```

   You should receive the following message, which confirms the secure upload port is available to secure Management Agents:

```
Http XML File receiver
Http Recceiver Servlet active!
```

To allow the Management Service to accept uploads from unsecure Management Agents, use the following command:

```
emctl secure unlock -upload
```

Note:

- The OMS need to be stopped before running 'secure unlock', and then restarted afterwards.
- To unlock the console and allow HTTP access to the console, enter the following command:

```
emctl secure unlock -console
```

- To unlock both, enter either of the following command:

```
emctl secure unlock
emctl secure unlock -console -upload
```

Example 11–8 Sample Output of the emctl secure lock Command

```
emctl secure lock
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
OMS Console is locked. Access the console over HTTPS ports.
Agent Upload is locked. Agents must be secure and upload over HTTPS port.
Restart OMS
```

Example 11–9 Sample Output of the emctl secure unlock Command

```
emctl secure unlock
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
OMS Console is unlocked. HTTP ports too can be used to access console.
Agent Upload is unlocked. Unsecure Agents may upload over HTTP.
Restart OMS
```

11.4.6 Managing Agent Registration Passwords

Enterprise Manager uses the Agent Registration password to validate that installations of Oracle Management Agents are authorized to load their data into the Oracle Management Service.

The Agent Registration password is created during installation when security is enabled for the Oracle Management Service. You can add/edit/delete registration passwords directly from the Enterprise Manager console.

Note: If you want to avoid new Agents from being registered with the OMS, delete all registration passwords.'

11.4.6.1 Using the Cloud Control Console to Manage Agent Registration Passwords

You can use the Cloud Control Console to manage your existing registration passwords or create additional registration passwords:

1. From the **Setup** menu, choose **Security** and then **Registration Passwords**.
2. Enterprise Manager displays the Registration Passwords page (Figure 11–3). The registration password you created when you ran the `emctl secure oms` command appears in the Registration Passwords table.
3. Use the Registration Passwords page to change your registration password, create additional registration passwords, or remove registration passwords associated with the current Management Repository.

Figure 11–3 Managing Registration Passwords in the Cloud Control Console

When you create or edit an Agent Registration Password on the Registration Passwords page, you can determine whether the password is persistent and available for multiple Management Agents or to be used only once or for a predefined period of time.

For example, if an administrator requests to install a Management Agent on a particular host, you can create a one-time-only password that the administrator can use to install and configure one Management Agent.

On the other hand, you can create a persistent password that an administrator can use for the next two weeks before it expires and the administrator must ask for a new password.

11.4.6.2 Using emctl to Add a New Agent Registration Password

To add a new Agent Registration Password, use the following `emctl` command on the machine on which the Management Service has been installed:

```
emctl secure setpwd [sysman pwd] [new registration pwd]
```

The `emctl secure setpwd` command requires that you provide the password of the Enterprise Manager super administrator user, `sysman`, to authorize the addition of the Agent Registration Password.

If you change the Agent Registration Password, you must communicate the new password to other Enterprise Manager administrators who need to install new Management Agents, enable Enterprise Manager Framework Security for existing Management Agents, or install additional Management Services.

As with other security passwords, you should change the Agent Registration Password on a regular and frequent basis to prevent it from becoming too widespread.

11.4.7 Configuring the OMS with Server Load Balance

When you deploy a Management Service that is available behind a Server Load Balancer (SLB), special attention must be given to the DNS host name over which the Management Service will be available. Although the Management Service may run on a particular local host, for example `myhost.mycompany.com`, your Management Agents will access the Management Service using the host name that has been assigned to the Server Load Balancer. For example, `oracleoms.mycompany.com`.

As a result, when you enable Enterprise Manager Framework Security for the Management Service, it is important to ensure that the Server Load Balancer host name is embedded into the Certificate that the Management Service uses for SSL communications. To do so, enter the following commands:

This may be done by using `emctl secure oms` and specifying the host name in the with an extra `-host` parameter as follows:

- Enable security on the Management Service by entering the following command:

  ```
  emctl secure oms -host <slb_hostname> [-slb_console_port <slb
  UI port>] [-slb_port <slb upload port>] [other params]
  ```

  Run this command on each OMS. You will need to restart each OMS after running the 'emctl secure oms' command.

- Create virtual servers and pools on the Server Load Balancer.
- Verify that the console can be accessed using the following URL:

  ```
  https://slbhost:slb_console_port/em
  ```

- Re-secure the Agents with Server Load Balancer by using the following command:

  ```
  emctl secure agent -emdWalletSrcUrl <SLB Upload or UI URL>
  ```

  For example:

  ```
  Agent_Home/bin/emctl secure agent -emdWalletSrcUrl
  https://slbost:slb_upload_port/em
  ```

11.4.8 Enabling Security for the Management Repository Database

This section describes how to enable Security for the Oracle Management Repository. This section includes the following topics:

- About Oracle Advanced Security and the sqlnet.ora Configuration File
- Configuring the Management Service to Connect to a Secure Management Repository Database
- Enabling Oracle Advanced Security for the Management Repository
- Enabling Security for a Management Agent Monitoring a Secure Management Repository or Database

11.4.8.1 About Oracle Advanced Security and the sqlnet.ora Configuration File

You enable security for the Management Repository by using Oracle Advanced Security. Oracle Advanced Security ensures the security of data transferred to and from an Oracle database.

See Also: *Oracle Database Advanced Security Administrator's Guide*

To enable Oracle Advanced Security for the Management Repository database, you must make modifications to the `sqlnet.ora` configuration file. The `sqlnet.ora` configuration file is used to define various database connection properties, including Oracle Advanced Security parameters.

The sqlnet.ora file is located in the following subdirectory of the Database home:

ORACLE_HOME`/network/admin`

After you have enabled Security for the Management Repository and the Management Services that communicate with the Management Repository, you must also configure Oracle Advanced Security for the Management Agent by modifying the `sqlnet.ora` configuration file in the Management Agent home directory.

> **See Also:** "Enabling Security for a Management Agent Monitoring a Secure Management Repository or Database"

It is important that both the Management Service and the Management Repository are configured to use Oracle Advanced Security. Otherwise, errors will occur when the Management Service attempts to connect to the Management Repository. For example, the Management Service might receive the following error:

```
ORA-12645: Parameter does not exist
```

To correct this problem, be sure both the Management Service and the Management Repository are configured as described in the following sections.

> **Note:** The procedures in this section describe how to manually modify the `sqlnet.ora` configuration file to enable Oracle Advanced Security. Alternatively, you can make these modifications using the administration tools described in the *Oracle Database Advanced Security Administrator's Guide*.

11.4.8.2 Configuring the Management Service to Connect to a Secure Management Repository Database

If you have enabled Oracle Advanced Security for the Management Service database—or if you plan to enable Oracle Advanced Security for the Management Repository database—use the following procedure to enable Oracle Advanced Security for the Management Service:

1. Stop the Management Service:

   ```
   ORACLE_HOME/bin/emctl stop oms
   ```

2. Set Enterprise Manager operational properties by using the `emctl set property` command.
3. Restart the Management Service.

   ```
   ORACLE_HOME/bin/emctl start oms
   ```

Table 11–4 Oracle Advanced Security Properties in the Enterprise Manager Properties File

Property	Description
oracle.sysman.emRep.dbConn.enableEncryption	Defines whether or not Enterprise Manager will use encryption between Management Service and Management Repository. Possible values are TRUE and FALSE. The default value is TRUE. For example: `oracle.sysman.emRep.dbConn. enableEncryption=true`

Table 11–4 (Cont.) Oracle Advanced Security Properties in the Enterprise Manager Properties File

Property	Description
oracle.net.encryption_client	Defines the Management Service encryption requirement. Possible values are REJECTED, ACCEPTED, REQUESTED, and REQUIRED. The default value is REQUESTED. In other words, if the database supports secure connections, then the Management Service uses secure connections, otherwise the Management Service uses insecure connections. For example: `oracle.net. encryption_client=REQUESTED`
oracle.net.encryption_types_client	Defines the different types of encryption algorithms the client supports. Possible values should be listed within parenthesis. The default value is `( DES40C ).` For example: `oracle.net. encryption_types_client= ( DES40C )`
oracle.net.crypto_checksum_client	Defines the Client's checksum requirements. Possible values are REJECTED, ACCEPTED, REQUESTED, and REQUIRED. The default value is REQUESTED. In other words, if the server supports checksum enabled connections, then the Management Service uses them, otherwise it uses normal connections. For example: `oracle.net. crypto_checksum_client=REQUESTED`
oracle.net.crypto_checksum_types_client	This property defines the different types of checksums algorithms the client supports. Possible values should be listed within parentheses. The default value is (MD5). For example: `oracle.net. crypto_checksum_types_client= ( MD5 )`

11.4.8.3 Enabling Oracle Advanced Security for the Management Repository

To be sure your database is secure and that only encrypted data is transferred between your database server and other sources, review the security documentation available in the Oracle Database documentation library.

See Also: *Oracle Database Advanced Security Administrator's Guide*

The following instructions provide an example of how you can confirm that Oracle Advanced Security is enabled for your Management Repository database and its connections with the Management Service:

1. Locate the `sqlnet.ora` configuration file in the following directory of the database Oracle Home:

   ```
   ORACLE_HOME/network/admin
   ```

2. Using a text editor, look for the following entries (or similar entries) in the `sqlnet.ora` file:

   ```
   SQLNET.ENCRYPTION_SERVER = REQUESTED
   SQLNET.CRYPTO_SEED = "abcdefg123456789"
   ```

   See Also: "Configuring Network Data Encryption and Integrity for Oracle Servers and Clients in the Oracle Application Server 10g Administrator's Guide

3. Save your changes and exit the text editor.

11.4.8.4 Enabling Security for a Management Agent Monitoring a Secure Management Repository or Database

After you have enabled Oracle Advanced Security for the Management Repository, you must also enable Advanced Security for the Management Agent that is monitoring the Management Repository:

1. Locate the `sqlnet.ora` configuration file in the following directory inside the home directory for the Management Agent that is monitoring the Management Repository:

   AGENT_HOME`/network/admin (UNIX)`
   AGENT_HOME`\network\admin (Windows)`

2. Using a text editor, add the following entry to the `sqlnet.ora` configuration file:

   ```
   SQLNET.CRYPTO_SEED = "abcdefg123456789"
   ```

   The `SQLNET.CRYPTO_SEED` can be any string between 10 to 70 characters.

   See Also: "Configuring Network Data Encryption and Integrity for Oracle Servers and Clients in the Oracle Application Server Administrator's Guide

3. Save your changes and exit the text editor.
4. Restart the Management Agent.

11.4.9 Configuring Third Party Certificates

You can configure third party certificates for:

- HTTPS Upload Virtual Host
- HTTPS Console Virtual Host

Note: Only Single Sign-On wallets are supported.

11.4.9.1 Configuring Third Party Certificate for HTTPS Upload Virtual Host

You can configure the third party certificate for the HTTPS Upload Virtual Host in two ways:

Method I

1. Create a wallet for each OMS in the Cloud.
2. While creating the wallet, specify the host name of the machine where the OMS is installed or the Load Balancer Name if the OMS is behind the Load Balancer for Common Name.
3. Write the certificates of all the Certificate Authorities in the certificate chain (like the Root Certificate Authority, Intermediate Certificate Authority) into a file named `trusted_certs.txt`.
4. Download or copy the `trusted_certs.txt` file to the host machines on which each Agent that is communicating with the OMS is running.
5. Run the add_trust_cert command on each Agent and then restart that Agent.

   ```
   emctl secure add_trust_cert -trust_certs_loc <location of the trusted_certs.txt
   file>
   ```

6. Secure the OMS and restart it.

   ```
   emctl secure oms -wallet <location of wallet> -trust_certs_loc <loc of trusted_
   certs.txt> [any other options]
   ```

Method 2

1. Create a wallet for each OMS in the Cloud.
2. Specify the host name of the machine where the OMS is installed or the Load Balancer Name if the OMS is behind the Server Load Balancer for Common Name (CN).
3. Write the certificates of all the Certificate Authorities in the certificate chain (like the Root Certificate Authority, Intermediate Certificate Authority) into a file named `trusted_certs.txt`.
4. Restart the OMS after it has been secured.

   ```
   emctl secure oms -wallet <location of wallet> -trust_certs_loc <loc of trusted_
   certs.txt> [any other options]
   ```

5. Either re-secure the Agent by running the `emctl secure agent` command (should be run on all Agents) or import the trust points by running the `emctl secure add_trust_cert -trust_certs_loc <location of the trusted_certs.txt file>` command. The `-trust_certs_loc` parameter must contain the path and the filename of the `trusted_certs.txt` file.

   Note: This file must only contain certificates in base64 format and no special characters or empty lines.

11.4.9.2 Configuring Third Party Certificate for HTTPS Console Virtual Host

To configure the third party certificate for HTTPS WebTier Virtual Host:

1. Create a wallet for each OMS in the Cloud. Specify the host name of the machine where the OMS is installed or the Load Balancer Name if the OMS is behind the Server Load Balancer for Common Name.

2. Run the following command on each OMS and the restart that OMS:

```
emctl secure console -wallet <location of wallet>
```

Note: Only single-sign-on wallets are supported.

11.5 Accessing Managed Targets

The following topics are discussed in this section:

- Credential Subsystem
- Pluggable Authentication Modules (PAM) Support
- Sudo and Powerbroker Support

11.5.1 Credential Subsystem

Credentials like user names and passwords are typically required to access targets such as databases, application servers, and hosts. Credentials are encrypted and stored in Enterprise Manager. Beginning with Enterprise Manager 12c, the credential subsystem supports, in addition to basic username-password, strong authentication schemes such as PKI, SSH-keys and Kerberos. SSH-key based host authentication, used by jobs, deployment procedures and other Enterprise Manger subsystems, is now supported.

By using appropriate credentials, you can:

- Collect metrics in the background as well as real-time
- Perform jobs such as backup, patching, and cloning
- Perform real-time target administration such as start, and stop
- Connect to My Oracle Support

Based on their usage, credentials can be classified into the following categories:

- Named Credential
- Job Credentials
- Monitoring Credentials
- Collection Credentials
- Preferred Credentials

11.5.1.1 Named Credential

Credentials are stored within Enterprise Manager as "named" entities. Administrators define and store credentials within Enterprise Manager and refer to the credential by a credential name. Named credentials can be a username/password, or a public key-private key pair. An Enterprise Manager administrator can then use the named credential for performing operations like running jobs, patching and other system management tasks. For example, an administrator can store the username and password they want to use for patching as "MyPatchingCreds". He can later submit a patching job that uses "MyPatchingCreds" to patch a production databases.

There are two categories of named credentials:

- **Global Named Credential**

A global named credential is an entity, which is not associated with any Enterprise Manager object. Global named credentials consist of the authentication scheme along with any authentication parameters. Because these are independent entities, an Enterprise Manger administrator can associate these credentials with objects at a later time.

- **Target Named Credentials**

  Target named credential is an entity which are associated with individual targets at the time of creation. This entity will also contain authentication scheme along with authentication parameters for a specific target.

Access Control for Named Credentials

The access control model for credentials adhere to the following rules:

- Only credential owners can grant privileges on their credential objects to other users.
- Enterprise Manager Super Administrators cannot obtain any privileges on a newly created credential until he is explicitly granted privileges on the credential object.
- Enterprise Manager administrators, regardless of privilege level, cannot see the sensitive fields such as passwords and private keys from the console UI.
- Credentials privileges cannot be assigned to a role. This eliminates back door entry by Enterprise Manager Super Administrators to grant themselves privileges on the credentials for which they do not have explicit access.
- An Enterprise Manager administrator cannot view other administrators' credentials unless an explicit grant is provided. Even Enterprise Manager Super Administrators cannot view other users' credentials.
- Any Enterprise Manager administrator can create his own credentials and have FULL privileges on the credentials owned.

Enterprise Manager Administrators will be able to grant privileges to other administrators while creating the credential or by granting the privileges when editing the credential.

All the credentials owned by an Enterprise Manager administrator will be deleted if that administrator is deleted from Enterprise Manager. Since access to shared credentials is not automatically granted to Super Administrators, re-assigning named credentials belonging to a regular Enterprise Manager administrator by a Super Administrator is not allowed.

Credential Privilege Levels

The following privilege levels are available for all credentials:

- **VIEW**: An administrator with VIEW privileges on other administrator's credentials will be able to view the structure and username of the credential. Sensitive information of the credential such as the password will never be shown. Administrators with VIEW privilege on a credential will also be able to use the credentials for running jobs, patching and other system management operations within Enterprise Manager.
- **EDIT**: Allows an Enterprise Manager administrator to change a sensitive information such as the password, or the public/private key pair of the credential. The administrator will not be able to change the Authentication Scheme of the credential. The username for the credential cannot be changed.
- **FULL**: Allows an Enterprise Manager administrator to change the credential username, sensitive information such as the password or the public/private key

pair, and authentication scheme. An administrator with FULL privilege on a named credential will be able to delete the named credential.

To create or edit a named credential, from the **Setup** menu, choose **Security** and then **Named Credential**. The Named Credential page displays as shown in the following figure.

Figure 11–4 Named Credentials Page

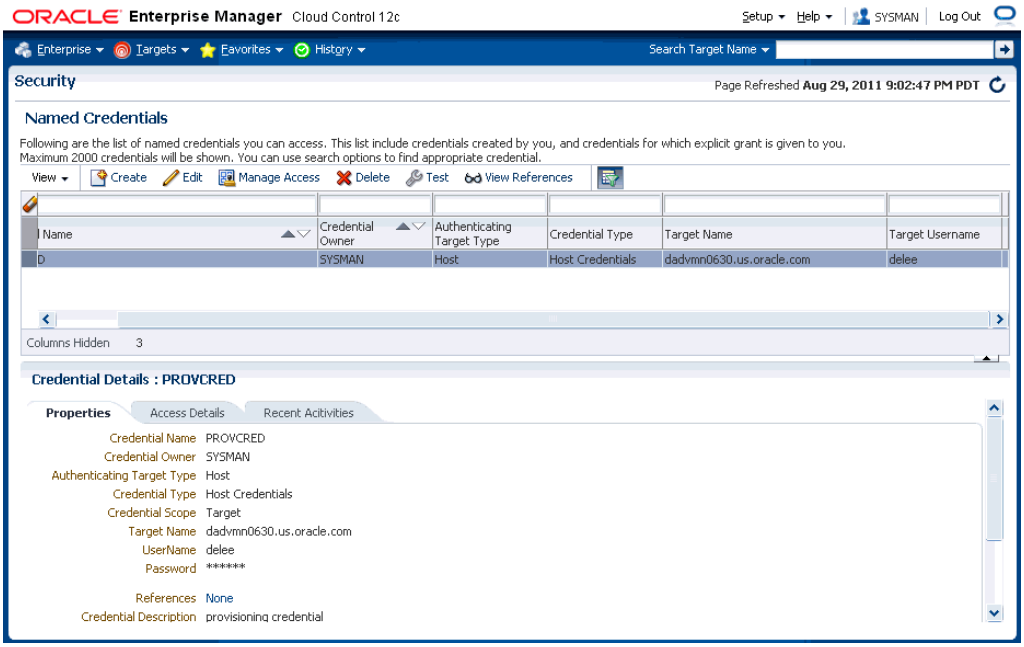

From the Named Credential page, you can **Create** a new named credential, **Edit** an existing credential, **Manage Access** (grant/revoke privileges), **Delete**, **Test**, **View References**, or click the *Query by Example* icon to filter the list of named credentials.

11.5.1.2 Job Credentials

The job system uses the credential subsystem to retrieve the credentials required to submit a job on the targets. The administrator can define their preferred and default credentials from the Setup -> Security -> Preferred Credentials page. As an administrator, you can:

1. Use Preferred Credentials

2. Use Named Credentials

3. Create new credentials

while submitting the job.

Note: If the user chooses to use preferred credentials, these credentials will be used when the user submits the job. If the preferred credentials are not available, the default credentials will be used. If default credentials are not present, the job cannot be submitted.

11.5.1.3 Monitoring Credentials

These credentials are used by the Management Agent to monitor certain types of targets. For example, most database monitoring involves connecting to the database, which requires a username, password, and optionally, a role. Monitoring credentials, if stored in the repository, can also be potentially used by management applications to connect directly to the target from the OMS.

To create or edit a monitoring credentials, from the **Setup** menu, choose **Security** and then **Monitoring Credentials**. The Monitoring Credentials page displays as shown in the following figure.

Figure 11–5 Monitoring Credentials

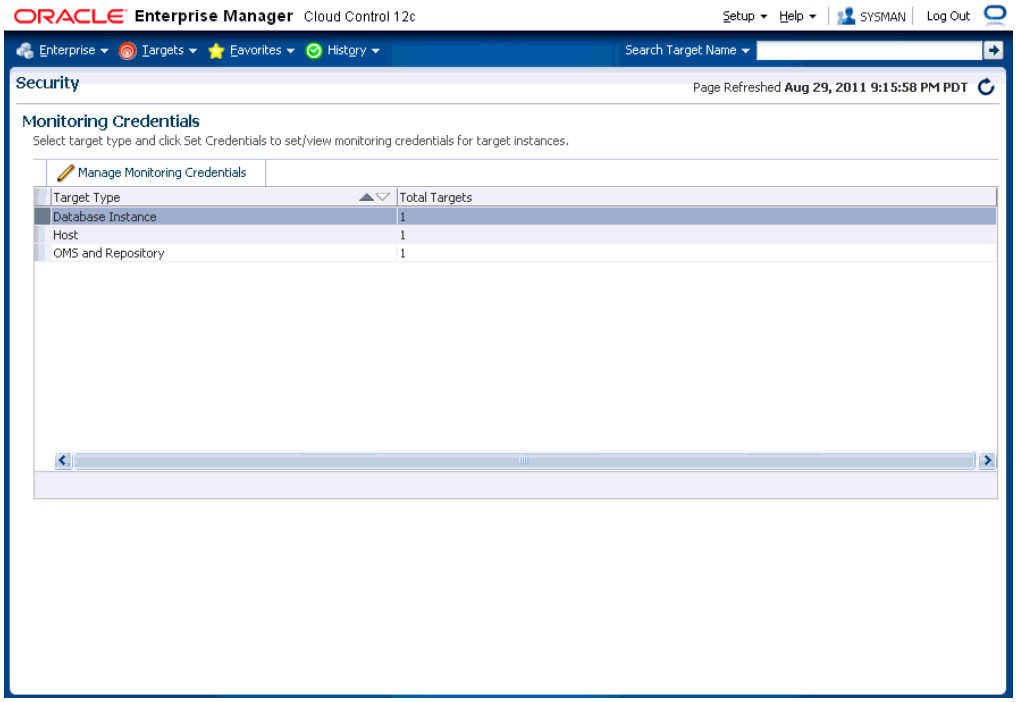

To modify monitoring credentials, select the desired target type and click **Manage Monitoring Credentials**. The monitoring credentials page for the selected target type displays.

11.5.1.4 Collection Credentials

These credentials are associated with metric extensions and older user-defined metrics.

11.5.1.5 Preferred Credentials

Preferred credentials are used to simplify access to managed targets by storing target login credentials in the Management Repository. With preferred credentials set, users can access an Enterprise Manager target that recognizes those credentials without being prompted to log into the target. Preferred credentials are set on a per user basis, thus ensuring the security of the managed enterprise environment.

- **Default Credentials**: Default credentials can be set for a particular target type and will be available for all the targets of the target type. It will be overridden by target preferred credentials.
- **Target Credentials**: Target credentials are preferred credentials set for a particular target. They could be used by applications such as the job system, notifications, or patching. For example, if the user chooses to use preferred credentials while submitting a job, then the preferred credentials set for the target (target credentials) will be used. If the target credentials are not present, the default credentials (for the target type) will be used. If the default credentials are not present, the job will fail. If not specified, by default, preferred credentials refer to preferred target credentials"

For example, to set the host preferred credentials, from the **Setup** menu, choose **Security** and then **Preferred Credential**. In the Preferred Credentials page, select the **Host** target type from the table and click **Manage Preferred Credentials**. The Host Preferred Credentials are displayed.

Figure 11–6 Host Preferred Credentials

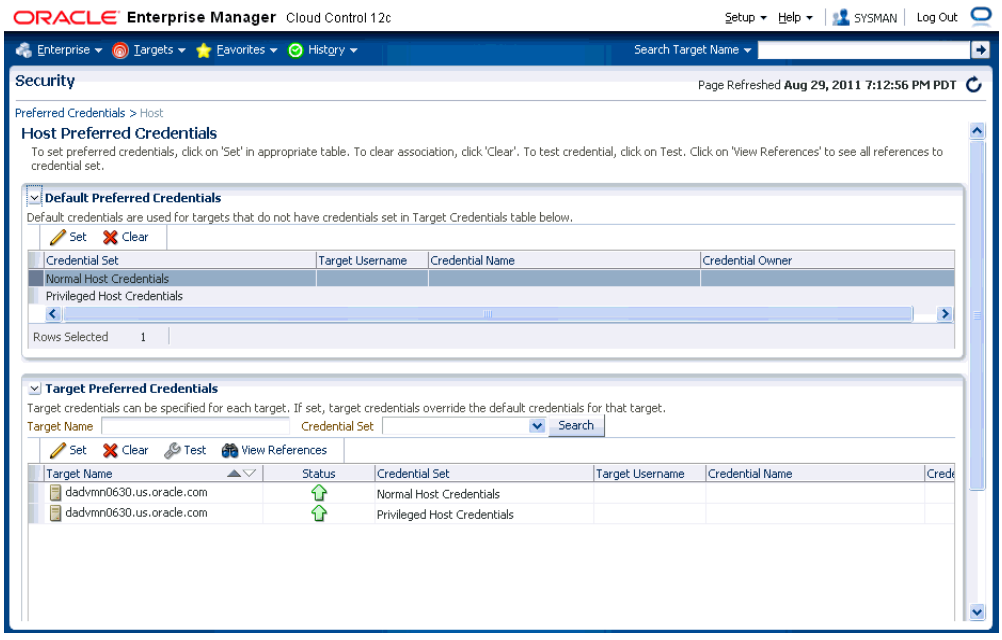

On this page, you can set both default and explicit preferred credentials for the host target types.

11.5.1.6 Managing Credentials Using EMCLI

You can manage passwords using EMCLI verbs. Using EMCLI, you can:

- Change the database user password in both the target database and Enterprise Manager.

```
emcli update_db_password -change_at_target=Yes|No -change_all_reference=Yes|No
```

- Update a password which has already been changed at the host target.

```
emcli update_host_password -change_all_reference=Yes|No
```

- Set preferred credentials for given users.

```
emcli set_preferred_credential
-set_name="set_name"
-target_name="target_name"
-target_type="ttype"
-credential_name="cred_name"
[-credential_owner ="owner]"
```

And

```
emcli set_preferred_credential
-set_name="set_name"
-target_name="target_name"
-target_type="ttype"
-credential_name="cred_name"
[-credential_owner ="owner]"
```

For detailed descriptions of these verbs, refer to *Enterprise Manager Command Line Interface* guide.

11.5.2 Setting Up SSH Key-based Host Authentication

Secure Shell or SSH allows data to be exchanged over the network using a secure channel between two devices. SSH is used primarily on Linux and Unix based systems. SSH was designed as a replacement for FTP, telnet and other unsecure remote shells, which send information, notably passwords in plaintext, leaving them open for interception. The encryption used by SSH provides confidentiality and integrity of data over an insecure network. SSH also protects the system against DNS spoofing attacks. This makes SSH a better choice in production environments over telnet/FTP and other username/password based authentications.

You can configure Enterprise Manager to use SSH while performing management operations, thus allowing Enterprise Manager administrators to leverage the security features provided by SSH along with the management capabilities of Enterprise Manager. When authenticating in this mode, the Agent acts as a Java SSH client and connect to the host using the username/password provided in the credential.

Enterprise Manager allows you to store a public-private key pair for administrators and allows them to view and install the public key on the hosts. Administrators can then submit jobs/patching operations in which they specify the credential that refers to the private key to perform the operation. The OMS passes the private key to the Agent along with the commands and the command parameters. Agent invokes the Java SSH client and attempts to connect to the host using the private key. Since the host already has the public key installed, it identifies the private key and successfully authenticates the Agent's Java SSH client. The Agent can now run the commands via the SSH client on the host to perform the requested operations.

Setup Example Session

> **Note:** The procedure shown in this example assumes that you have a firm understanding of SSH setup procedures and user and host equivalence using public private key pair using SSH.

To generate, manage, or convert SSH authentication keys, you use the *SSH-keygen* utility available on UNIX systems. This utility SSH-keygen tool provides different options to create with different strengths RSA keys for SSH protocol version 1 and RSA or DSA keys for use by SSH protocol version 2.

Example 11–10 Setting Up SSH key-based Authentication

```
$ ssh-keygen -t rsa
```
The command options instruct the utility to generate SSH keys (RSA key pair).

```
Generating public/private rsa key pair.
Enter file in which to save the key (/home/myhome/.ssh/id_rsa):
```
The path specified is the standard path to the location where SSH keys are stored ($HOME/.ssh).

```
Enter passphrase (empty for no passphrase)
Enter same passphrase again:
Your identification has been saved in /home/admin1/.ssh/id_rsa.
Your public key has been saved in /home/admin1/.ssh/id_rsa.pub.
The key fingerprint is:
bb:da:59:7a:fc:24:c6:9a:ee:dd:af:da:1b:1b:ed:7f admin1@myhost2170474
```

The ssh-regkey utility has now generated two files in the .ssh directory.

```
$ ls
id_rsa  id_rsa.pub
```
To permit access to the host without having SSH prompt for a password, copy the public key to the authorized_keys file on that system.

```
$ cp id_rsa.pub  authorized_keys
```

From this point, all keys listed in that file are allowed access.

Next, perform a remote login using SSH. The system will not prompt you for a password.

```
$ ssh myhost
The authenticity of host 'myhost (10.229.147.184)' can't be established.
RSA key fingerprint is de:a0:2a:d5:23:f0:8a:72:98:74:2c:6f:bf:ad:5b:2b.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'myhost,10.229.147.184' (RSA) to the list of known
hosts.
Last login: Mon Aug 29 16:48:45 2011 from anotherhost.us.oracle.com
$
```

You are now ready to add the credential to Enterprise Manager.

1. From the **Setup** menu, choose **Security** and then **Named Credentials**.
2. On the Named Credentials page, click Create.
3. Select **Host** from the **Authenticating Target Type** drop-down menu.

4. Select **SSH Key Credentials** from the **Credential Type** drop-down menu as shown in the following figure.

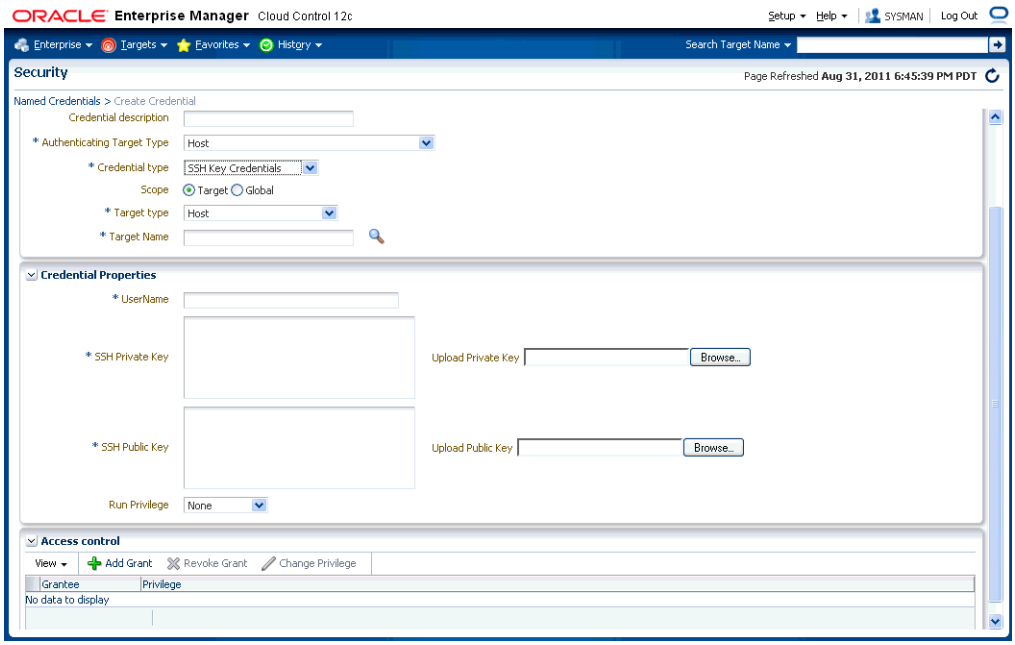

5. Ensure that the SSH private key/public key files have been copied to the host on which the browser is running.
6. From the **Credential Properties** region, click **Browse** for **Public Key** and **Private Key** to upload the generated public key/private key files.
7. Click **Test and Save** to verify the credentials and save them.

11.5.3 Pluggable Authentication Modules (PAM) Support for Hosts

Pluggable authentication modules, or PAM, is a mechanism to integrate multiple low-level authentication schemes into a high-level application programming interface (API). It allows programs that rely on authentication to be written independently of the underlying authentication scheme. By using PAM, instead of using the local password file to authenticate the user accessing the host, you can take advantage of other authentication mechanisms such as LDAP, RADIUS and Kerberors. If your host authentication is configured over PAM, the Management Agent needs to be configured accordingly to enable PAM Authentication. Refer to note 422073.1 for deployment details.

> **Note:** The local password file (usually `/etc/passwd`) will be checked and used first. This should be synchronized with the LDAP password if it is being used. If this fails, the Management Agent will switch to the external authentication module.

11.5.3.1 Configuring PAM for RHEL4 Users

For users on RHEL4, the PAM file configuration is as follows:

```
#%PAM-1.0
auth required pam_ldap.so
account required pam_ldap.so
```

```
password required pam_ldap.so
session required pam_ldap.so
```

For more details, see https://www.redhat.com/docs/manuals/enterprise/RHEL-4-Manual/ref-guide/s1-pam-format.html

11.5.3.2 Configuring PAM for AIX Users

For AIX users, use the `edit/etc/pam.conf` file and add the following lines:

```
emagent auth    required        /usr/lib/security/pam_aix
emagent account required        /usr/lib/security/pam_aix
emagent password  required      /usr/lib/security/pam_aix
emagent session required        /usr/lib/security/pam_aix
```

After editing the file, apply patch **5527130** and run `root.sh`

11.5.4 Sudo and PowerBroker Support

Privilege delegation allows a logged-in user to perform an activity with the privileges of another user. Sudo and PowerBroker are privilege delegation tools that allow a logged-in user to be assigned these privileges. Typically, the privileges that are granted to a specific user are administered centrally. For example, the sudo command can be used to run a script that requires root access:

```
sudo root root.sh
```

In the invocation of sudo in the example above, an administrator can use the sudo command to run a script as root provided he has been granted the appropriate privileges by the system administrator.

Enterprise Manager preferred credentials allow you to use two types of privilege delegation tools: Sudo and PowerBroker. You can use EMCLI or the Manage Privilege Delegation Settings page to set/edit privilege delegation settings for a host. See the *Enterprise Manager Command Line Interface* guide for more information on using the command line.

Sudo: sudo allows a permitted user to execute a command as the super user or another user, as specified in the sudoers file. If the invoking user is root or if the target user is the same as the invoking user, no password is required. Otherwise, sudo requires that users authenticate themselves with a password by default. Once a user has been authenticated, a timestamp is updated and the user may then use sudo without a password for a short period of time (5 minutes unless overridden in sudoers). sudo determines who is an authorized user by consulting the file `/etc/sudoers` file. For more information, see the manual page on sudo (man sudo) on Unix. Enterprise Manager authenticates the user using sudo, and executes the script as sudo. For example, if the command to be executed is foo -arg1 -arg2, it will be executed as sudo -S foo -arg1 -arg2.

PowerBroker: BeyondTrust Powerbroker enables UNIX system administrators to specify the circumstances under which other people may run certain programs such as root (or other important accounts). The result is that responsibility for such actions as adding user accounts, fixing line printer queues, and so on, can be safely assigned to the appropriate people, without disclosing the root password. The full power of root is thus protected from potential misuse or abuse-for example, modifying databases or file permissions, erasing disks, or more subtle damage.

BeyondTrust PowerBroker can access existing programs as well as its own set of utilities that execute common system administration tasks. Utilities being developed to run on top of BeyondTrust PowerBroker can manage passwords, accounts, backups, line printers, file ownership or removal, rebooting, logging people out, killing their programs, deciding who can log in to where from where, and so on. They can also provide TCP/IP, Load Balancer, cron, NIS, NFS, FTP, rlogin, and accounting subsystem management. Users can work from within a restricted shell or editor to access certain programs or files as root.

See your Sudo or PowerBroker documentation for detailed setup and configuration information.

11.5.4.1 Creating a Privilege Delegation Setting

Enterprise Manager allows you to create privilege delegation settings either by creating the setting directly on a host target, or by creating a Privilege Delegation Setting Template that you can apply to multiple hosts.

To create a privilege delegation setting directly on a host:

1. From the **Setup** menu, choose **Security** and then **Privilege Delegation**. The following screen is displayed:

Figure 11–7 Manage Privilege Delegation Settings

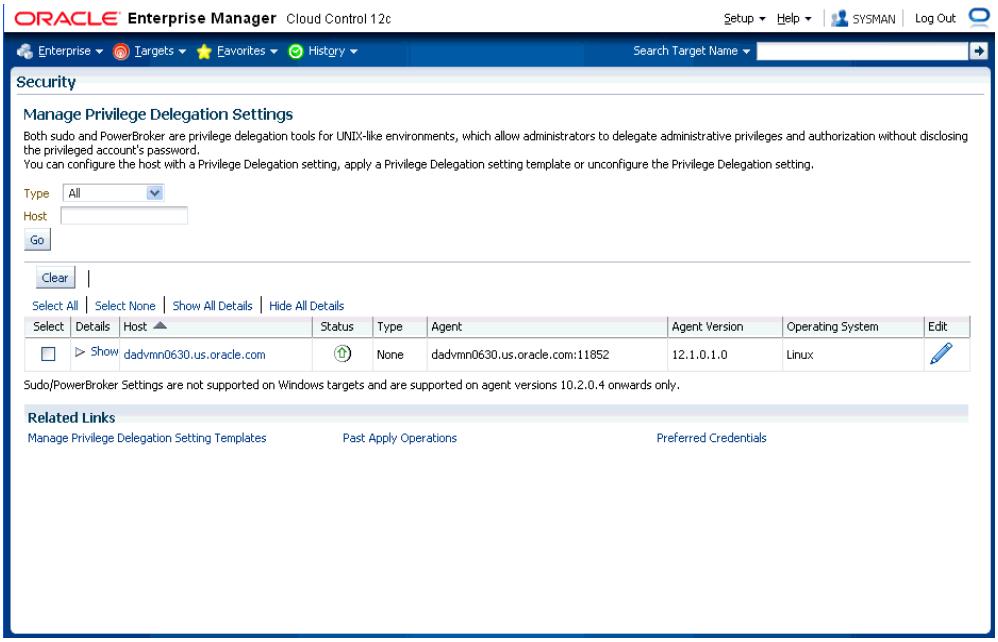

2. For any host target appearing in the table, click **Edit**. Enterprise Manager takes you to the Host Privilege Delegation Setting page.
3. Select a privilege delegation type (Sudo or PowerBroker).
4. Enter the privilege delegation command to be used and, in the case of PowerBroker, the optional Password Prompt.
5. Click **Update** to apply the settings to the host. The following figure shows the Host Privilege Delegation Setting window that you can use to create a PowerBroker setting.

Figure 11–8 Host Privilege Delegation Setting - PowerBroker

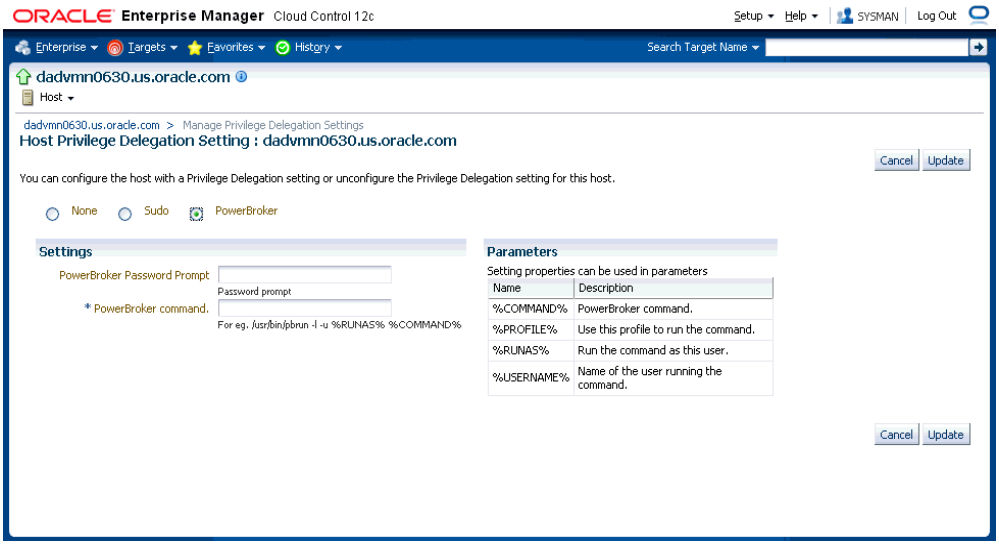

Once you have created a privilege delegation setting, you must apply this setting to selected targets. This setting can be applied to one more hosts or to a composite (Group) target (the group must contain at least one host target). You can apply a Privilege Delegation setting using the Cloud Control console. From the **Setup** menu, choose **Security** and then **Privilege Delegation**.

11.6 Cryptographic Support

To protect the integrity of sensitive data in Enterprise Manager, a signing on verification method known as the `emkey` is used. Encryption key is the master key that is used to encrypt/decrypt sensitive data, such as passwords and preferred credentials that are stored in the Repository. The key is originally in stored in repository. It is removed from repository and copied to the Credential Store during installation of the first OMS. (the `emkey` is secured out-of-the-box). A backup is created in `OMS_ORACLE_HOME/sysman/config/emkey.ora`. It is recommended to create a backup of this file on some other machine. When starting up, OMS reads the `emkey` from Credential Store and repository. If the `emkey` is not found or is corrupted, it fails to start. By storing the key separately from Enterprise Manager schema, we ensure that the sensitive data such as Named Credentials in the Repository remain inaccessible to the schema owner and other SYSDBA users (Privileged users who can perform maintenance tasks on the database) in the Repository. Moreover, keeping the key from the schema will ensure that sensitive data remain inaccessible while Repository backups are accessed. Further, the schema owner should not have access to the OMS/Repository Oracle homes.

11.6.1 Configuring the emkey

The `emkey` is an encryption key that is used to encrypt and decrypt sensitive data in Enterprise Manager such as host passwords, database passwords and others.

WARNING: If the `emkey.ora` file is lost or corrupted, all the encrypted data in the Management Repository becomes unusable. Maintain a backup copy of this file on another system.

During startup, the Oracle Management Service checks the status of the emkey. If the emkey has been properly configured, it uses it encrypting and decrypting data. If the emkey has not been configured properly, the following error message is displayed.

Example 11–11 emctl start oms Command

```
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
emctl start oms
Starting HTTP Server ...
Starting Oracle Management Server ...
Checking Oracle Management Server Status ...
Oracle Management Server is not functioning because of the following reason:
The Enterprise Manager Key is not configured properly. Run "emctl status emkey"
for more details.
```

11.6.2 emctl Commands

The emctl commands related to emkey are given below:

- emctl status emkey [-sysman_pwd <pwd>]
- emctl config emkey -copy_to_credstore [-sysman_pwd <pwd>]
- emctl config emkey -copy_to_repos [-sysman_pwd <pwd>]
- emctl config emkey -remove_from_repos [-sysman_pwd <pwd>]
- emctl config emkey -copy_to_file_from_credstore -admin_host <host> -admin_port <port> -admin_user <username> [-admin_pwd <pwd>] [-repos_pwd <pwd>] -emkey_file <emkey file>
- emctl config emkey -copy_to_file_from_repos (-repos_host <host> -repos_port <port> -repos_sid <sid> | -repos_conndesc <conn desc>) -repos_user <username> [-repos_pwd <pwd>] [-admin_pwd <pwd>] -emkey_file <emkey file>
- emctl config emkey -copy_to_credstore_from_file -admin_host <host> -admin_port <port> -admin_user <username> [-admin_pwd <pwd>] [-repos_pwd <pwd>] -emkey_file <emkey file>
- emctl config emkey -copy_to_repos_from_file (-repos_host <host> -repos_port <port> -repos_sid <sid> | -repos_conndesc <conn desc>) -repos_user <username> [-repos_pwd <pwd>] [-admin_pwd <pwd>] -emkey_file <emkey file>

11.6.2.1 emctl status emkey

This command shows the health or status of the emkey. Depending on the status of the emkey, the following messages are displayed:

- When the emkey has been correctly configured in the Credential Store, the following message is displayed.

Example 11–12 emctl status emkey - Example 1

```
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EmKey is configured properly, but is not secure. Secure the EmKey by running
"emctl config emkey -remove_from_repos"
```

- When the emkey has been correctly configured in the Credential Store and has been removed from the Management Repository, the following message is displayed.

Example 11–13 emctl status emkey - Example 2

```
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey is configured properly.
```

- When the emkey is corrupted in the Credential Store and removed from the Management Repository, the following message is displayed.

Example 11–14 emctl status emkey - Example 3

```
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey is not configured properly or is corrupted in the credential store and
does not exist in the Management Repository. To correct the problem:
1) Get the backed up emkey.ora file.
2) Configure the emkey by running "emctl config emkey -copy_to_credstore_from_
file"
```

11.6.2.2 emctl config emkey -copy_to_credstore

This command copies the emkey from the Management Repository to the Credential Store.

Example 11–15 Sample Output of the emctl config emkey -copy_to_credstore Command

```
emctl config emkey -copy_to_credstore [-sysman_pwd <pwd>]
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey has been copied to the Credential Store.
```

11.6.2.3 emctl config emkey -copy_to_repos

This command copies the emkey from the Credential Store to Management Repository.

Example 11–16 Sample Output of the emctl config emkey -copy_to_repos Command

```
emctl config emkey -copy_to_repos [-sysman_pwd <pwd>]
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey has been copied to the Management Repository. This operation will cause
the EMKey to become unsecure.
After the required operation has been completed, secure the EMKey by running
"emctl config emkey -remove_from_repos".
```

11.6.2.4 emctl config emkey -copy_to_file_from_credstore

This command copies the emkey from the Credential Store to a specified file.

Example 11–17 Sample Output of the emctl config emkey -copy_to_file_from_credstore Command

```
emctl config emkey -copy_to_file_from_credstore -admin_host <host> -admin_port
<port> -admin_user <username> [-admin_pwd <pwd>] [-repos_pwd <pwd>] -emkey_file
<emkey file>
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey has been copied to file.
```

11.6.2.5 emctl config emkey -copy_to_file_from_repos

This command copies the emkey from the Management Repository to a specified file.

Example 11–18 Sample Output of the emctl config emkey -copy_to_file_from_repos Command

```
emctl config emkey -copy_to_file_from_repos (-repos_host <host> -repos_port <port>
-repos_sid <sid> | -repos_conndesc <conn desc>) -repos_user <username> [-repos_pwd
<pwd>] [-admin_pwd <pwd>] -emkey_file <emkey file>
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey has been copied to file.
```

11.6.2.6 emctl config emkey -copy_to_credstore_from_file

This command copies the emkey from a specified file to the Credential Store.

Example 11–19 Sample Output of the emctl config emkey -copy_to_credstore_from_file Command

```
emctl config emkey -copy_to_credstore_from_file -admin_host <host> -admin_port
<port> -admin_user <username> [-admin_pwd <pwd>] [-repos_pwd <pwd>] -emkey_file
<emkey file>
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey has been copied to the Credential Store.
```

11.6.2.7 emctl config emkey -copy_to_repos_from_file

This command copies the emkey from a specified file to the repository.

Example 11–20 Sample Output of the emctl config emkey -copy_to_repos_from_file Command

```
emctl config emkey -copy_to_repos_from_file (-repos_host <host> -repos_port <port>
-repos_sid <sid> | -repos_conndesc <conn desc>) -repos_user <username> [-repos_pwd
<pwd>] [-admin_pwd <pwd>] -emkey_file <emkey file>
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey has been copied to the Management Repository. This operation will cause
the EMKey to become unsecure.
After the required operation has been completed, secure the EMKey by running
"emctl config emkey -remove_from_repos".
```

11.6.2.8 emctl config emkey -remove_from_repos

This command removes the emkey from the repository.

Example 11–21 Sample Output of emctl config emkey -remove_from_repos Command

```
emctl config emkey -remove_from_repos [-sysman_pwd <pwd>]
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey has been removed from the Management Repository.
```

Note: If the emkey is corrupted in the Credential Store , you cannot remove it from the Management Repository.

11.6.3 Install and Upgrade Scenarios

This section explains the install and upgrade scenarios for emkey.

11.6.3.1 Installing the Management Repository

A new emkey is generated as a strong random number when the Management Repository is installed.

11.6.3.2 Installing the First Oracle Management Service

When the Oracle Management Service is installed, the Installer copies the emkey to Credential Store and removes it from repository (emkey is secured out-of-box).

11.6.3.3 Upgrading from 10.2 or 11.1 to 12.1

The Management Repository is upgraded as usual. When upgrading the OMS, the `omsca` (OMS Configuration Assistant) copies the emkey to Credential Store and removes from repository. If the emkey is already secured before upgrade or has been removed from repository, then `omsca` reads the emkey from emkey.ora file present in old OMS Oracle Home and copies it to Credential Store.

Note: After all the Oracle Management Service have been upgraded, you can secure the emkey, that is, remove it from the Management Repository by running the following command:

```
emctl config emkey -remove_from_repos
```

11.6.3.4 Recreating the Management Repository

When the Management Repository is recreated, a new emkey is generated. This new key will not be in synchronization with the existing emkey in the Credential Store.

1. Copy the new emkey to Credential Store by using the `emctl config emkey -copy_to_credstore` command.
2. Take a backup by entering the `emctl config emkey -copy_to_file_from_repos` command or the `emctl config emkey -copy_to_file_from_credstore` command.
3. Secure the emkey by using the `emctl config emkey -remove_from_repos` command.

11.7 Setting Up the Auditing System for Enterprise Manager

All operations performed by Enterprise Manager users such as creating users, granting privileges, starting a remote job like patching or cloning, need to be audited to ensure compliance with the Sarbanes-Oxley Act of 2002 (SAS 70). This act defines standards an auditor must use to assess the contracted internal controls of a service organization. Auditing an operation enables an administrator to monitor, detect, and investigate problems and enforce enterprise wide security policies.

Irrespective of how the user has logged into Enterprise Manager, when auditing is enabled, each user action is audited and the audit details are stored in a record.

For Enterprise Manager 12c, BASIC auditing is enabled by default, thus creating an audit trail of credentials being created, edited, accessed, associated and deleted. Named credentials are first-class security objects on which privileges can be granted or revoked privileges. This means that multiple Enterprise Manager administrators will be able to use and modify the credential objects. Because credentials are sensitive data that can be used to perform various operations on the systems, there is a need to audit the operations on credentials.

Enterprise Manger audits all the operations performed on credentials. The auditing information includes, but is not be limited to, the current username, credential name, operation performed, operation status success or failure. The audit logs contain information about the credential owner, action initiator, credential name, user name, and target name, job names along with the date time of the operation. Credential fields like password, private keys are never logged.

The following operations are audited:

- **Creating a Named Credential**: Creating new Enterprise Manager credentials will be audited.
- **Editing a Named Credential**: Editing a credential may consist of changing the username and/or the sensitive credential attributes. Credential edits may also include changing the authentication scheme for the credential.
- **Delete a Named Credential**: Deleting a credential from Enterprise Manager will be audited.
- **Associating a Named Credential**: A named credential can be set as a preferred credential for a credential set at the target level or at target type level. The named credential can also be reference directly from a job. All operations involving the setting of the named credentials as preferred credentials and using it in a job or deployment procedure will be audited.
- **Accessing a Named Credential**: Enterprise Manager subsystems request credentials from the credential store to perform various system management tasks

11.7.1 Configuring the Enterprise Manager Audit System

You can configure the Enterprise Manager Audit System by using the following emcli commands:

- `enable_audit`: Enables auditing for all user operations.
- `disable_audit`: Disables auditing for all user operations.
- `show_operations_list`: Shows a list of the user operations being audited.
- `show_audit_settings`: Shows the audit status, operation list, externalization service details, and purge period details.

11.7.2 Configuring the Audit Data Export Service

Audit data needs to be protected and maintained for several years. The volume of audit data may become very large and impact the performance of the system. To limit the amount of data stored in the repository, the audit data must be externalized or archived at regular intervals. The archived audit data is stored in an XML file complying with the ODL format. To externalize the audit data, the `EM_AUDIT_EXTERNALIZATION` API is used. Records of the format <file-prefix>.NNNNN.xml, where NNNN is a number are generated. The numbers start with 00001 and continue to 99999.

You can set up the audit externalization service for exporting audit data into the file system by using the `update_audit_setting -externalization_switch` command.

11.7.3 Updating the Audit Settings

The `update_audit_settings` command updates the current audit settings in the repository and restarts the Management Service.

Example 11–22 Usage of the update_audit_setting command

```
emcli update_audit_settings
      -audit_switch="ENABLE/DISABLE"
      -operations_to_enable="name of the operations to enable, for all oprtations
       use ALL"
      -operations_to_disable="name of the operations to disable, for all
       oprtations use ALL"
      -externalization_switch="ENABLE/DISABLE"
      -directory_name="directory_name (DB Directory)"
      -file_prefix="file_prefix"
      -file_size="file_size (Bytes)"
      -data_retention_period="data_retention_period (Days)"
```

- `-audit_switch`: Enables auditing across Enterprise Manager. The possible values are `ENABLE/DISABLE`. Default value is `DISABLE`.
- `-operations_to_disable`: Enables auditing for specified operations. Enter **All** to enable all operations.
- `-operations_to_disable`: Disables auditing for specified operations. Enter **All** to disable all operations.
- `-externalization_switch`: Enables the audit data export service. The possible values are `ENABLE/DISABLE`. Default value is `DISABLE`.
- `-directory`: The database directory that is mapped to the OS directory where the export service archives the audit data files.
- `-file_prefix`: The file prefix to be used by the export service to create the file in which audit data is to be stored.
- `-file_size`: The size of the file on which the audit data is to be stored. The default value is 5000000 bytes.
- `data_retention_period`: The period for which the audit data is to be retained inside the repository. The default value is 365 days.

11.7.4 Searching the Audit Data

You can search for audit data that has been generated over a specified period. You can also search for the following:

- Audit details of a specific user operation or all user operations.
- Audit details of operations with a Success or Failure status or All operations.

From the **Setup** menu, choose **Management Services and Repository**. The Overview page is displayed. Click the **Audit Data** link under the Audit section. The Audit Data page is displayed. Specify the search criteria in the fields and click **Go**. The results are displayed in the Summary table.

Figure 11–9 Audit Data Search Page

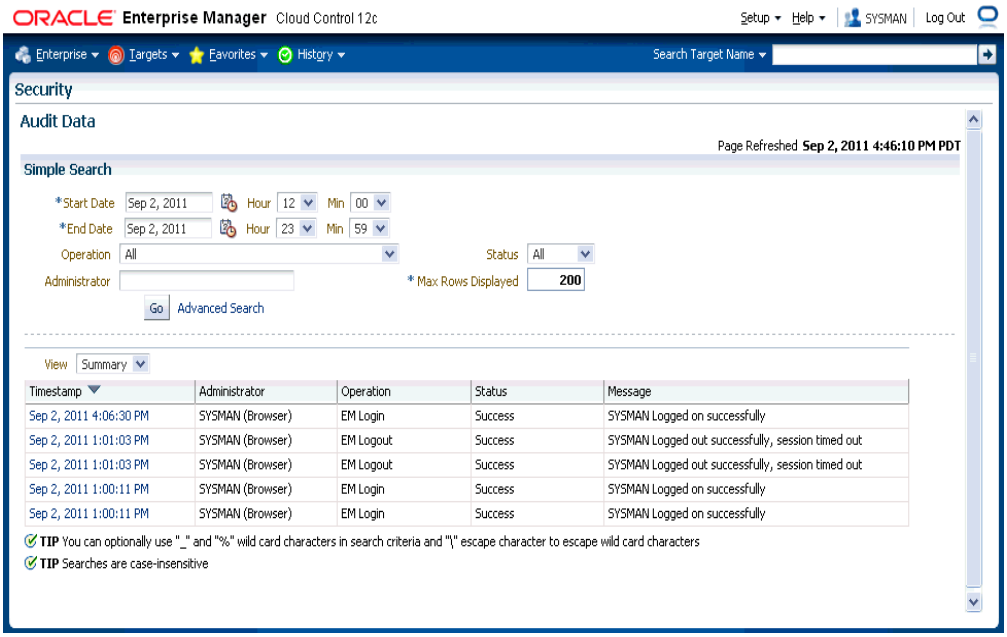

To view the details of each record that meets the search criteria, select Detailed in the View drop-down list. To drill down to the full record details, click on the **Timestamp**. The Audit Record page is displayed.

Figure 11–10 Audit Record Details Page

Field Name	Description
General	

Field Name	Description
Operation Timestamp	The date and time on which the operation took place.
Administrator	The id of the administrator who has logged into Enterprise Manager.
Operation	The type of operation being audited.
Status	The status of the operation which can be success or failure.
Message	A descriptive message indicating the status of the operation.
Normalized Timestamp	This is the UTC timestamp.
Client Information	
Session	This can either be the HTTP Session ID or the DBMS Session ID.
IP Address	The IP address of the client's host machine.
Hostname	The name of the client's host machine.
Upstream Component Type	The type of client, Console, Web Service, EMCLI, being used.
Authentication Type	The nature of the session (HTTP Session, DB Session).
Upstream Component Name	The name of the client being used.
OMS Information	
Hostname	The host name of the Oracle Management Service.
IP Address	The IP address of the Oracle Management Service.
Instance ID	The Instance ID of the Oracle Management Service.
Operation Specific Information	
Object Name	The operation being performed on an object

11.7.5 List of Operations Audited

The following table lists the names of operation and their description.

Table 11–5 List of Operations Audited

Operation Name	Description
TCAUD_ADD_TEMPLATE_ENTITY	Add entity to Template Collection
ADD_AGENT_REGISTRATION_ PASSWORD	Add Registration Password
SWLIBADDLOCATION	Configuring a new storage location in Software Library
ADD_CS_TARGET_ASSOC	Add Standard-Target Association
APPLY_TEMPLATE	APPLY_TEMPLATE
APPLY_UPDATE	Apply the update
TCAUD_ASSOC_TO_AG	Associate Template Collection to AG
ATTACH_MEXT	Attach Metric Extension

Table 11–5 (Cont.) List of Operations Audited

Operation Name	Description
AUDIT_EXPORT_SETTINGS	Audit Export Settings to externalize audit data
AUDIT_SETTINGS	Audit Settings to enable or Disable auditing
CHANGE_CONNECTOR_SETTINGS	enable/disable a Connector
CHANGE_PASSWORD	Change Password
CHANGE_PREFERRED_CREDENTIAL	change_pref_cred
CONFIG_CONNECTOR	Configure a Connector Instance
CREATE_CHANGE_MANAGEMENT_SETTING	Create the change management settings for the Real-time Monitoring rule.
CREATE_CONNECTOR	Create a Connector Instance
CCS_CREATE_MD	Create (or import) Custom Configuration Specification
CCS_CREATE_PARSER	Create Custom Configuration Specification Parser
CCS_CREATE_CUSTOM_TARGET_TYPE	Create Custom Target Type
CREATE_FACET	Create a new facet.
CREATE_FACET_PARAMETER	Create a new facet parameter.
CREATE_FACET_PATTERN	Create a new facet pattern.
CREATE_CSG	Create Framework
CREATE_MEXT	Create Metric Extension
CREATE_TEMPLATE	CREATE_TEMPLATE
CREATE_NAMED_CREDENTIAL	Create Named Credential
CREATE_PG_SCHED	Create Policy Group Schedule
CREATE_CCC_RULE	Create a Real-time Monitoring rule.
RES_STATE_CREATE_OP	Resolution State created
CREATE_ROLE	Create Role
CREATE_RULE	Create Rule
CREATE_CS	Create Standard
TCAUD_CREATE	Create Template Collection
CREATE_USER	Create User
CREATE_UDP	Create User Defined Policy
CREATE_UDPG	Create User Defined Policy Group
DB_LOGIN	Audit Database user Login
DB_LOGOUT	Audit Database user Logout
DELETE_CONNECTOR	delete a Connector Instance
CCS_DELETE_MD	Delete Custom Configuration Specification
CCS_DELETE_PARSER	Delete Custom Configuration Specification Parser
DELETE_FACET	Delete a facet.

Table 11–5 (Cont.) List of Operations Audited

Operation Name	Description
DELETE_FACET_PARAMETER	Delete a facet parameter.
DELETE_FACET_PATTERN	Delete a facet pattern.
DELETE_CSG	Delete Framework
DELETE_JOB	delete_job
DELETE_MEXT	Delete Metric Extension
DELETE_TEMPLATE	DELETE_TEMPLATE
DELETE_NAMED_CREDENTIAL	Delete Named Credential
DELETE_PG_EVAL	Delete Policy Group Evaluation Results
DELETE_PG_SCHED	Delete Policy Group Schedule
DELETE_CCC_RULE	Delete a Real-time Monitoring rule.
DELETE_AGENT_REGISTRATION_ PASSWORD	Delete Registration Password
RES_STATE_DELETE_OP	Resolution State deleted
DELETE_ROLE	Drop Role
DELETE_RULE	Delete Rule
SWLIBDELETEFOLDER	Deleting a directory in Software Library
SWLIBDELETEENTITY	Deleting an entity in Software Library
DELETE_CS	Delete Standard
TCAUD_DELETE	Delete Template Collection
DELETE_UPDATE	Delete the update
DELETE_USER	Delete User
DELETE_UDP	Delete User Defined Policy
DELETE_UDPG	Delete User Defined Policy Group
CCS_DEPLOY	Deploy Custom Configuration Specification
DETACH_MEXT	Detach Metric Extension
DISABLE_CS_TARGET_ASSOC	Disable Standard-Target Association
TCAUD_DEASSOC_FROM_AG	Disassociate Template Collection from AG
DOWNLOAD_UPDATE	Download an available update
EDIT_CSG	Edit Framework
EDIT_JOB	edit_job
EDIT_TEMPLATE	EDIT_TEMPLATE
EDIT_PG_SCHED	Edit Policy Group Schedule
EDIT_AGENT_REGISTRATION_ PASSWORD	Edit Registration Password
EDIT_RULE	Edit Rule
EDIT_CS	Edit Standard
EDIT_CS_TARGET_ASSOC	Edit Standard-Target Association
TCAUD_EDIT	Edit Template Collection

Table 11–5 (Cont.) List of Operations Audited

Operation Name	Description
EDIT_UDP	Edit User Defined Policy
EDIT_UDPG	Edit User Defined Policy Group
LOGIN	logon
LOGOUT	logoff
ENABLE_CS_TARGET_ASSOC	Enable Standard-Target Association
EVALUATE_UDP	Evaluate User Defined Policy
PERFORM_OPERATION_AS_AGENT	Execute any OS Command as the Agent User (uncredentialed)
FILE_TRANSFER	file_transfer
GET_FILE	get_file
GET_NAMED_CREDENTIAL	Get Named Credential
GRANT_JOB_PRIVILEGE	Grant Job Privilege
GRANT_PRIVILEGE	Grant Privilege
GRANT_ROLE	Grant Role
GRANT_SYSTEM_PRIVILEGE	Grant System Privilege
GRANT_TARGET_PRIVILEGE	Grant Target Privilege
IMPORT_FACET	Import a facet.
IMPORT_CSG	Import Framework
IMPORT_CCC_RULE	Import a Real-time Monitoring rule.
IMPORT_RULE	Import Rule
IMPORT_CS	Import Standard
IMPORT_UDP	Import User Defined Policy
INCLUDE_ACTION_TO_MONITOR	Include an action to monitor for the Real-time Monitoring rule.
INCLUDE_FILTER_FACET	Include a filter facet into the Real-time Monitoring rule.
INCLUDE_MONITORING_FACET	Include a monitoring facet into the Real-time Monitoring rule.
JOB_OUTPUT	Job output obtained after job execution
MODIFY_CHANGE_MANAGEMENT_SETTING	Modify the change management settings for the Real-time Monitoring rule.
MODIFY_FACET	Update a facet.
MODIFY_FACET_CONTENT	Update the basic facet information.
MODIFY_FACET_PARAMETER	Update a facet parameter.
MODIFY_FACET_PATTERN	Update a facet pattern.
MODIFY_METRIC_SETTINGS	MODIFY_METRIC_SETTINGS
UPDATE_NAMED_CREDENTIAL	Modify Named Credential
MODIFY_POLICY_SETTINGS	Modify Policy Settings
MODIFY_CCC_RULE	Update a Real-time Monitoring rule.

Table 11–5 (Cont.) List of Operations Audited

Operation Name	Description
RES_STATE_MODIFY_OP	Resolution State modified
MODIFY_ROLE	Modify Role
MODIFY_USER	Modify User
SWLIBMOVEENTITY	Moving all revisions of an entity in Software Library to another directory
PUBLISH_MEXT	Publish Metric Extension
SWLIBPURGELOCATION	Purging a storage location in Software Library
PUT_FILE_AS_AGENT	Put any File to the Agent's Filesystem as the Agent User (uncredentialed)
PUT_FILE	put_file
REFRESH_UPDATE	Refresh from EM Store
AGENT_REGISTRATION_PASSWORD_USAGE	Registration Password Usage
REMOTE_OPERATION_JOB	remote_op
REMOVE_ACTION_FROM_MONITOR	Remove an action from monitor for the Real-time Monitoring rule.
REMOVE_CHANGE_MANAGEMENT_SETTING	Remove the change management settings for the Real-time Monitoring rule.
TCAUD_REMOVE_TEMPLATE_ENTITY	Remove entity from Template Collection
REMOVE_FILTER_FACET	Remove a filter facet from the Real-time Monitoring rule.
REMOVE_MONITORING_FACET	Remove a monitoring facet from the Real-time Monitoring rule.
REMOVE_PRIVILEGE_DELEGATION_SETTING	Remove Privilege Delegation Setting
SWLIBDELETELOCATION	Removing a storage location in Software Library
REMOVE_CS_TARGET_ASSOC	Remove Standard-Target Association
REMOVE_UPDATE	Remove the update
TCAUD_RENAME	Rename Template Collection
AGENT_RESYNC	Agent resynchronization operation
REPOSITORY_RESYNC	Repository resynchronization operation
RETRY_JOB	
REVOKE_JOB_PRIVILEGE	Revoke Job Privilege
REVOKE_PRIVILEGE	Revoke Privilege
REVOKE_ROLE	Revoke Role
REVOKE_SYSTEM_PRIVILEGE	Revoke System Privilege
REVOKE_TARGET_PRIVILEGE	Revoke Target Privilege
SAVE_MONITORING_SETTINGS	SAVE_MONITORING_SETTINGS
SET_PRIVILEGE_DELEGATION_SETTING	Set Privilege Delegation Setting

Table 11–5 (Cont.) List of Operations Audited

Operation Name	Description
STOP_JOB	
SUBMIT_JOB	submit_job
SUBSCRIBE_UPDATE	Subscribe to an Update Type
SUSPEND_JOB	Suspend job
CCS_UNDEPLOY	Undeploy Custom Configuration Specification
UNSUBSCRIBE_UPDATE	Unsubscribe an Update Type
CCS_UPDATE_MD	Update Custom Configuration Specification
UPDATE_DB_PASSWORD	Update Database Password
INSERT_UPDATE	Show the update on self update home
UPDATE_MEXT	Update Metric Extension
UPDATE_PASSWORD	Update Password

11.8 Additional Security Considerations

After you enable security for the Enterprise Manager components and framework, there are additional security considerations. This section provides the following topics:

- Changing the SYSMAN and MGMT_VIEW Passwords
- Configuring Beacons to Monitor Web Applications Over HTTPS
- Patching Oracle Homes When the User is Locked
- Cloning Oracle Homes

11.8.1 Changing the SYSMAN and MGMT_VIEW Passwords

This section describes the commands used to change the SYSMAN and MGMT_VIEW passwords.

11.8.1.1 Changing the SYSMAN User Password

To change the password of the `SYSMAN` user, you use the following command:

```
emctl config oms -change_repos_pwd [-old_pwd <old_pwd>] [-new_pwd <new_pwd>]
[-use_sys_pwd [-sys_pwd <sys_pwd>]]
```

Parameter	Description
-change_repos_pwd	.
`-old_pwd`	This is the current SYSMAN password.
`-new_pwd`	This is the new password.
`-use_sys_pwd`	This parameter is optional and is used to connect to the database as a SYS user. Use this option if SYSMAN account on the database has expired/locked.
`-sys_pwd`	This is the password for the SYS user.

1. Stop all the OMSs.

   ```
   emctl stop oms
   ```

2. For each OMS, run the following command:

   ```
   emctl config oms -change_repos_pwd'
   ```

3. Restart the AdminServer and all the OMSs.

   ```
   emctl stop oms -all
   ```

   ```
   emctl start oms
   ```

11.8.1.2 Changing the MGMT_VIEW User Password

To change the password of the `MGMT_VIEW` user, you use the following command:

```
emctl config oms -change_view_user_pwd [-sysman_pwd <sysman_pwd>] [-user_pwd
<user_pwd>] [-auto_generate]
```

Parameter	Description
`-change_view_user_pwd`	Used to change MGMT_VIEW user's password.
`-sysman_pwd`	The password for the SYSMAN user.
`-user_pwd`	The new password for theMGMT_VIEW user..
`-auto_generate`	If this option is specified, the password is auto generated.

Important: In order to change the MGMT_VIEW password, you must ensure that the password of the WebLogic administrative user in the credential store matches the actual password of the user SYSMAN. If the credentials do not match, the connection to the Repository Database fails and the SYSMAN password cannot be modified

1. Stop all the OMSs.

   ```
   emctl stop oms -all
   ```

2. Execute the following command to update the WebLogic and nodemanager passwords in the Credential store:

   ```
   cd <OMS_HOME>/bin
   emctl secure create_admin_creds_wallet -admin_pwd <existing weblogic pwd>
   -nodemgr_pwd <existing nodemanager pwd>
   ```

3. Log into the Repository Database as a DBA user and execute the following to manually modify the password of the sysman_mds schema to the new password that will be set for the sysman user:

   ```
   SQL> alter user sysman_mds identified by <new_pwd of sysman user>;
   ```

4. For **ONE** of the OMSs, run the following command to modify the SYSMAN password::

   ```
   cd <OMS_HOME>/bin
   emctl config oms -change_repos_pwd -change_in_db -old_pwd <new_pwd> -new_pwd
   <new_pwd>
   ```

5. Restart the AdminServer and all the OMSs.

   ```
   emctl stop oms -all
   ```

```
emctl start oms
```

11.8.2 Configuring Beacons to Monitor Web Applications Over HTTPS

Oracle Beacons provide application performance availability and performance monitoring. They are part of the Application Service Level Management features of Enterprise Manager.

> **See Also:** "About Application Service Level Management" in the Enterprise Manager Online Help

When a Beacon is used to monitor a URL over Secure Sockets Layer (SSL) using an HTTPS URL, the Beacon must be configured to recognize the Certificate Authority that has been used by the Web site where that URL resides.

> **See Also:** "The Public Key Infrastructure Approach to Security" in the *Oracle Security Overview* for an overview of Public Key Infrastructure features, such as Certificate Authorities

The Beacon software is preconfigured to recognize most commercial Certificate Authorities that are likely to be used by a secure Internet Web Site. However, you may encounter Web Sites that, although available over HTTPS, do not have a Certificate that has been signed by a commercial Certificate Authority recognized by the Beacon. The following are out-of-box certificates recognized by Beacons:

- Class 1 Public Primary Certification Authority by VeriSign, Inc.
- Class 2 Public Primary Certification Authority by VeriSign, Inc.
- Class 3 Public Primary Certification Authority by VeriSign, Inc.
- Secure Server Certification Authority by RSA Data Security, Inc.
- GTE CyberTrust Root by GTE Corporation
- GTE CyberTrust Global Root by GTE CyberTrust Solutions, Inc.
- Entrust.net Secure Server Certification Authority by Entrust.net ((c) 1999
- Entrust.net Limited, www.entrust.net/CPS incorp. by ref. (limits liab.))
- Entrust.net Certification Authority (2048) by Entrust.net ((c) 1999
- Entrust.net Limited, www.entrust.net/CPS_2048 incorp. by ref. (limits liab.))
- Entrust.net Secure Server Certification Authority by Entrust.net ((c) 2000
- Entrust.net Limited, www.entrust.net/SSL_CPS incorp. by ref. (limits liab.))

In those cases, for example, if you attempt to use the Test section of the Beacon Performance page to test the HTTP Response of the secure URL, the following error appears in the **Status Description** column of the Response Metrics table on the URL Test Page:

```
javax.net.ssl.SSLException: SSL handshake failed:
X509CertChainIncompleteErr--https://mgmtsys.acme.com/OracleMyPage.Home
```

> **See Also:** "Using Beacons to Monitor Remote URL Availability" in the Enterprise Manager online help

To correct this problem, you must allow the Beacon to recognize the Certificate Authority that was used by the Web Site to support HTTPS. You must add the

Certificate of that Certificate Authority to the list of Certificate Authorities recognized by Beacon.

To configure the Beacon to recognize the Certificate Authority:

1. Obtain the Certificate of the Web Site's Certificate Authority, as follows:
   a. In Microsoft Internet Explorer, connect to the HTTPS URL of the Web Site you are attempting to monitor.
   b. Double-click the lock icon at the bottom of the browser screen, which indicates that you have connected to a secure Web site.

      The browser displays the Certificate dialog box, which describes the Certificate used for this Web site. Other browsers offer a similar mechanism to view the Certificate detail of a Web Site.
   c. Click the **Certificate Path** tab and select the first entry in the list of certificates.
   d. Click **View Certificate** to display a second Certificate dialog box.
   e. Click the **Details** tab on the Certificate window.
   f. Click **Copy to File** to display the Certificate Manager Export wizard.
   g. In the Certificate Manager Export wizard, select **Base64 encoded X.509 (.CER)** as the format you want to export and save the certificate to a text file with an easily-identifiable name, such as `beacon_certificate.cer`.
   h. Open the certificate file using a text editor.

      The content of the certificate file will look similar to the content shown in Example 11–23.
2. Update the list of Beacon Certificate Authorities as follows:
   a. Locate the `b64InternetCertificate.txt` file in the following directory of Agent Home of the Beacon host:

      agent_instance_home`/sysman/config/`

      This file contains a list of Base64 Certificates.
   b. Edit the `b64InternetCertificate.txt` file and add the contents of the Certificate file you just exported to the end of the file, taking care to include all the Base64 text of the Certificate including the BEGIN and END lines.
3. Restart the Management Agent.

   After you restart the Management Agent, the Beacon detects your addition to the list of Certificate Authorities recognized by Beacon and you can successfully monitor the availability and performance of the secure Web site URL.

Example 11–23 Sample Content of an Exported Certificate

```
-----BEGIN CERTIFICATE-----
MIIDBzCCAnCgAwIBAgIQTs4NcImNY3JAs5edi/5RkTANBgkqhkiG9w0BAQQFADCB
... base64 certificate content...
-----END CERTIFICATE-----
```

11.8.3 Patching Oracle Homes When the User is Locked

To patch an Oracle Home used by a user "Oracle" and the user is locked:

1. Edit the default patching script and prepend sudo or sudo -u or pbrun -u to the default patching step. You need to set a policy (by editing the sudoers file) to allow

the user submitting the job (who must be a valid operating system user) to be able to run sudo or pbrun without being prompted for password.

Note: You cannot patch Oracle Homes without targets. This must be done by using the Patching wizard.

11.8.4 Cloning Oracle Homes

The cloning application is wizard-driven. The source of the Oracle Home being cloned may be either an installed Oracle Home or a Software Library. Following are the steps in the cloning process:

1. If the source is an installed Oracle Home, then, after selecting the Oracle Home, a user will need to specify the Oracle Home credentials. These credentials once specified for an Oracle Home are stored in the repository. The next time a user clones the same Oracle Home, these credentials are automatically populated. Other parameters queried from the user at this point is a temporary location (on the source computer) and the list of files to be excluded from the Oracle Home. If the cloning source is a Software Library, the source Oracle Home credentials will not be queried for.
2. The user needs to specify the target location and provide the required credentials for each target location. These credentials will be the Oracle Home credentials for each of these target locations. Subsequently, if a user selects any of these cloned Oracle Homes as a source, the Oracle Home credentials are automatically populated.
3. Depending on the product being cloned, the user can view the Enterprise Manager page where query parameters required for the particular product being cloned are displayed.
4. The user can, then, view the execution of user-supplied pre-cloning and post-cloning scripts and the root.sh script. The root.sh script will always be run with sudo privileges, but the user has the option to decide if the pre-cloning and post-cloning scripts run with sudo privileges.
5. Finally, the user can schedule the cloning job at a convenient time.

For more information about cloning, refer to the Enterprise Manager Online Help.

12

Sizing Your Enterprise Manager Deployment

Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1 has the ability to scale for hundreds of users and thousands of systems and services on a single Enterprise Manager implementation.

This chapter describes techniques for achieving optimal performance using the Oracle Enterprise Manager application. It can also help you with capacity planning, sizing and maximizing Enterprise Manager performance in a large scale environment. By maintaining routine housekeeping and monitoring performance regularly, you insure that you will have the required data to make accurate forecasts of future sizing requirements. Receiving good baseline values for the Enterprise Manager Cloud Control vital signs and setting reasonable warning and critical thresholds on baselines allows Enterprise Manager to monitor itself for you.

This chapter also provides practical approaches to backup, recovery, and disaster recovery topics while addressing different strategies when practical for each tier of Enterprise Manager.

This chapter contains the following sections:

12.1 Oracle Enterprise Manager Cloud Control Architecture Overview

The architecture for Oracle Enterprise Manager Cloud Control exemplifies two key concepts in application performance tuning: distribution and parallelization of processing. Each component of Cloud Control can be configured to apply both these concepts.

The components of Enterprise Manager Cloud Control include:

- The Management Agent - A process that is deployed on each monitored host and that is responsible for monitoring all services and components on the host. The Management Agent is also responsible for communicating that information to the middle-tier Management Service and for managing and maintaining the system and its services.
- The Management Service - A J2EE Web application that renders the user interface for the Cloud Control Console, works with all Management Agents to process monitoring and jobs information, and uses the Management Repository as its data store.
- The Management Repository - The schema is an Oracle Database that contains all available information about administrators, services, and applications managed within Enterprise Manager.

For more information about the Cloud Control architecture, see the Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1 documentation:

- *Oracle Enterprise Manager Cloud Control Installation and Basic Configuration*
- *Oracle Enterprise Manager Concepts*

The Oracle Enterprise Manager 12*g* documentation is available at the following location on the Oracle Technology Network (OTN):

http://otn.oracle.com/documentation/oem.html

12.2 Enterprise Manager Cloud Control Sizing and Performance Methodology

An accurate predictor of capacity at scale is the actual metric trend information from each individual Enterprise Manager Cloud Control deployment. This information, combined with an established, rough, starting host system size and iterative tuning and maintenance, produces the most effective means of predicting capacity for your Enterprise Manager Cloud Control deployment. It also assists in keeping your deployment performing at an optimal level.

Here are the steps to follow to enact the Enterprise Manager Cloud Control sizing methodology:

1. If you have not already installed Enterprise Manager Cloud Control, choose a rough starting host configuration as listed in Table 12–1.
2. Periodically evaluate your site's vital signs (detailed later).
3. Eliminate bottlenecks using routine DBA/Enterprise Manager administration housekeeping.
4. Eliminate bottlenecks using tuning.
5. Extrapolate linearly into the future to plan for future sizing requirements.

Step one need only be done once for a given deployment. Steps two, three, and four must be done, regardless of whether you plan to grow your Enterprise Manager Cloud Control site, for the life of the deployment on a regular basis. These steps are essential to an efficient Enterprise Manager Cloud Control site regardless of its size or workload. You must complete steps two, three, and four before you continue on to step five. This is critical. Step five is only required if you intend to grow the deployment size in terms of monitored targets. However, evaluating these trends regularly can be helpful in evaluating any other changes to the deployment.

12.2.1 Step 1: Choosing a Starting Platform Cloud Control Deployment

If you have not yet installed Enterprise Manager Cloud Control on an initial platform, this step helps you choose a rough approximation based on experiences with real world Enterprise Manager Cloud Control deployments. **If you have already installed Enterprise Manager Cloud Control, proceed to Step 2.** Three typical deployment sizes are defined: small, medium, and large. The number and type of systems (or targets) it monitors largely defines the size of an Enterprise Manager Cloud Control deployment. This table represents Intel-based platforms.

Table 12–1 Management Server

Deployment Size	Hosts	Cores/Hosts	Memory/Host (GB)
Small (Agents < 100 and Targets < 1,000)	1	Minimum 2 cores (3 GHz)	*6
Medium (Agents < 1000 and Targets < 10,000)	1	4 cores (3 GHz)	*6
Large (Agents > 1000 or Targets > 10,000)	2	4 cores per host (3 GHz) 2	*8 per host

*If you plan on integrating BI Publisher version 11.1.1.5 with Enterprise Manager Release 12c Cloud Control, which is required for BI Publisher reports to function, add 1.5 GB to the memory requirements stated above, based on the size of your Enterprise Manager installation.

In any OMS host box, OPMN processes, Admin Server Process, Node Manager processes, will be running, so the minimum memory requirement is 4 GB per OMS host.

Table 12–2 Management Repository

Deployment Size	Hosts	Cores	Memory/Host (GB)
Small	1	2	6
Medium	1	4	6
Large	2	4	8

Table 12–3 Total Management Repository Storage

Deployment Size	Minimum Tablespace Sizes*				
	SYSTEM**	MGMT_ TABLESPACE	MGMT_ECM_ DEPOT_TS	MGMT_AD4J_ TS	TEMP
Small	600 MB	50 GB	1 GB	100 MB	10 GB
Medium	600 MB	200 GB	4 GB	200 MB	20 GB
Large	600 MB	300 GB	Greater than 4 GB	400 MB	40 GB

*These are strictly minimum values and are intended as rough guidelines only. The actual size of the MGMT_TABLESPACE could vary widely from deployment to deployment due to variations in target type distribution, user customization, and several other factors. These tablespaces are defined with AUTOEXTEND set to ON by default to help mitigate space constraint issues. On raw file systems Oracle recommends using more than the minimum size to help prevent space constraint issues.

**The SYSTEM and TEMP tablespace sizes are minimums for Enterprise Manager only repositories. If Enterprise Manager is sharing the repository database with other application(s), these minimums may be too low.

Note: You cannot monitor tablespaces through the use of alerts with auto extended files in version 11g of Enterprise Manager. You can either set up TABLESPACE FULL alerts generate if you want to have greater control over the management of your tablespaces, or you can allow Oracle to grow your database and not alert you through the AUTOEXTEND feature. Therefore to exercise greater control of the TABLESPACE FULL alerts, you can turn off autoextend.

The previous tables show the estimated minimum hardware requirements for each deployment size. Management Servers running on more than one host, as portrayed in the large deployment above, will divide work amongst themselves.

Deploying multiple Management Servers also provides basic fail-over capabilities, with the remaining servers continuing to operate in the event of the failure of one. Use of a Server Load Balancer, or SLB, provides failover for Enterprise Manager UI clients

in the event of a Management Server host failure, and it also balances the request load between the available Management Servers. SLBs are host machines dedicated for load balancing purposes. SLBs can be clustered to provide fail-over capability.

Using multiple hosts for the Management Repository assumes the use of Oracle Real Application Clusters (RAC). Doing so allows the same Oracle database to be accessible on more than one host system. Beyond the storage required for the Management Server, Management Repository storage may also be required. Management Server storage is less impacted by the number of management targets. The numbers suggested in the Enterprise Manager Cloud Control documentation should be sufficient in this regard.

12.2.1.1 Network Topology Considerations

A critical consideration when deploying Enterprise Manager Cloud Control is network performance between tiers. Enterprise Manager Cloud Control ensures tolerance of network glitches, failures, and outages between application tiers through error tolerance and recovery. The Management Agent in particular is able to handle a less performant or reliable network link to the Management Service without severe impact to the performance of Enterprise Manager as a whole. The scope of the impact, as far as a single Management Agent's data being delayed due to network issues, is not likely to be noticed at the Enterprise Manager Cloud Control system wide level.

The impact of slightly higher network latencies between the Management Service and Management Repository will be substantial, however. Implementations of Enterprise Manager Cloud Control have experienced significant performance issues when the network link between the Management Service and Management Repository is not of sufficient quality. The following diagram that displays the Enterprise Manager components and their connecting network link performance requirements. These are minimum requirements based on larger real world Enterprise Manager Cloud Control deployments and testing.

The latency between repositories and management servers should be half of that required in Release 11 of Enterprise Manger (15ms). This has a much larger impact on the user interface performance in Release 12 of Enterprise Manager.

12.2.2 Step 2: Periodically Evaluate the Vital Signs of Your Site

This is the most important step of the five. Without some degree of monitoring and understanding of trends or dramatic changes in the vital signs of your Enterprise Manager Cloud Control site, you are placing site performance at serious risk. Every monitored target sends data to the Management Repository for loading and aggregation through its associated Management Agent. This adds up to a considerable volume of activity that requires the same level of management and maintenance as any other enterprise application.

Enterprise Manager has "vital signs" that reflect its health. These vital signs should be monitored for trends over time as well as against established baseline thresholds. You must establish realistic baselines for the vital signs when performance is acceptable. Once baselines are established, you can use built-in Oracle Enterprise Manager Cloud Control functionality to set baseline warning and critical thresholds. This allows you to be notified automatically when something significant changes on your Enterprise Manager site. The following table is a point-in-time snapshot of the Enterprise Manager Cloud Control vital signs for two sites:

Module	Metrics	EM Site 1	EM Site 2
Site URL		emsite1.acme.com	emsite2.acme.com

Module	Metrics	EM Site 1	EM Site 2
Target Counts	Database Targets	192 (45 not up)	1218 (634 not up)
	Host Targets	833 (12 not up)	1042 (236 not up)
	Total Targets	2580 (306 not up)	12293 (6668 not up)
Loader Statistics	Loader Threads	6	16
	Total Rows/Hour	1,692,000	2,736,000
	Rows/hour/load/thread	282,000	171,000
	Rows/second/load thread	475	187
	Percent of Hour Run	15	44
Rollup Statistics	Rows per Second	2,267	417
	Percent of Hour Run	5	19
Job Statistics	Job Dispatchers	2	4
	Job Steps/second/dispatcher	32	10
Event Statistics	Events Processed (last hour)	536	1,100
Management Service Host Statistics	Average % CPU (Host 1)	9 (emhost01)	13 (emhost01)
	Average % CPU (Host 2)	6 (emhost02)	17 (emhost02)
	Average % CPU (Host 3)	N/A	38 (em6003)
	Average % CPU (Host 4)	N/A	12 (em6004)
	Number of CPUs per host	2 X 2.8 (Xeon)	4 X 2.4 (Xeon)
	Memory per Host (GB)	6	6
Management Repository Host Statistics	Average % CPU (Host 1)	12 (db01rac)	32 (em6001rac)
	Average % CPU (Host 2)		
	Average % CPU (Host 3)		
	Average % CPU (Host 4)		
	Number of CPUs per host		
	Buffer Cache Size (MB)		
	Memory per Host (GB)	6	12
	Total Management Repository Size (GB)	56	98
	RAC Interconnect Traffic (MB/s)	1	4
	Management Server Traffic (MB/s)	4	4
	Total Management Repository I/O (MB/s)	6	27
Enterprise Manager UI Page Response/Sec	Home Page	3	6

Module	Metrics	EM Site 1	EM Site 2
	All Host Page	3	30+
	All Database Page	6	30+
	Database Home Page	2	2
	Host Home Page	2	2

The two Enterprise Manager sites are at the opposite ends of the scale for performance.

EM Site 1 is performing very well with high loader rows/sec/thread and high rollup rows/sec. It also has a very low percentage of hours run for the loader and the rollup. The CPU utilization on both the Management Server and Management Repository Server hosts are low. Most importantly, the UI Page Response times are excellent. To summarize, Site 1 is doing substantial work with minimal effort. This is how a well configured, tuned and maintained Oracle Enterprise Manager Cloud Control site should look.

Conversely, EM Site 2 is having difficulty. The loader and rollup are working hard and not moving many rows. Worst of all are the UI page response times. There is clearly a bottleneck on Site 2, possibly more than one.

The following table outlines metric guidelines for the different modules based on tests that were run with the configurations outlined. It can serve as a reference point for you to extrapolate information and data based on the metrics and test environment used in the specified environment.

Table 12–4 Metric Guidelines for Modules Based On Test Environments

Module	Metrics	Value	Test Environment
Loader Statistics	NA	NA	**OMS Details**
	Total Rows/Hour	4,270,652	# of OMS Hosts = 2 # of CPU Per Host = 4 Intel Xeon Memory = 6 GB
	Rows/Hour/loaderthread	427,065	**Repository Details**
	Rows/second/loaderthread	120	# of Repository Nodes = 2 # of CPU per host = 4 Intel Xeon Memory = 6 GB
			EM Details
			Shared Recv Directory = Yes # of Agents = 867 # of Hosts = 867 Total Targets = 1803
			The Metrics are collected for 5 hours after 2 OMS instances were started and each agent had 50 MB of upload backlog files.
Rollup Statistics	Rows per second		
Job Statistics	Job Dispatchers	1 x Number of OMS instances	
	Job Steps/second/dispatcher		

Table 12–4 (Cont.) Metric Guidelines for Modules Based On Test Environments

Module	Metrics	Value	Test Environment
Notification Statistics	Notifications per second	16	**OMS Details** # of OMS Hosts = 1 # of CPU Per Host = 4 Intel Xeon Memory = 6 GB **Repository Details** # of Repository Nodes = 1 # of CPU per host = 4 Intel Xeon Memory = 6 GB **EM Details** # of OMS instances = 1 # of Repository Nodes = 1 # of Agents = 2474 # of Hosts = 2474 DB Total Targets = 8361
Alert Statistics	Alerts per hour	7200	**OMS Details** # of OMS Hosts = 1 # of CPU Per Host = 4 Intel Xeon Memory = 6 GB **Repository Details** # of Repository Nodes = 1 # of CPU per host = 4 Intel Xeon Memory = 6 GB **EM Details** # of OMS instances = 1 # of Repository Nodes = 1 # of Agents = 2474 # of Hosts = 2474 DB Total Targets = 8361
Management Service Host Statistics	Average % CPU (Host 1) Average % CPU (Host 2) Number of CPUs per host Memory per Host (GB)	31% 34% 4 (Xeon) 6	**OMS Details** # of OMS Hosts = 2 # of CPU Per Host = 4 Intel Xeon Memory = 6 GB **Repository Details** # of Repository Nodes = 2 # of CPU per host = 4 Intel Xeon Memory = 6 GB **EM Details** Shared Recv Directory = Yes # of Agents = 867 # of Hosts = 867 Total Targets = 1803 The Metrics are collected for 5 hours after 2 OMS instances were started and each agent had 50 MB of upload backlog files.

Table 12–4 (Cont.) Metric Guidelines for Modules Based On Test Environments

Module	Metrics	Value	Test Environment
Management Repository Host Statistics	Average % CPU (Host 1)	32%	**OMS Details** # of OMS Hosts = 2 # of CPU Per Host = 4 Intel Xeon Memory = 6 GB **Repository Details** # of Repository Nodes = 2 # of CPU per host = 4 Intel Xeon Memory = 6 GB **EM Details** Shared Recv Directory = Yes # of Agents = 867 # of Hosts = 867 Total Targets = 1803 The Metrics are collected for 5 hours after 2 OMS instances were started and each agent had 50 MB of upload backlog files.
	Average % CPU (Host 2)	26%	
	Number of CPUs per host	4	
	SGA Target	2 GB	
	Memory per Host (GB)	6	
	Total Management Repository Size (GB)	94	
	RAC Interconnect Traffic (MB/s)	1	
	Management Server Traffic (MB/s)		
	Total Management Repository I/O (MB/s)		
Enterprise Manager UI Page Response/Sec	Home Page	9.1 secs	**OMS Details** # of OMS Hosts = 1 # of CPU Per Host = 4 Intel Xeon Memory = 6 GB **Repository Details** # of Repository Nodes = 1 # of CPU per host = 4 Intel Xeon Memory = 6 GB **EM Details** # of OMS instances = 1 # of Repository Nodes = 1 # of Agents = 2474 # of Hosts = 2474 DB Total Targets = 8361
	All Host Page	9.8 secs	
	All Database Page	5.7 secs	
	Database Home Page	1.7 secs	
	Host Home Page	< 1 sec	

These vital signs are all available from within the Enterprise Manager interface. Most values can be found on the All Metrics page for each host, or the All Metrics page for Management Server. Keeping an eye on the trends over time for these vital signs, in addition to assigning thresholds for warning and critical alerts, allows you to maintain good performance and anticipate future resource needs. You should plan to monitor these vital signs as follows:

- Take a baseline measurement of the vital sign values seen in the previous table when the Enterprise Manager Cloud Control site is running well.
- Set reasonable thresholds and notifications based on these baseline values so you can be notified automatically if they deviate substantially. This may require some iteration to fine-tune the thresholds for your site. Receiving too many notifications is not useful.
- On a daily (or weekly at a minimum) basis, watch for trends in the 7-day graphs for these values. This will not only help you spot impending trouble, but it will also allow you to plan for future resource needs.

The next step provides some guidance of what to do when the vital sign values are not within established thresholds. Also, it explains how to maintain your site's performance through routine housekeeping.

12.2.3 Step 3: Use DBA and Enterprise Manager Tasks To Eliminate Bottlenecks Through Housekeeping

It is critical to note that routine housekeeping helps keep your Enterprise Manager Cloud Control site running well. The following are lists of housekeeping tasks and the interval on which they should be done.

12.2.3.1 Offline Monthly Tasks

Enterprise Manager Administrators should monitor the database built-in Segment Advisor for recommendations on Enterprise Manager Repository segment health. The Segment Advisor advises administrators which segments need to be rebuilt/reorganized and provides the commands to do so.

For more information about Segment Advisor and issues related to system health, refer to notes 242736.1 and 314112.1 in the My Oracle Support Knowledge Base.

12.2.4 Step 4: Eliminate Bottlenecks Through Tuning

The most common causes of performance bottlenecks in the Enterprise Manager Cloud Control application are listed below (in order of most to least common):

1. Housekeeping that is not being done (far and away the biggest source of performance problems)
2. Hardware or software that is incorrectly configured
3. Hardware resource exhaustion

When the vital signs are routinely outside of an established threshold, or are trending that way over time, you must address two areas. First, you must ensure that all previously listed housekeeping is up to date. Secondly, you must address resource utilization of the Enterprise Manager Cloud Control application. The vital signs listed in the previous table reflect key points of resource utilization and throughput in Enterprise Manager Cloud Control. The following sections cover some of the key vital signs along with possible options for dealing with vital signs that have crossed thresholds established from baseline values.

12.2.4.1 High CPU Utilization

When you are asked to evaluate a site for performance and notice high CPU utilization, there are a few common steps you should follow to determine what resources are being used and where.

1. Use the Processes display on the Enterprise Manager Host home page to determine which processes are consuming the most CPU on any Management Service or Management Repository host that has crossed a CPU threshold.
2. Once you have established that Enterprise Manager is consuming the most CPU, use Enterprise Manager to identify what activity is the highest CPU consumer. Typically this manifests itself on a Management Repository host where most of the Management Service's work is performed. Here are a few typical spots to investigate when the Management Repository appears to be using too many resources.

a. Click on the CPU Used database resource listed on the Management Repository's Database Performance page to examine the SQL that is using the most CPU at the Management Repository.

b. Check the Database Locks on the Management Repository's Database Performance page looking for any contention issues.

High CPU utilization is probably the most common symptom of any performance bottleneck. Typically, the Management Repository is the biggest consumer of CPU, which is where you should focus. A properly configured and maintained Management Repository host system that is not otherwise hardware resource constrained should average roughly 40 percent or less total CPU utilization. A Management Server host system should average roughly 20 percent or less total CPU utilization. These relatively low average values should allow sufficient headroom for spikes in activity. Allowing for activity spikes helps keep your page performance more consistent over time. If your Enterprise Manager Cloud Control site interface pages happen to be responding well (approximately 3 seconds) while there is no significant (constant) loader backlog, and it is using more CPU than recommended, you may not have to address it unless you are concerned it is part of a larger upward trend.

The recommended path for tracking down the root cause of high Management Repository CPU utilization is captured under step 3.b above. This allows you to start at the Management Repository Performance page and work your way down to the SQL that is consuming the most CPU in its processing. This approach has been used very successfully on several real world sites.

If you are running Enterprise Manager on Intel based hosts, the Enterprise Manager Cloud Control Management Service and Management Repository will both benefit from Hyper-Threading (HT) being enabled on the host or hosts on which they are deployed. HT is a function of certain late models of Intel processors, which allows the execution of some amount of CPU instructions in parallel. This gives the appearance of double the number of CPUs physically available on the system. Testing has proven that HT provides approximately 1.5 times the CPU processing power as the same system without HT enabled. This can significantly improve system performance. The Management Service and Management Repository both frequently have more than one process executing simultaneously, so they can benefit greatly from HT.

12.2.4.2 Loader Vital Signs

The vital signs for the loader indicate exactly how much data is continuously coming into the system from all the Enterprise Manager Agents. The most important items here are the percent of hour runs and rows/second/thread. The (Loader) % of hour run indicates whether the loader threads configured are able to keep pace with the incoming data volume. As this value approaches 100%, it becomes apparent that the loading process is failing to keep pace with the incoming data volume. The lower this value, the more efficiently your loader is running and the less resources it requires from the Management Service host. Adding more loader threads to your Management Server can help reduce the percent of hour run for the loader.

Rows/Second/Thread is a precise measure of each loader thread's throughput per second. The higher this number, the better. Rows/Second/Thread as high as 1200 have been observed on some smaller, well configured and maintained Enterprise Manager Cloud Control sites. If you have not increased the number of loader threads and this number is trending down, it may indicate a problem later. One way to overcome a decreasing rows/second/thread is to add more loader threads.

The number of Loader Threads is always set to one by default in the Management Server configuration file. Each Management Server can have a maximum of 10 loader threads. Adding loader threads to a Management Server typically increases the overall

host CPU utilization by 2% to 5% on a Enterprise Manager Cloud Control site with many Management Agents configured. Customers can change this value as their site requires. Most medium size and smaller configurations will never need more than one loader thread. Here is a simple guideline for adding loader threads:

Max total (across all Management Servers) loader threads = 2 X number of Management Repository host CPUs

There is a diminishing return when adding loader threads. You will not yield 100% capacity from the second, or higher, thread. There should be a positive benefit, however. As you add loader threads, you should see rows/second/thread decrease, but total rows/hour throughput should increase. If you are not seeing significant improvement in total rows/hour, and there is a constantly growing loader file backlog, it may not be worth the cost of the increase in loader threads. You should explore other tuning or housekeeping opportunities in this case.

To add more loader threads, you can change the following configuration parameter where *n* is a positive integer [1-10]:

em.loader.threadPoolSize=n

The default is 1 and any value other than [1-10] will result in the thread pool size defaulting to 1. This property file is located in the *{ORACLE_HOME}/sysman/config* directory. Changing this parameter will require a restart of the Management Service to be reloaded with the new value.

The following two parameters are set for the Receiver module which receives files from agents.

1. *em.loader.maxDataRecvThreads=n (Default 75)*

   Where *n* is a positive integer and default value is 75. This is used to limit the maximum number of concurrent data file receiver threads. So at the peak time only 75 receiver threads will be receiving files and an extra request will be rejected with a *Server Busy* error. These rejected requests will be resent by the agent after the default retry time.

   Care should be taken while setting this value as too high a value will put an increased load on the OMS machine and shared receiver directory box. If too low a value is set then data file receive throughput will be low.

2. *oracle.sysman.emRep.dbConn.maxConnForReceiver=n (Default 25)*

   Where *n* is a positive integer and default value is 25. This is used to set the maximum number of Repository Database connections for the receive threads. Oracle recommends you set this value equal to *em.loader.maxDataRecvThreads*, as each Receiver thread gets one DB session and there will be no wait for DB connections.

12.2.4.3 Rollup Vital Signs

The rollup process is the aggregation mechanism for Enterprise Manager Cloud Control. The two vital signs for the rollup are the rows/second and % of hour run. Due to the large volume of data rows processed by the rollup, it tends to be the largest consumer of Management Repository buffer cache space. Because of this, the rollup vital signs can be great indicators of the benefit of increasing buffer cache size.

Rollup rows/second shows exactly how many rows are being processed, or aggregated and stored, every second. This value is usually around 2,000 (+/- 500) rows per second on a site with a decent size buffer cache and reasonable speedy I/O. A downward trend over time for this value may indicate a future problem, but as long as % of hour run is under 100 your site is probably fine.

If rollup % of hour run is trending up (or is higher than your baseline), and you have not yet set the Management Repository buffer cache to its maximum, it may be advantageous to increase the buffer cache setting. Usually, if there is going to be a benefit from increasing buffer cache, you will see an overall improvement in resource utilization and throughput on the Management Repository host. The loader statistics will appear a little better. CPU utilization on the host will be reduced and I/O will decrease. The most telling improvement will be in the rollup statistics. There should be a noticeable improvement in both rollup rows/second and % of hour run. If you do not see any improvement in any of these vital signs, you can revert the buffer cache to its previous size. The old Buffer Cache Hit Ratio metric can be misleading. It has been observed in testing that Buffer Cache Hit Ratio will appear high when the buffer cache is significantly undersized and Enterprise Manager Cloud Control performance is struggling because of it. There will be times when increasing buffer cache will not help improve performance for Cloud Control. This is typically due to resource constraints or contention elsewhere in the application. Consider using the steps listed in the High CPU Utilization section to identify the point of contention. Cloud Control also provides advice on buffer cache sizing from the database itself. This is available on the database Memory Parameters page.

One important thing to note when considering increasing buffer cache is that there may be operating system mechanisms that can help improve Enterprise Manager Cloud Control performance. One example of this is the "large memory" option available on Red Hat Linux. The Linux OS Red Hat Advanced Server™ 2.1 (RHAS) has a feature called big pages. In RHAS 2.1, bigpages is a boot up parameter that can be used to pre-allocate large shared memory segments. Use of this feature, in conjunction with a large Management Repository SGA, can significantly improve overall Cloud Control application performance. Starting in Red Hat Enterprise Linux™ 3, big pages functionality is replaced with a new feature called huge pages, which no longer requires a boot-up parameter.

12.2.4.4 Rollup Process

The Rollup process introduces the concept of rollup participating instance; where rollup processing will be distributed among all participating instances. To add a candidate instance to the participating EMROLLUP group, the parameter instance_groups should be set on the instance level as follows:

- Add *EMROLLUP_1* to the *instance_group* parameter for node 1

  Add *EMROLLUP_2* to the *instance_group* parameter for node 2
- Introduce the *PQ* and *PW* parallel processing modes where:
  - PQ is the parallel query/parallel dml mode. In this mode, each participating instance will have one worker utilizing the parallel degree specified.
  - PW is the parallel worker mode. In this mode, each participating instance will have a number of worker jobs equal to the parallel level specified
- Distribute the work load for all participating RAC instances as follows:
  - Each participating instance will be allocated equal number of targets. So for (n) number of participating instances with total workload (tl), each instance will be allocated *(tl/n).*
  - Each worker on any participating instance will be allocated equal number of targets of that instance workload. So for (il) number of targets per instance with (w) number of workers, each worker will be allocated *(il/w).*
  - For each worker, the load is further divided into batches to control the number of times the rollup SQL is executed. The number of rows per batch will be the

total number of rows allocated for the worker divided by the number of batches.

Use the following recommendations as guidelines during the Rollup process:

- Use the parallel worker (PW) mode, and utilize the participating EMROLLUP_xx instance group.
- The recommendation is to use the parallel worker mode.
- Splitting the work among more workers will improve the performance and scalability until a certain point where the diminishing returns rule will apply. This is dependent on the number of CPUs available on each RAC node. In this test case, running with 10 workers was the optimal configuration, balancing the response time, machine CPU and IO utilization.
- It is important to set a proper batch size (10 recommended). The optimal run was the one with 10 batches, attributed to balancing the number of executions of the main SQL (calling EMD_1HOUR_ROLLUP) and the sort space needed for each individual execution.
- Start by setting the number of batches to 10 bearing in mind the number of batches can be changed based on the data distribution.

The recommendations above will yield the following results. Using the multi-instance parallel worker (8 PW) mode (with the redesigned code described earlier) improves the performance by a factor of 9-13 when utilizing two participating RAC instances.

Rollup row count (in millions) in MGMT_METRICS_1HOUR	Time (min)	Workers	Batch Size
29.5	30	8	1
9.4	5	8	10

** For the entire test there were 15779 distinct TARGET_GUID

** The test produced “29.5 Million” new rollup rows in MGMT_METRICS_1HOUR

Run **	Rows/Workers	Batches/Workers	Rows/Batch	Time (min)
8 PW /1 instance	3945	3945	1	40
8 PW /2 instances	1973	1973	1	30

12.2.4.5 Job, Notification, and Alert Vital Signs

Jobs, notifications, and alerts are indicators of the processing efficiency of the Management Service(s) on your Enterprise Manager Cloud Control site. Any negative trends in these values are usually a symptom of contention elsewhere in the application. The best use of these values is to measure the benefit of running with more than one Management Server. There is one job dispatcher in each Management Server. Adding Management Servers will not always improve these values. In general, adding Management Servers will improve overall throughput for Cloud Control when the application is not otherwise experiencing resource contention issues. Job, Notification, and Alert vital signs can help measure that improvement.

12.2.4.6 I/O Vital Signs

Monitoring the I/O throughput of the different channels in your Enterprise Manager Cloud Control deployment is essential to ensuring good performance. At minimum, there are three different I/O channels on which you should have a baseline and alert thresholds defined:

- Disk I/O from the Management Repository instance to its data files
- Network I/O between the Management Server and Management Repository
- RAC interconnect (network) I/O (on RAC systems only)

You should understand the potential peak and sustained throughput I/O capabilities for each of these channels. Based on these and the baseline values you establish, you can derive reasonable thresholds for warning and critical alerts on them in Cloud Control. You will then be notified automatically if you approach these thresholds on your site. Some Cloud Control site administrators can be unaware or mistaken about what these I/O channels can handle on their sites. This can lead to Enterprise Manager Cloud Control saturating these channels, which in turn cripples performance on the site. In such an unfortunate situation, you would see that many vital signs would be impacted negatively.

To discover whether the Management Repository is involved, you can use Cloud Control to check the Database Performance page. On the Performance page for the Management Repository, click on the wait graph showing the largest amount of time spent. From this you can continue to drill down into the actual SQL code or sessions that are waiting. This should help you to understand where the bottleneck is originating.

Another area to check is unexpected I/O load from non-Enterprise Manager Cloud Control sources like backups, another application, or a possible data-mining co-worker who engages in complex SQL queries, multiple Cartesian products, and so on.

Total Repository I/O trouble can be caused by two factors. The first is a lack of regular housekeeping. Some of the Cloud Control segments can be very badly fragmented causing a severe I/O drain. Second, there can be some poorly tuned SQL statements consuming much of the site I/O bandwidth. These two main contributors can cause most of the Cloud Control vital signs to plummet. In addition, the lax housekeeping can cause the Management Repository's allocated size to increase dramatically.

One important feature of which to take advantage is asynchronous I/O. Enabling asynchronous I/O can dramatically improve overall performance of the Cloud Control application. The Sun Solaris™ and Linux operating systems have this capability, but may be disabled by default. The Microsoft Windows™ operating system uses asynchronous I/O by default. Oracle strongly recommends enabling of this operating system feature on the Management Repository hosts and on Management Service hosts as well.

Automatic Storage Management (ASM) is recommended for Enterprise Manager Cloud Control repository database storage.

12.2.4.7 The Oracle Enterprise Manager Performance Page

There may be occasions when Enterprise Manager user interface pages are slow in the absence of any other performance degradation. The typical cause for these slow downs will be an area of Enterprise Manager housekeeping that has been overlooked. The first line of monitoring for Enterprise Manger page performance is the use of Enterprise Manager Beacons. These functionalities are also useful for web applications other than Enterprise Manager.

Beacons are designed to be lightweight page performance monitoring targets. After defining a Beacon target on an Management Agent, you can then define UI performance transactions using the Beacon. These transactions are a series of UI page hits that you will manually walk through once. Thereafter, the Beacon will automatically repeat your UI transaction on a specified interval. Each time the Beacon transaction is run, Enterprise Manager will calculate its performance and store it for historical purposes. In addition, alerts can be generated when page performance degrades below thresholds you specify.

When you configure the Enterprise Manager Beacon, you begin with a single predefined transaction that monitors the home page you specify during this process. You can then add as many transactions as are appropriate. You can also set up additional Beacons from different points on your network against the same web application to measure the impact of WAN latency on application performance. This same functionality is available for all Web applications monitored by Enterprise Manager Cloud Control.

After you are alerted to a UI page that is performing poorly, you can then use the second line of page performance monitoring in Enterprise Manager Cloud Control. This new end-to-end (or E2E) monitoring functionality in Cloud Control is designed to allow you to break down processing time of a page into its basic parts. This will allow you to pinpoint when maintenance may be required to enhance page performance. E2E monitoring in Cloud Control lets you break down both the client side processing and the server side processing of a single page hit.

The next page down in the Middle Tier Performance section will break out the processing time by tier for the page. By clicking on the largest slice of the Processing Time Breakdown pie chart, which is JDBC time above, you can get the SQL details. By clicking on the SQL statement, you break out the performance of its execution over time.

The JDBC page displays the SQL calls the system is spending most of its page time executing. This SQL call could be an individual DML statement or a PL/SQL procedure call. In the case of an individual SQL statement, you should examine the segments (tables and their indexes) accessed by the statement to determine their housekeeping (rebuild and reorg) needs. The PL/SQL procedure case is slightly more involved because you must look at the procedure's source code in the Management Repository to identify the tables and associated indexes accessed by the call.

Once you have identified the segments, you can then run the necessary rebuild and reorganization statements for them with the Management Server down. This should dramatically improve page performance. There are cases where page performance will not be helped by rebuild and reorganization alone, such as when excessive numbers of open alerts, system errors, and metric errors exist. The only way to improve these calls is to address (for example, clean up or remove) the numbers of these issues. After these numbers are reduced, then the segment rebuild and reorganization should be completed to optimize performance. These scenarios are covered in Section 12.2.3. If you stay current, you should not need to analyze UI page performance as often, if at all.

12.2.4.8 Determining the Optimum Number of Middle Tier OMS Servers

Determining the optimum number of middle tier OMS servers is not a trivial task. A number of data points must be considered for an informed, justified and acceptable decision for introducing additional OMS instances. The number of monitored targets is one of the first considerations, but its weight in decision making is normally not substantial.

The following items should be considered and examined as part of this exercise:

- The volume of job automation and scheduling used
- The number of administrators working simultaneously in the Console
- Network bandwidth and data channel robustness from agents to the OMS servers
- Number of triggered violations and notifications
- Speed and stability of the IO system the OMS servers use

Careful investigation of each category is essential to making an informed decision. In some cases, just adding an OMS server or providing more CPU or memory to the same host may not make any difference in performance enhancement. You can use the current running OMS instances to collect accurate statistics on current OMS performance to calculate the number of required OMS servers for current or future deployments. Enterprise Manager has "vital signs" that reflect its health. These vital signs should be monitored for trends over time as well as against established baseline thresholds.

12.2.5 Step 5: Extrapolating Linearly Into the Future for Sizing Requirements

Determining future storage requirements is an excellent example of effectively using vital sign trends. You can use two built-in Cloud Control charts to forecast this: the total number of targets over time and the Management Repository size over time.

Both of the graphs are available on the All Metrics page for the Management Service. It should be obvious that there is a correlation between the two graphs. A straight line applied to both curves would reveal a fairly similar growth rate. After a target is added to Enterprise Manager Cloud Control for monitoring, there is a 31-day period where Management Repository growth will be seen because most of the data that will consume Management Repository space for a target requires approximately 31 days to be fully represented in the Management Repository. A small amount of growth will continue for that target for the next year because that is the longest default data retention time at the highest level of data aggregation. This should be negligible compared with the growth over the first 31 days.

When you stop adding targets, the graphs will level off in about 31 days. When the graphs level off, you should see a correlation between the number of targets added and the amount of additional space used in the Management Repository. Tracking these values from early on in your Enterprise Manager Cloud Control deployment process helps you to manage your site's storage capacity proactively. This history is an invaluable tool.

The same type of correlation can be made between CPU utilization and total targets to determine those requirements. There is a more immediate leveling off of CPU utilization as targets are added. There should be no significant increase in CPU cost over time after adding the targets beyond the relatively immediate increase. Introducing new monitoring to existing targets, whether new metrics or increased collections, would most likely lead to increased CPU utilization.

12.2.6 Using Returning Query Safeguards to Improve Performance

On the All Targets page, Enterprise Manager uses a safeguard that prevents a flood of data from slowing performance and consuming excessive resources within the OMS by limiting the number of rows that can be returned from a query. By default, the limit is set to 2000, but an Enterprise Manager administrator can modify the limit with the following command:

```
emctl set property -name oracle.sysman.emSDK.eml.maxRows -value
2000
```

Providing a value equal to 0 will turn off the safeguard and fetch all rows. The new value takes immediate effect; no OMS restart is required. If the value is less than 0, the default value (2000) will be used instead. The only way to indicate that no limiting should be performed is to set the value to exactly 0.

When there are too many results returned from a query and this limit comes into effect, the following message appears under the results table:

"This table of search results is limited to 2000 targets. Narrow the results by using Refine Search or Search Target Name. See the tuning guide for how to modify this limit."

Similar behaviors (and messages) are applied to other large tables throughout Enterprise Manager. The same OMS property (`oracle.sysman.emSDK.eml.maxRows`) controls the maximum limit for all of them together. This matches the behavior (and reuses the existing property) from previous Enterprise Manager releases.

13

Enterprise Configuration Management

This chapter explains how Enterprise Manager Cloud Control simplifies the monitoring and management of the deployments in your enterprise. The chapter contains the following sections:

- Targets and Configuration Collections
- Cloud Control and Configuration Management

13.1 Targets and Configuration Collections

Enterprise Manager collects configuration information for all managed targets that have a running Management Agent. The agent periodically sends the configuration information to the Management Repository over HTTP or HTTPS, allowing you to access up-to-date configuration information for your entire enterprise through Cloud Control.

Cloud Control enables you to view, save, track, compare, search, and customize collected configuration information for all managed targets. Supported target types include:

- Databases (a database instance, for example)
- Groups, Systems, Services (a database system, for example)
- Middleware (application deployments, WebLogic Servers, and WebLogic Domains, for example)
- Servers, Storage, Network (hosts and virtual machines, for example)
- Others, such as Oracle Home
- Internal, such as agents and OMS

Table 13–1lists examples of configuration information collected for a cross-section of target types.

Table 13–1 Collected Configurations for Various Targets

Target Type	Collected Configuration Information
Host[1]	▪ Hardware (includes memory, CPU, I/O device, and network information) ▪ Operating system (includes installed patches and patch sets) ▪ Oracle software (includes installed products and their components, patch sets, and interim patches applied using OPatch) ▪ Other software (includes all software registered with the operating system)

Table 13–1 (Cont.) Collected Configurations for Various Targets

Target Type	Collected Configuration Information
Database[2]	■ Database and instance properties ■ Initialization and System Global Area parameters ■ Tablespace, datafile, and control file information ■ Redo logs, rollback segments, and high availability information ■ Licensing information
Middleware such as WebLogic Server	■ Node Manager, machine, Web service, and Web service port configurations ■ Resource Adapter, including outbound ■ Web and EJB modules ■ Server information ■ JDBC Datasource and Multi Datasource ■ Resource usage ■ Virtual hosts ■ Startup Shutdown classes ■ Jolt Connection Pool ■ Work Manager ■ JMS Topic, Queue and Connection Factory ■ Network channels
Client[3]	■ Hardware ■ Operating system (includes properties, file systems, patches) ■ Software registered with the operating system ■ Network data (includes latency and bandwidth to the Web server) ■ Client-specific data that describes configuration for the browser used to access the client configuration collection applet ■ Other client-oriented data items
Non-Oracle Systems	■ Hardware details including vendor, architecture, CPU, and I/O device information. ■ Operating system details including name, version, software and package lists, kernel parameters, and file system information. ■ OS Registered software including product name, vendor, location, and installation time.

[1] The default collection period for host configuration information is 24 hours.
[2] The default collection period for database configuration information is 12 hours.
[3] Refer to "Client Configurations" in this chapter for more information.

13.2 Cloud Control and Configuration Management

Use Cloud Control to manage enterprise configurations:

- Search collected configuration data
- Compare configurations
- View latest and saved configurations as well as inventory and usage details
- Monitor configuration history for changes
- Build custom configurations and introduce custom target types

- Collect and analyze external client configurations

13.2.1 Configuration Search

Use Configuration Search to search configuration data across the enterprise. Enterprise Manager ships with a set of out-of-box configuration searches, which you can use as a starting point to explore the volume of configuration data collected. As you work with a provided search, you can tailor the search criteria to refine or broaden the results, and save the altered search under a new name.

Perform powerful searches across the enterprise using sophisticated combinations of search filters, options, and relationships. Consider these search examples:

- Show all hosts with dual core CPUs
- Show all targets in a specific geographic area
- Show all tablespaces having at least one 10MB datafile
- Show all Oracle Homes installed on Linux 5.6 hosts
- Show all targets monitored by a particular agent
- Find all database instances with initialization parameters p1 and p2 having values v1 and v2, respectively

Enhance the search filtering criteria by adding your own SQL query statements. Save interesting search results by printing a report or exporting to a file.

To access the search capability, from the **Enterprise** menu select **Configuration**, then select **Search**. See the Cloud Control online help for information on setting up and executing configuration searches.

13.2.2 Configuration Comparisons

Enterprise Configuration Management deals with the collection, storage, and monitoring of configuration data tied to managed entities within the enterprise. A host, for example, has configuration item types related to its hardware and software components—number of CPUs, memory, IO devices, OS platform and version, installed software products, and so forth.

Changes to configuration data invariably happen, typically because of common events like patches and upgrades. At some point a change to one component can affect the overall system in a negative way. Detecting the root cause becomes paramount.

The comparison tool enables you to compare configurations of a target with configurations of another target of the same type. The comparisons can be done on the current configuration or configurations previously saved (perhaps, for example, just before applying a patch or doing an upgrade).

A comparison template is an exemplar for fine-tuning a comparison of like configurations. A template is associated with a specific target type, which determines the configuration item types, items, and properties to be compared. A set of default templates ships out-of-box to support various target types.

A template enables you to establish certain constants to take into account when comparing configurations of the given target type; for example, which property differences to ignore, and which property differences trigger an alert. You also can use constraints to establish acceptable values for specific properties. A configuration being compared that does not comply with the constraint constitutes a difference.

A template can invoke rules, or expressions, to be evaluated in determining when there is a match for comparison purposes, and when to disregard differences detected in a comparison.

For systems, you design a system template that references member templates, based on the target types that make up the system. Create the member templates before you create the system template.

Comparisons allow you to do the following:

- Ignore certain attributes during a comparison
- Notify key personnel when differences are detected
- Design and share comparison templates with other administrators
- Schedule a comparison to run on a recurring basis
- Compare complete target systems; match target system members automatically or manually
- Compare configuration file data as raw file content or in a parsed format
- Drill down in comparison results to address detected differences
- Execute a comparison again with a single click on the Jobs page

To access the comparison feature, from the **Enterprise** menu select **Configuration**, then select:

- **Compare** to access the comparison wizard.
- **Comparison Templates** to manage templates.
- **Comparison Job Activity** to view and rerun comparison jobs and drill down to comparison results.

Note: You must have the CREATE_JOB privilege to see these menu items.

For information on managing comparison templates and setting up configuration comparisons, see Chapter 14, "Configuration Comparisons, Templates, and Rules."

13.2.3 Configuration Views

Use the Configuration Browser to view configuration data in the context of a single managed entity. Configuration data can include:

- Configuration item types and properties
- System configuration data as well as all system members and their configuration data
- System and target relationships (immediate, member of, uses, used by, and so forth)
- Custom configuration collection data

The browser window consists of left and right panes. The left pane is a tree hierarchy. The right pane consists of tabs that display information in tables. As you navigate in the tree, your selection dictates the contents in the right pane. Depending on the selection, tabs appear containing data such as properties and values, relationships, a

hierarchical structure of a system and its members, and file contents in both a parsed and raw text format.

The viewed data can be the latest collected or previously saved. Saved configurations are snapshots in time of collected data preserved for future reference. You may simply want to view the saved data, or you may want to use it for some other purpose such as the basis of a comparison.

To view a target's configuration data, right-click the target of interest and select **Configuration**, then select **Saved** to search for a previous collection, or **Last Collected** to view the most recent data collection.

Cloud Control offers several summary views to monitor the state of the enterprise. From these views, you can drill down to the details and take various actions. Use the Inventory and Usage Details page to:

- View inventory summaries for deployments such as hosts, database installations, and fusion middleware installations on an enterprise basis or for specific targets.
- View inventory summary information in the context of different dimensions. For example, for host inventory summary, you can view by platform, vendor, or OS version.
- Drill down multiple levels of inventory details.
- See trends in inventory counts charted across a time line. Chart bars are color-coded to match the view selection.
- Switch to a pie chart to break down the inventory data for the rollup option by color-coded percentages.
- Click a patch indicator to link to patch details.
- Repeatedly revise selections to refresh chart and details based on new selections.
- Export deployment and details tables toa CSV file.

To view these summaries and drill down to details, from the **Enterprise** menu select **Configuration**, then select **Inventory and Usage Details**. Alternatively, click **See Details** in the **Inventory and Usage** region of the **Enterprise Summary** page.

See the Cloud Control online help for information on working within Configuration Browser and Inventory and Usage Details.

13.2.4 Configuration History

Use Configuration History to monitor change activity across the enterprise. The history is a log of changes to a managed entity (target) recorded over a period of one year; it includes changes both to configurations and to relationships. Relationships are the associations that exist among managed entities.

While viewing a configuration history you can:

- Track changes to targets over time by specifying and refining search criteria.
- View change history and manipulate how the information is presented.
- Annotate change records with comments that become part of the change history.
- Schedule a history search to capture future changes based on the same criteria.
- View the status of scheduled history jobs.
- Notify others of future change detection.
- Save change history details to a file.

You can also view a list of all current and past history searches. Use search criteria to filter the list of history jobs. For example, show all scheduled history searches started over the past 24 hours; or, show all successful history searches involving hosts started over the past 31 days. The jobs engine purges history jobs older than 31 days

To access Configuration History, from the **Enterprise** menu, select **Configuration**, then select **History**. See the Cloud Control online help for information on performing tasks related to Configuration History.

13.2.5 Custom Configuration Specifications

Custom configurations provide end users the means to define configurations to collect that Enterprise Manager has no way of knowing about. These customized configurations can be collected on well-known target types or on target types introduced as part of the custom configuration definition.

A custom configuration is a specification intended for deployment to an agent-monitored target where the agent uses it to gather configuration data about target instances. A custom configuration can be a combination of the following:

- File specifications—configuration files in a specified directory on the target to collect and upload to the repository
- Command specifications—commands and scripts to run against the target, given appropriate credentials, and upload command/script output as configuration data to the repository
- Query specifications—SQL database queries to run against a database on the target, given appropriate credentials, and upload query results to the repository

The configuration data that the agent collects and uploads is stored in both raw and parsed form. The custom configurations application has a host of out-of-box parsers that you can use to convert collected configuration data into a standard format for storing in the repository. The format is a tree of ordered nodes, also known as containers. Containers have names and possibly an additional identifier in square brackets determined by rules that you specify for the purpose of distinguishing the container from other containers. This is important when comparing configurations and tracking change history to avoid flagging "false positive" differences. It also aids in specifying search criteria and crafting SQL queries used in compliance rules.You can view collected configuration data in both raw and parsed form.

To work with custom configurations, from the **Enterprise** menu, select **Configuration**, then select **Custom**. For information on creating, maintaining, and deploying custom configuration specifications, including detailed parser reference information, see Chapter 15, "Custom Configurations, Parsers, and Rules."

13.2.6 Client Configurations

A "client" represents an end-user or customer system—a system that is not part of your own IT infrastructure. A "client configuration" represents the configuration data collected about the end-user's system. These configurations differ from the internal deployments that you manage using Cloud Control.

The Client System Analyzer (CSA) application allows Web server administrators to collect and analyze data from end-user systems. The client data is collected by an applet, diagnosed, and sent back to the CSA application. You can either use the CSA application that comes pre-installed with Cloud Control, or you can deploy CSA independently to your Web server.

To access client configurations, from the **Enterprise** menu, select **Configuration**, then select **Client Configurations**. See the Cloud Control online help for information on performing tasks related to client configurations.

13.2.6.1 Client System Analyzer in Cloud Control

Using the pre-installed application allows you to collect client data without having to set up a separate Web server. The Management Agents collect, analyze, and upload the client data to the Management Repository. End users do not need login credentials to access Cloud Control. Example usage scenarios include:

- End users who call the Help Desk may be asked to navigate to the out-of-box CSA page so that their system information is uploaded. The Technical Support Representative can then review the system information and offer solutions.
- The client's application can be changed to provide an "Upload my system information" link to the Client System Analyzer in the Cloud Control application. The link can specify certain configuration parameters, such as the URL to return to after the Client System Analyzer runs.
- The client's application can be modified to redirect its users to the Client System Analyzer in the Cloud Control page during login or at other points in the application. Collected information can then be used from within Cloud Control to obtain various bits of information about the client systems. Examples include most popular browser versions, or systems that do not have a necessary Operating System patch applied or do not have enough RAM.

To access the CSA application, from the **Enterprise** menu, select **Configuration**, then select **Client System Analyzer**. See the Cloud Control online help for information on working with the CSA application.

13.2.6.2 Client System Analyzer Deployed Independently

CSA can be deployed independently to any J2EE-capable Web server. This deployment strategy is appropriate when:

- Clients accessing CSA cannot reach or have limited access to a Cloud Control deployment; for example, due to a firewall.
- Further customization to the CSA application is required, such as:
  - Custom rules can be supplied to the CSA application so that the end users have immediate feedback as to whether their systems satisfy certain constraints.
  - The behavior of the applet can be changed to collect additional information or to present end users with additional or different user interfaces.
  - The load on the Management Service Web servers needs to be reduced.

Both pre-installed and standalone types of deployments assign a configurable identifier called a Client Configuration Collection Tag to every client configuration collection. After the client configuration data has been collected by the client configuration collection applet and written to the Web server directory specified by the CSA application, you must configure Cloud Control to collect the client configuration data and upload it to the Management Repository.

See the Cloud Control online help for information on collecting and viewing client configurations.

14

Configuration Comparisons, Templates, and Rules

This chapter describes the template creation process and the use of rules in the process. It also provides information on setting up comparisons and managing comparison templates.

14.1 Comparison Templates

A comparison template enables you to establish certain constants to take into account when comparing configurations of the given target type; for example, which property differences to ignore, and which property differences trigger an alert. You can use constraints to establish acceptable values for specific properties. A configuration being compared that does not comply with the constraint constitutes a difference.

A template can invoke rules, or expressions, to be evaluated in determining when there is a match for comparison purposes, and when to disregard differences detected in a comparison.

14.1.1 Create or Edit a Comparison Template

Use these instructions when creating a new template or editing an existing template; this includes create-like.

1. From the **Enterprise** menu, select **Configuration**, then select **Comparison Templates**.

   For a new template, click **Create** and provide a name and target type. To base a template on an existing one, select the template row, click **Create Like**, and provide a name. In either case, the action creates a new template row.

2. Select the appropriate template row in the table and click the **Edit** button. The **Template Settings** tab appears.

   The compared configurations' target type drives the hierarchy of configuration item types and configuration items on the left. The settings in play for the respective properties on the right derive from the selected template, unless you are creating a new template from scratch, in which case there are no settings.

   A system comparison takes an overall template and a template for each system member. Thus there is an additional tab for **Member Settings**. Edit the tab as follows:

   - Optionally select the member template to use for each system member type.

- For any given member type, you can elect not to compare configurations by clearing the check box.
- For member types that you are comparing, select a target property to use as a matching key. The default is target name, but typically you would want to use a distinguishable property to align comparison entities, such as department or location.

3. In the **Template Settings** tab, select a configuration item type or item in the left pane to expose its properties in the right pane. A key icon denotes a property that is defined as a key column in the configuration item type's metadata.
4. Click the **Property Settings** tab and check boxes for property differences to be ignored and alerted as appropriate. They are mutually exclusive. When you ignore differences in a property value in this fashion, you are doing so unconditionally for all differences detected in the property value for the configuration item type.

   Use a value constraint rule to filter the property value. In this case, the comparison engine compares the property value in the configurations being compared (the second through n configurations) to the constrained value. A property value that fails to satisfy the constraint constitutes a difference. For example, test for a version equal to or greater than 6. Any instance in the compared configurations of a version property value under 6 constitutes a difference. Clearly, you would not set a value constraint if you checked ignore differences. See "Specifying Rules" for details.
5. Repeat the preceding steps to set additional property settings on other configuration items.
6. Optionally, select an item in the left pane and click the **Rules for Matching Instances** tab. For a given property, specify a rule expression to be evaluated to determine when a match exists between configuration instances. In other words, if the expression resolves to true, compare the instances. See "Specifying Rules" for details.

   Match rules are column-based; they apply an AND logical operator. If you specify rules for multiple properties, they must all resolve to true to constitute a match.
7. Optionally, select an item in the left pane and click the **Rules for Ignoring Instances** tab. For a given property, specify a rule expression to be evaluated to determine when to ignore differences in configuration instances. In other words, if the expression resolves to true, disregard whatever differences the comparison detects. See "Specifying Rules" for details.

   Ignore rules are row-based; they apply an AND logical operator within a subset of rules and an OR logical operator between rule subsets. So, if you specify two rules for property A and two rules for property B, either both rules set on property A OR both rules set on property B must resolve to true to constitute a match.

14.1.2 Managing Templates

In addition to creating and editing comparison templates, you manage them by doing the following:

- View a template's settings and composition; this is read-only
- Delete a template (requires the proper permissions)
- Share templates by exporting them in XML file format and importing them into other Enterprise Manager systems

View a Comparison Template

You can view out-of-box templates and other users' templates to which you have access. Viewing a template is read-only: you see its makeup, but you cannot change anything, even temporarily.

1. Select a template in the Comparison Templates manager and click the **View** button.
2. Expand items in the tree on the left and peruse the settings and rules on the various tabs.
3. Notice that the **Save** button is disabled. Click **Cancel** to return to the Comparison Templates manager.

Delete a Comparison Template

Deleting a template is subject to the following constraints:

- You cannot delete a comparison template unless you have the proper permissions.
- You cannot delete a default comparison template.
- You cannot delete a comparison template currently in use.

To delete a template, select it in the Comparison Templates manager, click **Delete**, and confirm the operation.

Export a Comparison Template

Use the export feature to save a template as an external file that can be imported into another Enterprise Manager system.

1. Select a template in the Comparison Templates manager and click the **Export** button.

   A platform-specific file dialog opens. For example, if you are using Firefox, the dialog notes that you have chosen to open the named template, which it identifies as an XML file. The dialog asks what you want Firefox to do, open the file in an XML editor or save the file.
2. Select the save radio button and click **OK**.
3. Browse to the desired location in the file system and save the file, changing the name if applicable. You cannot change the name of a default (out-of-box) template on export.

Import a Comparison Template

Any comparison template import must comply with the comparison template .xsd. So, for all intents and purposes, the import should be a previously exported template to ensure compliance.

1. In the Comparison Templates manager click the **Import** button.
2. Browse to the template file location and click **Import**.

   The imported template appears as a new row in the template table.

An exported template is associated with its owner. A template whose owner is not the same as the login ID of the person importing the template retains its original ownership. If you want to be the owner of the imported template, you have to edit the `owner` attribute in the template XML file prior to import, changing the value to your login ID. Or, you can simply remove the attribute, in which case the default owner will be set to the ID of the person initiating the import operation.

The Template Manager disallows import of a default (out-of-box) template of the same name. Similarly, you could change the `name` attribute in the template XML file prior to import to allow the import to occur.

14.2 Comparison Wizard

Comparisons are an important factor in managing the enterprise. The comparison wizard walks you through the process of setting up a comparison. Setting up a comparison involves five steps, six if you are comparing systems:

- Select the first configuration in the comparison (the one to compare against)
- Select additional configurations (the one or more configurations to compare to the first configuration)
- Select a comparison template to frame the comparison (or no template)
- When comparing systems, map system members in the first configuration to members of the other configurations
- Schedule the comparison job and set up e-mail notifications
- Review your work and submit the job

A follow-on step would be to review the results and drill down to differences details.

14.2.1 Select a Configuration to Compare Against

The first step in setting up a comparison is to select a configuration against which to compare one or more other configurations. When you open the Comparison Wizard (from the **Enterprise** menu, select **Configuration**, then select **Compare**), available configuration collections appear in a table at the bottom.

1. Choose between the latest and a saved configuration. You can "mix and match" in that you can compare latest to saved and vice versa. When you choose saved, filtering criteria on the right become active so that you can enter dates and a description.
2. Specify filtering criteria to narrow the search. Minimally, you would probably want to select a target type, and for a saved configuration, a before or after date. Click **Search**.
3. In the list of matching configurations that appears at the bottom of the page, select the one to be the benchmark and click **Next** or **Comparison Configurations** on the workflow train.

14.2.2 Select Configurations to Compare

The next step is to select one or more configurations to compare to the first configuration you selected.

1. Click **Add Configurations** to open the Search and Select Configurations dialog.
2. As for the first configuration, choose between latest and saved, and enter filtering criteria to narrow the search. Click **Search**.
3. In the results list, select one or more configurations (you can multiselect) and then click **OK**.
4. Click **Next** or **Comparison Template** on the workflow train.

14.2.3 Select a Template to Use in the Comparison

Optionally, you can elect not to use a template (the default), in which case you can skip this step. You must use a template you select as is; that is, you cannot alter any settings for this particular comparison. If you want to change certain settings in a template, use the create-like feature to create a new template, based on an existing one.

Depending on the comparison target type, there may be a default template available. When done selecting a template, click **Next** or **Schedule and Notify** or **Mapping**, as appropriate, on the workflow train.

14.2.4 Map Members in a System Comparison

Mapping pertains exclusively to systems. It's a way to selectively indicate how members of respective systems should match up in a comparison. The mappings you see initially are system-defined, based on template selection. If you are comparing multiple systems, select a different system target to see the mappings for that system.

You can control which mapped members to compare by deselecting the **Compare Configurations** check box in the table header and then selecting table rows individually, or by deselecting rows within the table.

Template-based mappings may not account for all situations. Click **Create Mapping** to manually pair member configurations to compare. Select the system member type from the drop-down list and map a member of the first system to a member of the second system by selecting members in the respective table. Click **OK**.

Note that you can remove only those mappings that you manually create. When done with mapping, click **Next** or **Schedule and Notify** on the workflow train.

14.2.5 Schedule the Comparison and Create a Notification List

On the Schedule and Notify page, you schedule the comparison to run as a background job. The comparison can be one-time-only or run on a recurring basis. You can run the comparison immediately or at some later date. This also is where you supply e-mail addresses to which to send differences alerts.

Note: If you schedule a recurring job, you can subsequently change the job-related settings when viewing the results on the Jobs page. Click the **Edit** button and then go to the **Schedule** tab.

---

1. The comparison job must have a name. Although the system supplies a default name that identifies it by date and time as a configuration comparison, you may want to enter a meaningful name for the job.
2. Specify the job schedule:
   - **If not now, when.** Click **Later** to activate the calendar widget where you can select a date and time.
   - **How often.** Select report frequency in the drop-down list. Default is one-time-only.
   - **Wait how long.** If the job fails to run as scheduled, cancel within a specified time frame.
   - **Keep going.** Maintain the job schedule for the specified period.

3. Enter the e-mail addresses of those who are to be notified when the comparison detects a difference. Use a comma to separate addresses. Remember that the properties for which differences are alerted were specifically selected in the comparison template.

   Addressees entered here must have been properly set up for notifications (from the **Setup** menu, select **Notifications** on the console home page).

4. Click **Next** or **Review and Submit** on the workflow train.

14.2.6 Review the Comparison Parameters and Submit the Job

Review the comparison parameters before submitting the job:

- Is the benchmark configuration, that is, the first one, correct?
- Are the configurations you are comparing against the benchmark correct?
- Are you using the template you want?
- Is the job name suitable?
- Is the job scheduling as intended?
- Have you entered properly formatted e-mail addresses for differences alerts?

If you need to change anything, go back to the appropriate page; otherwise, click **Submit** to schedule the comparison job.

The Jobs page opens, showing a summary of the submitted job. Eventually the Jobs page reveals the results of the comparison; that is, whether the comparison detected differences or found the configurations to be the same. In either case (different or the same), the reported result is a link to the details of the comparison, although presumably different is the more interesting result. A separate entry appears for each compared configuration. In other words, if configurations B and C are being compared to A, there is a result for A compared to B, and a result for A compared to C.

14.3 Working with Comparison Results

This section covers comparison results from the following perspectives:

- Comparisons and Job Activity
- System Comparisons
- Single Target Comparisons

14.3.1 Comparisons and Job Activity

Comparisons run as scheduled jobs. Part of the process of setting up a comparison is to define the schedule: run once-only or on a recurring basis; run immediately or at some later time; retry after failure; and so forth. Given the permutations, you can have many jobs to keep track of.

To view comparison job activity, from the **Enterprise** menu, select **Configuration**, then select **Comparison Job Activity**.

Note that the comparison job activity page also affords the opportunity to resubmit comparisons already run; that is, you do not have to go through the entire comparison workflow to run a subsequent execution of the same comparison.

View a list of all current and past comparisons. Use search criteria to filter the list of comparison jobs. For example, show all failed comparisons started over the past 24

hours; or, show all successful comparisons involving hosts started over the past 31 days. The jobs engine purges comparison jobs older than 31 days.

Select a table row and click **View Result** to go to the Jobs page that reports the comparison result. From there you can drill down to results details. The job name is a hyperlink that takes you to the same place. Use the bread crumb on the Jobs page to navigate back to the list. The comparison jobs you can view beyond your own depend on your role and access level granted.

Select a table row and click **Resubmit Comparison Now** to run a new submission of a previously executed comparison, characterized by the following.

- As this is a new job, a date and time stamp distinguishes it from its predecessor.
- The job is scheduled to execute immediately.
- The job will execute only once.

An informational message confirms the job, which appears as a new row in the table.

If you are the job owner or otherwise have the proper access level, you can perform list maintenance by deleting comparison jobs that no longer have relevance.

14.3.2 System Comparisons

Results of a system comparison appear when the scheduled comparison job completes and you click the Different link on the Jobs page. This view summarizes the comparison and job details in the two regions at the top.

The bottom region summarizes the comparison results as follows:

- The left pane displays a hierarchy of system and member target types.
- The right pane reports comparison results based on the mappings established as part of comparison setup. A boxed 1 (left only) or 2 (right only) means there was nothing to compare to the first or second target, respectively. When this is the case, you can return to the mapping setup step and either create a mapping between the targets or clear the configuration check box for the member type.
- Click the Different link to see the differences between member targets. Use the bread crumb link at the top of the details page to return to the system comparison results page.

14.3.3 Single Target Comparisons

Difference details display when you click the Different link in the comparison results view on the Jobs page. The details view summarizes the comparison and job details. You may want to collapse these regions to provide more page real estate for the difference details.

Difference details break down as follows:

- The left pane displays a hierarchy of configuration items for the target type being compared, and, if applicable, custom configurations. Refine the scope of comparison results as follows:
  - Select the **Show Differences Only** check box to eliminate the "noise" of same and ignored results.
  - Select the **Show Ignored Properties** check box to display properties the comparison template ignores. Key properties and properties that are the same or different display be default. This option is disabled if you selected **Show Differences Only** in the left pane.

- The display on the right depends on the selection on the left.
  - Select a configuration item type on the left to display tables on the right where the top table contains rows of select (key) configuration properties and the bottom table contains rows of configuration properties and values. Use the filter feature to narrow the scope of configuration item types.
  - Select a row in the top table to move its configuration property to the top row of the bottom table.
  - Select a custom configuration file on the left to display file contents on the right. File contents that include property values for the compared configurations display both in raw and parsed form on separate tabs.

The icons that appear in this view are mostly intuitive: equal–same, not equal–different. The key icon denotes the key properties of the configuration item type. An indication of Out of Range means that the property value failed a value constraint set on the property. A boxed 1 (left only) or 2 (right only) means that the comparison did not find a matching item to compare to the first or second configuration, respectively. When this is the case, you can return to template setup and invoke a rule to create a match, and then rerun the comparison.

Configuration Key Properties

You might wonder about the column names that appear in the top table on the right. These columns represent configuration item type key properties. If the configuration item type does not have declared key properties, the comparison engine takes the top four properties in the CI type database table to serve as key stand-ins for purposes of matching up configurations. The comparison engine upholds the same precedence (top four properties in the database table) if for some reason the comparison is set up to ignore key properties (not recommended).

14.4 Specifying Rules

Specify rules in the context of creating or editing a comparison template (see"Create or Edit a Comparison Template").

Rules enable you to parse configuration data in order to fine-tune comparisons. In terms of the comparison, a rule applies the expression to the value of the selected item in the configuration instance that is being compared to the benchmark configuration. Matching rules are intended to devise a comparison key that aligns the instances being compared. Ignore rules are intended to establish a basis for disregarding any differences detected between instances being compared.

Value Constraint Rules

Specify value constraint rules as follows:v

1. Select a configuration item in the left pane.
2. Click the **Property Settings** tab in the right pane and select the property on which you want to set a value constraint.
3. Click the **Edit Rule** button in the toolbar. In the dialog that opens:
   a. Select an operator from the drop-down list.
   b. Type an operands expression and then click **OK**.

   To clear a rule, select the table row and click the **Remove Rule** button in the toolbar.

See "Rules Expression Reference" for details on the formation of a rules expression.

Matching Rules

Specify matching rules as follows:

1. Select a configuration item in the left pane.
2. Click the **Rules for Matching Instances** tab in the right pane and then click **New**.

   Click **Show Key Properties** to see which properties are defined as key columns in the selected configuration item type's metadata.
3. Select a property in the drop-down list that appears under **Property Name**.
4. To create the rule, select the table row and click the **Edit Rule** button in the toolbar. In the dialog that opens:

   a. Select an operator from the drop-down list.

   b. Type an operands expression and then click **OK**.

   c. To specify additional rules, click **New** and repeat Steps a and b

   To clear a rule, select the table row and click the **Remove Rule** button in the toolbar.

   See "Rules Expression Reference" for details on the formation of a rules expression.

   You can enter additional rules for the same or for a different configuration item. When there are multiple rules, they resolve in the order specified. Matching rules take an AND logical operator, which means all conditions must resolve to true to constitute a match.

Ignore Rules

Specify ignore rules as follows:

1. Select a configuration item in the left pane.
2. Click the **Rules for Ignoring Instances** tab in the right pane. and then click **New**.

   Click **Show Key Properties** to see which properties are defined as key columns in the selected configuration item type's metadata.
3. Select a property in the drop-down list that appears under **Property Name**.
4. To create the rule, select the table row and click the **Edit Rule** button in the toolbar. In the dialog that opens:

   a. Select an operator from the drop-down list.

   b. Type an operands expression and then click **OK**.

   c. To specify additional rules, click **New** and repeat Steps a and b

   To clear a rule, select the table row and click the **Remove Rule** button in the toolbar.

   See "Rules Expression Reference" for details on the formation of a rules expression.

   You can enter additional rules for the same or for a different configuration item. When there are multiple rules, they resolve in the order specified. Ignore rules take an AND logical operator for rules within a subset, and an OR logical operator between subsets. So, for two subsets, each with multiple rules, all rules in the first subset OR all rules in the second subset must resolve to true to constitute a match.
5. Select **New Or** to indicate the end of one rule subset and the beginning of another.

14.5 Rules Expression Reference

A rule consists of an operator and operands. Taken together, they form an expression that resolves to a value that is then compared to the value of the selected item. A true condition satisfies the rule.

Operands can be literals (string literals are enclosed in single quotes), legal numbers, or dates of the form `YYYY-MM-DD HH24:MI:SS.FF`. Operands that directly reference the value of a configuration item must be of the same date type as that value. Operands in square brackets in the syntax are optional.

Operator	Operands
is equal to*	An optional literal value to match; string values are case-sensitive; if unspecified, expression evaluates value of the property to which the rule applies Note that a matching rule compares the values of the configuration items in the respective configuration to one another, not to a third specified value, so the operator does not take an operand in this case. `[match-literal]`
is case-insensitive equal to*	An optional case-insensitive string literal; if unspecified, expression evaluates value of the property to which the rule applies Note that a matching rule compares the values of the configuration items in the respective configuration to one another, not to a third specified value, so the operator does not take an operand in this case. `['match-literal']`
is greater than or equal†	A literal value to match; required `match-literal`
is greater than †	A literal value to match; required `match-literal`
is less than or equal to†	A literal value to match; required `match-literal`
is less than†	A literal value to match; required `match-literal`
is one of†	A comma-separated list of literal values, at least one of which must be specified, but only one of which need match `match-literal-1[,match-literal-n,...]`
is between†	A range specified as start and end literal values; both must be specified; range is inclusive `start-range-literal , end-range-literal`
contains†	A string literal on which to perform pattern matching; required `[FALSE|TRUE,] 'pattern-literal'` `FALSE` (default) means string must comply with Oracle `LIKE` operator syntax; `TRUE` means string must comply with Posix regular expression syntax

Operator	Operands
replace‡	A string literal to match and replace with a second string literal `[FALSE|TRUE,]'pattern-literal'[,'replacement-literal'][,position-integer][,occurrence-integer]` `FALSE` (default) means string must comply with Oracle `LIKE` operator syntax; `TRUE` means string must comply with Posix regular expression syntax `TRUE` enables optional positional integer argument to indicate where within the column value to extract the string, and optional occurrence integer argument to indicate the position count to replace Mandatory pattern literal represents the string value to match If the replacement string literal is unspecified, replace the matched string literal with nothing
substring‡	Extract specified segment of string value `[FALSE|TRUE,]position-integer[,length-integer][,'pattern-literal'[,occurrence-integer]]` `FALSE` (default) means string must comply with Oracle `LIKE` operator syntax; `TRUE` means string must comply with Posix regular expression syntax Mandatory positional integer argument indicates where to begin string extraction: ■ If 0 or 1, returns all characters ■ If positive integer, starts extraction from beginning ■ If negative integer, starts extraction counting backwards from end Optional length integer argument indicates character count starting at position integer pattern literal represents the value to match; optional if the first argument is `FALSE`; required if `TRUE` occurrence integer argument indicates character count to match; valid only if pattern literal is specified

Notations are as follows:

*–Enabled for value constraints, matching rules, and ignore rules
†–Enabled for value constraints and ignore rules only
‡–Enabled for matching rules only

14.6 Rule Examples

These rule examples assume that you are in the process of creating or editing a template and are at the point where you have selected the configuration item in the tree on the left.

Matching Rule Examples

Suppose, when comparing the hardware of host configurations, you want, for matching purposes, to ignore case in respective vendor names. Here's a simple rule to make the comparison case-insensitive.

1. In the **Rules for Matching** tab, click **New**.
2. Select **Vendor Name** in the drop-down list.

3. Select the table row and click the **Edit Rule** button in the toolbar to open the rule dialog.
   - Set **Operator** to `is-case-insensitive-equal-to`. As this operator takes no operands for a matching rule, you are done.
   - Click **OK**.

You want to compare WebLogic Servers, aligning on server name, where the names are different: ManagedServer1 and ManagedServer2, for example. To ensure the comparison occurs, you need to fashion a match on server name.

1. In the **Rules for Matching Instances** tab, click **New**.
2. Select **Machine Name** in the drop-down list.
3. Select the table row and click the **Edit Rule** button in the toolbar to open the rule dialog.
   - Set **Operator** to `substring`.
   - Set **Operands** to `1, 13`.
   - Click **OK**.

   Effectively, the rule says use the first 13 characters of the name (ManagedServer), thus excluding the qualifying integer.
4. Another way to achieve the same result:
   - Set **Operator** to `replace`.
   - Set **Operands** to `true, '(*)(\d*)', '\1'`.
   - Click **OK**.

   This example uses a regular expression (`TRUE`) to resolve all characters prior to the qualifying integer.

For a more advanced example, consider a database instance comparison that requires a match on Datafiles filenames within a Tablespace, where filenames are of the form:

```
/u01/jblack_abc2d/oracle/dbs/dabc2/mgmt_ad4j.dbf
```

1. In the **Rules for Matching Instances** tab, click **New**.
2. Select **File Name** in the drop-down list.
3. Select the table row and click the **Edit Rule** button in the toolbar to open the rule dialog.
   - Set **Operator** to `replace`.
   - Set **Operands** to `true, '(/u01/)(.*)(oracle.*/dabc[0-9]+.*/)(.*)', '\2\4'`.
   - Click **OK**.

Effectively, the rule says use a regular expression (`TRUE`) to construct a matching key from the value between `/u01/` and `oracle`, combined with what remains of the original filename after `dabc2 /`, or `jblack_abc2d/mgmt_ad4j.dbf`.

Ignore Rule Examples

Generally, you use ignore rules to ignore differences in collections that are row-oriented, as opposed to column-oriented. Custom Configuration snapshots, for example, are row-oriented data collections.

Say, for example, you wanted to ignore in Custom Configuration parsed data, any row where the property `Attribute` identifies an internal ID or checksum.

1. In the **Rules for Ignoring Instances** tab, click **New**.
2. Select **Attribute** in the drop-down list.
3. Select the table row and click the **Edit Rule** button in the toolbar to open the rule dialog.
   - Set **Operator** to `is one of`.
   - Set **Operands** to `'id', 'checksum'`.
   - Click **OK**.

The rule ensures that the comparison ignores any row in the collection data that contains either of the specified values.

Now consider an ignore rule that demonstrates how the comparison engine applies the logical operators AND and OR against the same configuration item type. In this example the objective is to ignore rows in Custom Configuration parsed data when any of three rule sets satisfies the following conditions:

Data Source = `'sqlnet.ora'` AND **Attribute** = `'ADR_BASE'`
OR
Data Source = `'tnsnames.ora'` AND **Attribute** = `'HOST'`
OR
Data Source = `'resources.xml'` AND **Attribute** = `'authMechanismPreference'`

Notice that the comparison engine applies the AND operator to rules within a set and the OR operator between rule sets. Rules for ignoring instances support inheritance; thus, in this case, the Data Source property is available in rules creation, as demonstrated in the example.

1. In the **Rules for Ignoring Instances** tab, click **New**.
2. Select **Data Source** in the drop-down list.
3. Select the table row and click the **Edit Rule** button in the toolbar to open the rule dialog.
   - Set **Operator** to `is equal to`.
   - Set **Operands** to `'sqlnet.ora'`.
   - Click **OK**.
4. Click **New** and select **Attribute** in the drop-down list.
5. Select the table row and click the **Edit Rule** button in the toolbar to open the rule dialog.
   - Set **Operator** to `is equal to`.
   - Set **Operands** to `'ADR_BASE'`.
   - Click **OK**.
6. Click **New Or** to insert a logical OR operator to signal the end of the first rule set.
7. Add two new rules where **Data Source** is equal to `'tnsnames.ora'` and **Attribute** is equal to `'HOST'`.
8. Click **New Or** to insert a logical OR operator to signal the end of the second rule set.

9. Add two new rules where **Data Source** is equal to `'resources.xml'` and **Attribute** is equal to `'authMechanismPreference'`.

The comparison ignores any row in the collection data that satisfies any of the three rule sets.

15

Custom Configurations, Parsers, and Rules

This chapter provides detailed information on the management and implementation of custom configuration specifications, and how they use parsers to transform and store configuration data. The chapter serves as a reference for the suite of parsers that ships with the application, and includes examples of how parsers and rules work together to align configuration data in operations such as comparisons.

15.1 Manage Custom Configurations

Custom configurations are stored in and managed from the Custom Configurations library. To access the library, from the **Enterprise** menu, select **Configuration**, then select **Custom**. A list of all available specifications appears by default. Use search criteria to find a specific custom configuration.

Use the following actions to manage custom configurations:

Action	Description
Create	Create a new custom configuration. See "Create or Edit a Custom Configuration" for more information.
Create Like	Create a new custom configuration based on the selected specification. See "Create or Edit a Custom Configuration" for more information.
Edit	Make edits to the selected specification. See "Create or Edit a Custom Configuration" for more information.
View Details	View the make-up of the selected specification. See "View a Custom Configuration" for more information.
Delete	Delete the selected specification. See "Delete a Custom Configuration" for more information.
Enable Facet Synchronization	Synchronize the selected specification with facets in the Compliance Library for real-time facet monitoring. See "Enable Facet Synchronization" for more information.
Export	Save the selected specification as an XML file. See "Export a Custom Configuration" for more information.
Import	Import a specification as an XML file from the local file system. See "Import a Custom Configuration" for more information.
Deploy	Indicate the monitored target from which to collect configuration data based on the specification. See "Deploy and Undeploy Custom Configurations" for more information.

To perform these actions, click the respective button, or select a specification in the table and click the button, as appropriate.

15.1.1 Create or Edit a Custom Configuration

Use the instructions below to create, create like, or edit a custom configuration.

Given appropriate privileges, you can edit a custom configuration and save the edited version, in which case, the version number increases. You might also edit and save as a draft, or edit a draft for publishing. Note that when you edit a custom configuration, you cannot change the target type, as this would cause the underlying metadata to be incompatible with existing deployments of the custom configuration.

See the section "About Custom Configurations and Privileges" for information on privileges required to perform various actions on custom configurations.

Note: When you edit a deployed custom configuration, it is automatically redeployed upon saving. This does not apply to saving as draft.

---

1. In the Custom Configurations library, click the **Create** button; or, select an existing specification in the library and click **Create Like** or **Edit**.
2. On the Create Custom Configuration page, enter a name for the custom configuration and an optional description. The create like action requires minimally that you rename the specification.
3. Select a target type. If no currently defined target type satisfies your requirements, click the **Create Custom Target Type** button to the right of the target type drop-down list. Type a name and click **OK**. The new type now appears in the drop-down list of target types.

   To create a new target type, ensure that the administrator has installed a software library (**Setup > Provisioning and Patching > Software Library**). This must be done once, after Enterprise Manager installation.
4. Optionally, set up a sample target. A sample target resides on the host from which you intend to collect configuration data. If you do not set up a sample target, you cannot browse the file system or use the preview feature in entering your specifications. You can select an existing target instance as a sample, or add a new one. Typically, you would add a new target instance to match a custom target type added in Step 4.
   - To select an existing target instance, click the appropriate link. A dialog opens containing known instances of the target type. Use the filtering criteria as necessary to locate the instance you want and then click **Select**.
   - To add a new target instance, click the appropriate link. As instructed in the dialog that opens, you must first select an agent to monitor the target you are adding. Next, click **Add Target**. In the dialog that opens, provide target properties appropriate to the instance target type. Minimally, provide values for required properties (denoted by an asterisk). For a new target instance that matches a custom target type, the pertinent target property is the path to install home, as this is the likely location of configuration files relevant to the custom target type.
5. See "Files & Commands Tab" for instructions on how to complete the tab.
6. See "SQL Tab" for instructions on how to complete the tab.
7. After you complete the specification definition and have mapped credentials to the target type, use the preview feature to validate your entries, in particular, to ensure the parsed view is what you expect.

8. Save the new or edited specification. Remember that custom configurations are in the public domain. Use the save-as-draft feature to keep the specification private while you test and refine it. See "About Save, Save As, and Versioning" for more information on the ramifications of save actions.

   If you are editing a draft, the buttons change as follows:

   - **Publish** implies that you are making the draft public.
   - **Save** implies that you are creating the next version of the draft.

When done, you can begin collecting configuration data by deploying the custom configuration to target instances. See "Deploy and Undeploy Custom Configurations" for more information.

Files & Commands Tab

Create file and command specifications as follows:

1. Click the search icon to browse to a default base directory location. This is where the configuration files reside, or where the commands you specify are to execute.

   Click the **Use Property** button to open a dialog where you can select a target property to include as part of the directory path. These properties serve as variables, denoted by curly braces, to be substituted with actual values at runtime. You can type additional text in the box to supplement your selection. So, for example, you might select OracleHome and append a directory–`{OracleHome}/config`–to collect files on the target located in the config subdirectory under the Oracle Home path. Note that the target type definition determines available target properties. User-defined properties do not appear in the list, as they are not available at the agent.

2. Click **Advanced Settings** to specify the following:

   - An alternate base directory for the sample target.
   - The encoding to use in collecting the data at the agent. Configuration data is stored in UTF-8 format in the repository. Oracle Default means use UTF-8 for XML files and the locale encoding of the target for all other file types; Target Locale means store all file types including XML in the locale encoding of the target; otherwise, select an encoding from the drop-down list. Selecting directly from the list automatically selects the accompanying radio button.
   - Whether to use the agent credentials (file specification only) or some other predefined credential set to access data on the target. If the customized credential set does not appear in the drop-down list, click **Create** to identify the credential set to use. Note that you must then specify the credentials that map to the credential set name you create. If you don't know a mapped name, you can specify a credential set when you open the Remote File Browser to add files as described in Step 3. See "Set Up Credentials" for more information.

3. Click **Add** and select file or command as the specification type.

   For a **file specification**, enter a file name in the space provided or browse the base directory to select a file on the target. Use of wildcards (* and **) is allowed, where ** indicates 0 or more subdirectories. In using wildcards (and as a general caveat), ensure that collections do not result in too many (or too large) files, and that the files collected be configuration-related, that is, files under administrative control that change relatively rarely, so as not to overload Enterprise Manager.

   For a **command specification**, enter command syntax in the space provided or browse the base directory to a script. You must assign a unique alias to the

command. The alias you assign appears in the configuration browser as a link when viewing the custom configuration hierarchy. When you click the link, it opens the command specification in the tab on the right. The same caveats as mentioned for files apply to command output; that is, that their results are constrained in number and size, and to configuration-related data.

Select a parser to convert the configuration file or command output into a standard format for storing in the repository. There is no default. If you do not specify a parser, only the raw data format is stored and available for viewing. See "About Parsers" for more information.

Optional. Specify post-parser rules to align tree nodes. See"Set Up Rules"for information on entering rules.

4. Repeat Step 3 to specify additional files or commands.

Return to the section "Create or Edit a Custom Configuration" and resume with Step 8.

SQL Tab

Create SQL query specifications as follows:

1. Select credentials to use to connect to the database. If the customized credential set does not appear in the drop-down list, click **Create** to identify the credential set to use. Note that you must then specify the credentials that map to the credential set name you create. Custom configurations only support database credentials with NORMAL roles, not those with SYSDBA, SYSOPER, or other roles.
2. Specify a JDBC connection to an Oracle database from which to extract data via an SQL query. The connection string can be either a URL or an abstraction of database target properties. It cannot be a combination of the two; that is, partial URL and some target properties.

   The URL must contain the name of the target database host, applicable port number, and the Oracle Service name (SID); for example, `mydatabase.us.oracle.com:1521:ORCL`.

   If you want to use target properties, leave the field blank. At runtime the application will substitute values for these target properties—{MachineName}{Port}{SID}—to make the connection.
3. Click **Add** and type or paste a SQL query in the provided text box. Ensure that the query is sufficiently selective to return only pertinent configuration-related data of manageable size and scope.

   You must assign a unique alias to the query. The alias you assign appears in the configuration browser as a link when viewing the custom configuration hierarchy. When you click the link, it opens the SQL query in the tab on the right.

   Database Query Parser should be preselected in the drop-down list.

   Optional. Specify post-parser rules to align tree nodes. See"Set Up Rules"for information on entering rules.
4. Repeat Step 3 to specify additional SQL queries.

Return to the section "Create or Edit a Custom Configuration" and resume with Step 8.

Set Up Credentials

If you create a credential set while creating a custom configuration, you have to specify the credentials that make up the credential set. To do this, you have to return to the Custom Configurations library and proceed as follows:

1. From the **Setup** menu (top right of the page next to the Help menu), select **Security**, then select **Monitoring Credentials**.
2. Select the applicable target type in the table and click **Manage Monitoring Credentials**.
3. Select the row with the credential set name you created during the custom configuration definition for the given target type and click **Set Credentials**.
4. Enter the username and password for the credential set and click **Save** (or **Test and Save** for database credentials).
5. Return to the "Files & Commands Tab" or "SQL Tab" description.

Set Up Rules

Use rules to differentiate nodes in the parsed representation that have the same name. This is particularly important in comparisons and change history when trying to match nodes in the parsed tree, or when expressing SQL queries to verify compliance. Rules resolve to an identifier that is appended in square brackets to node text in the tree as a way of uniquely identifying the node. An operation such as a comparison will then use the combination of node text and bracketed identifier for evaluation purposes.

A rule consists of a condition and an expression, both of which must be valid XPath expressions. The condition resolves to a node that requires the identifier. The expression resolves to a string computation for the identifier. You can use a special case `SKIP` expression to bypass the node specified in the condition; this is a convenient way to eliminate "noise." In other words, for purposes of comparison, ignore the node the condition resolves to.

Some parsers have default parser rules already defined. They execute automatically on the parsed representation. You can elect to use a subset of default rules, edit them, or override them with custom rules that you define.

The number in the **Rules** column is significant. Initially, the number is zero (0). A whole number greater than zero indicates the number of custom rules defined. Zero also appears for a parser that has default parser rules. So the appearance of a whole number in the column stipulates an override of default parser rules, if any, with the custom rules the number represents.

Set up rules as follows:

1. Click the **Parser Rules** button. The Edit Parser Rules page displays.
2. To define a custom rule, click **Add**. In the table row that appears, enter a condition and an expression as valid XPath expressions.

   You can define multiple rules; they are applied to the parsed content in the order specified. Click **Return** when you are done.

   Select a table row to delete a custom rule.
3. To manipulate default rules, click **Add Default Rules**.

   Rules appear in table rows, provided the parser you selected has default parser rules. Edit and delete default rules as appropriate to your purposes. Remember that you are working with a copy of these rules; the originals remain safely intact.

   Note that if you delete all rules, you are merely removing the copies you imported. Default parser rules will still fire unless overridden by custom rules.

Return to the "Files & Commands Tab" or "SQL Tab" description.

15.1.2 View a Custom Configuration

You can view a custom configuration in read-only mode to get an idea of the make-up of a specification. Perhaps, for example, to see if it is a likely candidate on which to base a new specification.

1. In the Custom Configurations library, select the specification table row and click **View Details**.
2. Peruse the settings and rules on the various tabs.

15.1.3 Enable Facet Synchronization

You can synchronize a custom configuration specification with real-time monitoring facets to monitor real-time changes to the configuration files and queries that make up the custom configuration. Real-time monitoring enables you to know such things as when files and database settings change, who made the change, whether observations were automatically reconciled, whether the actions observed were authorized, and so forth.

When you synchronize custom configurations with real-time monitoring facets, future changes to custom configurations automatically propagate to corresponding facets, which means configurations are not only collected, compared, tracked, and so forth, but also are monitored for authorized real-time changes. Note that to associate a custom configuration with a facet and to subsequently edit a custom configuration synchronized with a facet requires the additional role of EM_COMPLIANCE_DESIGNER.

1. In the Custom Configurations library, select the specification table row and click **Enable Facet Synchronization**.
2. The **Facet Synchronization** column displays a **Use Facet** link in the custom configuration table row. Click the link to go to the **Real-time Monitoring Facets** tab in the Compliance Library where you can manage the synchronization of facets with the custom configuration.

About Real-time Monitoring Rules and Facets

A real-time monitoring rule monitors operating system and database level entities that store configuration data. The rule defines the entities to monitor, user actions to watch, and any types of filters to apply to the monitoring, including when changes occurred, who made the changes, and what process made the changes. The real-time monitoring rule definition includes facets that are used to determine what is important to monitor for a given target type, target properties, and entity type.

A facet is a collection of patterns that make up one attribute of a target type. For example, you may choose to define a facet that lists all of the critical configuration files for the host target type. These are the files that, if changed, would most likely result in host instability. Another facet might be one that lists all DBA users.

For a given target type, you can create any number of facets. A target type might have a facet that identifies critical configuration files, another facet that identifies log files, another that identifies executable files, another that identifies database tables with sensitive configuration data, and so forth. The sum of all of these facets for a given target type constitutes what is important to monitor for the given target type in terms of compliance.

A facet that is specific to a target type can also include a combination of target type properties. A facet for a host target type on Windows is different from a facet for a host

target type on Linux. If no target type criteria are set, a facet is assumed to apply to all criteria (any target of this type).

15.1.4 Export a Custom Configuration

You can export a custom configuration as an XML file that can subsequently be imported into the same or another system.

1. In the Custom Configurations library, select the specification table row and click **Export**.
2. Browse to a file system location where you want to save the specification as an XML file. The saved file takes the name of the custom configuration by default.

15.1.5 Import a Custom Configuration

Given appropriate privileges, you can import a custom configuration that was previously exported as an XML file.

1. In the Custom Configurations library, select the specification table row and click **Import**.
2. Browse to the file location. Select the file and click the **Import** button on the dialog.

   The imported specification appears in the Custom Configurations library.

15.1.6 Delete a Custom Configuration

You must be the owner or otherwise have sufficient privileges to delete a custom configuration. Note that there are dependencies that potentially impact deletion, including deployments, job schedules, existing collections, and so forth.

1. In the Custom Configurations library, select the specification table row and click **Delete**.
2. The system validates permissions and otherwise checks for dependencies that might prevent the deletion, although some dependencies cannot be verified until a job submission involving the custom configuration.

15.1.7 Deploy and Undeploy Custom Configurations

Deployment of a custom configuration means to direct the specification to a target where a monitoring agent will collect configuration data based on the specification's definition. A custom configuration can be deployed to multiple targets. You must have sufficient privileges to deploy and undeploy custom configurations.

To deploy a custom configuration:

1. In the Custom Configurations library, select the specification table row and click **Deploy**.
2. On the Deployments page, click **Add**. In the dialog that opens, search for and select targets of the specified target type where you want to deploy the custom configuration.
3. When you close the dialog (click **Select**), a new column appears denoting a pending action of **Deploy** and the status becomes **Selected for deployment**.
4. Proceed as follows:

- Click **Apply** to confirm the action while remaining on the Deployments page. The action column disappears, and the status becomes **Deployment job in progress**.
- Click **OK** to schedule the deployment and return to the library.
- Click **Cancel** to void the request and return to the library.

5. Click **Refresh Status** on the Deployments page to confirm a successful outcome.

If you update a deployed custom configuration, redeployment occurs automatically.

To undeploy a custom configuration:

1. Select the deployment in the table.
2. Click **Remove**. A new column appears denoting a pending action of **Undeploy**; status remains **Deployed**.
3. Proceed as follows:
   - Click **Apply** to confirm the action while remaining on the Deployments page. The action column disappears, and the status becomes **Undeployment job in progress**.
   - Click **OK** to schedule the undeployment and return to the library.
   - Click **Cancel** to void the request and return to the library.
4. Click **Refresh Status** on the Deployments page to confirm a successful outcome.

When viewing custom configurations in the library, a green check mark in the Deployments column denotes a currently deployed custom configuration. Click the check mark to open the Deployments page.

Edit a Deployment

To edit a deployment:

1. In the Custom Configurations library, locate the appropriate table row and click the deployments link.
2. On the Deployments page, select the deployment in the table and click **Edit**.
3. The type of custom configuration, that is, file/command-based or SQL-based, determines the make-up of the dialog that opens. Specify a base directory to override the default base directory currently in effect, or change the JDBC URL, as appropriate.
4. Proceed as follows:
   - Click **Apply** to confirm the action while remaining on the Deployments page. The action column disappears, and the status becomes **Redeployment job in progress**.
   - Click **OK** to schedule the redeployment and return to the library.
   - Click **Cancel** to void the request and return to the library.
5. Click **Refresh Status** on the Deployments page to confirm a successful outcome.

Note that the edit applies to the deployment of the specification; it does change the custom configuration definition.

View a Configuration Collection

You must have sufficient privileges to view a custom configuration's collected data.

1. In the Custom Configurations library, locate the appropriate table row and click the deployments link.
2. On the Deployments page, select the deployment in the table and click **View Configuration**.
3. In the configuration browser popup window, peruse details of the custom configuration by selecting nodes in the tree hierarchy on the left:
   - The root node represents the target instance being monitored. The right pane displays target properties and immediate relationships.
   - The next level down in the tree represents a template for the specification. The right pane displays specification details such as configurations being collected and the base directory from which they are collected.
   - The remaining leaf nodes in the tree represent the configuration data collected. The right pane displays the configuration data in both parsed and raw format.

15.2 About Save, Save As, and Versioning

When you create a custom configuration, you have options to save or save as draft. A normal save action makes the specification publicly available to the general user community. A save as draft action keeps the specification private to you. How you use these actions when creating and editing specifications influences the mechanics of versioning. Consider the following scenarios:

- You create and save a custom configuration; this is public version 1. You subsequently edit public1 and save as a draft; this becomes draft1. Public1 is still generally available. You edit draft1 and publish; this becomes public2. Note that in parallel, someone else with the proper permissions can also edit public1 and save as a draft to create version 1 of draft2.
- You create and save a custom configuration as a draft; this is version1 of draft1. You edit and save again; this becomes version 2 of draft1. Repeat the edit-and-save operation; this becomes version 3 of draft1. Edit version 3 of draft1 and publish; this becomes public version 1.

15.3 About Custom Configurations and Privileges

Working with custom configurations requires privileges specific to the given operation you want to perform.

Operation	Required Privilege (Role)
Create new target type	EM_PLUGIN_OMS_ADMIN To create a new target type, ensure that the administrator has installed a software library (**Setup > Provisioning and Patching > Software Library**). This must be done once, after Enterprise Manager installation.
Create new target instance	EM_PLUGIN_AGENT_ADMIN
Create or import custom configuration	"Manage custom configurations owned by user" (or the more powerful "Manage custom configurations owned by any user")
Associate custom configuration with an auto-synchronized real-time monitoring facet	EM_COMPLIANCE_DESIGNER

Operation	Required Privilege (Role)
Edit or delete custom configuration	Differs, depending on the specific activity within the realm of editing: ■ Custom configuration owner requires "Manage custom configurations owned by user"; nonowner requires "Manage custom configurations owned by any user" ■ Schedule redeployment jobs for already deployed targets requires "Create" privilege for Job System resource type ■ For custom configurations associated with real-time monitoring facet, EM_COMPLIANCE_DESIGNER
Deploy or undeploy custom configuration on a target	"Manage target metrics" privilege on the target instance; "Create" privilege for Job System resource type (to schedule deployment/undeployment;) EM_PLUGIN_AGENT_ADMIN (to deploy a plug-in to an agent)
Create a new credential set	Superuser
View custom configuration definition	None
View custom configuration collected data	Regular "target instance view" privilege

Note that editing an imported custom configuration may be restricted to edits that do not change the version, depending on options set during export. One such permissible edit would be to credential set information.

15.4 About Parsers

A Parser takes raw configuration data and parses it into a nested attribute structure. This structure is a tree hierarchy where nodes are containers and leaves are name value pairs of attributes, or properties. The structure is well-suited to representation in a graphical user interface.

Cloud Control includes a host of parsers out-of-box. Each parser consists of a base parser and, in many cases, parser parameters. Some parsers also contain post-parsing rules. A base parser refers to a parser category that is capable of parsing data of a particular format. Parser parameters provide a way to tailor the base format to accommodate variations in data formatting. Post-parsing rules are a mechanism for aligning nodes in the tree that otherwise have no distinct identity. This is important when comparing configurations and tracking change history to avoid flagging "false positive" differences. It also aids in specifying search criteria and crafting SQL queries used in compliance rules.

There are four base parser varieties:

- XML
- Format-specific
- Columnar
- Properties

Some parsers have out-of-box default rules. These rules address well-known instances where nodes need additional clarifying information, or, in some cases, to be skipped altogether (ignored) . Specifically, the WebLogic and WebSphere parsers contain default rules to address such instances. You can leave these rules as is, execute a subset of them, or replace them with your own custom rules.

Note: Parser parameters described in the following sections that take regular expressions as values conform to the syntax and semantics of the Java package `java.util.regex`. Visit `http://download.oracle.com/javase/tutorial/essential/regex/` for more information on Java regular expressions.

15.5 Manage Parsers

While creating, editing, or viewing custom configurations, you can peruse the list of available parers, their default parameters, and post-parser rules, if applicable. Parser parameters dictate formatting such as comment character, delimiters, start and end characters, and so forth. You cannot edit these parameters, but you can export a parser as an XML file, edit the file, and import it back into the application under a new name. Some parsers also have default rules that serve to align nodes in the parsed tree for purposes of comparison, for example.

1. While working with a custom configuration, click **Manage Parsers**. A list of available parsers appears in a table. The column on the right (Base Parsers) denotes a general parser category, Properties for example, which implies file types that contain name/value pairs.
2. Select a parser and click **Details**. This dialog also shows default rules, if any.
   - Click the **Parameters** tab to see the parameter defaults in effect. You can then judge if you need to edit the parser to conform with your file format conventions.
   - Click the **Default Rules** tab to see the post-parsing rules that ship with certain parsers. This is a good way to get exposure to rules formation.
3. Assume you want to change the delimiter character in a given parser.
   a. With the parser selected in the table, click **Export**.
   b. In the dialog that opens click **Save** and navigate to a file system location. Save the XML file with an appropriate name.
   c. In making your edits, be sure to change the parser ID and parser name in the XML, as you are creating a customized version of an out-of-box parser.
4. Assume you edited a shipped parser and now want to import it for use in creating custom configurations.
   a. With the Parsers table open, click **Import**.
   b. In the dialog that opens, browse to the file location where you edited the exported parser file. Select it and click **Import** on the dialog.

   The edited parser file now appears in the Parsers table where it can be used in custom configuration creation.

Note: The parser you edit must be one of the base parsers. For a complete list of supported parsers and details about parser parameters, see Enterprise Manager reference documentation.

15.6 XML Parsers

Cloud Control has two XML parsers: a default (attribute-keyed) XML parser and a generic XML parser.

15.6.1 Default XML Parser

Parsing occurs as follows:

- XML elements with no XML attributes or child elements become parsed attributes; all other elements become containers.
- XML attributes become parsed attributes.
- Element text content becomes a parsed attribute, with its name dependent on whether or not the tag contains any XML attributes. If the tag contains XML attributes, the parsed attribute name takes the value specified in the `STORE_CONTENT_AS` parameter; otherwise, the parsed attribute name takes the tag name.

The default XML parser accepts the following parameters:

Parameter	Description
`MULTIKEY_DELIMITER`	Delimiter that separates a list of XML attribute names in the `CONTAINER_NAME` parameter; default is tilde (~)
`STORE_CONTENT_AS`	Name to give to parsed attributes derived from element text content, where the element contains XML attributes; default is `text_value`

`CONTAINER_NAME`	A list of XML attribute names delimited by the value of the `MULTIKEY_DELIMITER` parameter. If an attribute name in this list appears in a tag in the original file, the tag becomes a container named for the value of the XML attribute. All other XML attributes become parsed attributes as usual. The tag name itself is discarded.
	For example, the list includes attribute names Moe and Larry in this order. The original file contains an XML tag Stooges that has attributes Moe, Larry, and Curly. As Moe appears first in the delimited list, its value, leader, becomes the parsed container name; Larry and Curly become parsed attributes. The tag name Stooges is discarded. The original XML fragment might be as follows:
	`<?xml version="1.0" encoding="UTF-8"?>`
	`<Comedy>`
	`<Stooges Moe="leader", Larry="zany", Curly="bald">`
	`</Stooges>`
	`</Comedy>`

15.6.1.1 WebLogic Attribute-keyed Parser

Cloud Control provides an out-of-box attribute-keyed parser specifically designed to parse the WebLogic `config.xml` file. It has the same parameters as the default XML parser and comes with 26 default post-parsing rules to uniquely identify nodes with the same name.

15.6.1.2 WebSphere Attribute-keyed Parsers

Cloud Control provides several out-of-box attribute-keyed parsers designed to parse specific WebSphere configuration files. Each parser has the same parameters as the default XML parser and comes with a set of default post-parsing rules to uniquely identify nodes with the same name. There are parsers for the following WebSphere configuration files:

- `node.xml` (1 default post-parsing rule)
- `plugin-cfg.xml` (7 default post-parsing rules)

- `resource.xml` (9 default post-parsing rules)
- `server.xml` (13 default post-parsing rules)
- `variables.xml` (1 default post-parsing rule)

15.6.2 Generic XML Parser

Parsing occurs as follows:

- All XML elements become containers.
- All XML attributes become parsed attributes.
- Element text content becomes a parsed attribute that takes the name `text_value`, where the text content becomes the parsed attribute value.

The generic XML parser accepts no parameters.

15.6.2.1 WebSphere Generic Parser

Cloud Control provides one out-of-box generic parser designed to parse the WebSphere `serverindex.xml` configuration file. It comes with three default post-parsing rules to uniquely identify nodes with the same name.

15.6.3 XML Parser Examples

This section contains three XML parser examples:

- As parsed using the default XML parser, with out-of-box parameter values
- As parsed using the default XML parser, with modified parameter values
- As parsed using the generic XML parser

Parsed examples derive from the following original XML file:

```
<?xml version="1.0" encoding="UTF-8"?>
<Application>
   <AppName>foo</AppName>
   <Server name="ajax" os="linux">production</Server>
</Application>
```

Default XML Parser (Out-of-Box Parameter Values)

When parsed using the default XML parser with out-of-box parameter values, the parsed version appears as follows:

```
Application
   AppName = foo
   Server
      name = ajax
      os = linux
      text_value = production
```

Note the following about this parsed version:

- The element contents of the AppName and Server tags become parsed attributes.
- Since the AppName tag contains no XML attributes, the parsed attribute name takes the tag name.
- Contrast with the Server tag, which has XML attributes (name and os). This results in a container named for the tag (Server), with three parsed attributes, one for each

of the XML attributes, and a third for the text content of the Server tag, which is set to the value of the STORE_CONTENT_AS parameter (text_value).

Default XML Parser (Modified Parameter Values)

To modify parameter values, you have to create a new parser by exporting the default XML parser, modifying the exported XML file, and importing the modified parser, using a new name and parser ID.

Assume you followed this process, making the following modifications:

- Set the STORE_CONTENT_AS parameter to the value myVal
- Set the CONTAINER_NAME parameter to the value name

When parsed using the default XML parser with modified parameter values, the parsed version appears as follows:

```
Application
   AppName = foo
   ajax
      os = linux
      myVal = production
```

Note the following about this parsed version:

- The AppName tag remains the same; that is, it has no XML attributes so it becomes a parsed attribute.
- Since the Server tag has an XML attribute that matches the value of CONTAINER_NAME, the container takes the value of the attribute (ajax), obviating the name=ajax parsed attribute. Remember that the out-of-box CONTAINER_NAME parameter has a placeholder but no actual default value; thus, the difference in this version of the parsed representation.
- The remaining Server tag attribute (os) becomes a parsed attribute as usual, and the text content associated with the tag becomes the value of the attribute myVal, per the edited STORE_CONTENT_AS parameter.

Generic XML Parser

When parsed using the generic XML parser (the one that takes no parameters), the parsed version appears as follows:

```
Application
   AppName
      text_value = foo
   Server
      name = ajax
      os = linux
      text_value = production
```

Refer to the generic XML parser description for a reminder of how parsing occurs.

15.7 Format-specific Parsers

Format-specific base parsers are applicable only to a particular data format. Format-specific parsers run the gamut from having no parameters to a few to many with which to tailor formatting.

Parser	Description
Blue Martini DNA	Parser for Blue Martini DNA files (no parameters).
Connect:Direct	Parser for Connect:Direct `.cfg` files (no parameters).
Database Query (see Section 15.10.3 for an example)	Parser for custom configuration database query output. Enterprise Manager automatically transforms query results into a format the parser accepts, organizing results into sections similar to a Windows `.ini` file. Each section represents one record; each line in a section contains a table column name and a value. See Database Query Parser Parameters.
Db2	Parser for the output of the DB2 `GET DATABASE CONFIGURATION` command (no parameters).
Directory	Parser for files containing multiple name value pairs on the same line, where each line may have varying numbers of pairs. For example, the first line might be a=b j=k, the second line c=d m=n y=z, and so forth. See Directory Parser Parameters.
E-Business Suite	Parser for E-Business Suite `.drv` files. The parser converts `IF...THEN...ELSE` structures in the file into containers in the parsed representation, and the rest of the lines into a container with a fixed number of parsed attributes. These lines can be of two types: directory specifications, whose parsed attribute names are specified in the `DIR_HEADER` parser parameter; configuration file specifications, whose parsed attribute names are specified in the `HEADER` parser parameter. See E-Business Suite Parser Parameters.
Galaxy CFG	Parser for Galaxy `.cfg` files. See Galaxy CFG Parser Parameters.
Introscope	Parser for Introscope files (no parameters).
MQ-Series	Parser for MQ-Series files. See MQ-Series Parser Parameters.
Odin	Parser for Odin files (no parameters).
Oracle ORA	Parser for Oracle `.ora` files, such as `tnsnames.ora` (no parameters).
Siebel	Parser for Siebel `siebns` files. The parser creates a container for each unique path in the file, and attributes for name value pairs, except where a line contains the string `Type=empty`, in which case the parser does not create a parsed attribute for the line. See Siebel Parser Parameters.
UbbConfig	Parser for BEA Tuxedo files (no parameters). The parser converts sections prefixed with an asterisk (*), and names in double quotes at the start of a new line, into containers. It converts all other data into attributes.
Unix Installed Patches	Parser for Unix installed patches data. The parser creates one container per (non-comment) line of the file. It treats every field ending with a colon (:) on each line as a property name field and the value following, if any, as the property value. Note that a property does not have to have a value. See Unix Installed Patches Parser Parameters.
Unix Recursive Directory List	Parser for output of Unix recursive directory listing (`ls -l -R`). The parser converts each subdirectory line into a container, and each file information line into a container with a fixed set of attributes. SeeUnix Recursive Directory List Parser Parameters.

Remember, to modify a format-specific parser, you have to create a new parser by exporting the particular parser, modifying the exported XML file, and importing the modified parser, using a new name and parser ID.

Database Query Parser Parameters

The following table describes the parameters with which you can customize the Database Query parser:

Parameter	Description
CELL_DELIMITER	Character that separates name value pairs; default is =.
PROPERTY_DELIMITER	Character that separates the length of a name or value from the value itself; default is _.
COMMENT	Character that tells the parser to ignore the line that follows; default is #.
SECTION_START	Character that denotes the start of a section; default is \[(backslash is escape character).
SECTION_END	Character that denotes the end of a section; default is \] (backslash is escape character).
USE_INI_SECTION	Flag that tells the parser to use Windows .ini type sections; default is true.

Directory Parser Parameters

The following table describes the parameters with which you can customize the Directory parser:

Parameter	Description
CELL_DELIMITER	Character that separates one property from another; default is a space.
EXTRA_DELIMITER	Character that separates a property name from its value; default is =.
COMMENT	Character that tells the parser to ignore the line that follows; default is #.

E-Business Suite Parser Parameters

The following table describes the parameters with which you can customize the E-Business Suite parser:

Parameter	Description
DIR_HEADER	A tilde-delimited list of attribute names for directory specifications.
STRUCTURE_START	A tilde-delimited list of regular expressions denoting the start of a structure.
CELL_DELIMITER	A tilde-delimited list of regular expressions denoting name value pair delimiters.
HEADER	A tilde-delimited list of attribute names for file specifications.
COMMENT	A tilde-delimited list of regular expressions denoting comments.
STRUCTURE_END	A tilde-delimited list of regular expressions denoting the end of a structure.
LAST_FREE_FORM	Flag that tells the parser to ignore cell delimiters in the last value of a directory or file specification; default is true.
ELEMENT_FIELD	A tilde-delimited list of file specification attribute names. The parser concatenates values of the specified attributes to form the name of the container associated with the file specification.

Parameter	Description
DIR_ELEMENT_FIELD	A tilde-delimited list of directory specification attribute names the parser uses to determine the name of the container associated with the directory specification.

Galaxy CFG Parser Parameters

The following table describes the parameters with which you can customize the Galaxy CFG parser:

Parameter	Description
COMMENT	Character that tells the parser to ignore the line that follows; default is !.
ADD_SUFFIX	Names of attributes whose values to append to a container name.
MONO_PROP_SECTION	Names of sections that have a single property.
MULTI_PROP_SECTION	Names of sections that have multiple properties.
NODES_SECTION	Names of section start and end elements

MQ-Series Parser Parameters

The MQ-Series parser has a single parameter that you can customize: COMMENT, which defaults to *.

Siebel Parser Parameters

The following table describes the parameters with which you can customize the Siebel parser:

Parameter	Description
LINES_TO_SKIP	Tells the parser the number of lines to ignore at the beginning of the file; default is 4.
CELL_DELIMITER	A tilde-delimited list of regular expressions denoting name value pair delimiters.
COMMENT	A tilde -delimited list of regular expressions denoting comments.
SECTION_START	A tilde-delimited list of regular expressions denoting the start of a unique path specification section.
SECTION_END	A tilde-delimited list of regular expressions denoting the end of a unique path specification section.
USE_INI_SECTION	Flag that tells the parser to use Windows .ini type sections; default is true.

Unix Installed Patches Parser Parameters

The following table describes the parameters with which you can customize the Unix Installed Patches parser:

Parameter	Description
CELL_DELIMITER	Character that separates name value pairs; default is a space.
EXTRA_DELIMITER	Character that separates a property name from its value; default is :.

Parameter	Description
COMMENT	Character that tells the parser to ignore the line that follows; default is #.

Unix Recursive Directory List Parser Parameters

The following table describes the parameters with which you can customize the Unix Recursive Directory List parser:

Parameter	Description
LINES_TO_SKIP	Tells the parser the number of lines to ignore at the beginning of the file; default is 4.
CELL_DELIMITER	A tilde-delimited list of regular expressions denoting name value pair delimiters.
COMMENT	A tilde-delimited list of regular expressions denoting comments.
HEADER	A tilde-delimited list of attribute names.
LAST_FREE_FORM	Flag that tells the parser to ignore cell delimiters in the last value of a line; default is true.
SECTION_START	A tilde-delimited list of regular expressions denoting the start of a subdirectory section.
SECTION_END	A tilde-delimited list of regular expressions denoting the end of a subdirectory section.
ELEMENT_FIELD	A tilde-delimited list of attribute names. The parser concatenates values of the specified attributes to form the name of the container associated with the line.

15.8 Columnar Parsers

Columnar parsers are inherently flexible owing to the parameters they can accept to tailor formatting. All columnar parsers use a subset of the same parameters.

Parser	Description
Cron Access	Parser for `cron.allow` and `cron.deny` files.
Cron Directory	Parser for Unix `etc` and `cron.d` files.
CSV	Parser for comma-separated-value data. The out-of-box parameter values support CSV files with these characteristics: ■ Each line has the same number of values ■ The first parsed (that is, non-comment) line is a header line whose content is a comma-separated list of column names ■ Commas in double quotes are considered part of the value, not value delimiters ■ One of the column names is "name" whose value becomes the container name associated with each line Text inside double quotes is considered part of a value; to specify a value that contains a double quote, escape the double quote with a backslash (\). Use a backslash to escape the backslash character itself (\\).
Hosts Access	Parser for `hosts.allow` and `hosts.deny` files.
Kernel Modules	Parser for `kernel modules` files.

Parser	Description
Linux Directory List	Parser for Linux directory listing data format (for example, output of a `ls -l` command).
PAM Configuration	Parser for `pam.conf` files.
PAM Directory	Parser for Unix `etc/pam.d` files.
Process Local	Parser for `process.local` files.
Secure TTY	Parser for Unix `etc/securetty` files.
Solaris Installed Packages	Parser for Solaris installed packages files.
Unix Crontab	Parser for Unix crontab files.
Unix Directory List	Parser for Unix directory listing data format for example, the output of a `ls -l` command).
Unix Groups	Parser for Unix `etc/group` files. The parser ignores group name and password information.
Unix GShadow	Parser for Unix `etc/gshadow` files.
Unix Hosts	Parser for Unix `etc/hosts` files.
Unix INETD	Parser for Unix `etc/inetd.conf` files.
Unix Passwd	Parser for Unix `etc/passwd` files. The parser ignores password values.
Unix Protocols	Parser for Unix `etc/hosts` files.
Unix Services	Parser for Unix `etc/services.conf` files.
Unix Shadow	Parser for Unix `etc/shadow` files.
Unix System Crontab	Parser for Unix system crontab files. System crontab files are very similar to crontab files, but may contain name value pairs such as `PATH=/a/b`.

15.8.1 Columnar Parser Parameters

This section describes all columnar base parser parameters. Although the base parser can accept values for any of these parameters, a given parser specification does not necessarily need to provide values for all of them. All parameters have default values, which are used in the absence of a specified value, although in some cases, parameters have explicit values.

Use quotes when delimiters or other special text such as comment characters or new lines are part of some value. The `QUOTE_DELIMITER` determines the character value to use. Prefix the quote delimiter with a backslash (\) if you need to escape the character. Use a backslash to escape the backslash character itself (\\) in quoted strings.

Parameter	Description
`COMMENT`	A tilde-delimited list of regular expressions that denote comment characters or sequences. For example, `#[^\r\n]*` specifies that everything on a line following the # character is a comment. Default is an empty list; that is, parse all file contents.
`LINES_TO_SKIP`	The number of initial lines (excluding blank or comment lines) to ignore for parsing purposes, treating them in effect as comments. Default is 0; that is, skip no lines.

Parameter	Description
CELL_DELIMITER	A tilde-delimited list of regular expressions that delimit line values. Default is an empty list; that is, no delimiters (it is unusual to use the default).
QUOTE_DELIMITER	A tilde-delimited list of regular expressions that define how quoted values begin and end (usually either a single or double quote character). The beginning and end quote delimiter must be the same. Default is an empty list; that is, parser does not recognize quoted values.
PROPERTY_DELIMITER	A tilde-delimited list of regular expressions that delimit property names and values. Default is an empty list; that is, no property delimiters. Rarely, a columnar file may contain name value pairs of the syntax a=b.
RESERVED_ DIRECTIVES	A tilde-delimited list of property keywords. Some crontab files contain lines of simple name value pairs, separated by a delimiter (foo=bar), thus violating the requirement that each line have the same number of fields. This parameter provides a workaround to specify property keywords. In the example, the property keyword would be foo. This says, in effect, parse any line beginning with this keyword as a parsed attribute name value pair under the root container. Default is an empty list; that is, no property keywords.
ALTERNATE_ DELIMITER	An alternate delimiter for property names and values. Default is '/' (used only if ALTERNATE_FIELD parameter is nonempty).
ALTERNATE_FIELD	A tilde-delimited list of fields separated by alternate delimiters. Default is an empty list; that is, no alternate delimiters.
HEADER_FLAG	A flag specifying whether or not the file has a header line that specifies the column names. Default is false.
HEADER	A tilde-delimited list of column names to use if there is no header line. Default is an empty list; that is, no column names (it is unusual to use the default).
ELEMENT_FIELD	A tilde-delimited list of column names whose values the parser concatenates to create the container name associated with a line. Default is an empty list; that is, no column names (it is unusual to use the default).
IGNORE_FIELD	A tilde-delimited list of column names to ignore. No parsing of values in these columns occurs. Default is an empty list; that is, ignore nothing.
LAST_FREE_FORM	A flag that specifies whether the last column is free form. The parser ignores all delimiters in a free form column value. Default is false.
USE_LINE_COMMENT	A flag that specifies whether to treat end of line comments as a value to appear in the parsed representation of the data. Default is false.

15.9 Properties Parsers

Properties parsers are inherently flexible owing to the parameters they can accept to tailor formatting and handle disparate organizational elements. All properties parsers use the same set of basic and advanced parameters, as well as advanced constructs.

Parser	Description
AIX Installed Packages	Parser for AIX installed packages files.
Apache HTTPD	Parser for Apache `HTTPD.conf` files.
Autosys	Parser for `Autosys.jil` files.
Custom CFG	Parser for custom `.cfg` files. This syntax defines an element with `E = {}` syntax, where the brackets may contain name value pairs, nested elements, or both.
Java Policy	Parser for `java.policy` files.
Java Properties	Parser for `java.properties` files.
LDAP	Parser for LDAP `.cfg` files.
Mime Types	Parser for `mime.types` files.
Radia	Parser for Radia `.cfg` files.
Sectioned Properties	Parser for files containing name value pairs organized into sections, such as a Windows `.ini` file.
SiteMinder Agent	Parser for SiteMinder agent files.
SiteMinder Registry	Parser for SiteMinder `.registry` files.
SiteMinder Report	Parser for SiteMinder `SmReport.txt` files.
SmWalker	Parser for SiteMinder `SmWalker.dat` files.
Sun ONE Magnus	Parser for Sun ONE `magnus.conf` files.
Sun ONE Obj	Parser for Sun ONE `obj.conf` files.
Tuxedo	Parser for Tuxedo files.
Unix Config	Parser for Unix `etc/config` files.
Unix Login	Parser for Unix `etc/login.defs` files.
Unix PROFTPD	Parser for Unix `etc/proftpd.conf` files.
Unix Resolve	Parser for Unix `etc/resolve.conf` files.
Unix SSH Config	Parser for Unix `etc/ssh/sshd.conf` files.
Unix System	Parser for Unix `etc/system` files.
Unix VSFTPD	Parser for Unix `etc/vsftpd.conf` files.
Unix XINETD	Parser for Unix `etc/xinetd.conf` files.
WebAgent	Parser for WebAgent files.
Windows Checksum	Parser for Windows checksum output generated with `fciv.exe`.

15.9.1 Basic Properties Parser Parameters

This section describes basic properties parser parameters that are required to parse simple property data formats. Simple property data formats specify a property as a name value pair, usually with a defined delimiter separating the name and the value: foo=bar. The basic data format is a list of properties, one property to a line, together with optional comments; a `java.properties` file, for example. All parameters have default values, which are used in the absence of a specified value.

Use quotes when delimiters or other special text such as comment characters or new lines are part of some value. The `QUOTE_DELIMITER` determines the character value

to use. Prefix the quote delimiter with a backslash (\) if you need to escape the character. Use a backslash to escape the backslash character itself (\\) in quoted strings.

A comment character such as the pound sign (#), or a particular character sequence (//) usually denotes a comment. Special sequences such as a C style comment (/*….*/) might denote the beginning and end of a comment. Some files might have generic informational content in the first couple of lines. In this case, a parameter is available to tell the parser to ignore these lines.

Parameter	Description
COMMENT	A tilde-delimited list of regular expressions that denote comment characters or sequences. For example, #[^\r\n]* specifies that everything on a line following the # character is a comment. Default is an empty list; that is, parse all file contents.
LINES_TO_SKIP	The number of initial lines (excluding blank or comment lines) to ignore for parsing purposes, treating them in effect as comments. Default is 0; that is, skip no lines.
CELL_DELIMITER	A tilde-delimited list of regular expressions that delimit line values. Default is an empty list; that is, no delimiters (it is unusual to use the default).
QUOTE_DELIMITER	A tilde-delimited list of regular expressions that define how quoted values begin and end (usually either a single or double quote character). The beginning and end quote delimiter must be the same. Default is an empty list; that is, parser does not recognize quoted values.
ALLOW_NAME_ONLY_ PROPERTIES	A flag that indicates whether the parser allows property names without a delimiter or a value. Default: false.
REVERSE_PROPERTY	A flag that indicates whether the parser allows the value to come before the delimiter and property name. Default: false.

15.9.2 Advanced Properties Parser Parameters

This section describes advanced properties parser parameters that are required to parse more complex property data formats. All parameters have default values, which are used in the absence of a specified value.

Parameter	Description
PROPERTY_DELIMITER	A tilde-delimited list of regular expressions denoting property delimiters. For example, the text "a=b : x=y" could be interpreted in either of two ways: ■ As a single property "a" with value "b : x=y" ■ As two separate properties," a=b" and "x=y" If a colon (:) is the property delimiter, the parsing engine interprets the text as containing two separate properties. Default is an empty list; that is, parser does not recognize property delimiters.
LINE_END_DELIMITER	A tilde-delimited list of regular expressions denoting line end sequences. When the parser encounters a line end delimiter, it assumes a new property or construct starts on the next line. Default is an empty list; that is, parser does not recognize line end delimiters.

Parameter	Description
CONTINUE_LINE	A tilde-delimited list of regular expressions denoting continue line sequences. When the parser encounters a continue line pattern, it interprets data on the following line as a continuation of the construct or property on the previous line, as opposed to interpreting the new line as the beginning of a new property or construct. For example, the parser must encounter a line continuation pattern to recognize property values that span multiple lines. Default is an empty list; that is, parser does not recognize line continuation patterns.
SECTION_START	A tilde-delimited list of regular expressions denoting the beginning of a section. Sections cannot be nested. Default is an empty list; that is, parser does not recognize sections.
SECTION_END	A tilde-delimited list of regular expressions denoting the end of a section. Default is an empty list.
STRUCTURE_START	A tilde-delimited list of regular expressions denoting the beginning of a structure. Structures can be nested. Default is an empty list; that is, parser does not recognize structures.
STRUCTURE_END	A tilde-delimited list of regular expressions denoting the end of a structure. Default is an empty list.
XML_STYLE_TAG	A flag that indicates whether structures in the file are XML style tags. Default: false.
USE_INI_SECTION	A flag that indicates whether INI style sections are present. Default: false.
RESERVED_ DIRECTIVES	A tilde-delimited list of reserved names indicating the start of a reserved directive. Default is an empty list; that is, parser does not recognize reserved directives.
RESERVED_FUNCTIONS	A tilde-delimited list of reserved names indicating the start of a reserved function. Default is an empty list; that is, parser does not recognize reserved functions.
DIRECTIVE_ PROPERTIES	A tilde-delimited list of reserved directive- implicit property names. Default is an empty list.
FUNCTION_ PROPERTIES	A tilde-delimited list of required reserved function-explicit property names. Default is an empty list.
SECTION_PROPERTIES	A tilde-delimited list of section-implicit property names. Default is an empty list.
STRUCTURE_ PROPERTIES	A tilde-delimited list of structure-implicit property names. Default is an empty list.
ELEMENT_FIELD	A keyword to be ignored by the parser when parsing properties. This typically applies to data formats that specify a keyword before a name value pair; "set a=b" for example. Default is an empty list; that is, parser ignores nothing.
ALLOW_ELEMENT_CELL	A flag that indicates whether the file format supports element cell structures. Default: false.
SECTION_EXPLICIT_ PROPERTIES	A flag that indicates whether sections support explicit properties. Default: false.
STRUCTURE_ EXPLICIT_ PROPERTIES	A flag that indicates whether structures support explicit properties. Default: false.
NEWLINE_CONTINUE_ LIN	A flag that indicates whether newlines can be line continuation sequences. Default: false.

Parameter	Description
KEYWORD_FIELD	A tilde-delimited list of regular expressions denoting keywords that precede properties that use a whitespace delimiter. Default is an empty list; that is, parser does not recognize keywords.

15.9.3 Advanced Properties Parser Constructs

Properties files come in variety of file formats. To accommodate the widest possible range of formats, the generic properties base parser uses combinations of constructs found in most files.

The constructs fall into two categories:

- Container constructs, which can be reserved functions, reserved directives, XML structures, structures, delimited structures, INI sections, delimited sections, sections, and element cells
- Property constructs, which can be simple properties, reverse properties, keyword properties, keyword name properties, bracket properties, implicit properties and explicit properties

Of the element constructs, section constructs cannot be nested, but can contain any other construct. Structure constructs can be nested, and can contain any construct except a section. Element cells can be nested, but can only contain element cells and simple properties. Reserved directives and reserved functions cannot be nested, nor can they contain any other constructs.

The rest of this section describes the constructs the base properties parser supports.

Simple Property

A simple property consists of a property name, cell delimiter, property value, and newline sequence, in that order. A simple property may take up more than one line, although properties that span multiple lines usually contain a line continuation character or sequence. The parser ignores whitespace such as tabs and spaces, unless a parameter specifies whitespace as having some significance (cell delimiter, for example).

Example: `name=value_that_wraps_to_next_line_/`, where the forward slash serves as a line continuation character. A Java Properties file typifies this data format.

Keyword Property

This construct is the same as a simple property, only with a keyword in front, which the parser ignores.

Example: `set name=value`, where `set` is the ignored keyword. A Unix System file typifies this data format.

Keyword Name Property

This construct is a simple property where the property name matches a regular expression specified in the `KEYWORD_FIELD` parser parameter. This is a special case property type specific to the Unix XINETD parser. The XINETD file uses an equal sign (=) as a cell delimiter except when the property begins with the keyword "include" or "includedir", in which case the cell delimiter is whitespace.

While added specifically for XINETD files, the property can be used for other file types where appropriate.

Example: `includedir /etc`, where `includedir` is the parser parameter regular expression and whitespace is the cell delimiter.

Explicit Property

An explicit property consists of a property name, a delimiter, and a property value. Unlike a simple or keyword property, an explicit property is bound to a container construct such as a section or a structure; an XML tag attribute, for example.

Examples:

```
[SectionName p1=v1 p2=v2]

<StructureName p1=v1 p2=v2>
...
</StructureName>
```

In these constructs, the name value pairs p1 v1 and p2 v2 are explicit properties. A Sun ONE Obj file typifies this data format.

Implicit Property

An implicit property is a property value without an associated property name. Like an explicit property, an implicit property is bound to a container construct, usually a reserved directive. The `DIRECTIVE_PROPERTIES` parser parameter contains the property names of implicit properties.

Examples:

```
[SectionName myName myPath]

<StructureName myName myPath>
...
</StructureName>
```

In these constructs, the implicit properties have the values `myName` and `myPath`, with the presumed property names `name` and `path`, as declared in the `DIRECTIVE_PROPERTIES` parser parameter. An Apache HTTPD file typifies this data format.

Reserved Function

A reserved function is a keyword followed by one or more explicit properties. The `RESERVED_FUNCTIONS` parser parameter specifies keywords that denote reserved functions.

Example: `Error fn="query-handler" type="forbidden"`, where `Error` is the reserved function keyword specified in the `RESERVED_FUNCTIONS` parser parameter. A Sun ONE Magnus file typifies this data format.

Reserved Directive

A reserved directive is a keyword followed by one or more implicit properties. The `RESERVED_DIRECTIVES` parser parameter specifies keywords that denote reserved directives.

Example: `LoadModule cgi_module "/bin/modules/std/cgi"`, where `LoadModule` is the reserved function keyword specified in the `RESERVED_DIRECTIVES` parser parameter. An Apache HTTPD file typifies this data format.

XML Structure

An XML structure is a standard XML tag that can contain a name only, a name followed by explicit properties, or a name followed by implicit properties.

Examples:

```
<Name>
...
</Name>

<Name p1=v1 p2=v2>
...
</Name>
<Name "implicit_property1" "implicit_property2">
...
</Name>
```

A WebAgent file typifies this data format.

Delimited Structure

A delimited structure consists of the following (in the specified order):

- Structure name
- Delimiter
- Start structure character or character sequence
- Structure contents
- End structure character or character sequence

Example:

```
StructureName = {
...
}
```

Explicit and implicit properties are not allowed. Java Policy and Custom CFG files typify this data format.

Structure

A structure consists of the following (in the specified order):

- Structure name
- Start structure character or character sequence
- Structure contents
- End structure character or character sequence

The only difference between a delimited structure and a structure is the delimiter; that is, a structure does not require a delimiter between the structure name and the start structure indicator.

Example:

```
StructureName {
...
}
```

Explicit and implicit properties are not allowed. A Unix XINETD file typifies this data format.

INI Section

And INI section resembles a section heading in a Windows `.ini` file, characterized by:

- Section start character or character sequence
- Section name
- Optional (explicit and implicit) properties
- Section end character or character sequence

Examples:

```
[SectionName]

[SectionName p1=v1 p2=v2]

[SectionName "implicit_property1" "implicit_ property2"]
```

SmWalker and Sectioned Properties files typify this data format.

Delimited Section

A delimited section is a line that begins with a common pattern, but otherwise resembles a simple property.

Examples:

```
HKEY_LOCAL_MACHINE\SOFTWARE\A\B\C=789

HKEY_LOCAL_MACHINE\SOFTWARE\X\Y\Z=123
```

These are two delimited section headings where the common pattern is `HKEY_`. SiteMinder Registry and LDAP files typify this data format.

Element Cell

An element cell consists of an element cell name and a property name value pair of the form `A = B = C`. Element cells typically use line continuation sequences and nesting to clarify the structure. An element cell that has multiple properties uses a property delimiter to separate them.

Example 1:

```
EC = \
    B = C, D = F
```

This example is an element cell named `EC` with two property name value pairs, `B = C` and `D = F`, separated by a comma. The structure uses the backslash character (\) to indicate line continuation. The advanced properties parser parameters `PROPERTY_DELIMITER` and `CONTINUE_LINE` define the respective format characters.

Example 2:

```
EC = \
              EC2 = \
                              A = B, \
                              C = D
```

This example is an element cell named `EC` that has a nested element cell named `EC2` that contains two property name value pairs, `A = B` and `C = D`. This example uses the same delimiter and line continuation characters.

15.10 Parsed Files and Rules

A collected configuration file is stored in raw form and, if a parser is specified, in a tree structure of nodes, or containers, and attributes, or properties. The file also is generated internally in XML format for the purpose of applying post-parsing rules, which consist of XPath conditions and expressions. Note that even non-XML files are generated in this internal format. Since the internal format must accommodate other file types, it introduces an additional root node in the XML to compensate for files such as Java properties files that have only attribute names and values.

Examples of how files are parsed and displayed, and the effects of post-parsing rules help to clarify:

- Sample XML File Parsing and Rule Application
- Sample Non-XML File Parsing and Rule Application
- Sample SQL Query Parsing and Rule Application

15.10.1 Sample XML File Parsing and Rule Application

Consider the following simple XML file:

```
<dir name="/a/b/c">

  <file name="file1" size=120/>
  <file name="file2" size=350/>
</dir>
```

Its parsed form, using the default XML parser, appears in the user interface in the following tree structure:

```
dir
   name    = /a/b/c
      file
              name = file1
              size = 120
      file
              name = file2
              size = 350
```

Notice that two containers have the same name (file), which makes it impossible to distinguish between the two, at the container level, at least. Thus, this file is a candidate for a post-parsing rule. As mentioned, there is a special internal XML format against which to apply a rule's XPath condition and expression. This format treats nodes and attributes as XML elements, and converts attribute values into corresponding element text content. It also adds a root element that doesn't appear in the original file:

```
<root>

  <dir>
      <name>/a/b/c</name>
      <file>
            <name>file1</name>
            <size>120</size>
      </file>
      <file>
            <name>file2</name>
            <size>350</size>
      </file>
  </dir>
```

```
</root>
```

Given the problem in the parsed form of having two containers with the same name, a rule resolution might consist of the following:

Condition: `/root/dir/file`
Expression: `name/text()`

Effectively, this says: for each file evaluate `name/text()` to produce an identifier that distinguishes one file from another within the dir node.

After applying the post-parsing rule, the parsed tree structure appears as follows:

```
dir
      name = /a/b/c
      file[file1]
            name = file1
            size = 120
      file[file2]
            name = file2
            size = 350
```

The rule resolves to an identifier appended in square brackets to the container name. The combination (`file [file1]`, for example) enables various operations such as compare, search, change history, and so forth, to distinguish between file containers.

15.10.2 Sample Non-XML File Parsing and Rule Application

Consider the following simple ORA file:

```
acme=
   (DESCRIPTION=
      (SOURCE_ROUTE=yes)
      (ADDRESS=(PROTOCOL=tcp)(HOST=host1)(PORT=1630))
      (ADDRESS_LIST=
         (FAILOVER=on)
         (LOAD_BALANCE=off)
   (ADDRESS=(PROTOCOL=tcp)(HOST=host2a)(PORT=1630))
         (ADDRESS=(PROTOCOL=tcp)(HOST=host2b)(PORT=1630)))
      (ADDRESS=(PROTOCOL=tcp)(HOST=host3)(PORT=1630))
      (CONNECT_DATA=(SERVICE_NAME=Sales.us.acme.com)))
```

Its parsed form, using the Oracle ORA parser, appears in the user interface in the following tree structure:

```
acme
        DESCRIPTION
             SOURCE_ROUTE    yes
             ADDRESS
                     PROTOCOL         tcp
                     HOST             host1
                     PORT             1630
             ADDRESS_LIST
                     FAILOVER         on
                     LOAD_BALANCE     off
                     ADDRESS
                             PROTOCOL          tcp
                             HOST              host2a
                             PORT              1630
                     ADDRESS
                             PROTOCOL          tcp
```

```
                        HOST            host2b
                        PORT            1630
                ADDRESS
                        PROTOCOL        tcp
                        HOST            host3
                        PORT            1630
                CONNECT_DATA
                        SERVICE_NAME    Sales.us.acme.com
```

Notice that the ADDRESS_LIST address containers are indistinguishable. Thus, this file is a candidate for a post-parsing rule. As mentioned, there is a special internal XML format against which to apply a rule's XPath condition and expression. This format treats nodes and attributes as XML elements, and converts attribute values into corresponding element text content. It also adds a root element that doesn't appear in the original file:

```
<root>
        <acme>
                <DESCRIPTION>
                        <SOURCE_ROUTE>yes</SOURCE_ROUTE>
                        <ADDRESS>
                                <PROTOCOL>tcp</PROTOCOL>
                                <HOST>host1</HOST>
                                <PORT>1630</PORT>
                        </ADDRESS>
                        <ADDRESS_LIST>
                                <FAILOVER>on</FAILOVER>
                                <LOAD_BALANCE>off</LOAD_BALANCE>
                                <ADDRESS>
                                        <PROTOCOL>tcp</PROTOCOL>
                                        <HOST>host2a</HOST>
                                        <PORT>1630</PORT>
                                </ADDRESS>
                                <ADDRESS>
                                        <PROTOCOL>tcp</PROTOCOL>
                                        <HOST>host2b</HOST>
                                        <PORT>1630</PORT>
                                </ADDRESS>
                        </ADDRESS_LIST>
                        <ADDRESS>
                                <PROTOCOL>tcp</PROTOCOL>
                                <HOST>host3</HOST>
                                <PORT>1630</PORT>
                        </ADDRESS>
                        <CONNECT_DATA>
                                <SERVICE_NAME>Sales.us.acme.com</SERVICE_NAME>
                        </CONNECT_DATA>
                </DESCRIPTION>
        </acme>
</root>
```

Given the problem in the parsed form of having two containers with the same name, a rule resolution might consist of the following:

Condition: `//ADDRESS_LIST/ADDRESS`
Expression: `/HOST/text()`

Effectively, this says: for each address in the address list evaluate `/HOST/text()` to extract the host name as the address identifier.

After applying the post-parsing rule, the parsed tree structure appears as follows:

```
acme
        DESCRIPTION
             SOURCE_ROUTE    yes
             ADDRESS
                     PROTOCOL         tcp
                     HOST             host1
                     PORT             1630
             ADDRESS_LIST
                     FAILOVER         on
                     LOAD_BALANCE     off
                     ADDRESS[host2a]
                             PROTOCOL          tcp
                             HOST              host2a
                             PORT              1630
                     ADDRESS[host2b]
                             PROTOCOL          tcp
                             HOST              host2b
                             PORT              1630
             ADDRESS
                     PROTOCOL         tcp
                     HOST             host3
                     PORT             1630
             CONNECT_DATA
                     SERVICE_NAME     Sales.us.acme.com
```

The rule resolves to an identifier appended in square brackets to the container name. The combination (`ADDRESS [host2a]`, for example) enables various operations such as compare, search, change history, and so forth, to distinguish between address containers.

15.10.3 Sample SQL Query Parsing and Rule Application

Consider the following three-column database table `SERVER_DETAILS`:

SERVER_NAME	ENVIRONMENT	HOSTED_APPLICATIONS
webserver-100	QA	5
webserver-200	PERFORMANCE	6
webserver-500	PRODUCTION	3

The SQL query expressed as part of the custom configuration creation is as follows:

```
select * from SERVER_DETAILS
```

This query returns the following raw output:

```
[row]
11_SERVER_NAME=13_ webserver-100
11_ENVIRONMENT=2_ QA
19_HOSTED_APPLICATIONS=1_5
[row]
11_SERVER_NAME=13_ webserver-200
11_ENVIRONMENT=11_ PERFORMANCE
19_HOSTED_APPLICATIONS=1_6
[row]
11_SERVER_NAME=13_ webserver-500
11_ENVIRONMENT=10_ PRODUCTION
```

```
19_HOSTED_APPLICATIONS=1_3
```

The Configuration Browser Source tab renders the data the same way.

Its parsed form, using the Database Query parser, appears in the user interface in the following tree structure:

```
row
       SERVER_NAME=webserver-100
       ENVIRONMENT=QA
       HOSTED_APPLICATIONS=5
row
       SERVER_NAME=webserver-200
       ENVIRONMENT=PERFORMANCE
       HOSTED_APPLICATIONS=6
row
       SERVER_NAME=webserver-500
       ENVIRONMENT=PRODUCTION
       HOSTED_APPLICATIONS=3
```

Notice that the `row` containers are indistinguishable. Thus, this query result is a candidate for a post-parsing rule. As mentioned, there is a special internal XML format against which to apply a rule's XPath condition and expression. This format treats nodes and attributes as XML elements, and converts attribute values into corresponding element text content. It also adds a root element that doesn't appear in the original file:

```
<root>
            <row>
                        <SERVER_NAME>webserver-100</SERVER_NAME>
                        <ENVIRONMENT>QA</ENVIRONMENT>
                        <HOSTED_APPLICATIONS>5</HOSTED_APPLICATIONS>
            </row>
            <row>
                        <SERVER_NAME>webserver-200</SERVER_NAME>
                        <ENVIRONMENT>PERFORMANCE</ENVIRONMENT>
                        <HOSTED_APPLICATIONS>6</HOSTED_APPLICATIONS>
            </row>
            <row>
                        <SERVER_NAME>webserver-500</SERVER_NAME>
                        <ENVIRONMENT>PRODUCTION</ENVIRONMENT>
                        <HOSTED_APPLICATIONS>3</HOSTED_APPLICATIONS>
            </row>
</root>
```

Given the problem in the parsed form of having three containers with the same name, a rule resolution might consist of the following:

Condition:`/root/row/SERVER_NAME`
Expression:`SERVER_NAME/text()`

Effectively, this says: for each row evaluate `SERVER_NAME/text()` to produce an identifier that distinguishes one row from another within the tree structure.

After applying the post-parsing rule, the parsed tree structure appears as follows:

```
row[webserver-100]
       SERVER_NAME=webserver-100
       ENVIRONMENT=QA
       HOSTED_APPLICATIONS=5
row[webserver-200]
```

```
        SERVER_NAME=webserver-200
        ENVIRONMENT=PERFORMANCE
        HOSTED_APPLICATIONS=6
row[webserver-500]
        SERVER_NAME=webserver-500
        ENVIRONMENT=PRODUCTION
        HOSTED_APPLICATIONS=3
```

The rule resolves to an identifier appended in square brackets to the container name. The combination (`row[webserver-100]`, for example) enables various operations such as compare, search, change history, and so forth, to distinguish between row containers.

16

Enterprise Manager Diagnosability

This chapter introduces diagnostic capabilities in Enterprise Manager that extend to Oracle Management Service (OMS) and Management Agents.

16.1 Fault Diagnostics Framework

Enterprise Manager includes a fault diagnostics framework for collecting and managing diagnostic data. Diagnostic data includes trace files, dumps, and core files as well as other information that enables customers and Oracle Support to identify, investigate, track, and resolve problems quickly and effectively.

The diagnostics framework offers the following benefits:

- Automatic capture of diagnostic data upon first failure contributes to quick problem resolution thereby reducing downtime
- Integration with Incident Manager and Support Workbench, and access to My Oracle Support (MOS) from Enterprise Manager ensure simplified customer interaction with Oracle Support
- Ability to initiate proactive health checks and availability of the Enterprise Manager Diagnostics Kit enhance problem prevention and resolution

16.2 Automatic Diagnostic Workflow

The general fault diagnostic workflow is as follows:

- A critical error occurs in OMS or an Agent
- The diagnostic framework automatically creates an incident and organizes the diagnostic information in the Automatic Diagnostic Repository (ADR)
- Administrators use the Incident Manager to manage the complete life cycle of the incident
- Administrators use the Support Workbench to view and process the contents of the ADR
- Administrators perform health checks and run the Enterprise Manager Diagnostics Kit for finer grained analysis
- Administrators use the Incident Packaging System (IPS) to package incident data and diagnostic results in a zip file for upload to My Oracle Support (MOS)
- Administrators receive confirmation that an SR was created with which to track problem resolution

16.3 Fault Diagnosability Infrastructure

For critical errors, the ability to capture error information at first-failure greatly increases the chance of a quick problem resolution and reduced downtime. An always-on, memory-based tracing system proactively collects diagnostic data from many Enterprise Manager components, and can help isolate root causes of problems. The system of data collection is similar to that of airplane "black box" flight recorders. When a problem is detected, alerts are generated and the fault diagnosability infrastructure is activated to capture and store diagnostic data.

The fault diagnosability infrastructure aids in preventing, detecting, diagnosing, and resolving problems. The problems that are targeted in particular are critical errors such as those caused by code bugs, metadata corruption, and customer data corruption.

When a critical error occurs, it is assigned an incident number, and diagnostic data for the error (such as trace files) are immediately captured and tagged with this number. The data is then stored in the Automatic Diagnostic Repository (ADR), where it can later be retrieved by incident number and analyzed.

16.3.1 Automatic Diagnostic Repository

The ADR is a file-based hierarchical data store for depositing diagnostic information produced by diagnostic framework clients. The repository contains data describing incidents, traces, dumps, alert messages, data repair records, health check records, SQL Trace information, core dumps, and other information essential for problem diagnosis.

You can view and process ADR contents through the Support Workbench. There also is a command line interface, the ADR Command Interpreter, with which you can manipulate the contents.

The default ADR home for OMS is:

```
<MiddlewareHome>/gc_inst/user_projects/domains/<DOMAIN_ NAME>/servers/<SERVER_
NAME>/adr
```

The default ADR home for Agents is:

```
<Middleware Home>/agent/agent_inst
```

16.3.2 Incident Manager

The Incident Manager provides a central point of control for managing events, incidents and problems detected within Enterprise Manager.

The Incident Manager gives you in-context access to diagnostic and resolution capabilities. You also have in-context access to My Oracle Support, where you can research knowledge base articles and create service requests.

The Guided Resolution region offers recommendations and provides links to diagnostics and resolutions.

16.3.3 Support Workbench

The Enterprise Manager Support Workbench (Support Workbench) is a facility that enables you to investigate, report, and in some cases, repair problems (critical errors), all with an easy-to-use graphical interface. The Support Workbench provides a self-service means for you to gather first-failure diagnostic data, obtain a support request number, and upload diagnostic data to Oracle Support quickly and with a minimum of effort, thereby reducing time-to-resolution for problems.

The Support Workbench allows you to view and process the contents of ADRs. From the Home and Problem Details pages you can do the following:

- View recent and historical problems
- View and create diagnostic packages
- Create user-reported problems
- Review checker findings
- Search MOS knowledge base

16.3.4 Perform Health Checks and Run Diagnostics Kit

Perform Health Checks and Run Diagnostics Kit

Health checks test the viability of various system components. Health checks run automatically in response to an incident. You also can perform targeted checks proactively. The diagnostic framework includes a comprehensive set of 26 out-of-box health checks to test components such as Jobs, Credential, Event, Loader, Plugin, ASLM, and so forth. Health check results are stored in the ADR.

The Enterprise Manager Diagnostics Kit is a set of Oracle-supplied scripts specifically designed to identify inconsistencies in Enterprise Manager that are known to contribute to errors. In some cases, the script may be able to resolve the issue.

The scripts run repository diagnostics against system modules. You can run diagnostics against all or selected modules. The kit is accessible via a link in the Support Workbench. Diagnostic output is stored in the ADR with other dump files.

16.3.5 Incident Packaging Service (IPS)

The IPS enables you to automatically and easily gather the diagnostic data (traces, dumps, health check reports, and so forth) pertaining to a critical error and package the data into a zip file for transmission to Oracle Support.

Because all diagnostic data and files related to a critical error are tagged with that error's incident number, you do not have to search through all the stored information to determine the files required for analysis. The IPS identifies the required files automatically and adds them to the zip file.

Before creating the zip file, the IPS first collects diagnostic data into an intermediate logical structure called an incident package (package) and stores it in the ADR, where you can view the package and modify its contents. For example, you may want to add additional diagnostic data or remove existing data before uploading the zip file to Oracle Support.

16.3.6 Draft Service Request Acknowledgment

After you upload the zip file, you receive a confirmation that the file was successfully generated and that a Draft Service Request has been created. You are advised to go to My Oracle Support to finalize and submit the SR to Oracle Support.

16.4 Using the Support Workbench to Investigate Problems

When Enterprise Manager encounters a critical error that prevents you from completing a task, Enterprise Manager logs an error and generates an incident for this critical error, which then generates an alert. Enterprise Manager stores incident details,

including dump and trace files where applicable, in the Automatic Diagnostic Repository so that Support Workbench can access this information and display it.

After receiving an alert notifying you of a problem or incident, take the following action:

1. From the **Enterprise** menu, select **Monitoring**, then select **Support Workbench**.
2. In the list of targets that support ADR, locate the target about which you were notified and click the target link.
3. On the Support Workbench page, perform any of the following actions as appropriate:
   - View problem or incident details
   - View, create, or modify incident packages
   - View health checker findings
   - Close resolved problems

For details on performing these actions, see the Cloud Control online help.

17

Updating Cloud Control

The Self Update feature allows you to expand Enterprise Manager's capabilities by updating Enterprise Manager components whenever new/updated features become available between official releases. Oracle makes functional updates available between releases by publishing them to the Enterprise Manager Store, an external site that is periodically checked by Cloud Control to obtain information about available updates.

Cloud Control also provides support for plug-in and connector management. The ability to update plug-ins is particularly important because core Enterprise Manager features - such as Oracle Database management functionality - is now made available via plug-ins.

This chapter covers the following topics:

17.1 Using Self Update

The Self Update feature is accessed via the Self Update home page, a common dashboard used to obtain information about new updates and a common workflow to review, download and apply the updates. The Self Update console frees you from having to monitor multiple channels to get informed about new updates that are available from Oracle. The Self Update console automatically informs you whenever new updates are made available by Oracle. Only those updates that are applicable to your site are shown, eliminating the need to wade through unrelated updates.

17.1.1 What Can Be Updated?

Specific updates authored by Oracle that are usually bundled with specific Cloud Control releases can be updated via Self Update. Some examples are Oracle authored Management Plug-ins or Deployment Procedures. In general, Oracle-supplied entities are read-only. You can create a copy and customize the copy as per your needs but you cannot modify the original Oracle-supplied entity.

These entities can also be published on Oracle Web sites such as Oracle Technology Network (OTN) and My Oracle Support (MOS). You can download and import the entity archive into their Cloud Control deployment using specific import features provided by the update-able entity.

Entity Types That Can Be Updated

Examples of update-able entity types are:

- Management Agents
- Management Plug-ins
- Management Connectors
- Database Profiles and Gold Images
- Application Server Profiles and Gold Images
- Provisioning Bundles
- Enterprise Manager Deployment Pre-requisite Checks
- Compliance Content
- Diagnostic Checks

17.2 Setting Up Self Update

Before the Self Update feature can be used, a few prerequisites must be met.

- My Oracle Support credentials have been set up. This is required to enable entities to be downloaded from the My Oracle Support site.
- The Software Library (also known as the local store) has been configured. Updates are downloaded to this local store before being deployed into Cloud Control.

Review the following sections for instructions on setting up Self Update:

- Setting Up Enterprise Manager Self Update Mode
- Assigning Self Update Privileges to Users
- Setting Up the Software Library
- Setting Up the EMCLI Utility (Optional)

17.2.1 Setting Up Enterprise Manager Self Update Mode

In order to setup/modify the Enterprise Manager Self Update feature, you must have Enterprise Manager Super Administrator privileges.

1. Log into on to Enterprise Manager as an administrator with Super Administrator privileges.
2. From the **Setup** menu, choose **Extensibility>Self Update**. The Self Update console appears with the default setup displayed.
3. From the **General** status area, click on the **Connection Mode** status to set either offline or online mode. Enterprise Manager takes you to the Patching Setup page to specify online and offline settings.
4. Once the desired connection mode has been selected, return to the Self Update console.

   From here you can select entity types, schedule updates from the Enterprise Manager Update Store.

17.2.2 Assigning Self Update Privileges to Users

In order for Enterprise Administrators to use the Self Update feature, they must have the requisite privileges. The Enterprise Manager Super Administrator must assign the following Self Update roles to these administrators:

- VIEW_SELF_UPDATE - User can view the Self Update console and can monitor the status of download and apply jobs.
- MANAGE_SELF_UPDATE - User can schedule download and apply jobs. User can also suppress/unsuppress updates. This privilege implicitly contains VIEW_SELF_UPDATE.

By default, the Super Administrator will have MANAGE_SELF_UPDATE privilege granted to him.

To assign Self Update privileges to regular Enterprise Manager administrators:

1. From the **Setup** menu, choose **Security**>**Administrators**.
2. Select an administrator and click **Edit**.
3. From the Roles page, assign the appropriate Self Update roles.

17.2.3 Setting Up the Software Library

The Software Library is a repository that stores software patches, virtual appliance images, reference gold images, application software and their associated directive scripts. It allows maintaining versions, maturity levels, and states of entities.

In the context of applying updates, it is the "local store" that entities are downloaded to before deployment.

Follow these steps to set up the Software Library.

1. Create a folder in the system where Enterprise Manger is installed. For example, /net/*hostname*/scratch/aime/swlib1.
2. From the **Enterprise** menu, select **Provisioning and Patching**, then **Software Library**.
3. Click **Actions**, then **Administration**.
4. Click **Add**.
5. In the pop up window, enter a name and location for the folder you want to use as the Software Library. For example, `swlib1` and `/net/`*`hostname`*`/scratch/aime/swlib1`. This should be the folder that you created in step 1.
6. Wait for the processing to complete.

17.2.4 Setting Up the EMCLI Utility (Optional)

If you plan to apply software updates in offline mode, or to deploy non-Oracle plug-ins, you will need to use the Enterprise Manager Command Line Utility, or EMCLI, to import entity archives for deployment to Enterprise Manager.

A page is provided in the Cloud Control console with instructions on setting up EMCLI. Access the page by appending `/console/emcli/download` to the URL used to access the Cloud Control console:

```
https://<emcc_host>:<emcc_port>/em/console/emcli/download
```

For example:

```
https://server001.acme.com:7801/em/console/emcli/download
```

17.3 Applying an Update

The process for applying updates is essentially as follows: an update to Cloud Control, you must first download the

- Check for the latest updates available from Oracle.
- Download the updates you want to apply to the Software Library.
- Apply the update.

Review the following sections to learn how to apply an update:

- Applying an Update in Online Mode
- Applying an Update in Offline Mode

17.3.1 Applying an Update in Online Mode

Updates must be downloaded to the Software Library (the local store) before they can be applied. You can review the latest available updates from the Self Update console.

Note that Enterprise Manager must have access to the Enterprise Manager Store via the Internet to download available updates. If this access is not possible, you can download entities in offline mode. See Section 17.3.2, "Applying an Update in Offline Mode" for details.

1. From the **Setup** menu, choose **Extensibility>Self Update**.
2. Click **Self Update** to get the complete list of available updates.
3. Select the desired entity type and choose **Open** from the **Action** menu. The entity type page appears.
4. Select an update from the list of available updates.
5. Click **Download**. The **Schedule Download** dialog appears.
6. Select when to download the update. Note that multiple downloads can be scheduled simultaneously.

   The following options are available:

   - Immediately
   - Later (specified time)
   - Whether or not to send a notification when the download is complete.
7. Click **Select**. An Enterprise Manager job is created to download the update to the Software Library.

   Enterprise Manager starts downloading the archive from the Oracle Enterprise Manager store. Wait for the download to complete. (When in offline mode the system starts reading from the specified location.)

   When the download is complete, Enterprise Manager displays the Confirmation page, and the downloaded plug-in is shown in the local Oracle Enterprise Manager Store.

> **Note:** The page is not refreshed automatically. Click the refresh icon to view the updated download status.

8. Once an entity has been downloaded to the Software Library, it is ready to be applied to your installation. Select an update from the list whose status is **Downloaded**, then click **Apply**.

   Note that the application process varies according to the entity type:

   - For connectors, diagnostic checks, and compliance content, clicking **Apply** will install the update to Enterprise Manager. No further action is required.
   - For plug-ins, you will be redirected to the plug-in deployment page.
   - For provisioning bundles, you will need to exit the Enterprise Manager console, run Opatch and other commands via a terminal, and then restart the OMS.

17.3.2 Applying an Update in Offline Mode

Under certain circumstances, such as in high security environments, an active Internet connection between Enterprise Manager and the Enterprise Manager Update Store may not be available. In such situations, the Self Update feature can be used in "offline mode".

The update process still requires that a computer exist at your site that has an Internet access, as a connection to the Enterprise Manager Update Store is still required to obtain the updates. Update files from this computer can then be transferred to a computer behind your firewall.

The generic offline mode update procedure is as follows:

1. Ensure Cloud Control is set to offline mode. From the **Setup** menu, choose **Provisioning and Patching>Offline Patching**, then change the setting for Connection to **Offline**.
2. Click **Check for Updates** on the Self Update home page. A message is displayed that contains the URL to be accessed to download a catalog of all updates.
3. From an Internet-enabled computer, download the catalog file using the aforementioned URL.
4. Copy the downloaded file to the Oracle Management Service host or the Management Agent host you will deploy the update to.
5. Run the `emcli import_update` command to import the file into the Oracle Management Service instance or the Management Agent you want to update.

   See Section 17.5.1.2, "Importing an External Archive into Enterprise Manager" for instructions on using the `emcli import_update` command.
6. Review the update from Self Update Home and select and click download. A messages is displayed with URL and instructions.
7. Select and click **Apply** to apply the update.

17.4 Acquiring or Updating Management Agent Software

Management Agent software for the various platforms (operating systems) supported by Enterprise Manager Cloud Control can be downloaded to the Software Library using the Self Update console. Once an Agent is persisted to the Software Library, it can be installed on host machines that you want to bring under Cloud Control management using the Management Agent installation wizard.

Steps for obtaining Agent software in both online and offliine modes are discussed below.

- Acquiring Management Agent Software in Online Mode
- Acquiring Management Agent Software in Offline Mode

17.4.1 Acquiring Management Agent Software in Online Mode

Using Self Update in online mode requires Enterprise Manager to have access to My Oracle Support via the Internet.

1. From the **Setup** menu, choose **Extensibility>Self Update**.
2. Select the entity type *Agent Software* and choose **Open** from the **Action** menu. The entity type page appears to show agent software for different platforms.
3. Select an update from the list of available updates. All entries other than the one which matches the platform of the OMS host should show their status as *Available*.
4. Click **Download**. The **Schedule Download** dialog appears.
5. Select when to download the update. The following options are available:
   - Immediately
   - Later (specified time)
   - Whether or not to send a notification when the download is complete
6. Click **Select**. An Enterprise Manager job is created to download the Agent software to the Software Library.

   Enterprise Manager starts downloading the archive from the Oracle Enterprise Manager store. Wait for the download to complete. (When in offline mode the system starts reading from the specified location.)

   When the download is complete, Enterprise Manager displays the Confirmation page, and the downloaded plug-in is shown in the local Oracle Enterprise Manager Store.
7. Once the download is complete, select the Agent, then click **Apply**. This step will stage the Agent software in the Software Library and make it available to the Add Targets wizard, which you will use to install the Agent on host machines.
8. Click **Agent Software** to launch the Add Targets/Agent Installation wizard.

17.4.2 Acquiring Management Agent Software in Offline Mode

Follow this Self Update process only when Enterprise Manager is in offline mode.

1. Ensure Enterprise Manager is set to offline mode. From the Setup menu, choose **Provisioning and Patching>Offline Patchin**g, the change the Connection setting to Offline.
2. From the **Setup** menu, choose **Extensibility>Self Update**.
3. Click **Check for updates** on Self Update home page. A message is displayed that contains the URL to be accessed to download a catalog of all updates.

   Note that the archive containing the Management Agent software should also be available from the Oracle Technology Network (OTN) site.
4. From an Internet-enabled computer, download the catalog file using the aforementioned URL.

5. Copy the downloaded file to either of the following:
   - To any host that has a Management Agent and EMCLI installed
   - To the Oracle Management Service host
6. Run the `emcli import_update` command to import the archive into the Oracle Management Service instance or the Management Agent you want to update.

   See Section 17.5.1.2, "Importing an External Archive into Enterprise Manager" for instructions on using the `emcli import_update` command.
7. Select the entity type *Agent Software* and choose **Open** from the **Action** menu. The entity type page appears displaying agent software for different platforms.
8. Select an update from the list of available updates. All entries other than the one which matches the platform of the Oracle Management Service host will show their status as Available.
9. Click **Download**. A message is displayed with a URL and instructions.
10. From an Internet-enabled computer, download the file from the URL displayed in step 8. Do one of the following:
    - Copy the file to a Management Agent and follow the instructions displayed in step 8.
    - Copy the file to Oracle Management Service and follow the instructions displayed in step 8.

    At this stage, the update will show up in downloaded state in the Self Update home page.
11. Once the download is complete, select the Managment Agent, then click **Apply**. This step will stage the Management Agent software in the Software Library and make it available to the Add Targets wizard, which you will use to install the Management Agent on host machines.
12. Click **Agent Software** to launch the Add Targets/Agent installation wizard.

17.5 Deploying and Updating Plug-ins

A plug-in is a component (module) that can be plugged into an existing Enterprise Manager Cloud Control installation to extend its management and monitoring capabilities. The plug-in management features provided with Cloud Control simplify lifecycle management by allowing you to install and deploy plug-ins, as well as upgrade to newer versions as they become available.

With all releases of Enterprise Manager, external companies provide plug-ins that monitor specific types of targets, such as non-Oracle databases or applications. In a major architectural change, core Enterprise Manager Cloud Control features for managing and monitoring Oracle technologies - such as Oracle Database, Oracle Fusion Middleware and Oracle Fusion Applications - are also provided via plug-ins that can be downloaded and deployed.

This new “pluggable” framework enables Cloud Control to be updated with management support for the latest Oracle product releases, without having to wait for the next Cloud Control release to provide such functionality. For example, when a new version of Oracle Database is released, you can simply download and deploy the latest Oracle Database plug-in, which will include management support for the latest release.

- Importing an External Archive into Enterprise Manager

- Deploying a Plug-in
- Updating a Plug-in

17.5.1 Deploying a Plug-in

The process of enabling a plug-in to monitor targets in Enterprise Manager Cloud Control is essentially as follows:

- The plug-in archive is downloaded to the Software Library.
- The plug-in is deployed to the Oracle Management Service (OMS) instance that manages and monitors targets of the plug-in's type. As part of this process, the plug-in metadata is written to the Management Repository.
- Targets that the plug-in will monitor are added - or promoted to managed status - in Cloud Control.
- As part of the target promotion process, a Management Agent containing the required plug-in content is assigned to monitor the target. (If a Management Agent already exists on the target host, it will be updated with the plug-in content.)

See the following sections for details:

- Downloading a Plug-in from the Enterprise Manager Store
- Importing an External Archive into Enterprise Manager
- Deploying a Plug-in to Oracle Management Service (OMS)
- Adding Targets for the Plug-in to Monitor
- Important Details Regarding Plug-in Deployment

17.5.1.1 Downloading a Plug-in from the Enterprise Manager Store

Plug-in archives that are provided by Oracle will be made available through the Enterprise Manager Store, just like any other update. The plug-in must be downloaded to the Software Library, again like any other update, before it can be deployed.

See Section 17.3.1, "Applying an Update in Online Mode" for instructions on download plug-in archives.

17.5.1.2 Importing an External Archive into Enterprise Manager

External, non-Oracle plug-ins - that is, plug-ins that have been created by a company other than Oracle - must be imported into the Software Library using EMCLI before the plug-in deployment process can be initiated.

Note that this process can also be used to import other types of entity archives if Self Update is used in offline mode.

The plug-in archive must first be downloaded to an accessible location. Once downloaded, the plug-in can be imported into Enterprise Manager Cloud Control using the `emcli import_upate` command. See Section 17.2.4, "Setting Up the EMCLI Utility (Optional)" for instructions on setting up EMCLI.

You have two options for importing the plug-in archive, depending on where EMCLI is installed:

- If EMCLI is on the same system as the system as the location you downloaded the plug-in archive (*.opar file) to, run the following command.

  ```
  emcli import_update
  ```

```
-file="<path to *.opar file>"
-omslocal
```

The `-omslocal` flag indicates that the plug-in archive is on the same system where you are running this command and the path exists on this system.

- If EMCLI is on a different system than the plug-in archive, run the following command:

```
emcli import_update
       -file="<path to *.opar file you created>"
       -host="host1.example.com"
       -credential_name="host1_creds"
       -credential_owner="admin1"
```

The command syntax is as follows:

  - `-file`: The absolute path to the *.opar file on the system where you created the archive.
  - `-host`: The target name for a host target where the file is available.
  - `-credential_name`: The name of the credentials on the remote system you are connecting to.
  - `-credential_owner`: The owner of the credentials on the host system you are connecting to.

- As an alternative to the previous step, you can also run the following command:

```
emcli import_update
     -file="<path to *.opar file you created>"
     -host="hostname"
     -credential_set_name="setname"
```

`-credential_set_name`: The set name of the preferred credential stored in the Management Repository for the host target. It can be one of the following:

  - `HostCredsNormal`: The default unprivileged credential set.
  - `HostCredsPriv`: The privileged credential set.

Once the archive has been imported, the plug-in (or other entity downloaded in offline mode) can be deployed.

17.5.1.3 Deploying a Plug-in to Oracle Management Service (OMS)

A plug-in must be deployed on Oracle Management Service (OMS) before it can be used to monitor targets. Follow the steps below to deploy the plug-in on Enterprise Manager Cloud Control.

To deploy a plug-in on the Oracle Management Server, follow these steps:

1. From the **Setup** menu, select **Extensibility**, then **Plug-ins**. Enterprise Manager displays the list of plug-ins that have been downloaded and can be deployed on the Plug-ins page.
2. On the Plug-ins page, select the specific plug-in you want to deploy. Note that the plug-in archive must have already been downloaded to the Software Library.
3. Click **Deploy On>Management Servers**.

   Be sure that dependent plug-ins are deployed and that all existing Management Agents are compatible with the version of the specified plug-in. Enterprise Manager prompts for credentials if the agent is not available.

Note: Plug-ins must be deployed on Oracle Management Service prior to being deployed on Management Agents.

4. Specify the required details on the Deploy Plug-in dialog box. Note that you will need the Management Repository SYS user password to complete the deployment process.

   In the Version of Plug-in to Deploy section, select or choose the **Plug-in** version from the Plug-in drop-down. The **Target Type** information is displayed in the table. Enter the **Repository sys Password**, then click **Continue**.

5. Proceed through the steps in the Deploy Plug-in dialog box.

6. Click **Deploy** to deploy the selected plug-in on all Enterprise Manager servers.

   Enterprise Manager displays a page that allows you to monitor the deployment status. Enterprise Manager deploys the selected plug-in on all Enterprise Manager Servers.

   You can also monitor the deployment status by going to the Enterprise Manager Cloud Control console, then going to the plug-ins page as described in step 1, selecting the plug-in and select the **Recent Deployment Activities** tab at the bottom of the page for the selected plug-ins. This bottom section also lets you see details of your plug-in, which includes the plug-in ID, version, vendor, and so on.

   If any of the steps during plug-in deployment fails, the log is available in $ORACLE_HOME/cfgtoollogs/pluginca/*. Append these in while logging a support request for failure while deploying the plug-in. You can also use them to debug the problem.

17.5.1.4 Adding Targets for the Plug-in to Monitor

In the current Cloud Control release, deployment of a plug-in to a Management Agent that will monitor targets is no longer required. Instead, the plug-in for a specific target type is automatically deployed with the Management Agent that will monitor targets of that type.

This is a significant change from previous releases, in which plug-ins had to first be manually deployed to a Management Agent; a target instance then had to be manually added to the Management Agent.

You can add targets that the plug-in will monitor through Enterprise Manager Cloud Control by selecting **Add Targets** from the **Setup** menu. The process for adding targets - known in Cloud Control terminology as *target promotion* - will vary depending on the option you choose.

See XXX LINK TO ADDING TARGETS CHAPTER for details on adding targets to monitor.

17.5.1.5 Important Details Regarding Plug-in Deployment

- You can import multiple versions of the same plug-in. The version to deploy can be selected from a list if you are using Cloud Control to deploy the plug-in, or can be specified on the command line if using EMCLI.
- Only one version can be deployed on the Oracle Management Service (OMS) at any given time. If a later version has been deployed previously, it cannot be downgraded to an earlier version.

- Updating a plug-in to a new version does not remove the content of the older plug-in.
- The Management Agent can have the same or earlier version of the plug-in that is deployed on the OMS host. However a version later than the version on the OMS host is not allowed on the Management Agent host.
- The plug-in on the OMS host and the plug-in on the Management Agent host can be updated independently of each other.
- Available updates are visible on the Plug-ins page. They can be downloaded from the Enterprise Manager store or imported using `emcli` as described in Section 17.5.1.2, "Importing an External Archive into Enterprise Manager".

17.5.2 Updating a Plug-in

When a new version of a plug-ins is released, you can upgrade Oracle Management Service and Management Agents that use the plug-in with the latest version.

Updated versions of Oracle plug-ins are made available through the Oracle Enterprise Manager Store, where then can be downloaded for deployment. The Oracle Enterprise Manager Store is a central store maintained by Oracle, where plug-in metadata and archives are published. Plug-ins must be downloaded to a location (the Local Store) before being deployed.

Updated versions of external non-Oracle plug-ins are not available through the Oracle Enterprise Manager Store, and must first be imported into Oracle Management Service before being deployed. See Section 17.5.1.2, "Importing an External Archive into Enterprise Manager" for instructions.

Note that before a plug-in can be upgraded on a Management Agent, it must first be upgraded on the associated Oracle Management Service instance. In addition, you do not need to remove the existing version of a plug-in from OMS or Management Agents before upgrading to the latest version.

For more information about using plug-in lifecycle management, see the following sections:

- Downloading the Latest Plug-in Archive from the Oracle Enterprise Manager Store
- Upgrading a Plug-in on a Management Agent
- Removing a Plug-in

17.5.2.1 Downloading the Latest Plug-in Archive from the Oracle Enterprise Manager Store

Plug-ins must be downloaded to the Software Library (the local store) before being deployed.

To download the latest plug-in archive from the Oracle Enterprise Manager Store to the Software Library, follow the instructions in Section 17.5.1.1, "Downloading a Plug-in from the Enterprise Manager Store".

17.5.2.2 Upgrading a Plug-in on Oracle Management Service

The new plug-in must be deployed to the Oracle Management Service instance where the plug

Once the plug-in is available in the Local Store, you can deploy it to Oracle Management Service. See Section 17.5.1.3, "Deploying a Plug-in to Oracle Management Service (OMS)" for instructions.

The plug-in must be updated on the OMS instance managing relevant targets before it is updated on active Management Agents.

Note that you do not need to remove the existing version of a plug-in from the OMS instance before upgrading to the latest version.

17.5.2.3 Upgrading a Plug-in on a Management Agent

Once the updated version of the plug-ins has been deployed on OMS, all Management Agents that use the plug-in can be updated with the latest version.

Note that you do not need to remove the existing version of a plug-in from the Management Agents before upgrading to the latest version.

1. From the **Setup** menu, choose **Extensibility>Plug-ins**.
2. Select the row for the plug-in you want to update to in the table.
3. Click **Deploy On>Management Agent**.

> **Note:** Plug-ins must be deployed on a Managed Server prior to deploying on Management Agents.

4. On the Deploy Plug-in dialog box, use the Management Agents to Deploy section to **Add** or **Remove** the agents to which you want to deploy the plug-in.

   When you click **Add**, Enterprise Manager displays the Search and Select dialog box where you can select the agents to add. Click **OK** to return to the Deploy Plug-in dialog box. Only agents running the operating systems supported by the selected plug-in may be selected.

   To remove an agent, highlight it in the Management Agents to Deploy table and click **Remove**.

5. Specify other required details on the Deploy Plug-in on Agents dialog box.

   Select the **Plug-in** version from the Plug-in drop-down. Click **Next**.

6. Click **Deploy** to deploy the chosen plug-in on the selected Management Agents.
7. Enterprise Manager displays a page that monitors the deployment status and begins the deployment process with the Install option. Deployment occurs in parallel on all selected agents.

17.5.2.4 Removing a Plug-in

To remove a plug-in from Oracle Management Service or a Management Agent, follow the steps below. Removing a plug-in also removes all of its metadata from the Management Repository.

Note that default plug-ins provided by Oracle cannot be un-deployed.

1. From the **Setup** menu, choose **Extensibility>Plug-ins**.
2. Select the row for the plug-in you want to remove to in the table.
3. Click **Undeploy From**, then either **Management Servers** or **Management Agen**t. You can then select the OMS or Management Agent you want to remove the plug-in from.

4. Confirm the plug-in removal. Enterprise Manager notifies the connected and relevant Enterprise Manager users and begins the de-configuration process.

8

Patching Enterprise Manager

This chapter provides an overview of patching, and describes how you can patch the Enterprise Manager core components, mainly Oracle Management Service (OMS), Oracle Management Agent (Management Agent), and Oracle Management Repository (Management Repository).

18.1 Overview

A patch is an entity that contains one more bugs fixes. In order to transform a software product with a defect to a software product without a defect, you must apply patches, this process is called Patching. The patching cycle involves downloading patches, applying patches, and verifying the applied patch to ensure that the bug fixes present in the patch reflect appropriately.

Patching involves migrating from one version of the software product to another, within a particular release, unlike upgrading which involves moving from one release of a product to another newer release of the software product.

Oracle periodically releases the following types of patches to fix the bugs encountered in the core Enterprise Manager components:

- **Interim Patches** are released to fix a bug, or a collection of bugs.
- **Interim Patches (for Security bug fixes)** are released to provide customer specific security fixes.
- **Diagnostic Patches** mainly help diagnose and verify a fix, or a collection of bug fixes.
- **Bundle Patch Updates** are cumulative collection of fixes for a specific product or component.
- **Patch Set Updates** (PSU), are cumulative patch bundles that contain well-tested and proven bug fixes for critical issues. PSUs have limited new content, and do not include any changes that require re-certification.
- **Security Patch Updates** are cumulative collection of security bug fixes.

18.2 Patching OMS and Management Repository

Patching OMS and Management Repository follows the Manual patching approach that requires you to follow step-by-step instructions to patch a target. This mechanism of patching expects you to meet certain prerequisites, manually validate the patch for applicability and conflicts, and patch only one target at a time.

This section contains the following topics:

- OMS Patches
- Repository Patches
- Applying OMS and Repository Patches

18.2.1 OMS Patches

OMS patches typically fix one or more OMS errors encountered. These patches can be downloaded from *My Oracle Support* portal. For information about the critical OMS patches released by Oracle that apply to your environment, see the Patch Recommendation region available on the **Patches and Updates** tab in *My Oracle Support*. Alternately, if you know the patch number of the OMS patch that you need to apply, go to the Patch Search region available on the **Patches and Updates** tab, enter the patch number, and click **Search**. The OMS patches are applied using the `OPatch` utility, or the `emctl` utility. For information on applying the patch see "Applying OMS and Repository Patches"

In Enterprise Manager 12*c*, OMS patches are available as OMS Core Patches and OMS Plugin Patches, required to patch the core component and plugins respectively. Ensure that you navigate to the correct directory location under `<middleware_home>` to patch OMS core or OMS plugin:

```
<middleware_home>
    |_____oms
    |_____plugins
          |_____oracle.sysman.db.oms.plugin_12.1.0.1.0
          |_____oracle.sysman.emas.oms.plugin_12.1.0.1.0
          |_____oracle.sysman.mos.oms.plugin_12.1.0.1.0
```

For example, to patch OMS core, you must navigate to `<middleware_home>/oms` and for plugins, navigate to `<middleware_home>/plugins`.

18.2.2 Repository Patches

Enterprise Manager Repository patches typically update PL/SQL procedures, or other SQL content. They are patched using the `emctl` patching tool.

> **Note:** The Repository Patches are applied on the OMS targets. For location details of a core, and plugin patch, see "OMS Patches"

For information on applying the patch see "Applying OMS and Repository Patches"

18.2.3 Applying OMS and Repository Patches

To apply OMS or repository patches, follow these steps:

1. Log into *My Oracle Support* (`https://support.oracle.com`) console with the necessary credentials.

> **Note:** Check the Patch Recommendation region to view the patches recommended for your environment. You can also provide the recommended patch number in the patch search region to download the recommended patch.

2. On the *My Oracle Support* home page, click **Patches and Updates.**
3. Enter the patch number in the Patch Search region, and click **Search.**
4. Select the patch, and from the context menu, click **Download**.
5. After downloading the zip file, follow the instructions available in the `Readme.html` or `Readme.txt` to patch the target.

> **Note:** Depending on whether it is a core patch or a plugin patch you must log into the host machine, navigate to the directory location, and unzip the patch file.

18.3 Patching Enterprise Manager Agents

Oracle offers two approaches to apply the Agent patches: automated approach, and the manual approach. Oracle strongly recommends you to use the automated approach because it not only saves time and effort in mass-deploying patches but also reduces human intervention, thereby minimizing the errors involved while patching.

This section contains the following topics:

- Management Agent Patches
- Automated Agent Patching
- Manual Agent Patching

18.3.1 Management Agent Patches

Management Agent Patches are released to fix one or more errors encountered in the agent targets. In addition to the Management Agent running on OMS, you can patch all the agent targets that report to the OMS running on your host machine.

A GUI based utility called Patches and Updates is used to patch the Agent targets. For information about accessing the tool, see "Accessing Patches and Updates".

In Enterprise Manager 12*c*, there are separate Agent patches for core components and for plugins. Ensure that you navigate to the correct directory location under `<installation_base_directory>` to patch Agent core or Agent plugin:

```
<installation_base_directory>
    |_____core
          |_____12.1.0.1.0
    |_____plugins
    |_____plugins.txt
    |_____plugins.txt.status
    |_____agent_inst
    |_____sbin
    |_____agentimage.properties
```

For example, to patch Agent core component, you must navigate to `<installation_base_directory>/core/` and for plugins, navigate to `<installation_base_directory>/plugins`.

18.3.1.1 Patches and Updates Versus My Oracle Support

The Table 18–1 captures the advantages of using Enterprise Manager Cloud Console over *My Oracle Support* for Automated Patching:

Table 18–1 Advantages of using Enterprise Manager Cloud Console to My Oracle Support

Enterprise Manager Cloud Control	My Oracle Support
Enterprise Manager Cloud Control enables you to leverage existing Cloud Control Agents to perform configuration data collection without having to install individual configuration managers on each managed target.	You must install Oracle Configuration Manager 10.3.2 (or higher) or Enterprise Manager Cloud Control for My Oracle Support configuration collection (Enterprise Manager harvester) on each of the managed target for the latest patch recommendations.
Enterprise Manager Cloud console is inherently built with the intelligence to support the entire lifecycle of patching, it allows you to select the recommended patch for your environment, add to a plan, and deploy the patch on the selected targets after validating the patches for applicability and conflicts in your environment.	MOS only supports adding the recommended patches to a plan, and validating the patch for conflicts. You cannot deploy the patch from MOS.
In Enterprise Manager, once the patches are applied, they will not appear in the recommended patches list for the selected target.	In MOS this is not an automated process, depending on the data collection by the configuration manager present on the patched targets, the patch recommendation region is automated.

18.3.2 Automated Agent Patching

Automated Patching is a quick-and-easy, reliable, and a GUI-based patching mechanism that is facilitated using Patch Plans, a new concept introduced through the Patches and Updates functionality within the Enterprise Manager Cloud Control console (Cloud Control console).

Automated patching can be performed in the Online mode, and in the Offline mode. In the Online Mode, you can connect to *My Oracle Support* Website to download the patches. However, if you are patching in the offline mode, then you must ensure that the patches to be applied are already available on the Software Library.

This section includes the following:

- Accessing Patches and Updates
- Viewing Patch Recommendations
- Searching Patches
- Applying Agent Patches
- Verifying the Applied Agent Patches
- Validating Agent Patch Errors
- Deinstalling the Applied Agent Patches
- Troubleshooting Agent Deployment

18.3.2.1 Accessing Patches and Updates

To access the Patches and Updates page, in Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching,** and then click **Patches and Updates**.

The following page appears:

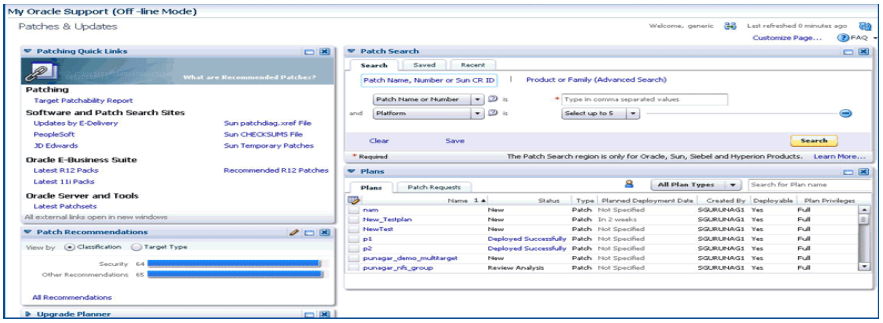

Some of the advantages of using the Patches and Updates page (Automated Patching approach) are:

- Patching operations are more organized, done through a single window, and is always initiated only from the OMS.
- Allow you to schedule jobs that will run periodically, and connect to *My Oracle Support*, check for the latest patches, and automatically download them. This relieves you of the manual effort involved in searching the latest patches and patch sets, and downloading them whenever they are available.
- One patch plan can be used to add multiple patches to multiple sets of homogeneous targets. For example, both core and plugin agent patches can be added to the same plan.

18.3.2.2 Viewing Patch Recommendations

The Patch Recommendation region is one of the regions in the patching domain that proactively communicates all the recommendations that are applicable to your environment. This region minimizes human effort in terms of searching for Oracle recommended patches, which may or may not apply to your environment.

Note: Keep the following points in mind:

- Recommendations are not available for custom plugins. Oracle only supports the default plugins that ship with the product, and releases timely updates for them.
- You must use the configuration manager release 10.3.2 or higher for Patch Recommendations to be enabled.

Figure 18–1 captures the Patch Recommendations region as it appears in the Patches and Updates page.

Figure 18–1 Patch Recommendations

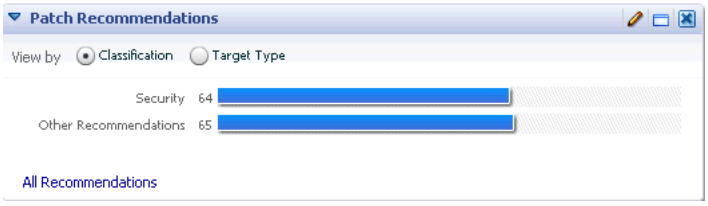

Patches are primarily classified as Security patches, and Other Recommended patches. For example, if you select **Security** from the graph, all the security related patches are displayed on the Patch Recommendation page. Alternately, you can also view the patches by their target types. Click the bar graph to drill down to a list of recommended patches, view details about those patches, download the patches, or add them to a patch plan. The bar graph summarizes the number of issues found (for example, if there is one issue, then there is one recommendation for one target).

Patch Recommendation region allows you to:

- Compare the patches installed in your configuration with what Oracle recommends, and identifies any missing patches.
- Identify and prioritize targets missing Critical Patch Updates.
- Identify missing recommended patches issued by Oracle.

18.3.2.3 Searching Patches

The Patch Search is a region in the patching domain that allows you to search for Oracle, Sun, Siebel, and Hyperion products. The primary purpose of searching a patch is to limit results to the exact patch name or number.

From Figure 18–2 you can see that the Patch Search region contains three tabs: **Search**, **Saved**, and **Recent**. The Search tab lets you apply the desired filters to drill down to the exact results. You can choose to save the search, and view it later. All the saved searches appear in the Saved region. A history of all the searches is registered in the Recent tab, and you can check the logs if you want to access them.

Figure 18–2 Patch Search

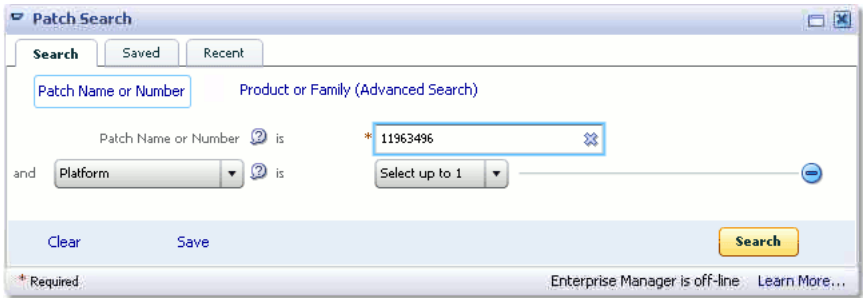

Use any of the following approaches to search for a patch:

- Basic Search
- Advanced Search

Basic Search

To perform a **Basic Search**, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Patching and Provisioning**, and then click **Patches and Updates**.
2. On Patches and Updates page, enter the **Patch Name or Number, or Sun CR ID,** if there are multiple values, then they must be separated by commas. You can select the platform name from the list in the Patch Search Region, and then click **Search**.
3. The Patch Search Results page displays the results based on the search criterion provided.

   From the Patch Search results page, you can do the following:

   – You can highlight one or more patches in the search results and, from the inline tool bar, add the patches to a patch plan (if you use the collector), download the patches, or copy patch details to the system clipboard.

   – If a single patch is highlighted, you can view the patch readme.

Advanced Search

To perform an **Advanced Search**, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Patching and Provisioning**, and then click **Patches and Updates**.
2. On the Patches and Updates page, click **Product or Family (Advanced Search)**.
3. Enter a **Product** name, and set the **Release** number to narrow down your search.

   Additionally, if you are accessing Patches and Updates page in the Offline mode, then click **(+)** icon to add more filters like **Type**, **Platform, or Language**. In the Online mode, in addition to all the filters available in the Offline mode, you can use advanced search categories like: **Classification, Description, Patch Target, or Updated** to drill down to the desired results.

   Note: Advanced Search allows you to search for recommended patches for your environment through the **Classification** search category available when you are accessing Patches and Updates page in the online mode.

4. Click **Search**.

   After updating the appropriate search filters, you can save the search combination by clicking **Save**.

5. The Patch Search Results page displays the results based on the search criterion provided.

   From the Patch Search results page, you can do the following:

   – You can highlight one or more patches in the search results and, from the inline tool bar, add the patches to a patch plan (if you use the collector), download the patches, or copy patch details to the system clipboard.

   – If a single patch is highlighted, you can view the patch readme.

18.3.2.4 Applying Agent Patches

To apply an agent patch using patch plans, follow these steps:

Note: Applying the core agent patch, or the plugin agent patch follows the same patching process (as listed in this section.)

Ensure that the patches selected are applied on homogeneous set of targets, which means that the patches selected should have the same platform and version, as the targets being patched. For example, 12.1.0.1.0 Linux x86 patches can only be applied on 12.1.0.1.0 LinuxX86 targets, any mismatch will result in a patching error.

---

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, and then click **Patches and Updates**.
2. On the Patches and Updates page, select Agent patches from one of the following regions:
   - In the Patch Recommendation region, click the graph to drill down to the list of recommended patches for your environment.

     For more information on Patch Recommendation, see "Viewing Patch Recommendations"
   - In the Patch Search region, enter the patch number of the Agent patch or perform an Advanced Search to search, and select the desired patch.

     For more information on Basic and Advanced search, see "Searching Patches"
3. Select one or more patches, and from the context menu, click **Add to a Plan**.

   Click **Add to New** to create a new plan, if not you can update an existing plan by selecting **Add to Existing** option.
4. If you select **Add to New**, then one of the following dialog box appears:
   - In the Create a New Plan dialog box, enter a unique name for your plan, and click **Create Plan**.

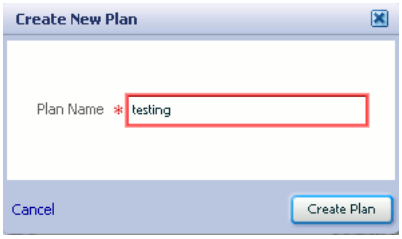

     The patch selected along with the associated target gets added to the plan.
   - In the Add Patch to Plan dialog box as shown in Figure 18–3, enter a unique name for the plan, and select the targets. To search and select targets, follow one of these approaches:
     - If you know the target name, then enter the name of the target in the search field.
     - Click the search icon to view all the targets reporting to the OMS running on your host machine, group them by type **Agent**, and select the desired Agent targets.
     - Create a group of Agent targets, and provide the group name in the search field to add all the targets in that particular group to the plan. To create a group of targets, from **Setup** menu, select **Add Target**, and then click

group. On the Create Group page, add all the Agent targets to the group, and create the group with a unique name.

After selecting the targets, click **Create Plan**. The selected patches and associated targets are added to the plan after validating for conflicts.

Figure 18–3 Add Patch To Plan

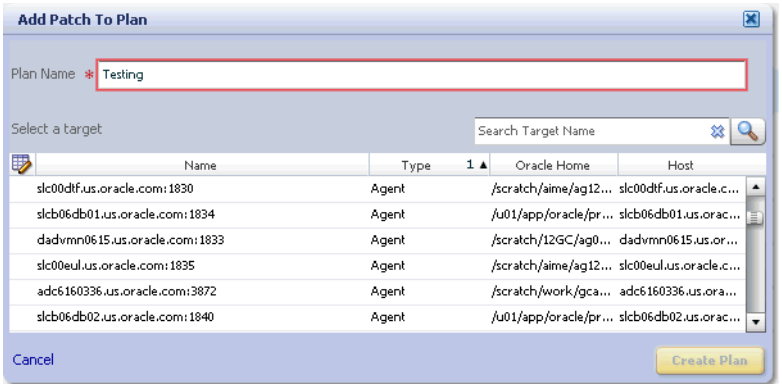

Note: If you want to add to an existing plan, click **Add to Existing** and, then click **More.** From the Add Selected Patch to Plan dialog box, select the plan name and click **Select Targets**. From the Add Selected Patch to Plan `<plan_name>` dialog box, select the targets and click **Add to Plan**.

5. If the selected patches are applied on homogeneous targets, then the plan gets successfully created with a link to **View Plan**. Click the link to view the plan details.

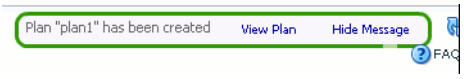

If any of the agent targets added to the patch plan are NFS-Agents, then you may see a warning message **Issues with Adding Patch** as shown in Figure 18–4. As a solution to this problem, a list of all the targets impacted appear, click **Add All To Plan** to add all the affected targets to the patch plan.

Figure 18–4 Issues with Adding Patch: Impacts Other Targets

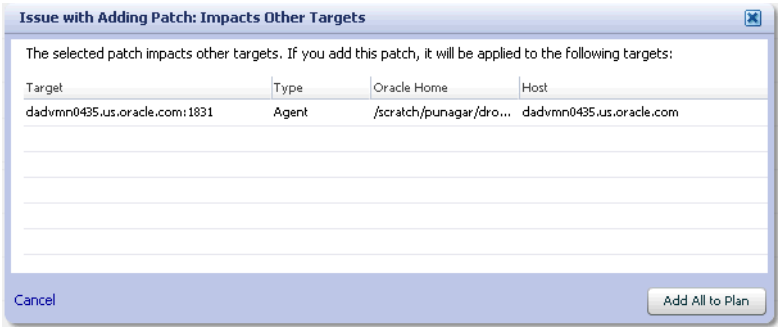

However, if there is a mismatch between the platform and version of the patch selected, and that of the target, you may see one of the following warnings:

Note: Oracle recommends that you fix the warning before proceeding as it may result in an error during the plan validation. However, if you still want to proceed, you can select the patches, and click **Ignore Warnings and Add** to proceed

---

a. **Target is Down**: This warning message appears when the target added is not up and running. All the other homogeneous targets get added to the plan, and the plan is created without this target. Click **View Plan** to see the details.

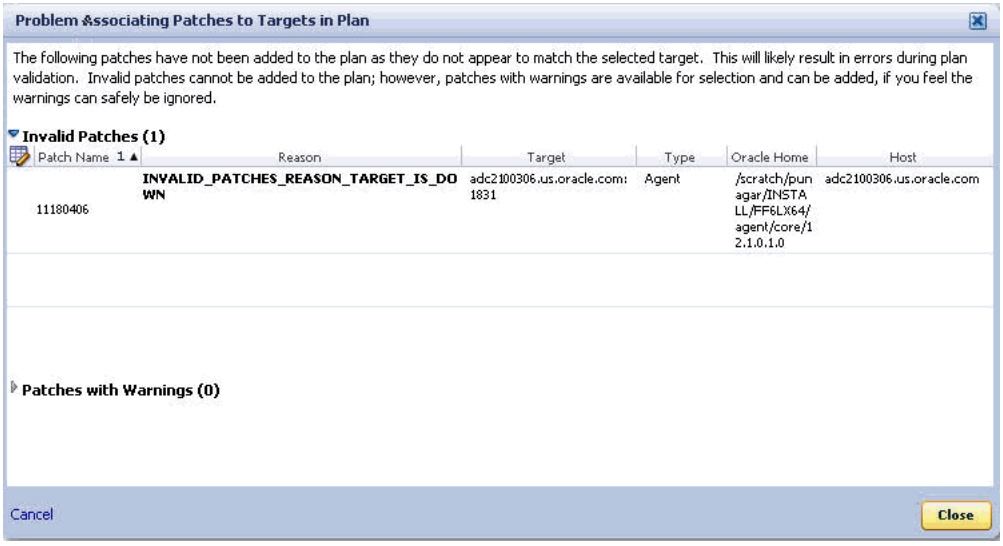

b. **Null Platform**: This error occurs when the target selected appears with a null platform, the validation fails as there is a mismatch in the platform of the patch and that of the target. This could happen when a target is down, in this case the plan is not created until the error is fixed.

6. On the Patches & Updates page, in the Plans region, click the plan name you want to view, and from the context menu, click **View**.

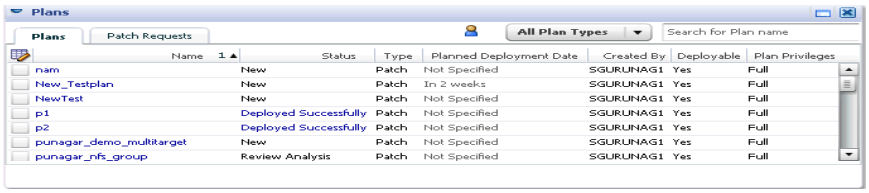

   To filter the plans table, select **All Plan Types** or **Patch** depending on your preference. To search for a plan, enter a plan name or partial plan name in the search box, then click the search button.

7. In the Create Plan Wizard, on the Plan Information page, in the Overview section, validate the Patch Plan name. You can choose to edit it if you want.

   (Optional) Select a date and time when you want to deploy the Patch Plan, and enter a short description to describe the Patch Plan.

8. Click **Next**.

9. On the Patches page, review the patches added to the Patch Plan.

   To associate additional targets to a patch that is already in your Patch Plan, click **Add Patch**. In the Edit Search dialog box, enter the patch number and click **Search**. Select the patch, and click **Add to This Plan**. From Add Patch To Plan dialog box, select the targets, and click **Add to This Plan.**

10. Click **Next**.

11. In the Deployment Options page, retain the default location (`%emd_emstagedir%`) available on the target machine or edit the **Stage Location** to provide a new location for staging the Agent patches.

    In the Credentials section, select Oracle Home Preferred Credentials if you have already set them earlier. You can otherwise click **Override Oracle Home Preferred credentials** and set the Normal Oracle Home Credentials and Privileged Oracle Home Credentials to access the Oracle home of the target.

12. Click **Next**.

13. On the Validation page, click **Analyze** to validate the patch before deploying it. A Validation job is submitted, that performs an exhaustive list of checks like: check for conflicts, check for the latest OPatch version, check if the version and platform of the targets and the patch match (homogeneity rule), and so on in the background.

    To track the progress of the job, 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**. If the status is **Succeeded**, then the job is valid. If not, then review the issues, and try to resolve it according to the corresponding problem description available on the page. After resolving the issue, click **Re-Analyze**.

    Upon validation, if there are conflicts between the two patches, then you might be recommended to request for replacement patches. In this case, click **Request Patch**.

    If there is a Merge Patch already available, you can directly opt to replace the conflicting patches with the Merge Patch. In this case, click **Replace Patch**.

    > **See Also:** For more information about the common errors during the Validation phase, see "Validating Agent Patch Errors"

14. Click **Next**.

15. On the Review & Deploy page, review the details you have provided for the patch plan, then click **Deploy**.

    A job is submitted, to track the progress of the job, 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**.

    > **Note:** For a demonstration about how to apply patches on Enterprise Manager 12*c* Agents using Cloud Control Console, see *My Oracle Support* note 1359221.1.

18.3.2.5 Verifying the Applied Agent Patches

To verify the applied patches using Enterprise Manager, perform the following steps:

1. In Cloud Control, click **Targets**, and then select **All Targets**
2. On the All Targets page, enter the **target name** in the Search Target Name field to search for the target you patched.

    For example, enter **adc2101818** in the search field, and click the search icon.

3. Click the target name to select Oracle home of the target that was patched.

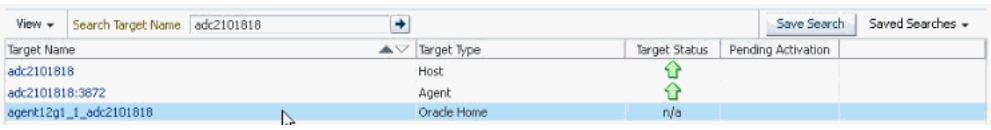

    A summary page with details about the target is displayed.

4. On the Summary page, in the Patch Advisories region, select **Patches Applied** tab to verify all the patches that have been successfully applied on the target.

18.3.2.6 Validating Agent Patch Errors

Here are some of the errors that may appear during the Validation phase of patch plans:

- **Problems Associating the Patches To the Plan**

  This error occurs when patches do not match the targets selected.

  For example if the patch is released for Linux x86 platform, and you are trying to apply the patch on a Linux x64 target, then the plan fails.

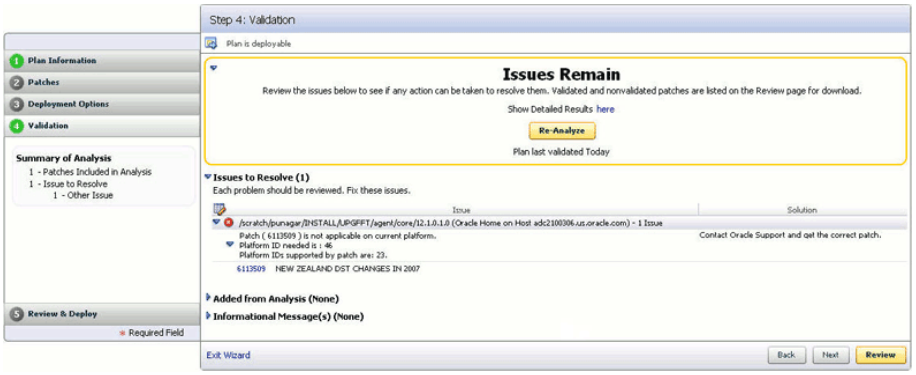

- **Oracle Home Credentials Not Set**

  This error occurs when the oracle Home credentials like Privileged Oracle Home credentials or the Normal Oracle Home credentials are not set.

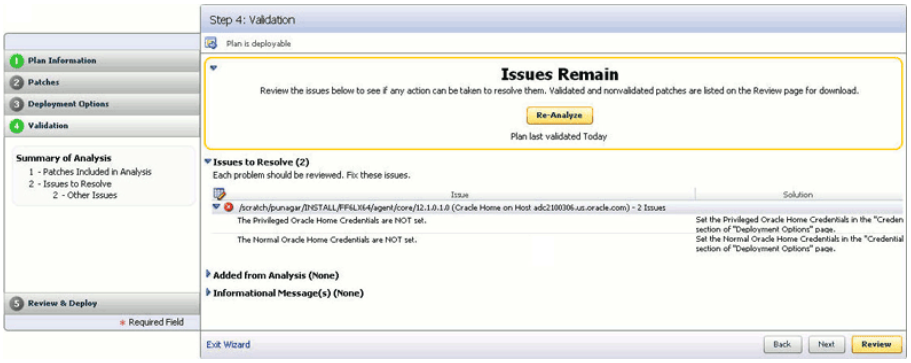

- **Conflict Check Analysis Failure**

  This error occurs when there is a conflict in the patches added.

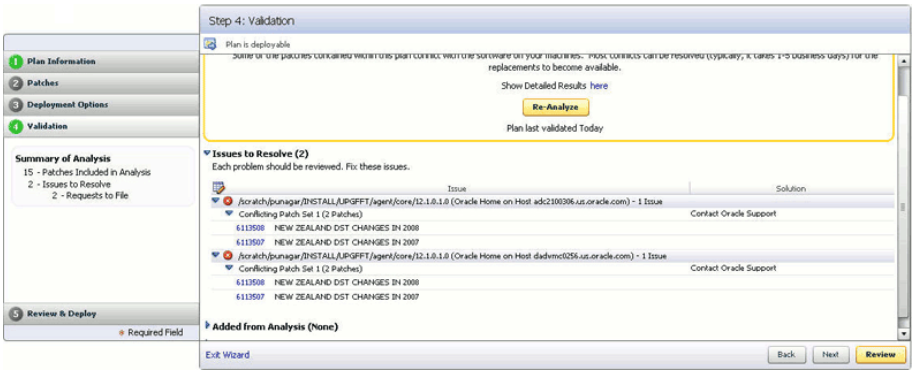

18.3.2.7 Deinstalling the Applied Agent Patches

To roll back applied Agent patches, follow the deinstallation steps available in the `Readme.html` or `Readme.txt` for the patch.

18.3.2.8 Troubleshooting Agent Deployment

The section describes how to troubleshoot issues that you might encounter while patching Agents through Patch Plans.

Error

Applying a Patch Set Update on an Agent target that has already been patched a couple of times earlier, can cause a conflict, which may result in an Analyze Job Failure.

Workaround

To roll back the patch applied, do the following:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, and then click **Software Library**.
2. On the Software Library home page, expand the folders under **Software Library** as follows (**Software Library >>Patching>>Agent>>All>>Generic>Directive**):

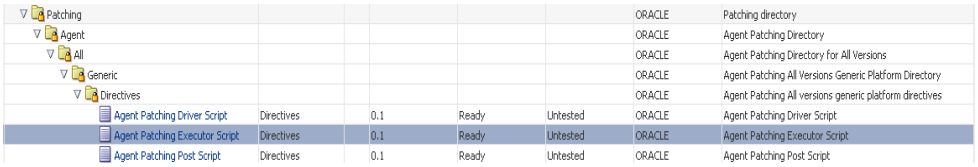

3. From the Directives folder, select **Agent Patching Executor Script**, and then click **View**.
4. On the View Entity page, select **Select Files** tab, and then click **emdpatch.pl** to download the file.

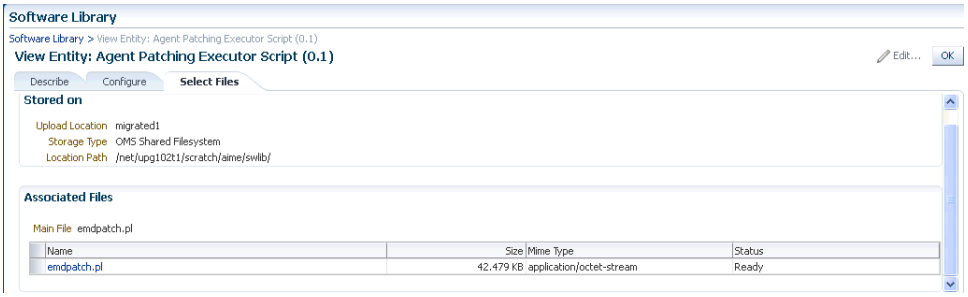

5. Use any editor to open the file, and do the following:
   a. Search for the text **# Step 4: apply the patch here** in the file.
   b. Add the following lines:

```
# Begin of Custom Rollback steps

my $custRollbackCmd =  File::Spec->catfile($ORACLE_HOME, 'OPatch',
'opatch');
my $custRollbackStatus = echodo($custRollbackCmd . " nrollback -id 6113507
-silent " . $inv_loc);
$custRollbackStatus = statusf($custRollbackStatus);
if($custRollbackStatus == 0)
```

```
  {
     logf("Patches rolled back successfully");
   }
else
  {
      logf("Rollback of the patches failed.");
   }

# End of custom rollback steps
```

6. Copy the updated file `emdpatch.pl`, edited in the preceding step to the following location:

   `$OMS_HOME/sysman/metadata/swlib/patch/directives`

7. Navigate to the following location:

   `$OMS_HOME/sysman/metadata/swlib/patch/`

8. Edit `swlib.xml` file to update the entity version of the Agent Patching Executor Script to **0.2**, as follows:

```
<Entity name="Agent Patching Executor Script">
  <LocalizableDescription default="Agent Patching Executor Script" nlsid =
"AGENTPATCH_EMDPATCH_SCRIPT_DESC"/>
  <Type>COMP_Directives</Type>
  <Directory>Patching/Agent/All/Generic/Directives</Directory>
  <Maturity>Production</Maturity>
  <Dictionary filename="entityMetadata/Up2date_dictionary.xml"/>
  <ConfigProps filename="entityMetadata/Up2date_properties.xml"/>
  <Fileset>
   <FileEntry mainfile="true">
      <sourcePath>directives/emdpatch.pl</sourcePath>
      <path>emdpatch.pl</path>
   </FileEntry>
</Fileset>
<Status>Status_Complete</Status>
<ExternalID>0.2</ExternalID>
</Entity>
```

9. Save the changes made to `swlib.xml`, and then run the following commands:

   ```
   $OMS_HOME/bin/emctl register oms metadata -service swlib
   -file
   ```

   ```
   $OMS_HOME/sysman/metadata/swlib -core
   ```

18.3.3 Manual Agent Patching

Manual patching is a patching mechanism that requires you to follow step-by-step instructions to patch a Management Agent manually. This mechanism of patching expects you to meet certain prerequisites, manually validate the patch for applicability and conflicts, and patch only one Management Agent at a time.

> **Note:** Oracle recommends you to use the automated patching mechanism because it not only saves time and effort in mass-deploying patches but also reduces human intervention, thereby minimizing the errors involved while patching.

To patch manually, you must perform the following steps:

1. Log into *My Oracle Support* (`https://support.oracle.com`) console with the necessary credentials.

   Note: Check the Patch Recommendation region to view the patches recommended for your environment. You can also provide the recommended patch number in the patch search region to download the recommended patch.

2. On the *My Oracle Support* home page, click **Patches and Updates.**
3. Enter the patch number in the Patch Search region, and click **Search.**
4. Select the patch, and from the context menu, click **Download**.
5. After downloading the zip file, follow the instructions available in the `Readme.html` or `Readme.txt` to install the patch.

19

Configuring Software Library

This chapter describes how to configure a new Software Library, and maintain an existing Software Library in the Enterprise Manager Cloud Control environment.

In particular, this chapter covers the following:

- Overview of Software Library
- Users, Roles, and Privileges
- Software Library Storage
- Prerequisites for Configuring Software Library
- Configuring Software Library Storage Location
- Maintaining Software Library

19.1 Overview of Software Library

Oracle Software Library (Software Library) is one of the core features offered by Enterprise Manager Cloud Control. Technically, it is a repository that stores software entities such as software patches, virtual appliance images, reference gold images, application software, and their associated directive scripts. In addition to storing them, it also enables you to maintain versions, maturity levels, and states of these software entities.

To access the Software Library console page, in Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching** and then, click **Software Library.** On the Software Library home page, as shown in Figure 19–1, there are two types of folders: Oracle-owned folders (marked by a lock symbol) and User-owned folders.

Oracle owned folders and their contents (including other subfolders and entities) ship with the product by default, and appear on the Software Library home page after software library is configured post Enterprise Manager installation. User-owned folders are logical top level folders that the user creates to organize the entities that he intends to create.

Figure 19–1 Software Library Console

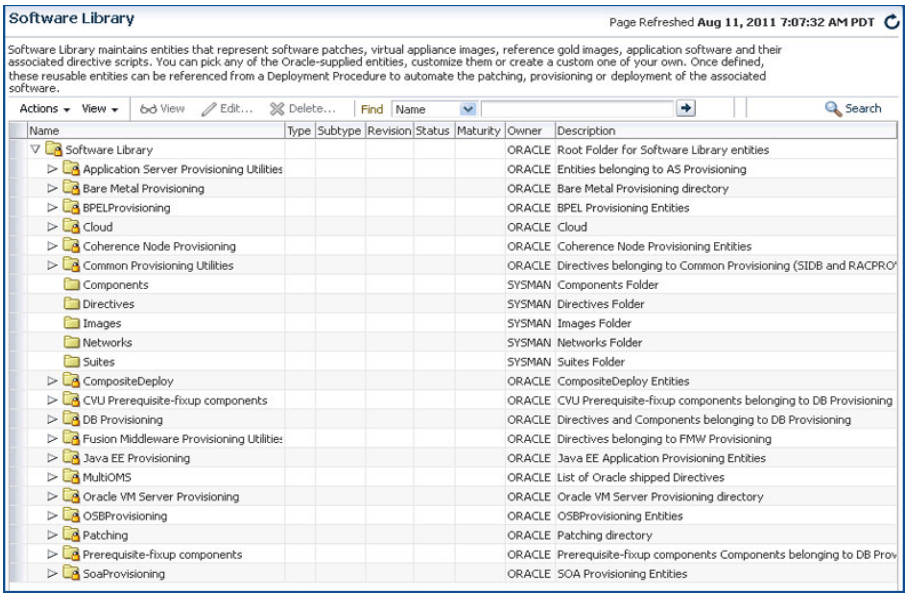

The Software Library Page is a GUI based screen that enables you to create and maintain entities that are used by automation framework to perform a number of Patching and Provisioning operations.

From the Software Library Console page, you can perform the following tasks:

- Configure Software Library Storage, see "Configuring Software Library Storage Location" for more information.
- Create Software Library Entities. For example, Creating a Generic Component, Creating Directives, and so on.
- Manage Software Library Entities. For example, Viewing Entities, Editing Entities, Deleting Entities, Searching Entities, and so on.

> **See Also:** For information about Creating and Managing Software Library Entities, see *Oracle Enterprise Manager Administrator's guide for Software and Server Provisioning and Patching*.

19.2 Users, Roles, and Privileges

Software Library can be accessed only if you are granted one or more Software Library privileges according to your roles and responsibilities. Fine grained privileges provide a way to control user access to the different entities in the Software Library.

> **Note:** All the folders and entities that ship with the product, by default are viewable by all the Enterprise Manager users.

Administrator by default do not have any Software Library privileges, it is for the Super Administrator, to grant access, privileges to an Administrator. Table 19–1 describes all the available Software Library privileges that can be granted to a user or role.

Users and roles can be granted privileges on specific entities by the owner of the entity or the Super Administrator. For more details, see *Oracle Enterprise Manager Administrator's Guide for Software and Server Provisioning and Patching*.

Table 19–1 Software Library Privileges for Administrators

Resource Type	Description
View any Template Entity	Ability to view any Template Entity
Export Any Software Library Entity	Ability to export any Software entity
Edit any Software Library Entity	Ability to edit any Software Library entity
Manage Any Software Library Entity	Ability to create, view, edit, and delete any Software Library entity
Import Any Software Library Entity	Ability to import any Software Library entity
Create Any Software Library Entity	Ability to create any Software Library entity
View Any Software Library Entity	Ability to view any Software Library entity
View Any Assembly Entity	Ability to view any Assembly entity
Grant Any Entity Privilege	Ability to grant view, edit, and delete privileges on any Software Library entity.This privilege is required if the user granting the privilege on any entity is not a Super Administrator or owner of the entity.

Table 19–2 describes all the primary users of Software Library, and their associated privileges:

Table 19–2 Roles and Privileges

Role	Software Library Privileges
Super Administrator	All Software Library Privileges
`EM_PROVISIONING_DESIGNER` (Designer)	Create Any Software Library Entity
`EM_PROVISIONING_OPERATOR` (Operator)	View Any Software Library Entity
`EM_PATCH_OPERATOR`	Create Any Software Library Entity View Any Software Library Entity
`EM_USER` (Administrator)	Access Enterprise Manager

Super Administrators have complete privileges on all the entities present in Software Library, and can exercise access control on the entities by granting one or more privileges, and later revoking the previously granted privilege to another user or role.

Designers by default are given create privileges, which allow them to create entities and manage them.

Operators by default are given view privileges, which allow them to view all the entities in Enterprise Manager Cloud Control.

Any Enterprise Manager user requires, at the very least, a view privilege on an entity for the entity to be visible on the Software Library console. The Super Administrator can choose to grant additional privileges described in Table 19–1 to the user or role. Users will not be able to see this entity till the Super Administrator grants them at least a view privilege on the entity.

19.3 Software Library Storage

The Software Library Administration console allows you to configure and administer Software Library. To start using the Software Library, you must add at least one upload file storage location (OMS Shared location, or OMS Agent location) on the host where the OMS is running. A storage location in Software Library represents a repository of files, these files are either uploaded by Software Library, or generated and saved by some user-owned process.

To access the administration console, log into Enterprise Manager Cloud Control with Administration access, and follow these steps:

In Cloud Control, from **Setup** menu, select **Provisioning and Patching**, and then click **Software Library**.

OR

In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching** and then, click **Software Library.** On the Software Library home page, from **Actions** menu, select **Administration**.

Figure 19–2 Software Library Administration

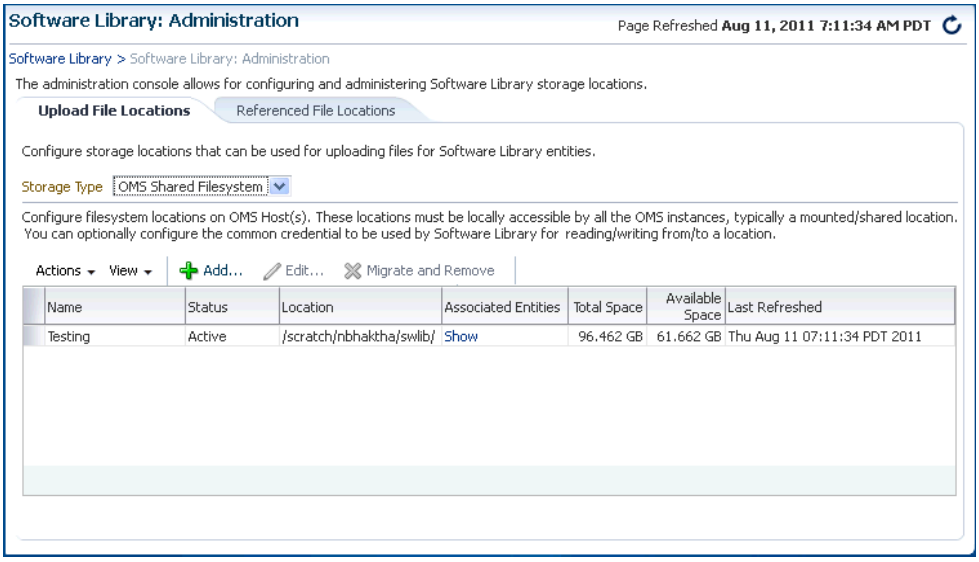

The Software Library Administration Page as shown in Figure 19–2 is a GUI based screen, that enables you to create one or more storage locations to store or refer to files that are associated with an entity. To view the entities present in the storage location, click **show** on the Administration page. You can create a storage location on the OMS or the agent running on the same host as the OMS. With Enterprise Manager 12*c*, a new feature called Referenced Storage location has been introduced, wherein Software Library entities can refer to files that are stored on another host. These locations are however read-only for Software Library, will not be used for uploading files.

The space requirements for configuring Software Library depends on the amount of space required for storing the software binaries, and its associated scripts. Understandably, this space requirement increases over a period of time as you create more entities. Depending on the features or software required for provisioning and patching, you must decide and configure the Software Library storage space.

> **Note:** When a storage location starts running out of space, then it is important to deactivate the configured storage location so that no new uploads can happen to this location. For more information about removing a storage location, see "Maintaining Software Library"

The following types of storage locations are available:

- Upload File Locations
- Referenced File Location

19.3.1 Upload File Locations

Upload File Locations are locations configured for storing files uploaded by Software Library as part of creating or updating an entity.

> **Note:** For Software Library to become usable, at least one upload file location must be configured. On adding the first upload file location, the default Software Library metadata for all installed plugins is imported from the OMS Oracle home, and a job is submitted. Ensure that you wait for this job to complete successfully, before performing other patching or provisioning operations.

As a prerequisite, before using Upload File Locations as storage option, you must set credentials for using an OMS Shared File System or OMS Agent File System:

- For multiple OMS environment, all the OMS hosts must have a preferred normal host credential set.

  > **Note:** When OMS instances are added, it is necessary to ensure that the configured locations are accessible from the designated host where the new OMS will be provisioned.
  >
  > For a OMS that will be provisioned using the Add OMS functionality, the shared location configured as upload location should be mounted on the designated host, and verified manually.

- For OMS Agent File System location configuration, a credential (preferred or named) has to be specified.

Upload File Locations support two storage options:

OMS Shared File System

An OMS Shared File System location is required to be shared (or mounted) across all the Oracle Management Server (OMS) hosts. This option is ideal for UNIX systems.

For single OMS environments, you can configure the Software Library either on the host where the OMS is running, or in a shared location. However, in future, if you plan to expand the single OMS setup to a multiple OMS setup, then local file system path is not recommended.

For multiple OMS environments, Oracle recommends you to configure the Software Library in a non-local, shared file system path that is accessible through NFS mount points to all Oracle Management Servers in the environment. Besides accessibility, it is

important to ensure that there is enough space available for the storage of software binaries, and associated scripts for the entities that you want to create and store.

OMS Agent File System

An OMS Agent File System location should be accessible to the agent running on the host machine where the OMS is deployed, and is recommended for multiple OMS setup on Windows. To use this storage option, ensure that you have a preferred, or a named credential for the OMS host. Click **Change Credential** to change the associated credential to be used to access this location. Ensure that credential associated with the location is viewable by all Software Library designers, so that other designers can upload, and stage the files associated with any entity.

19.3.2 Referenced File Location

Referenced File Locations are locations that allow you to leverage the organization's existing IT infrastructure (like file servers, web servers, or storage systems) for sourcing software binaries and scripts. Such locations allow entities to refer to files without having to upload them explicitly to a Software Library storage.

Referenced File Locations support three storage options:

- **HTTP**: An HTTP storage location represents a base URL which acts as the source of files that can be referenced.

  For example, the base URL `http://my.files.com/scripts` could be configured as an HTTP location for sourcing files such as `http://my.files.com/scripts/perl/installMyDB.pl` or `http://my.files.com/scripts/linux/stopMyDB.sh.`

- **NFS**: An NFS storage location represents an exported file system directory on a server. The server need not be an Enterprise Manager host target.

  For example, the directory `/exported/scripts` is exported on server `my.file.server` could be configured as an NFS location for sourcing files such as `/exported/scripts/generic/installMyDB.pl` or `/exported/scripts/linux/stopMyDB.sh` once mounted on a target host file system.

- **Agent**: An Agent storage location is similar to the OMS Agent File System option, but can be any host monitored by an Enterprise Manager Agent. The Agent can be configured to serve the files located on that host.

  For example, the directory `/u01/binaries` on the Enterprise Manager Host `my.em.file.server` could be configured as an Agent location for sourcing files such as `/u01/binaries/rpms/myCustomDB.rpm` or `/u01/binaries/templates/myTemplate.tar.gz`.

  These locations require a named credential to be associated which will be used to access the files from the base location on the host through the Enterprise Manager Agent.

19.4 Prerequisites for Configuring Software Library

To administer the different storage types, and to configure software library, keep the following points in mind:

- Depending on the features or software required for provisioning and patching, you must decide and configure the Software Library storage space.

  A minimum storage of 2 GB is required to store all the Oracle shipped files.

- On UNIX systems, Oracle recommends that you configure at least one OMS Shared File System location that is accessible from all the OMS hosts.

  On Windows systems, Oracle recommends configuring OMS Agent File System storage.
- Each OMS host must have a preferred normal host credential set before configuring the location. For OMS Agent File System location configuration, a credential (preferred or named) has to be specified.
- You (the user configuring the Software Library) must have view privilege on all the OMS, and the agent targets running on the host machine. As per the accessibility verification, you must be able to view, and edit these newly configured locations using the credentials described in the proceeding point.
- All the credentials used while configuring the locations, must be viewable by all the Software Library designers, as this will enable the designers to upload the files to the Software Library storage. Additionally, designers can also stage the uploaded files to other hosts for provisioning or patching activities.

19.5 Configuring Software Library Storage Location

System Administrators are responsible for configuring a storage location. Only after the storage location is configured, you can start uploading the entity files.

You can configure the Software Library in one of the following locations:

- Configuring an OMS Shared File System Storage Location
- Configuring an OMS Agent Storage Location
- Configuring a Referenced Storage Location

19.5.1 Configuring an OMS Shared File System Storage Location

To configure an OMS Shared File System storage location that can be used for uploading Software Library entity files, perform the following steps:

1. In Cloud Control, from **Setup** menu, select **Provisioning and Patching** and then, click **Software Library.**
2. On Software Library: Administration page, select **OMS Shared Filesystem**.
3. To add a new OMS Shared File System, click **+Add**.
4. In the Add OMS Shared File System location dialog box, provide a unique name, and location on the OMS host, where you want to set up the upload location.

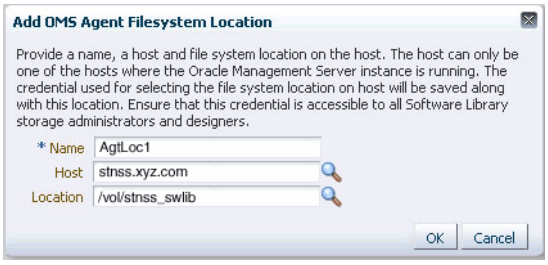

Ensure that the configured storage location is a shared location that is accessible by all the OMS instances. For a Multi OMS setup, set the Normal Preferred Credentials for all the OMS(s).

When you configure an upload location for the first time, a metadata registration job is submitted which imports all the metadata information of all the installed plugins from the Oracle home of the OMS.

To track the progress of the job, 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, choose **Targetless** as the Target Type, and then click **Search**. Typically, the name of the job starts with `SWLIBREGISTERMETADATA_*`.

If the Import job fails, see "Maintaining Software Library" for information on Re-importing metadata for Oracle-owned files.

5. Click **OK** to create a new entry for the storage location in the table, with details like **Name**, **Location**, **Host**, **Status**, and **Host Credentials.**

   In addition to this, you can click **Associated Entities** to view or search the entities present in the selected upload location.

19.5.2 Configuring an OMS Agent Storage Location

To configure an OMS Agent location, perform the following steps:

1. In Cloud Control, from **Setup** menu, select **Provisioning and Patching,** and then click **Software Library.**
2. On Software Library: Administration page, select **OMS Agent Filesystem**.
3. Click **+Add,** in the Add OMS Agent File System Location dialog box, enter the following details:

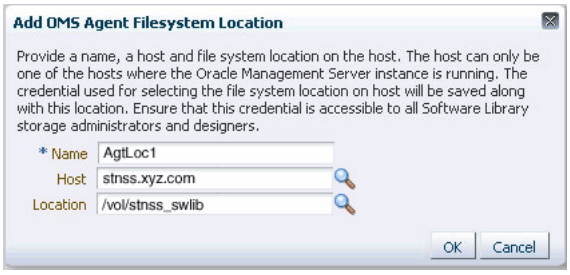

   a. In the **Name** field, enter a unique name for the storage.

   b. In the **Host** field, click the magnifier icon. From the Search and Select: Hosts dialog box, select a host where the OMS is running, and click **Select**.

      For example, `xyz.acme.com`

   c. In the **Location** field, click the magnifier icon. In the Remote File Browser dialog box, click **Login As** to log into the host machine with either Preferred, Named or New credentials.

      Navigate to the location on the host where you want to create the Agent File System, and click **OK**.

      The selected credential is saved along with the host and selected file system path. The saved credential is used to upload files and stage the uploaded files to a host target as part of some provisioning or patching activity.

      Note: The administrator configuring the Software Library must grant view privilege (at least) on the credential chosen to all designers uploading or staging the files to or from this location.

4. Click **OK** to create a new entry for the storage location in the table, with details like **Name**, **Location**, **Host**, **Status**, and **Host Credentials.**

   In addition to this, you can click **Associated Entities** to view or search the entities present in the selected upload location.

These newly configured OMS Agent locations are now available for storing entity files.

19.5.3 Configuring a Referenced Storage Location

To configure storage location that can be used for referring to files from the Software Library entities, perform the following steps:

1. In Cloud Control, from **Setup** menu, select **Provisioning and Patching,** and then click **Software Library.**
2. On Software Library: Administration page, click **Referenced File Locations** tab.
3. To add an HTTP location that can be accessed through a HTTP URL, select **HTTP** from the Storage Type list and click **+Add**.

   In the Add HTTP Location dialog box, enter a unique name and a HTTP location for the storage that you want to reference, and click **OK**.

   A new entry for the storage location is created, with details like **Name**, **Location**, and **Status**.

4. To add an NFS shared location, select **NFS** from the Storage Type list and click **+Add**.

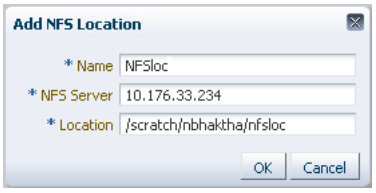

   In the Add NFS Location dialog box, do the following:

   a. Enter a unique name in the **Name** field for the storage.

   b. In **NFS server** field, provide a fully qualified domain name or the IP address of the hosted machine that has NFS services running on them.

   c. In the **Location** field, provide the shared location or directory path on the NFS server to define a storage location, then click **OK**.

   A new entry for the storage location is created in the table, with details like **Name**, **Location**, and **Status**.

5. To add an Agent location that has read-only privileges set on it, select **Agent** from the Storage Type list and click **+Add**.

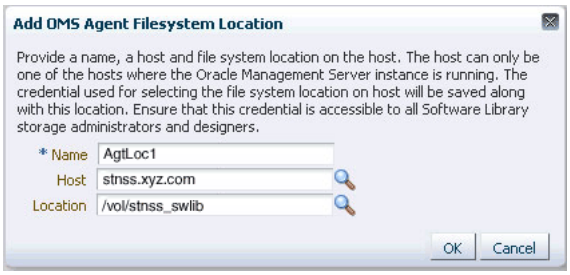

In the Add Agent Location dialog box, enter the following details:

a. In the **Name** field, enter a unique name for the storage.

b. In the **Host** field, click the magnifier icon to select a target from the list available.

For example, `xyz.company.com`

c. In the **Location** field, click **Login As** to select the credentials and browse the previously selected host.

The credential selected, either Preferred, Named or New, is saved along with the host and selected file system path. The saved credential is used for staging the files to a host target as part of some provisioning or patching activity.

Note: The administrator configuring the Software Library must grant view privilege (at least) on the credential chosen to all designers uploading or staging the files to or from this location.

Note: When you create a new entity, these newly configured Referenced File Locations are available as storage options.

19.6 Maintaining Software Library

To maintain the health and proper functionality of the Software Library, the administrator who configured the Software Library, or the Designer who has administration access on it must perform the tasks listed here.

This section includes:

- Periodic Maintenance Tasks
- Re-Importing Oracle Owned Entity Files
- Deleting Software Library Storage Location

19.6.1 Periodic Maintenance Tasks

Periodically, the Administrator must perform the following tasks for proper functioning of the Software Library:

- Refresh the Software Library regularly to compute the free and used disk space.
- Purge deleted entities to conserve disk space
- Check accessibility of the configured Software Library locations to ensure that they are accessible.

19.6.2 Re-Importing Oracle Owned Entity Files

> **Note:** Re-importing metadata applies only to the Oracle owned files, which means all the entity files that ship with the Enterprise Manager product by default. The metadata of User owned entity files cannot be recovered through the Re-import functionality.

Re-Importing the metadata of Oracle owned entity files is not a periodic activity. Re-import helps to recover the metadata files in one of the following situations:

- If you delete the file system location where the metadata was imported. For example, `/scratch/swlib1/`
- If the import job submitted while creating the first upload location fails.

To re-import the metadata of Oracle owned files, do the following:

1. In Cloud Control, from **Setup** menu, select **Provisioning and Patching,** and then click **Software Library.**
2. On the Software Library Administration page, in the Upload File Location tab, from **Actions** menu, select **Re-Import Metadata** option to submit a job that re-initiates the re-import process.

19.6.3 Deleting Software Library Storage Location

Software Library Storage Administrators have the required privileges to delete a storage location. Before removing a storage location, you are prompted to choose an alternate location where the files will be migrated. On selecting an alternate location, a migration job is submitted, and the location is marked **Inactive**. After successful migration of the entity files to the new location, the location configuration is deleted.

> **Note:** To remove a location from OMS Agent File System or Referenced Agent File System storage, you must have a view privilege on the credentials for the location being removed, and the alternate location where the files are migrated.

To delete a configured storage location, perform the following steps:

1. In Cloud Control, from **Setup** menu, select **Provisioning and Patching,** and then click **Software Library.**
2. On the Software Library Administration page, select the storage location, and click **Migrate and Remove**.
3. In the confirmation dialog box, click **Remove** to submit a job, which on successful completion deletes the storage entry from the table.

> **Note:** If you have only one Upload File storage location, either OMS Shared File System or OMS Agent File System, you cannot delete this location. Only if there are more than one configured storage locations, then the **Migrate and Remove** option is enabled to perform a delete operation.

20

Monitoring Using Web Services and JMX

You can extend Enterprise Manager to monitor Web Services and JMX-instrumented applications for critical events, performance problems, error conditions, and statistics.

Enterprise Manager's ability to monitor WSDL and JMX-enabled targets enables you to consolidate monitoring and management operations. When added to the Enterprise Manager framework, Enterprise Manager functionality, such as notifications, jobs, and reporting, is automatically extended to these targets.

This chapter covers the following topics:

Note: This chapter assumes knowledge of Management Plug-ins and the requisite target definition files. For information on Management Plug-in concepts, or developing and deploying Management Plug-ins, see the Enterprise Manager Programmer Developer's Kit documentation.

20.1 Overview

Using Enterprise Manager to monitor targets that expose a Web Services management interface, JMX-instrumented applications and servers, and standalone Java Virtual Machine (JVM) targets entails defining a new target type via Management Plug-ins.

Creating a new Management Plug-in consists of four basic steps:

1. Generate the target metadata and default collection files to be added to the Management Plug-in.
2. Create a Management Plug-in Archive containing the target definition files for one or more Management Plug-ins. A single archive may contain more than one Management Plug-in.
3. Import the Management Plug-in into Enterprise Manager.
4. Deploy the Management Plug-in to the appropriate Management Agent(s).

Procedural information for the monitoring targets can be found in the following sections:

- Section 20.2 discusses software components exposing an external interface that communicate across a network using a standard messaging protocol.
- Section 20.4 discusses J2EE applications running on an OC4J that are instrumented using JMX MBeans.
- Section 20.5 discusses standalone Java applications running on J2SE5.0 or higher that are instrumented using JMX MBeans.
- Section 20.6 discusses JMX applications running on Oracle WebLogic Application Servers 9.x or above.

Section 20.4 and Section 20.5 explain how to generate metadata and default collection files for your custom JMX-enabled application by guiding you through the MBeans for which you are interested in collecting data, and helping you define the MBeans as metrics in Enterprise Manager. Even if your standalone Java application is not instrumented through JMX, you can still monitor the JVMs it is running on by directly creating the built-in JVM target instances as defined in Section 20.10.3.

After the metadata and default collection files are created, you can follow the normal Management Plug-in mechanism to deploy your plug-in and create target instances of your Java application target type as discussed in Section 20.7.

20.2 Monitoring Using Web Services in Enterprise Manager

Web Services are loosely coupled software components that expose an external interface via the Web Service Definition Language (WSDL). These components communicate across a network using a standard messaging protocol called Simple Object Access Protocol (SOAP). The Management Agent's Web Service Fetchlet (with ID WSF) supports SOAP communication.

Note: For more information about the Web Services standard, see the World Wide Web Consortium (W3C) website:

`HTTP://www.w3.org`

Prerequisites

- Enterprise Manager Management Agent version 12.1.0.0.0 or greater installed on that host.
- Enterprise Manager Management Server (OMS) version 12.1.0.0.0 or greater with which the Management Agent communicates.

20.2.1 Creating Metadata and Default Collection Files

Defining a target type to be monitored via a Web Services interface entails creating the requisite target definition files, which are required to collect metrics from resources that support the WSDL interface:

- Target Metadata
- Default Collection

Enterprise Manager provides an easy-to-use Web Services CLI command-line tool that simplifies creating new Management Plug-ins by automatically generating these requisite files. Information retrieval is achieved via the Web Service Fetchlet that is integrated with the Management Agent.

The command-line tool works by parsing a specified WSDL file for all operations, and enables you to select one or more operations to be invoked. If multiple port types are specified in the WSDL file, the tool prompts you to select one of them. Operations are listed along with their parameters. A Web Service operation can be one of four types:

- One Way
- Request Response
- Solicit Response
- Notification

The Request Response operation type is particularly useful: The selected operation could have primitive or complex parameters and results. The result of Web Service invocation is displayed in a table (the tool prompts you to provide labels for the table columns). You can also filter result attributes by specifying an Xpath expression (see the `RowType` property in the generated target metadata, Example 20–3). Filter attributes can be useful for complex return types from which only few attributes are interesting.

The Web Services command-line tool supports Web Services with the following binding and encoding styles:

- DOC/literal
- DOC/Wrapped
- RPC/encoded

20.2.1.1 Web Services CLI Command-line Tool Syntax

The Web Services CLI command-line tool syntax is as follows:

```
emctl wscli [-metadata | -help] [-options]
```

The command accepts the following options:

- -wsdl=<file | URL> : WSDL file or URL [mandatory]
- -username=<user ID> : username if the WSDL is protected

The command-line tool requires a WSDL file name or URL to locate the WSDL for a Web Service. For example, for a Calculator service Web Service, a WSDL URL would be as follows:

```
http://localhost:44861/CalWS/CalculatorPort?WSDL
```

The command tool script requires access to the Enterprise Manager home directory (EM_HOME) to run. The tool defaults to ORACLE_HOME (ensure this environment variable is set properly before using this tool).

The tool parses specified WSDL for all the port types and binding (supported protocols such as HTTP get/post, SOAP) to list all the operations. If there are multiple port types in WSDL, you will first be prompted to choose a port type.

20.2.1.2 Web Services Command-line Tool Security

The command-line tool generates metadata required by Enterprise Manager for target monitoring purposes via the WSDL file. When you run this tool, you only need read permission on the WSDL file or URL and permission to save generated files to the appropriate directory.

20.2.1.3 Generating the Files

Example 20–1 shows a sample WSDL file passed to the command-line tool to generate the target metadata and collection files.

Example 20–1 Sample WSDL File CalculatorService.wsdl

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- Published by JAX-WS RI at http://jax-ws.dev.java.net. RI's version is Oracle
JAX-WS 2.1.5. -->
<wsdl:definitions xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/"
xmlns:tns="http://tests.jaxws.oracle.com/"
xmlns:ns0="http://www.oracle.com/jaxws/tests"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns:mime="http://schemas.xmlsoap.org/wsdl/mime/"
xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
xmlns:soap12="http://schemas.xmlsoap.org/wsdl/soap12/" name="CalculatorService"
targetNamespace="http://tests.jaxws.oracle.com/">
    <wsdl:types>
        <xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema" version="1.0"
targetNamespace="http://www.oracle.com/jaxws/tests/types">
            <xs:complexType name="calculatorFaultInfo">
                <xs:sequence>
                    <xs:element name="number" type="xs:int"/>
                    <xs:element name="reason" type="xs:string" minOccurs="0"/>
                </xs:sequence>
            </xs:complexType>
        </xs:schema>
        <xs:schema xmlns:ns1="http://www.oracle.com/jaxws/tests/types"
xmlns:tns="http://www.oracle.com/jaxws/tests"
xmlns:xs="http://www.w3.org/2001/XMLSchema" version="1.0"
targetNamespace="http://www.oracle.com/jaxws/tests">
            <xs:import namespace="http://www.oracle.com/jaxws/tests/types"/>
            <xs:element name="CalculatorException" nillable="true"
type="tns:CalculatorException"/>
            <xs:element name="CalculatorWrapperException" nillable="true"
type="ns1:calculatorFaultInfo"/>
            <xs:complexType name="CalculatorException">
                <xs:sequence>
```

```
                    <xs:element name="Message" type="xs:string"/>
                    <xs:element name="Number" type="xs:int"/>
                    <xs:element name="Reason" type="xs:string"/>
                </xs:sequence>
            </xs:complexType>
        </xs:schema>
        <schema xmlns="http://www.w3.org/2001/XMLSchema"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns:tns="http://tests.jaxws.oracle.com/"
targetNamespace="http://tests.jaxws.oracle.com/">
            <xsd:complexType name="add">
                <xsd:sequence>
                    <xsd:element name="arg0" type="xsd:int"/>
                    <xsd:element name="arg1" type="xsd:int"/>
                </xsd:sequence>
            </xsd:complexType>
            <xsd:element name="add" type="tns:add"/>
            <xsd:complexType name="addResponse">
                <xsd:sequence>
                    <xsd:element name="return" type="xsd:int"/>
                </xsd:sequence>
            </xsd:complexType>
            <xsd:element name="addResponse" type="tns:addResponse"/>
            <xsd:complexType name="square">
                <xsd:sequence>
                    <xsd:element name="arg0" type="xsd:int"/>
                </xsd:sequence>
            </xsd:complexType>
            <xsd:element name="square" type="tns:square"/>
            <xsd:complexType name="squareResponse">
                <xsd:sequence>
                    <xsd:element name="arg0" type="xsd:int"/>
                </xsd:sequence>
            </xsd:complexType>
            <xsd:element name="squareResponse" type="tns:squareResponse"/>
            <xsd:complexType name="checkNumber">
                <xsd:sequence>
                    <xsd:element name="arg0" type="xsd:int"/>
                </xsd:sequence>
            </xsd:complexType>
            <xsd:element name="checkNumber" type="tns:checkNumber"/>
            <xsd:complexType name="checkNumberResponse">
                <xsd:sequence>
                    <xsd:element name="return" type="xsd:boolean"/>
                </xsd:sequence>
            </xsd:complexType>
            <xsd:element name="checkNumberResponse"
type="tns:checkNumberResponse"/>
        </schema>
    </wsdl:types>
    <wsdl:message name="addInput">
        <wsdl:part name="parameters" element="tns:add"/>
    </wsdl:message>
    <wsdl:message name="addOutput">
        <wsdl:part name="parameters" element="tns:addResponse"/>
    </wsdl:message>
    <wsdl:message name="squareInput">
        <wsdl:part name="parameters" element="tns:square"/>
    </wsdl:message>
    <wsdl:message name="squareOutput">
```

```
        <wsdl:part name="parameters" element="tns:squareResponse"/>
    </wsdl:message>
    <wsdl:message name="checkNumberInput">
        <wsdl:part name="parameters" element="tns:checkNumber"/>
    </wsdl:message>
    <wsdl:message name="checkNumberOutput">
        <wsdl:part name="parameters" element="tns:checkNumberResponse"/>
    </wsdl:message>
    <wsdl:message name="CalculatorWrapperException">
        <wsdl:part name="CalculatorWrapperException"
element="ns0:CalculatorWrapperException"/>
    </wsdl:message>
    <wsdl:message name="CalculatorException">
        <wsdl:part name="CalculatorException" element="ns0:CalculatorException"/>
    </wsdl:message>
    <wsdl:portType name="Calculator">
        <wsdl:operation name="add">
            <wsdl:input xmlns:ns1="http://www.w3.org/2006/05/addressing/wsdl"
message="tns:addInput" ns1:Action=""/>
            <wsdl:output xmlns:ns1="http://www.w3.org/2006/05/addressing/wsdl"
message="tns:addOutput" ns1:Action=""/>
        </wsdl:operation>
        <wsdl:operation name="square">
            <wsdl:input xmlns:ns1="http://www.w3.org/2006/05/addressing/wsdl"
message="tns:squareInput" ns1:Action=""/>
            <wsdl:output xmlns:ns1="http://www.w3.org/2006/05/addressing/wsdl"
message="tns:squareOutput" ns1:Action=""/>
        </wsdl:operation>
        <wsdl:operation name="checkNumber">
            <wsdl:input xmlns:ns1="http://www.w3.org/2006/05/addressing/wsdl"
message="tns:checkNumberInput" ns1:Action=""/>
            <wsdl:output xmlns:ns1="http://www.w3.org/2006/05/addressing/wsdl"
message="tns:checkNumberOutput" ns1:Action=""/>
            <wsdl:fault name="CalculatorWrapperException"
message="tns:CalculatorWrapperException"/>
            <wsdl:fault name="CalculatorException"
message="tns:CalculatorException"/>
        </wsdl:operation>
    </wsdl:portType>
    <wsdl:binding name="CalculatorSoapHttp" type="tns:Calculator">
        <soap:binding style="document"
transport="http://schemas.xmlsoap.org/soap/http"/>
        <wsdl:operation name="add">
            <soap:operation soapAction=""/>
            <wsdl:input>
                <soap:body use="literal"/>
            </wsdl:input>
            <wsdl:output>
                <soap:body use="literal"/>
            </wsdl:output>
        </wsdl:operation>
        <wsdl:operation name="square">
            <soap:operation soapAction=""/>
            <wsdl:input>
                <soap:body use="literal"/>
            </wsdl:input>
            <wsdl:output>
                <soap:body use="literal"/>
            </wsdl:output>
        </wsdl:operation>
```

```
        <wsdl:operation name="checkNumber">
            <soap:operation soapAction=""/>
            <wsdl:input>
                <soap:body use="literal"/>
            </wsdl:input>
            <wsdl:output>
                <soap:body use="literal"/>
            </wsdl:output>
            <wsdl:fault name="CalculatorWrapperException">
                <soap:fault name="CalculatorWrapperException" use="literal"
encodingStyle=""/>
            </wsdl:fault>
            <wsdl:fault name="CalculatorException">
                <soap:fault name="CalculatorException" use="literal"
encodingStyle=""/>
            </wsdl:fault>
        </wsdl:operation>
    </wsdl:binding>
    <wsdl:service name="CalculatorService">
        <wsdl:port name="CalculatorPort" binding="tns:CalculatorSoapHttp">
            <soap:address
location="http://localhost:8888/CalWSBA/CalculatorPort"/>
        </wsdl:port>
    </wsdl:service>
</wsdl:definitions>
```

Example 20–2 uses the WSDL file shown in Example 20–1, "Sample WSDL File CalculatorService.wsdl". First, the tool parses the WSDL for all port types and bindings (supported protocols such as HTTP get/post or SOAP) to list all the operations. If there are multiple port types in the WSDL, the tool first prompts you to select a port type.

To start the command-line tool:

1. Go to the $AGENT_HOME/bin directory.
2. Run the following command:

```
$ emctl wscli -metadata -wsdl=/tmp/CalculatorWS.wsdl
```

Once invoked, the command-line tool automatically prompts you for the requisite information, as shown in Example 20–2, "Sample Web Services Command-Line Tool Session". If you need to abort a command-line tool session, you can press Ctrl+C at any point to exit. Session information will not be saved.

Example 20–2 Sample Web Services Command-Line Tool Session

```
Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.

OracleHome : /oracle/oms/agent
EMDROOT    : /oracle/oms/agent

Generate Metric Metadata for Web Service Monitoring

Reading WSDL Document at /tmp/CalculatorWS.wsdl...done.

==> Enter the metadata file name [/tmp/target/metadata/CalculatorService.xml] :
```

```
* Selected Service: CalculatorService

* Selected Port: CalculatorPort

All operations for the selected Port "CalculatorPort":
[1]   squareResponse square(int arg0)
[2]   checkNumberResponse checkNumber(int arg0)
[3]   addResponse add(int arg0, int arg1)

==> Enter the index [1-3] of operation to select: 1
* Selected Operation:
    squareResponse square(int arg0)

Define new metric group:
==> Enter the name for this metric group [square]:

Return value(s) for the selected operation:
[1]   //ns0:squareResponse/arg0 <int>

==> Enter the index [1-1] of metric to display: 1
==> Enter the name for this metric [arg0]: SquareResult
==> Enter the label for this metric [SquareResult]:
==> Is this a key metric <y/n>? [n] :
==> Do you want to create threshold for this item <y/n>? [n] :


Setup operation Argument: square.arg0 <type:int>
==>   Enter value [%square.arg00001%] :

==> Do you want to use jps-config-jse.xml <y/n>? [n] :

==> Do you want to add User/Password Credential <y/n>? [n] : y
==> Enter the name for User/Password credential set [UserCredentialSet01] :

==> Do you want to add SSL TrustStore Credential <y/n>? [n] :

==> Do you want to add SSL KeyStore Credential <y/n>? [n] :

==> Do you want to add KeyStore Credential <y/n>? [n] :

==> Do you want to add Encryption Key Credential <y/n>? [n] :

==> Do you want to add Signature Key Credential <y/n>? [n] :

==> Is this metric group for periodic collection <y/n>? [y] :
The following units are for collection frequency:
[1]   Min
[2]   Hr
[3]   Day

==> Enter the index [1-3] of unit for this collection: 1
==> Enter the frequency of collection in Min: 30

==> Do you want to add another metric group <y/n>? [n] :

Files Generated:
- Target Metadata file: /tmp/target/metadata/CalculatorService.xml
- Target Collection file: /tmp/target/metadata/CalculatorServiceCollection.xml
```

The command-line tool generates the metadata required to monitor the target type `CalculatorService` as shown in Example 20–3.

Example 20–3 CalculatorService Target Metadata File

```
<!DOCTYPE TargetMetadata SYSTEM "../dtds/TargetMetadata.dtd">
<TargetMetadata META_VER="1.0" TYPE="CalculatorService">
  <Display>
    <Label NLSID="NLSID_CALCULATOR_SERVICE">CalculatorService</Label>
    <ShortName NLSID="NLSID_CALCULATOR_SERVICE">CalculatorService</ShortName>
    <Description NLSID="NLSID_CALCULATOR_SERVICE">CalculatorService</Description>
  </Display>
  <Metric NAME="square" TYPE="TABLE">
    <Display>
      <Label NLSID="NLSID_SQUARE">square</Label>
    </Display>
    <TableDescriptor>
      <ColumnDescriptor IS_KEY="FALSE" NAME="SquareResult" TYPE="STRING">
        <Display>
          <Label NLSID="COL_SQUARE_RESULT">SquareResult</Label>
        </Display>
      </ColumnDescriptor>
    </TableDescriptor>
    <QueryDescriptor FETCHLET_ID="WSF">
      <Property NAME="ProxyHost" SCOPE="INSTANCE"
OPTIONAL="TRUE">ProxyHost</Property>
      <Property NAME="ProxyPort" SCOPE="INSTANCE"
OPTIONAL="TRUE">ProxyPort</Property>
      <Property NAME="SecurityPolicy" SCOPE="INSTANCE"
OPTIONAL="FALSE">square.SecurityPolicy</Property>
      <Property NAME="ServiceEndpoint" SCOPE="INSTANCE"
OPTIONAL="FALSE">square.ServiceEndpoint</Property>
      <Property NAME="ServiceName" SCOPE="GLOBAL"
OPTIONAL="FALSE">ns0:CalculatorService</Property>
      <Property NAME="PortName" SCOPE="GLOBAL"
OPTIONAL="FALSE">ns0:CalculatorPort</Property>
      <Property NAME="OperationName" SCOPE="GLOBAL"
OPTIONAL="FALSE">square</Property>
      <Property NAME="MessageType" SCOPE="GLOBAL" OPTIONAL="FALSE">SOAP</Property>
      <Property NAME="SOAPBindingStyle" SCOPE="GLOBAL"
OPTIONAL="FALSE">DOCUMENT</Property>
      <Property NAME="SOAPBindingUse" SCOPE="GLOBAL"
OPTIONAL="FALSE">LITERAL</Property>
      <Property NAME="ParameterStyle" SCOPE="GLOBAL"
OPTIONAL="FALSE">WRAPPED</Property>
      <Property NAME="SOAPVersion" SCOPE="GLOBAL" OPTIONAL="FALSE">SOAP_1_
1</Property>
      <Property NAME="Namespace" SCOPE="GLOBAL"
OPTIONAL="FALSE"><![CDATA[[ns0="http://tests.jaxws.oracle.com/"]]]></Property>
      <Property NAME="RowType" SCOPE="GLOBAL"
OPTIONAL="FALSE">//ns0:squareResponse/arg0</Property>
      <Property NAME="ColType" SCOPE="GLOBAL"
OPTIONAL="FALSE">SquareResult:STRING</Property>
      <Property NAME="Payload" SCOPE="GLOBAL"
OPTIONAL="FALSE"><![CDATA[<soap:Envelope
xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
            <soap:Body xmlns:ns1="http://tests.jaxws.oracle.com/">
              <ns1:square>
                <arg0>%square.arg00001%</arg0>
```

```
              </ns1:square>
            </soap:Body>
          </soap:Envelope>]]></Property>
      <Property NAME="UserCredential" SCOPE="GLOBAL"
OPTIONAL="FALSE">UserCredentialSet01</Property>
      <CredentialRef
NAME="UserCredentialSet01">UserCredentialSet01</CredentialRef>
    </QueryDescriptor>
  </Metric>
  <CredentialInfo>
    <CredentialType NAME="CSFKeyCredential">
      <Display>
        <Label NLSID="CRED_TYPE">CSF-Key Credential Type</Label>
      </Display>
      <CredentialTypeColumn NAME="CSFKey">
        <Display>
          <Label NLSID="CRED_C_S_F_KEY">Alias CSF Key</Label>
        </Display>
      </CredentialTypeColumn>
    </CredentialType>
    <CredentialType NAME="AliasCredential">
      <Display>
        <Label NLSID="CRED_TYPE">Alias Credential Type</Label>
      </Display>
      <CredentialTypeColumn NAME="Alias">
        <Display>
          <Label NLSID="CRED_ALIAS">Alias (i.e. username, encryption key,
signature key, etc)</Label>
        </Display>
      </CredentialTypeColumn>
      <CredentialTypeColumn NAME="Password">
        <Display>
          <Label NLSID="CRED_PASSWORD">Password for the alias</Label>
        </Display>
      </CredentialTypeColumn>
    </CredentialType>
    <CredentialSet NAME="UserCredentialSet01" USAGE="MONITORING">
      <AllowedCredType TYPE="CSFKeyCredential"/>
      <AllowedCredType TYPE="AliasCredential"/>
    </CredentialSet>
  </CredentialInfo>
  <InstanceProperties>
    <InstanceProperty NAME="ProxyHost" CREDENTIAL="FALSE" OPTIONAL="TRUE">
      <Display>
        <Label NLSID="PROP_PROXY_HOST">Proxy Server Name</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="ProxyPort" CREDENTIAL="FALSE" OPTIONAL="TRUE">
      <Display>
        <Label NLSID="PROP_PROXY_PORT">Proxy Server Port</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="square.SecurityPolicy" CREDENTIAL="FALSE"
OPTIONAL="FALSE">
      <Display>
        <Label NLSID="PROP_SQUARE_SECURITY_POLICY">[square] Authentication/Web
Service Policy</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="square.ServiceEndpoint" CREDENTIAL="FALSE"
```

```
OPTIONAL="FALSE">
      <Display>
        <Label NLSID="PROP_SQUARE_SERVICE_ENDPOINT">[square] Web Service Endpoint
URL</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="square.arg00001" CREDENTIAL="FALSE" OPTIONAL="FALSE">
      <Display>
        <Label NLSID="PROP_SQUARE_ARG00001">[square] square.arg0</Label>
      </Display>
    </InstanceProperty>
  </InstanceProperties>
</TargetMetadata>
```

The command-line tool also generates the requisite collection file as shown in Example 20–4.

Example 20–4 CalculatorService Default Collection File

```
<!DOCTYPE TargetCollection SYSTEM "../dtds/TargetCollection.dtd">
<TargetCollection TYPE="CalculatorService">
  <CollectionItem NAME="square">
    <Schedule>
      <IntervalSchedule TIME_UNIT="Min" INTERVAL="30"/>
    </Schedule>
  </CollectionItem>
</TargetCollection>
```

After the tool generates the target metadata and collection files, you can create the Management Plug-in archive. See Section 20.7, "Creating a Management Plug-in Archive".

20.3 Monitoring Using WS-Management in Enterprise Manager

Beginning with Enterprise Manager 12c, WS-Management (WS-MAN)-compliant resources can be monitored using the fetchlet WSManagementFetchlet..

The fetchlet communicates with the WS-MAN resources using WS-Transfer protocol, which defines a number of management operations that the managed resources should support. However, in the current release, the fetchlet only supports the operation WS-Transfer GET.

Note: For more information about the monitor WS-Management standard, see the DMTF Web Services Management website:

`http://www.dmtf.org/standards/wsman`

Prerequisites

- Enterprise Manager Cloud Control Management Agent version 12.1.0.0.0 or greater installed on that host.
- Enterprise Manager Cloud Control Management Server (OMS) version 12.1.0.0.0 or greater with which the Management Agent communicates.

20.3.1 Creating Metadata and Default Collection Files

Enterprise Manager provides an easy-to-use WS-Management CLI command-line tool that simplifies creating new Management Plug-ins by automatically generating the requisite target metadata and default collection files. Information retrieval is achieved via the WSManagementFetchlet that is integrated with the Management Agent.

Resources, which support WS-Management interface, should describe their model-specific elements using XML Schema Definition (XSD) representation and expose the XSD as a public accessible link just like WSDL for Web Services.

The command-line tool works by parsing a specified XSD file for the managed WS-MAN resource and then prompts you to select the interested resource properties to construct a monitoring metric.

20.3.1.1 WS-Management CLI Command-line Tool Syntax

The WS-Management CLI command-line tool syntax is as follows:

```
Usage: emctl wsmancli [-metadata | -help] [-options]
```

The command accepts the following options:

- *-schema=<file | URL>*: Resource XSD file or URL [mandatory]
- *-username=<user ID>* : Username if the schema is protected

The command-line tool requires a XSD file name or URL to locate the resource schema. For example, for a Traffic Light WS-Management service, a XSD URL would be as follows:

```
http://localhost:8888/TrafficLight?xsd
```

The command tool script requires access to the Enterprise Manager home directory (EM_HOME) to run. The tool defaults to ORACLE_HOME (ensure this environment variable is set properly before using this tool).

20.3.1.2 Command-line Tool Security

The command-line tool generates metadata required by Enterprise Manager for target monitoring purposes via the resource XSD. When you run this tool, you only need read permission on the XSD file or URL and permission to save generated files to the appropriate directory.

20.3.1.3 Generating Target Metadata and Collection Files

The following example shows a sample XSD file passed to the command-line tool to generate the target metadata and collection files.

Example 20–5 Sample XSD File TrafficLight.xsd

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema
targetNamespace="http://schemas.wiseman.dev.java.net/traffic/1/light.xsd"
elementFormDefault="qualified" blockDefault="#all"
xmlns:tl="http://schemas.wiseman.dev.java.net/traffic/1/light.xsd"
xmlns:xs="http://www.w3.org/2001/XMLSchema">
         <xs:complexType name="TrafficLightType">
               <xs:sequence>
                     <xs:element name="name" type="xs:string"/>
                     <xs:element name="color" type="xs:string"/>
                     <xs:element name="x" type="xs:int"/>
```

```
                    <xs:element name="y" type="xs:int"/>
                </xs:sequence>
            </xs:complexType>
            <xs:element name="trafficlight" type="tl:TrafficLightType"/>
</xs:schema>
```

To start the command-line tool:

1. Go to the $AGENT_HOME/bin directory.
2. Execute the following command:

```
$ emctl wsmancli -metadata -schema= http://localhost:8080/Traffic?xsd
```

Once invoked, the command-line tool automatically prompts you for the requisite information, as shown in Example 20–6, "Sample WS-Management CLI Command-Line Tool Session". If you need to abort a command-line tool session, you can press <Ctrl+C> at any point to exit. Session information will not be saved.

Example 20–6 Sample WS-Management CLI Command-Line Tool Session

```
Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.0.0
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.

OracleHome : /oracle/oms/agent
EMDROOT    : /oracle/oms/agent

Generate Metric Metadata for WS-Management Resource Monitoring

Reading Resource XSD Document at http://localhost:8080/Traffic?xsd...done.

==> Enter the name for this target type: TrafficLight

==> Enter the metadata file name [/tmp/target/metadata/TrafficLight.xml] :

Define new metric group name:
==> Enter the name for this metric group: trafficLight

WS-Addressing namespaces:
[1]   http://www.w3.org/2005/08/addressing
[2]   http://schemas.xmlsoap.org/ws/2004/08/addressing

==> Enter the index [1-2] to select: 1


SOAP Envelope namespaces:
[1]   http://www.w3.org/2003/05/soap-envelope
[2]   http://schemas.xmlsoap.org/soap/envelope/

==> Enter the index [1-2] to select: 1

Resource properties:
[1]   trafficlight:color
[2]   trafficlight:name
[3]   trafficlight:x
[4]   trafficlight:y

==> Enter the index [1-4] of property to display: 2
==> Enter the name for this metric [name]:
==> Enter the label for this metric [name]:
==> Is this a key metric <y/n>? [n] : y
```

```
==> Do you want to add another metric <y/n>? [n] : y

Resource properties:
[1]   trafficlight:color
[2]   trafficlight:x
[3]   trafficlight:y

==> Enter the index [1-3] of property to display: 1
==> Enter the name for this metric [color]:
==> Enter the label for this metric [color]:
==> Is this a key metric <y/n>? [n] :
==> Do you want to create threshold for this item <y/n>? [n] :
==> Do you want to add another metric <y/n>? [n] : y

Resource properties:
[1]   trafficlight:x
[2]   trafficlight:y

==> Enter the index [1-2] of property to display: 1
==> Enter the name for this metric [x]:
==> Enter the label for this metric [x]:
==> Is this a key metric <y/n>? [n] :
==> Do you want to create threshold for this item <y/n>? [n] :
==> Do you want to add another metric <y/n>? [n] : y

Resource properties:
[1]   trafficlight:y

==> Enter the index [1-1] of property to display: 1
==> Enter the name for this metric [y]:
==> Enter the label for this metric [y]:
==> Is this a key metric <y/n>? [n] :
==> Do you want to create threshold for this item <y/n>? [n] :


==> Enter comma-separated list of Selector elements: name

==> Do you want to add User/Password Credential <y/n>? [n] : y
==> Enter the name for User/Password credential set [UserCredentialSet01] :

==> Is this metric group for periodic collection <y/n>? [y] :
The following units are for collection frequency:
[1]   Min
[2]   Hr
[3]   Day

==> Enter the index [1-3] of unit for this collection: 1
==> Enter the frequency of collection in Min: 30

==> Do you want to add another metric group <y/n>? [n] :

Files Generated:
- Target Metadata file: /tmp/target/metadata/TrafficLight.xml
- Target Collection file: /tmp/target/metadata/TrafficLightCollection.xml
```

The command-line tool generates the metadata required to monitor the target type TrafficLight as shown in Example 6-7.

Example 20–7 TrafficLight Target Metadata File

```
<!DOCTYPE TargetMetadata SYSTEM "../dtds/TargetMetadata.dtd">
<TargetMetadata META_VER="1.0" TYPE="TrafficLight">
  <Display>
    <Label NLSID="NLSID_TRAFFIC_LIGHT">TrafficLight</Label>
    <ShortName NLSID="NLSID_TRAFFIC_LIGHT">TrafficLight</ShortName>
    <Description NLSID="NLSID_TRAFFIC_LIGHT">TrafficLight</Description>
  </Display>
  <Metric NAME="trafficLight" TYPE="TABLE">
    <Display>
      <Label NLSID="NLSID_TRAFFIC_LIGHT">trafficLight</Label>
    </Display>
    <TableDescriptor>
      <ColumnDescriptor IS_KEY="TRUE" NAME="name" TYPE="STRING">
        <Display>
          <Label NLSID="COL_NAME">name</Label>
        </Display>
      </ColumnDescriptor>
      <ColumnDescriptor IS_KEY="FALSE" NAME="color" TYPE="STRING">
        <Display>
          <Label NLSID="COL_COLOR">color</Label>
        </Display>
      </ColumnDescriptor>
      <ColumnDescriptor IS_KEY="FALSE" NAME="x" TYPE="STRING">
        <Display>
          <Label NLSID="COL_X">x</Label>
        </Display>
      </ColumnDescriptor>
      <ColumnDescriptor IS_KEY="FALSE" NAME="y" TYPE="STRING">
        <Display>
          <Label NLSID="COL_Y">y</Label>
        </Display>
      </ColumnDescriptor>
    </TableDescriptor>
    <QueryDescriptor FETCHLET_ID="WSManagementFetchlet">
      <Property NAME="ProxyHost" SCOPE="INSTANCE"
OPTIONAL="TRUE">ProxyHost</Property>
      <Property NAME="ProxyPort" SCOPE="INSTANCE"
OPTIONAL="TRUE">ProxyPort</Property>
      <Property NAME="SecurityPolicy" SCOPE="INSTANCE"
OPTIONAL="TRUE">trafficLight.SecurityPolicy</Property>
      <Property NAME="ResourceURL" SCOPE="INSTANCE"
OPTIONAL="FALSE">trafficLight.ResourceURL</Property>
      <Property NAME="To" SCOPE="INSTANCE"
OPTIONAL="FALSE">trafficLight.To</Property>
      <Property NAME="OptionSet" SCOPE="INSTANCE"
OPTIONAL="TRUE">trafficLight.OptionSet</Property>
      <Property NAME="Locale" SCOPE="INSTANCE"
OPTIONAL="TRUE">trafficLight.Locale</Property>
      <Property NAME="MaxEnvelopeSize" SCOPE="INSTANCE"
OPTIONAL="TRUE">trafficLight.MaxEnvelopeSize</Property>
      <Property NAME="OperationTimeout" SCOPE="INSTANCE"
OPTIONAL="TRUE">trafficLight.OperationTimeout</Property>
      <Property NAME="Namespace" SCOPE="GLOBAL"
OPTIONAL="FALSE"><![CDATA[[ns1="http://schemas.wiseman.dev.java.net/traffic/1/ligh
t.xsd"][ns0="http://www.w3.org/2001/XMLSchema"][wsa="http://www.w3.org/2005/08/add
ressing"][env="http://www.w3.org/2003/05/soap-envelope"]]]></Property>
      <Property NAME="RowType" SCOPE="GLOBAL" OPTIONAL="FALSE">
//ns1:trafficlight/ns1:name,//ns1:trafficlight/ns1:color,//ns1:trafficlight/ns1:x,
//ns1:trafficlight/ns1:y</Property>
```

```
      <Property NAME="ColType" SCOPE="GLOBAL"
OPTIONAL="FALSE">name:STRING,color:STRING,x:STRING,y:STRING</Property>
      <Property NAME="ReplyTo" SCOPE="GLOBAL"
OPTIONAL="FALSE">http://www.w3.org/2005/08/addressing/role/anonymous</Property>
      <Property NAME="Action" SCOPE="GLOBAL"
OPTIONAL="FALSE">http://schemas.xmlsoap.org/ws/2004/09/transfer/Get</Property>
      <Property NAME="TransferOperation" SCOPE="GLOBAL"
OPTIONAL="FALSE">GET</Property>
      <Property NAME="SelectorSet" SCOPE="GLOBAL"
OPTIONAL="FALSE">[name,%trafficLight.name%]</Property>
      <Property NAME="UserCredential" SCOPE="GLOBAL"
OPTIONAL="FALSE">UserCredentialSet01</Property>
      <CredentialRef
NAME="UserCredentialSet01">UserCredentialSet01</CredentialRef>
    </QueryDescriptor>
  </Metric>
  <CredentialInfo>
    <CredentialType NAME="CSFKeyCredential">
      <Display>
        <Label NLSID="CRED_TYPE">CSF-Key Credential Type</Label>
      </Display>
      <CredentialTypeColumn NAME="CSFKey">
        <Display>
          <Label NLSID="CRED_C_S_F_KEY">Alias CSF Key</Label>
        </Display>
      </CredentialTypeColumn>
    </CredentialType>
    <CredentialType NAME="AliasCredential">
      <Display>
        <Label NLSID="CRED_TYPE">Alias Credential Type</Label>
      </Display>
      <CredentialTypeColumn NAME="Alias">
        <Display>
          <Label NLSID="CRED_ALIAS">Alias (i.e. username, encryption key,
signature key, etc)</Label>
        </Display>
      </CredentialTypeColumn>
      <CredentialTypeColumn NAME="Password">
        <Display>
          <Label NLSID="CRED_PASSWORD">Password for the alias</Label>
        </Display>
      </CredentialTypeColumn>
    </CredentialType>
    <CredentialSet NAME="UserCredentialSet01" USAGE="MONITORING">
      <AllowedCredType TYPE="CSFKeyCredential"/>
      <AllowedCredType TYPE="AliasCredential"/>
    </CredentialSet>
  </CredentialInfo>
  <InstanceProperties>
    <InstanceProperty NAME="ProxyHost" CREDENTIAL="FALSE" OPTIONAL="TRUE">
      <Display>
        <Label NLSID="PROP_PROXY_HOST">Proxy Server Name</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="ProxyPort" CREDENTIAL="FALSE" OPTIONAL="TRUE">
      <Display>
        <Label NLSID="PROP_PROXY_PORT">Proxy Server Port</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="trafficLight.SecurityPolicy"
```

```
      CREDENTIAL="FALSE" OPTIONAL="TRUE">
      <Display>
        <Label NLSID="PROP_TRAFFIC_LIGHT_SECURITY_POLICY">[trafficLight]
Authentication/Web Service Policy</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="trafficLight.ResourceURL" CREDENTIAL="FALSE"
OPTIONAL="FALSE">
      <Display>
        <Label NLSID="PROP_TRAFFIC_LIGHT_RESOURCE_U_R_L">[trafficLight] Resource
URL (wsman:ResourceURL)</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="trafficLight.To" CREDENTIAL="FALSE" OPTIONAL="FALSE">
      <Display>
        <Label NLSID="PROP_TRAFFIC_LIGHT_TO">[trafficLight] Network Address of the
service (wsa:To)</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="trafficLight.OptionSet" CREDENTIAL="FALSE"
OPTIONAL="TRUE">
      <Display>
        <Label NLSID="PROP_TRAFFIC_LIGHT_OPTION_SET">[trafficLight] Set of
wsman:Option. Format: [&lt;OptionName1&gt;, value:&lt;value1&gt;,
type:&lt;type1&gt;, mustComply:&lt;true|false&gt;][&lt;OptionName2&gt;,
value:&lt;value2&gt;, type:&lt;type&gt;,
mustComply:&lt;true|false&gt;][...]</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="trafficLight.Locale" CREDENTIAL="FALSE"
OPTIONAL="TRUE">
      <Display>
        <Label NLSID="PROP_TRAFFIC_LIGHT_LOCALE">[trafficLight] wsman:Locale (RFC
3066 language code). Format: e.g. en-US</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="trafficLight.MaxEnvelopeSize"
      CREDENTIAL="FALSE" OPTIONAL="TRUE">
      <Display>
        <Label NLSID="PROP_TRAFFIC_LIGHT_MAX_ENVELOPE_SIZE">[trafficLight]
wsman:MaxEnvelopeSize in Octets. Format: e.g. 8192</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="trafficLight.OperationTimeout"
      CREDENTIAL="FALSE" OPTIONAL="TRUE">
      <Display>
        <Label NLSID="PROP_TRAFFIC_LIGHT_OPERATION_TIMEOUT">[trafficLight]
wsman:OperationTimeout. Format: e.g. PT30S</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="trafficLight.name" CREDENTIAL="FALSE"
OPTIONAL="FALSE">
      <Display>
        <Label NLSID="PROP_TRAFFIC_LIGHT_NAME">[trafficLight] Value for the
Selector "name"</Label>
      </Display>
    </InstanceProperty>
  </InstanceProperties>
</TargetMetadata>
```

The command-line tool also generates the requisite collection file as shown in Example 20–8, "TrafficLight Default Collection File".

Example 20–8 TrafficLight Default Collection File

```
<!DOCTYPE TargetCollection SYSTEM "../dtds/TargetCollection.dtd">
<TargetCollection TYPE="TrafficLight">
  <CollectionItem NAME="trafficLight">
    <Schedule>
      <IntervalSchedule TIME_UNIT="Min" INTERVAL="30"/>
    </Schedule>
  </CollectionItem>
</TargetCollection>
```

After the command-line tool generates the target metadata and collection files, you can create the Management Plug-in archive. See Creating a Management Plug-in Archive on page 20-46.

20.4 Monitoring JMX Applications Deployed on Oracle Application Servers (OC4J)

The Java Management Extensions (JMX) framework improves manageability of your JMX-instrumented applications by enabling you to see what is happening inside. You gain insight into your applications and infrastructure through modular plug-ins called Managed Beans (MBeans). MBeans integrate with your application, components (such as Enterprise Java-Beans), or other resources to expose attributes (values) and operations.

The OJMX fetchlet, supplied with 10.2 Management Agents, enables you to monitor key metrics in your JMX-instrumented applications deployed on Oracle Application Server 10.1.3 or above. The fetchlet extends monitoring capabilities via JMX to the J2EE 1.4-compliant Oracle containers for J2EE (OC4J) servers themselves.

Monitoring JMX-instrumented applications/servers with Enterprise Manager entails defining a new target type that Enterprise Manager can monitor via Management Plug-ins. As with the Web Services `wscli` command-line tool, Enterprise Manager provides a `jmxcli` command-line tool to automate the generation of the target metadata and collection files.

Prerequisites

- Oracle Application Server 10.1.3 instance running on a specific host with a JMX-enabled application deployed on it that needs to be monitored as a target in Enterprise Manager.
- Enterprise Manager Management Agent version 10.2.0.2 or greater installed on that host.
- Enterprise Manager Management Server (OMS) version 10.2.0.2 or greater with which the Management Agent communicates.

Known Limitations

Currently, the jmxcli tool and OJMX fetchlet only allow you to browse and monitor MBeans (system and application-defined) that are available on the default MBeanserver on the target OC4J instance. The jmxcli tool primarily handles attributes and parameter and return values for operations that are OpenTypes. Examples: SimpleTypes, CompositeTypes, TabularTypes, and arrays of SimpleTypes.

20.4.1 Creating Metadata and Default Collection Files

As with Web Services, the JMX command-line tool (jmxcli) simplifies creating the requisite target definition files: metadata and the default collection file. The tool is an offline configuration utility that connects you to an MBeanServer and enables you to browse available MBeans. It can also append metrics to an existing set of files during a subsequent invocation of the tool.

During a command-line tool session, you select specific MBeans and then choose the desired attributes/statistical values or operations Enterprise Manager needs to retrieve or invoke periodically on these MBeans to collect these values. The tool helps define packaging for these collected values as one or more Enterprise Manager metrics (with columns), and also enables you to specify a metric collection interval.

20.4.1.1 JMX Command-line Tool Syntax

The JMX command-line tool syntax is as follows for a JMX-enabled target on an OC4J: Note that usage has changed from prior releases. The `cli` is now integrated with the `emctl` utility on the Agent.

```
cd <Agent Instance Home>/bin
emctl jmxcli <TARGET_HOME>
               [  -h <hostname>
                  -p <port>
                  -u <username>
                  -c <credential/password>
                  -w <work directory>
                  -e <true/false>
                  [-m <MBeanName> | -d <jmx_domain> | -s <mBeanPattern>]
               ]
```

`<TARGET_HOME>` is an Oracle Home directory 10.1.3 or greater Oracle Application Server Container for J2EE (OC4J).

The `jmxcli` command accepts the following options:

- **-h** Hostname of the OC4J. Default: "localhost"
- **-p** RMI/RMIS port of the OC4J. Default: "23791"
  From the `ORACLE_HOME/opmn/bin` directory of your Application Server 10.1.3.0 or later instance, run `opmnctl status -l` to determine the RMI port for the OC4J for which MBeans were deployed.
- **-u** Valid username for the OC4J. Default: "oc4jadmin"
- **-c** Password associated with the OC4J user specified by the `-u` option. Default: None. If you do not specify a password, you are prompted for the password.
- **-w** Directory where the metadata and default collection files created by the JMX command-line tool are placed. Default: Current directory. When invoking the command-line tool, you must have write permission on this directory to create subdirectories and add files. If the metadata and default collection files already exist within that directory, you have the option of appending to or overwriting the original files.
- **-e** Whether or not the RMIS connection to the OC4J is enabled (true or false). Default: false

You can also specify ONE of the following three parameters (`-m`, `-d` or `-s`) to retrieve a subset of MBeans available on the MBeanServer. By default, all MBeans on the MBeanServer are displayed for you to select from if none of these parameters are specified.

- **-m** MBean ObjectName of the required MBean that needs to be retrieved and examined. If this is an ObjectName pattern-matching multiple MBeans, you are shown a list of all MBeans that match the pattern, and you can select one at a time to work on.
- **-d** MBean domain of the required application whose MBeans need to be retrieved and examined. For example, you want to browse all MBeans for an application (myApp). MBeans for this application would be available in the JMX domain "myApp".
- **-s** MBean pattern-matching an existing set of MBeans from which the metrics are to be defined. The `-s` parameter allows bulk retrieval of JMX Attributes/Statistics from multiple MBeans of a similar type.

  If you specify the `-s` parameter, the resulting metrics created during this `jmxcli` session appear as a table in the Enterprise Manager console with multiple rows — one row representing each MBean that matches the specified pattern, and with the MBean ObjectName as a key column. For example, if you specify `-s 'oc4j:j2eeType=Servlet,*'` the resulting metric will have multiple rows, one for each servlet that matches the ObjectName pattern. Besides the MBean ObjectName column, other columns would be the attributes or fields from the return object of the operation, selected during the `jmxcli` session.

20.4.1.2 Generating the Files

To start the JMX command-line tool:

1. Go to the `$AGENT_HOME/bin` directory.
2. Run the following command:

   ```
   emctl jmxcli <Oracle Home of the target 10.1.3 or greater OC4J> [OPTIONS]
   ```

   Once invoked, the command-line interface automatically prompts you for the requisite information, as shown in Example 20–10. If you need to abort a JMX command-line tool session, you can press Ctrl+C at any point to exit. Session information will not be saved.

Example 20–9 Sample JMXCLI Invocation

```
./emctl jmxcli /scratch/shiphomes/oc4j/1013_SOA_M1/ -h localhost -p 12404 -m
'oc4j:J2EEApplication=orabpel,name=\"ServerBean\",*'
```

Example 20–10 Sample JMXCLI Session

```
oracleHome=/ade/sparmesw_10202_ssm/oracle
targetHome=/scratch/shiphomes/oc4j/1013_SOA_M1/
The Port is 12404

Connecting to server: localhost:12404
Connecting as user: oc4jadmin
Enter the password:

Obtained 1 MBeans matching pattern
oc4j:J2EEApplication=orabpel,name="ServerBean",*.

Enter the target type for this metric: [myJ2EEApp] myBPELApp
```

This is the target type for the new J2EE application as it should show up within Enterprise Manager.

```
Enter the target version: [1.0]

Enter the target metadata file: [./metadata/myBPELApp.xml]
```

This is the location of the metadata file that jmxcli generates. You must have write permission on the directories where the target metadata and default collection files are to be created.

```
Enter the default collections file: [./default_collection/myBPELApp.xml]
The file ./metadata/myBPELApp.xml already exists.
Do you want to overwrite the existing file, append to it, or quit <o/a/q>? [a] a
Appending to existing file: ./metadata/myBPELApp.xml.
The available targets are:
0: Identifies a deployed stateless session bean
        (oc4j:EJBModule="ejb_ob_
engine",J2EEApplication=orabpel,J2EEServer=standalone,j2eeType=StatelessSessionBea
n,name="ServerBean")

Enter the index of target/MBean you wish to monitor or press <Ctrl-C> to quit: 0
```

If multiple MBeans matched the -m <MBean pattern> specified when jmxcli was invoked, all MBean ObjectNames matching the pattern are listed during this part of the command-line session, at which point you can select one among the list. You can choose another MBean from the above list after creating metrics for the first one without exiting this jmxcli session.

If you want to append metrics from another MBean that does not match the above -m pattern, you must exit and start another jmxcli session with the MBean ObjectName/Pattern of the latter MBean, and create metrics from this MBean which will be appended to the original target metadata and default collection files from the previous jmxcli session. Using this method, you can append metrics created from multiple jmxcli sessions to the same target metadata and default collection files, if necessary.

```
Following metric source types are available for selected target(s):
        0: JMX Attributes
        1: JMX Operations
        2: J2EE Statistics
Enter the index of your choice or press <Ctrl-C> to quit: 2

Statistics are:
        0: CreateCount
        1: ejbCreate()ClientActive
        2: ejbCreate()ClientTime
        3: ejbRemove()ClientActive
        4: ejbRemove()ClientTime
        5: MethodReadyCount
        6: RemoveCount
        7: setSessionContext(javax.ejb.SessionContext)ClientActive
        8: setSessionContext(javax.ejb.SessionContext)ClientTime
Select one or more items as comma-separated indices: 0,6

JavaBean is : CreateCount
        0: count
        1: description
        2: lastSampleTime
        3: name
        4: startTime
        5: unit
```

This indicates that the Statistic call CreateCount is not a simple data type, but has a JavaBean pattern with the above listed properties, of which some may interest you.

```
Select one or more items as comma-separated indices: 0
JavaBean is : RemoveCount
        0: count
        1: description
        2: lastSampleTime
        3: name
        4: startTime
        5: unit
Select one or more items as comma-separated indices: 0
Number of possible columns in the resultant metric are 2.

Enter the name for this metric column at index=0 : [countOfCreateCount]
createCount
```

You can specify any meaningful name here. If you do not specify a name, the JMX command-line tool generates a default name that may not be appropriate in all cases.

```
Is this column a KEY Column <y/n>? [n]
```

In situations where multiple rows can be returned, as might be the case when the Attribute or return value of the Operation is TabularData, you need to specify one or more of your chosen metrics as "Key" columns.

```
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [createCount]
Enter the NLSID for column: [createCount]
Enter the UNIT for column "createCount": [count]
Do you want to create a threshold for this column <y/n>? [n] y
Creating threshold!!
Following operators are available for creating thresholds:
 0: GT
 1: EQ
 2: LT
 3: LE
 4: GE
 5: CONTAINS
 6: NE
 7: MATCH
```

If you want to create a threshold on this column, you can specify an operator and then a value that would trigger a CRITICAL or WARNING alert.

```
Enter the index of your choice or press <Ctrl-C> to quit: 0
Enter the CRITICAL threshold: [NotDefined] 100
Enter the WARNING threshold: [NotDefined] 85
Enter the number of occurrences that trigger threshold: [6] 3
```

This is the number of consecutive occurrences of above CRITICAL or WARNING values that would trigger an alert.

```
Enter the message to be used when threshold is triggered: [createCount is %value%
and has crossed warning (%warning_threshold%) or critical (%critical_threshold%)
threshold.]
Enter NLSID for the message used when threshold is triggered: [createCount_cond]
Enter the name for this metric column at index=1 : [countOfRemoveCount]
removeCount
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [removeCount]
```

```
Enter the NLSID for column: [removeCount]
Enter the UNIT for column "removeCount": [count]
Do you want to create a threshold for this column <y/n>? [n]
Enter the name of this metric: ServerBeanStats
Enter the label for this metric: [ServerBeanStats]

Do you want periodic collection for this metric <y/n>? [n] y
```

If the metric does not have to be collected periodically, as would be the case with real-time-only metrics, you can specify "no".

```
Enter the collection interval in seconds: 60
Periodic collection interval is: 60 seconds.

Do you want to create another metric <y/n>? [n] y
The available targets are:
0: Identifies a deployed stateless session bean
        (oc4j:EJBModule="ejb_ob_
engine",J2EEApplication=orabpel,J2EEServer=standalone,j2eeType=StatelessSessionBea
n,name="ServerBean")

Enter the index of target/MBean you wish to monitor or press <Ctrl-C> to quit: 0
```

If multiple MBeans match the MBean pattern for the -m option (specified when jmxcli was invoked) you can select a different MBean from the above list for the next iteration of this command-line session.

```
Following metric source types are available for selected target(s):
        0: JMX Attributes
        1: JMX Operations
        2: J2EE Statistics

Enter the index of your choice or press <Ctrl-C> to quit: 0

Attributes are:
        0: activeInstances       Return Value: int
        1: activeInstancesHighWaterMark  Return Value: int
        2: eventProvider         Return Value: boolean
        3: maxInstances  Return Value: int
        4: minInstances  Return Value: int
        5: ObjectName    Return Value: javax.management.ObjectName
        6: stateManageable       Return Value: boolean
        7: statisticsProvider    Return Value: boolean
        8: stats         Return Value: javax.management.j2ee.statistics.Stats
        9: transactionTimeout    Return Value: int
Select one or more items as comma-separated indices: 0,3,4

Number of possible columns in the resultant metric are 3.
Enter the name for this metric column at index=0 : [activeInstances]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [activeInstances]
Enter the NLSID for column: [activeInstances]
Enter the UNIT for column "activeInstances": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]
Enter the name for this metric column at index=1 : [maxInstances]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [maxInstances]
Enter the NLSID for column: [maxInstances]
Enter the UNIT for column "maxInstances": [millisec, kb etc.. ]
```

```
Do you want to create a threshold for this column <y/n>? [n]

Enter the name for this metric column at index=2 : [minInstances]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [minInstances]
Enter the NLSID for column: [minInstances]
Enter the UNIT for column "minInstances": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]

Enter the name of this metric: ServerBeanCount
Enter the label for this metric: [ServerBeanCount]

Do you want periodic collection for this metric <y/n>? [n] y
Enter the collection interval in seconds: 300
Periodic collection interval is: 300 seconds.

Do you want to create another metric <y/n>? [n] n
Written the metadata xml file: ./metadata/myBPELApp.xml.
Updated the default collection file for myBPELApp at location ./default_
collection/myBPELApp.xml.
Exiting...
```

After the JMX command-line tool generates the target metadata and collection files, you can create the Management Plug-in archive. See Section 20.7, "Creating a Management Plug-in Archive". A sample of each generated file from the command-line tool session above is shown in Example 20–11 and Example 20–12.

Example 20–11 Generated Target Metadata File

```
<!DOCTYPE TargetMetadata SYSTEM "../dtds/TargetMetadata.dtd">
<TargetMetadata META_VER="1.0" TYPE="myBPELApp" CATEGORY
_PROPERTIES="VersionCategory">
   <Display>
      <Label NLSID="myBPELAppNLSID">myBPELApp</Label>
      <ShortName NLSID="myBPELAppShortName">myBPELApp</ShortName>
      <Description NLSID="myBPELAppDescription">myBPELApp</Description>
   </Display>

   <Metric NAME="ServerBeanStats" TYPE="TABLE">
      <Display>
         <Label NLSID="ServerBeanStats">ServerBeanStats</Label>
      </Display>
      <TableDescriptor>
         <ColumnDescriptor NAME="createCount" TYPE="NUMBER">
            <Display>
               <Label NLSID="createCount">createCount</Label>
               <Unit NLSID="count">count</Unit>
            </Display>
         </ColumnDescriptor>
         <ColumnDescriptor NAME="removeCount" TYPE="NUMBER">
            <Display>
               <Label NLSID="removeCount">removeCount</Label>
               <Unit NLSID="count">count</Unit>
            </Display>
         </ColumnDescriptor>
      </TableDescriptor>
      <QueryDescriptor FETCHLET_ID="OJMX">
         <Property NAME="machine" SCOPE="INSTANCE">HTTPMachine</Property>
         <Property NAME="OracleHome" SCOPE="INSTANCE">OracleHome</Property>
```

```
         <Property NAME="oc4jInstanceName" SCOPE="INSTANCE"
OPTIONAL="TRUE">OC4JInstanceName</Property>
         <Property NAME="jvmId" SCOPE="INSTANCE" OPTIONAL="TRUE">JVMId</Property>
         <Property NAME="mgmtWebSite" SCOPE="INSTANCE"
OPTIONAL="TRUE">MgmtWebSite</Property>
         <Property NAME="authuser" SCOPE="INSTANCE"
OPTIONAL="TRUE">authUser</Property>
         <Property NAME="authpwd" SCOPE="INSTANCE"
OPTIONAL="TRUE">authPasswd</Property>
         <Property NAME="metric" SCOPE="GLOBAL">ServerBeanStats</Property>
         <Property NAME="delimiter" SCOPE="GLOBAL">|</Property>
         <Property NAME="name" SCOPE="GLOBAL">getStatistics</Property>
         <Property NAME="signature"
SCOPE="GLOBAL">objectName,statNames,languageCode,countryCode</Property>
<Property NAME="returnType" SCOPE="GLOBAL">arrayOfComplexObjectBean</Property>
         <Property NAME="dontAddDefaultRowKey" SCOPE="GLOBAL">true</Property>
         <Property NAME="columnOrder"
SCOPE="GLOBAL">/CreateCount/count,/RemoveCount/count</Property>
         <Property NAME="arguments" SCOPE="GLOBAL">
            <![CDATA[<arguments>
   <argument>
      <value>oc4j:EJBModule="ejb_ob_
engine",J2EEApplication=orabpel,J2EEServer=standalone,j2eeType=StatelessSessionBea
n,name="ServerBean"</value>
   </argument>
   <argument>
      <value>CreateCount</value>
      <value>RemoveCount</value>
   </argument>
   <argument>
      <value>en</value>
   </argument>
   <argument>
      <value>US</value>
   </argument>
</arguments>]]>
         </Property>
      </QueryDescriptor>
   </Metric>

<Metric NAME="ServerBeanCount" TYPE="TABLE">
      <Display>
         <Label NLSID="ServerBeanCount">ServerBeanCount</Label>
      </Display>
      <TableDescriptor>
         <ColumnDescriptor NAME="activeInstances" TYPE="NUMBER">
            <Display>
               <Label NLSID="activeInstances">activeInstances</Label>
            </Display>
         </ColumnDescriptor>
         <ColumnDescriptor NAME="maxInstances" TYPE="NUMBER">
            <Display>
               <Label NLSID="maxInstances">maxInstances</Label>
            </Display>
         </ColumnDescriptor>
         <ColumnDescriptor NAME="minInstances" TYPE="NUMBER">
            <Display>
               <Label NLSID="minInstances">minInstances</Label>
            </Display>
         </ColumnDescriptor>
```

```
      </TableDescriptor>
      <QueryDescriptor FETCHLET_ID="OJMX">
         <Property NAME="machine" SCOPE="INSTANCE">HTTPMachine</Property>
         <Property NAME="OracleHome" SCOPE="INSTANCE">OracleHome</Property>
         <Property NAME="oc4jInstanceName" SCOPE="INSTANCE"
OPTIONAL="TRUE">OC4JInstanceName</Property>
         <Property NAME="jvmId" SCOPE="INSTANCE" OPTIONAL="TRUE">JVMId</Property>
         <Property NAME="mgmtWebSite" SCOPE="INSTANCE"
OPTIONAL="TRUE">MgmtWebSite</Property>
         <Property NAME="authuser" SCOPE="INSTANCE"
OPTIONAL="TRUE">authUser</Property>
         <Property NAME="authpwd" SCOPE="INSTANCE"
OPTIONAL="TRUE">authPasswd</Property>
         <Property NAME="metric" SCOPE="GLOBAL">ServerBeanCount</Property>
         <Property NAME="delimiter" SCOPE="GLOBAL">|</Property>
         <Property NAME="name" SCOPE="GLOBAL">getAttributes</Property>
         <Property NAME="signature"
SCOPE="GLOBAL">objectName,attributeNames,languageCode,countryCode</Property>
         <Property NAME="returnType"
SCOPE="GLOBAL">arrayOfComplexObjectBean</Property>
         <Property NAME="dontAddDefaultRowKey" SCOPE="GLOBAL">true</Property>
         <Property NAME="columnOrder"
SCOPE="GLOBAL">/activeInstances,/maxInstances,/minInstances</Property>
<Property NAME="arguments" SCOPE="GLOBAL">
            <![CDATA[<arguments>
   <argument>      <value>oc4j:EJBModule="ejb_ob_
engine",J2EEApplication=orabpel,J2EEServer=standalone,j2eeType=StatelessSessionBea
n,name="ServerBean"</value>
   </argument>
   <argument>
      <value>activeInstances</value>
      <value>maxInstances</value>
      <value>minInstances</value>
   </argument>
   <argument>
      <value>en</value>
   </argument>
   <argument>
      <value>US</value>
   </argument>
</arguments>]]>
         </Property>
      </QueryDescriptor>
   </Metric>

<Metric NAME="Response" TYPE="TABLE">
      <Display>
         <Label NLSID="Response">Response</Label>
      </Display>
      <TableDescriptor>
         <ColumnDescriptor NAME="Status" TYPE="NUMBER">
            <Display>
               <Label NLSID="Status">Status</Label>
            </Display>
         </ColumnDescriptor>
      </TableDescriptor>
      <QueryDescriptor FETCHLET_ID="OJMX">
         <Property NAME="machine" SCOPE="INSTANCE">HTTPMachine</Property>
         <Property NAME="OracleHome" SCOPE="INSTANCE">OracleHome</Property>
         <Property NAME="oc4jInstanceName" SCOPE="INSTANCE"
```

```
OPTIONAL="TRUE">OC4JInstanceName</Property>
         <Property NAME="jvmId" SCOPE="INSTANCE" OPTIONAL="TRUE">JVMId</Property>
         <Property NAME="mgmtWebSite" SCOPE="INSTANCE"
OPTIONAL="TRUE">MgmtWebSite</Property>
         <Property NAME="authuser" SCOPE="INSTANCE"
OPTIONAL="TRUE">authUser</Property>
         <Property NAME="authpwd" SCOPE="INSTANCE"
OPTIONAL="TRUE">authPasswd</Property>
         <Property NAME="metric" SCOPE="GLOBAL">Response</Property>
         <Property NAME="delimiter" SCOPE="GLOBAL">|</Property>
         <Property NAME="name" SCOPE="GLOBAL">getAttributes</Property>
         <Property NAME="signature"
SCOPE="GLOBAL">objectName,attributeNames,languageCode,countryCode</Property>
         <Property NAME="returnType"
SCOPE="GLOBAL">arrayOfComplexObjectBean</Property>
         <Property NAME="dontAddDefaultRowKey" SCOPE="GLOBAL">true</Property>
         <Property NAME="columnOrder" SCOPE="GLOBAL">/state</Property>
         <Property NAME="arguments" SCOPE="GLOBAL">
            <![CDATA[<arguments>
   <argument>

<value>oc4j:J2EEServer=standalone,j2eeType=J2EEApplication,name=orabpel</value>
   </argument>
   <argument>
      <value>state</value>
   </argument>
   <argument>
      <value>en</value>
   </argument>
   <argument>
      <value>US</value>
   </argument>
</arguments>]]>
         </Property>
      </QueryDescriptor>
   </Metric>

   <InstanceProperties>
      <InstanceProperty NAME="HTTPMachine">
         <Display>
            <Label NLSID="dms_HTTPMachine_iprop">Machine name</Label>
         </Display>
      </InstanceProperty>
      <InstanceProperty NAME="OracleHome">
         <Display>
            <Label NLSID="dms_OracleHome_iprop">Oracle home path</Label>
         </Display>
      </InstanceProperty>
<InstanceProperty NAME="OC4JInstanceName" OPTIONAL="TRUE"><Display><Label
NLSID="OC4JInstanceNameiprop">OC4JInstanceName</Label></Display>home</InstanceProp
erty>
      <InstanceProperty NAME="JVMId" OPTIONAL="TRUE"><Display><Label NLSID="JVMId_
iprop">JVMId</Label></Display>1</InstanceProperty>
      <InstanceProperty NAME="MgmtWebSite" OPTIONAL="TRUE"><Display><Label
NLSID="MgmtWebSite_
iprop">MgmtWebSite</Label></Display>default-web-site</InstanceProperty>
      <InstanceProperty NAME="URI" OPTIONAL="TRUE"><Display><Label
NLSID="URI">URI</Label></Display>/JMXSoapAdapter/JMXSoapAdapter</InstanceProperty>
      <InstanceProperty NAME="authUser" OPTIONAL="TRUE">
         <Display>
```

```
            <Label NLSID="dms_authUser_iprop">Username for Basic
authorization</Label>
         </Display>
      </InstanceProperty>
      <InstanceProperty NAME="authPasswd" OPTIONAL="TRUE" CREDENTIAL="TRUE">
         <Display>
            <Label NLSID="dms_authPasswd_iprop">Password for Basic
authorization</Label>
         </Display>
      </InstanceProperty>
      <InstanceProperty NAME="Version" OPTIONAL="TRUE"><Display><Label
NLSID="oc4j_version_iprop">Version of
myBPELApp</Label></Display>1.0</InstanceProperty>
   </InstanceProperties>
</TargetMetadata>
```

Example 20–12 Generated Metric Collection File

```
<!DOCTYPE TargetCollection SYSTEM "../dtds/TargetCollection.dtd">
<!-- This file is generated by Collector at 2011-04-28 12:11:55 -->

<TargetCollection TYPE="myBPELApp" INCLUDE_DEFAULT="TRUE">
        <CollectionItem NAME="ServerBeanStats" UPLOAD="YES">
                <Schedule>
                        <IntervalSchedule INTERVAL="60" TIME_UNIT="Sec"/>
                </Schedule>
                <MetricColl NAME="ServerBeanStats">
                      <Condition COLUMN_NAME="createCount" CRITICAL="100"
WARNING="85" OPERATOR="GT" OCCURRENCES="3" MESSAGE="createCount is %value% and has
crossed warning (%warning_threshold%) or critical (%critical_threshold%)
threshold." MESSAGE_NLSID="createCount_cond"/>
                </MetricColl>
        </CollectionItem>
        <CollectionItem NAME="ServerBeanCount" UPLOAD="YES">
                <Schedule>
                        <IntervalSchedule INTERVAL="300" TIME_UNIT="Sec"/>
                </Schedule>
                <MetricColl NAME="ServerBeanCount">
                </MetricColl>
        </CollectionItem>

        <CollectionItem NAME="Response" UPLOAD="YES">
                <Schedule>
                        <IntervalSchedule INTERVAL="30" TIME_UNIT="Sec"/>
                </Schedule>
                <MetricColl NAME="Response">
                        <Condition COLUMN_NAME="Status" CRITICAL="1"
WARNING="NotDefined" OPERATOR="NE" OCCURRENCES="2" MESSAGE="Status is %value% and
has crossed warning (%warning_threshold%) or critical (%critical_threshold%)
threshold." MESSAGE_NLSID="Status_cond"/>
                </MetricColl>
        </CollectionItem>
</TargetCollection>
```

20.4.2 Displaying Target Status Information

For the status information of your targets to appear correctly within the Enterprise Manager console, you must define a metric, called Response, that has a column, named Status, with a critical threshold set. The status of target instances of this type appears in the console as "Up" (available) if the metric value is below the critical

threshold. When the threshold is exceeded, the target status appears as "Down" in the console.

You can create the Response metric in another `jmxcli` session (append the metric to the metadata and collection files created in an earlier session). Example 20–13 illustrates adding a Response metric to previously generated metadata and collection files from a new command-line session.

Example 20–13 Adding a Response Metric

```
./emctl jmxcli /scratch/shiphomes
//oc4j/1013_PRODUCTION/ -p 12403 -c welcome1 -m 'oc4j:j2eeType=J2EEApplication,n
ame=orabpel,*'

oracleHome=/ade/sparmesw_10202_ssm/oracle
targetHome=/scratch/shiphomes//oc4j/1013_PRODUCTION/
The Port is 12403

Connecting to server: localhost:12403
Connecting as user: oc4jadmin

Obtained 1 MBeans matching pattern oc4j:j2eeType=J2EEApplication,name=orabpel,*.

Enter the target type for this metric: [myJ2EEApp] myBPELApp

Enter the target version: [1.0]

Enter the target metadata file: [./metadata/myBPELApp.xml]

Enter the default collections file: [./default_collection/myBPELApp.xml]
The file ./metadata/myBPELApp.xml already exists.

Do you want to overwrite the existing file, append to it, or quit <o/a/q>? [a] a
Appending to existing file: ./metadata/myBPELApp.xml.
The available targets are:
0: Identifies a J2EE application EAR that has been deployed
        (oc4j:J2EEServer=standalone,j2eeType=J2EEApplication,name=orabpel)

Enter the index of target/MBean you wish to monitor or press <Ctrl-C> to quit: 0
Following metric source types are available for selected target(s):
        0: JMX Attributes
        1: JMX Operations

Enter the index of your choice or press <Ctrl-C> to quit: 0

Attributes are:
        0: allAccessibleGroups   Return Value: java.util.Set
        1: allAccessibleUsers    Return Value: java.util.Set
        2: applicationRootDirectoryPath  Return Value: java.lang.String
        3: archivePath   Return Value: java.lang.String
        4: childApplicationNames         Return Value: [Ljava.lang.String;
        5: childApplications     Return Value: [Ljavax.management.ObjectName;
        6: dataSourcesDescriptor         Return Value: java.lang.String
        7: deploymentDescriptor  Return Value: java.lang.String
        8: ejbClassLoaderPath    Return Value: java.lang.String
        9: eventProvider         Return Value: boolean
        10: groups       Return Value: java.util.Set
        11: iiopStubs    Return Value: [B
        12: metricRulesDescriptor        Return Value: java.lang.String
        13: Modules      Return Value: [Ljavax.management.ObjectName;
```

```
        14: objectName   Return Value: java.lang.String
        15: ohsRouting   Return Value: boolean
        16: parentApplication    Return Value: javax.management.ObjectName
        17: parentApplicationName        Return Value: java.lang.String
        18: properties   Return Value: java.util.Properties
        19: proprietaryDeploymentDescriptor      Return Value: java.lang.String
        20: proxyInterfaceSQLObjects     Return Value: [Ljava.lang.String;
        21: routingId    Return Value: java.lang.String
        22: Server       Return Value: javax.management.ObjectName
        23: sharedLibraryImports   Return Value:
[Loracle.oc4j.admin.management.shared.SharedLibraryImport;
        24: startTime    Return Value: long
        25: state        Return Value: int
        26: stateManageable      Return Value: boolean
        27: statisticsProvider   Return Value: boolean
        28: syntheticWebModules  Return Value:
oracle.oc4j.admin.management.shared.WebModule
        29: users        Return Value: java.util.Set
        30: webSite      Return Value: java.lang.String
        31: webSiteBindings      Return Value: java.util.Map
Select one or more items as comma-separated indices: 25

Number of possible columns in the resultant metric are 1.

Enter the name for this metric column at index=0 : [state] Status
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [Status]
Enter the NLSID for column: [Status]
Enter the UNIT for column "Status": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n] y
Creating threshold!!
Following operators are available for creating thresholds:
 0: GT
 1: EQ
 2: LT
 3: LE
 4: GE
 5: CONTAINS
 6: NE
 7: MATCH
Enter the index of your choice or press <Ctrl-C> to quit: 6
Enter the CRITICAL threshold: [NotDefined] 1
Enter the WARNING threshold: [NotDefined]
Enter the number of occurrences that trigger threshold: [6] 2
Enter the message to be used when threshold is triggered: [Status is %value% and
has crossed warning (%warning_threshold%) or critical (%critical_threshold%)
threshold.]
Enter NLSID for the message used when threshold is triggered: [Status_cond]

Enter the name of this metric: Response
Enter the label for this metric: [Response]

Do you want periodic collection for this metric <y/n>? [n] y
Enter the collection interval in seconds: 30
Periodic collection interval is: 30 seconds.

Do you want to create another metric <y/n>? [n] n
Written the metadata xml file: ./metadata/myBPELApp.xml.
Updated the default collection file for myBPELApp at location ./default_collecti
```

```
on/myBPELApp.xml.
Exiting...

Please note that the Response metric collected in this jmxcli session would be
appended to the metadata and default_collection file created in an earlier session
of the tool. (User can chose to overwrite the earlier file as well if they
specific  the "o" option to the following prompt)

Do you want to overwrite the existing file, append to it, or quit <o/a/q>? [a] a
```

20.5 Monitoring a Standalone JMX-instrumented Java Application or Java Virtual Machine (JVM) Target

> **Note:** If your Java application is not JMX-instrumented, but you want to monitor the J2SE 5.0 JVM on which it is running, go directly to Section 20.10.3, "Configuring a Standalone Java Application or JVM Target" to create target instances of type JVM. This enables you to monitor these JVMs in Enterprise Manager, preferably from an Enterprise Manager Agent installed on the same host as your JVM. However, the prerequisites and known limitations discussed below still apply.

Enterprise Manager provides an out-of-box JVM target type. This enables you to add and configure metrics from standalone J2SE1.5 JVMs that are enabled for remote management in Enterprise Manager version 10.2.0.3 or greater.

If your standalone Java application exposes data through JMX MBeans as for a J2EE application deployed on an Oracle Container for J2EE, you can use the JMX command-line tool to define such an application as an Enterprise Manager target type and generate a metadata and default collection file for this target type. You can monitor your standalone application targets from an Enterprise Manager Agent, preferably installed on the same host as your JVM. Multiple JVMs running on that host can be monitored by the same Enterprise Manager Agent.

You can collect metrics from user-defined MBeans on a standalone (J2SE1.5-based) JVM and place them into Enterprise Manager using the JMX fetchlet. The fetchlet is designed for a standalone Sun J2SE1.5-based (or higher) JVM containing user-defined MBeans that use JMX OpenTypes as arguments and return values.

Prerequisites

- Java virtual machine J2SE 1.5 or higher instance running on a specific host. This JVM could be running a JMX-enabled application that exposes metrics via MBeans that need to be monitored as a target in Enterprise Manager. If the application does not expose MBeans, the JVM itself could be monitored using the built-in JVM target type provided in Enterprise Manager. See Section 20.10.3, "Configuring a Standalone Java Application or JVM Target" for more information.
- JMX agent enabled for local access. Set this system property when you start the JVM or Java application:

  `com.sun.management.jmxremote`
- Monitoring and management from remote systems enabled. Set this system property when you start the JVM:

  `com.sun.management.jmxremot.port=portNum`

For additional information about enabling the JVM for remote management, see the following document:

`http://java.sun.com/j2se/1.5.0/docs/guide/management/agent.html#remote`

- Enterprise `Manager` Management Agent version 10.2.0.3 or greater installed on that host.
- Enterprise Manager Management Server (OMS) version 10.2.0.3 or greater with which the Management Agent communicates.

Known Limitations

Currently, the jmxcli tool only allows you to browse and monitor MBeans (platform and application-defined) that are available on the default platform MBeanserver on the target JVM instance. The tool does not support monitoring a custom MBeanServer on the target JVM instance. The jmxcli tool primarily handles attributes as well as parameter and return values for operations that are OpenTypes, such as SimpleTypes, CompositeTypes, TabularTypes, and arrays of SimpleTypes.

20.5.1 Generating Metadata and Default Collection Files

As with Web Services and the J2EE application on OC4J, the command-line tool (jmxcli) simplifies creating the requisite target definition files: metadata and the default collection file for a standalone JMX-instrumented Java application. The tool is an offline configuration utility that connects you to an MBeanServer on a J2SE1.5 or higher JVM and enables you to browse available MBeans. It can also append metrics to an existing set of files during a subsequent invocation of the tool.

During a command-line tool session, you select specific MBeans and then choose the desired attributes/statistical values or operations Enterprise Manager needs to retrieve or invoke periodically on these MBeans to collect these values. The tool helps define packaging for these collected values as one or more Enterprise Manager metrics (with columns), and also enables you to specify a metric collection interval.

20.5.1.1 JMX Command-line Tool Syntax

The JMX command-line tool syntax is as follows:

```
cd <Agent Instance dir>/bin
emctl jmxcli -t JVM
               [  -l <JMXServiceURL>
                  -h <hostname>
                  -p <port>
                  -u <username>
                  -c <credential/password>
                  -w <work directory>
                  -e <true/false>
                  [-m <MBeanName> | -d <jmx_domain> | -s <mBeanPattern>]
               ]
```

The `jmxcli` command accepts the following options:

- **-t JVM** Indicates that the MBeanServer is on a standalone JVM
- **-l** JMXServiceURL of the target JVM
- **-h** Hostname of the JVM. Default: "localhost" if the -l option is not specified
- **-p** RMI/RMIS port of the JVM. Default: "23791" if the -l option is not specified. From the `ORACLE_HOME/opmn/bin` directory of your Application Server 10.1.3.0

or later instance, run `opmnctl status -l` to determine the RMI port for the OC4J for which MBeans were deployed.

- **-u** Valid username for the JVM. Default: None
- **-c** Password for the above user. Default: None. The password is only used to retrieve data and is not stored anywhere.
- **-w** Work directory where the metadata and default collection files are created. Default: Current directory. When invoking the command-line tool, you must have write permission on this directory to create subdirectories and add files. If the metadata and default collection files already exist within that directory, you have the option of appending to or overwriting the original files.
- **-e** True for enabling the SSL connection to the JVM. Default: false

You can also specify ONE of the following three parameters (`-m`, `-d` or `-s`) to retrieve a subset of MBeans available on the MBeanServer. By default, all MBeans on the MBeanServer are displayed for you to select from if none of these parameters are specified.

- **-m** MBean ObjectName of the required MBean that needs to be retrieved and examined. If this is an ObjectName pattern-matching multiple MBeans, you are shown a list of all MBeans that match the pattern, and you can select one at a time to work on.
- **-d** MBean domain of the required application whose MBeans need to be retrieved and examined. For example, you want to browse all MBeans for an application (myApp). MBeans for this application would be available in the JMX domain "myApp".
- **-s** MBean pattern matching an set of similar MBeans from which the metrics are to be defined. The `-s` parameter allows bulk retrieval of JMX Attributes/Statistics from multiple MBeans of a similar type.

  If you specify the `-s` parameter, the resulting metrics created during this `jmxcli` session appear as a table in the Enterprise Manager console with multiple rows — one row representing each MBean that matches the specified pattern, and with the MBean ObjectName as a key column. For example, if you specify `-s 'oc4j:j2eeType=Servlet,*'` the resulting metric will have multiple rows, one for each servlet that matches the ObjectName pattern. Besides the MBean ObjectName column, other columns would be the attributes or fields from the return object of the operation, selected during the `jmxcli` session.

20.5.1.2 Generating the Files

The following steps explain how to prepare for and then use the JMX command-line tool to generate the files.

1. Bring up the standalone JVM instance with the MBeans. The following example shows an invocation of the JVM:

   ```
   JDK15/bin/java -Dcom.sun.management.jmxremote
   -Dcom.sun.management.jmxremote.port=6789
   -Dcom.sun.management.jmxremote
   -Dcom.sun.management.jmxremote.authenticate=false
   -Dcom.sun.management.jmxremote.ssl=false MyJMXEnabledApp $*
   ```

   The `jmxcli` tool connects to the port number above as a JSR-160 client.

2. Go to the `$ORACLE_HOME/bin` directory of the 10.2.0.3 or higher version of the Enterprise Manager Agent.

3. Set the environment variable as follows:

```
setenv USER_JARS /myAppHome/<myJar1.jar>;/myAppHome/<myJar2.jar>
```

This step is needed if custom classes are being returned in attributes and/or operations in any of the MBeans registered with the target MBeanServer. The Enterprise Manager Agent (fetchlet) can only effectively monitor attributes and/or operations that return JMX OpenTypes, but it could also handle Java Bean properties (through getters and setters) on any custom classes.

> **Note:** If the application-defined MBeans are returning custom classes, you need to also set up the corresponding user jar file in the `CLASSPATH` of the Enterprise Manager Agent monitoring this application. To do this, manually insert the location of this jar into the `$ORACLE_HOME/sysman/config/classpath.lst` file on the Enterprise Manager Agent, then restart the Agent.

4. Run the following command:

```
./emctl jmxcli -t JVM -h localhost -p 6789 u <user> -c <password>
```

where:

- **-t** JVM indicates that the MBeanServer is running on a standard JVM
- **-h** Hostname where the JVM is running
- **-p** Port number that enables the JVM for JSR-160 remote access

You can also specify an `-l <JMXServiceURL>` option instead of `-h <host>` and `-p <port>` options.

You can invoke `jmxcli` with a `-w <work directory>` option to create the metadata and default collection files in the specified work directory. If you do not specify `-w` when you start `jmxcli`, it defaults to the current directory, which is the directory where you start `jmxcli`.

Once invoked, the command-line interface automatically prompts you for the requisite information, as shown in Example 20–14. For most of the prompts, you can just press enter to use defaults. If you need to abort a JMX command-line tool session, you can press Ctrl+C at any point to exit. Session information will not be saved.

When the session concludes after you exit, the result will be a `myJ2EEApp.xml` file (or whatever target type you specified) as `metadata/myJ2EEApp.xml`, and a `default_collection/myJ2EEApp.xml` file if you specified periodic collection.

Sample JMXCLI Invocations

The following sample enables you to browse all MBeans on a remote MBeanServer:

```
./emctl jmxcli -t JVM -p 6789 (the host defaults to "localhost")
```

The following sample invokes the command-line interface and filters MBeans based on the MBeanPattern specified as the argument for the `-m` option:

```
./emctl jmxcli -t JVM -p 6789 -m "java.lang:*"
```

Example 20–14 Sample JMXCLI Session

```
oracleHome=/ade/sparmesw_emas_ml/oracle
userJars=
```

```
Connecting to server: localhost:6789
Connecting without authentication. For specifying username and password use
the
-u and -c options.

Obtained 14 MBeans matching pattern java.lang:*.

Enter the target type for this metric: [myJ2EEApp] myJavaApp

Enter the target version: [1.0]

Enter the target metadata file: [./metadata/myJavaApp.xml]

Enter the default collections file: [./default_collection/myJavaApp.xml]

Enter a label for this target type: [myJavaApp]

Enter a description for this target type: [myJavaApp]
The available targets are:
0: sun.management.CompilationImpl
        (java.lang:type=Compilation)
1: sun.management.MemoryManagerImpl
        (java.lang:name=CodeCacheManager,type=MemoryManager)

2: sun.management.GarbageCollectorImpl
        (java.lang:name=Copy,type=GarbageCollector)

3: sun.management.MemoryPoolImpl
        (java.lang:name=Eden Space,type=MemoryPool)

4: sun.management.RuntimeImpl
        (java.lang:type=Runtime)

5: sun.management.ClassLoadingImpl
        (java.lang:type=ClassLoading)

6: sun.management.MemoryPoolImpl
        (java.lang:name=Survivor Space,type=MemoryPool)

7: sun.management.ThreadImpl
        (java.lang:type=Threading)

8: sun.management.GarbageCollectorImpl
        (java.lang:name=MarkSweepCompact,type=GarbageCollector)

9: com.sun.management.UnixOperatingSystem
        (java.lang:type=OperatingSystem)

10: sun.management.MemoryImpl
        (java.lang:type=Memory)

11: sun.management.MemoryPoolImpl
        (java.lang:name=Code Cache,type=MemoryPool)

12: sun.management.MemoryPoolImpl
        (java.lang:name=Tenured Gen,type=MemoryPool)

13: sun.management.MemoryPoolImpl
        (java.lang:name=Perm Gen,type=MemoryPool)
```

```
Enter the index of target/MBean you wish to monitor or press <Ctrl-C> to quit: 4

Following metric source types are available for selected target(s):
        0: JMX Attributes

Enter the index of your choice or press <Ctrl-C> to quit: 0

Attributes are:
        0: BootClassPath         Return Value: java.lang.String
        1: BootClassPathSupported        Return Value: boolean
        2: ClassPath     Return Value: java.lang.String
        3: InputArguments        Return Value: [Ljava.lang.String;
        4: LibraryPath   Return Value: java.lang.String
        5: ManagementSpecVersion         Return Value: java.lang.String
        6: Name  Return Value: java.lang.String
        7: SpecName      Return Value: java.lang.String
        8: SpecVendor    Return Value: java.lang.String
        9: SpecVersion   Return Value: java.lang.String
        10: StartTime    Return Value: long
        11: SystemProperties     Return Value:
javax.management.openmbean.TabularData
        12: Uptime       Return Value: long
        13: VmName       Return Value: java.lang.String
        14: VmVendor     Return Value: java.lang.String
        15: VmVersion    Return Value: java.lang.String
Select one or more items as comma-separated indices: 6,7,8

Number of possible columns in the resultant metric are 3.

Enter the name for this metric column at index=0 : [Name]
Is this column a KEY Column <y/n>? [n] y
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [Name]
Enter the NLSID for column: [Name]
Enter the UNIT for column "Name": [millisec, kb etc.. ]

Enter the name for this metric column at index=1 : [SpecName]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [SpecName]
Enter the NLSID for column: [SpecName]
Enter the UNIT for column "SpecName": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]

Enter the name for this metric column at index=2 : [SpecVendor]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [SpecVendor]
Enter the NLSID for column: [SpecVendor]
Enter the UNIT for column "SpecVendor": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]

Enter the name of this metric: RuntimeMetric
Enter the label for this metric: [RuntimeMetric]

Do you want periodic collection for this metric <y/n>? [n] y
Enter the collection interval in seconds: 300
Periodic collection interval is: 300 seconds.
```

```
Do you want to create another metric <y/n>? [n]
Written the metadata xml file: ./metadata/myJavaApp.xml.
Creating new file: ./default_collection/myJavaApp.xml.
Updated the default collection file for myJavaApp at location
./default_collection/myJavaApp.xml.
Exiting...
```

20.5.2 Using the Metadata and Default Collection Files

Look at the `<currentDir>/metadata` and `currentDir/default_collection` directories to see the `myTarget.xml` files (for the target type you specified earlier).

You can use these files as follows:

- Convert the files to a Management Plug-in Archive (MPA) and push them from OMS to the Agent and target instances created from OMS. See Section 20.7, "Creating a Management Plug-in Archive" and Section 20.10.3, "Configuring a Standalone Java Application or JVM Target".
- Edit the files, extract the metric definitions and `QueryDescriptors`, move them to other metadata and default collection files, and post-process them by creating `ExecutionDescriptors` as needed.

If you want the status information of your targets to appear correctly in the Enterprise Manager console, you need to define a Response metric. See Section 20.4.2, "Displaying Target Status Information" for more information.

20.6 Monitoring JMX Applications Deployed on Oracle WebLogic Application Servers

The JMX fetchlet, supplied with 11.1 Management Agents, enables you to monitor key metrics in your JMX-instrumented applications deployed on Oracle WebLogic Application Server 9.x or above.

Monitoring JMX-instrumented applications with Enterprise Manager entails defining a new target type that Enterprise Manager can monitor via Management Plug-ins. As with the Web Services wscli command-line tool (and as was possible for Oracle Application Servers (OC4J), Enterprise Manager provides an jmxcli command-line tool to automate the generation of the target metadata and collection files for custom JMX instrumented applications on weblogic servers..

Prerequisites

- Oracle WebLogic Server 9.x instance running on a specific host with a JMX-enabled application deployed on it that needs to be monitored as a target in Enterprise Manager.
- Enterprise Manager Management Agent version 11.1 or greater installed (preferably) on that host.
- Enterprise Manager Management Server (OMS) version 10.2.0.4 or greater with which the Management Agent communicates.
- The `jmxcli` tool primarily handles attributes and parameter and return values for operations that are OpenTypes. Examples: SimpleTypes, CompositeTypes, TabularTypes, and arrays of SimpleTypes.

20.6.1 Creating Metadata and Default Collection Files using jmxcli

As with Web Services, the JMX command-line tool (jmxcli) simplifies creating the requisite target definition files: metadata and the default collection file. The tool is an offline configuration utility that connects you to an MBeanServer and enables you to browse available MBeans. It can also append metrics to an existing file during a subsequent invocation of the tool.During a command-line tool session, you select specific MBeans and then choose the desired JMX attributes/statistical values the Enterprise Manager needs to retrieve or JMX operations that need to be invoked periodically on these MBeans to collect these values. The tool helps define packaging for these collected values as one or more Enterprise Manager metrics (with columns), and also enables you to specify a metric collection interval and thresholds for the columns.

20.6.1.1 JMX Command-line Tool Syntax

The JMX command-line tool syntax is as follows for a JMX-enabled target on an Oracle WebLogic Application Server:

```
./emctl jmxcli -t WebLogic [help|options]
  where options are:
           [  -l <JMX ServiceURL>
              -u <username>
              -c <credential/password>
              -w <work directory>
              [-m <MBeanName> | -d <jmx_domain> | -s mBeanPattern>]
                ]
```

The `jmxcli` command accepts the following options:

- l - JMXServiceURL to the WebLogic managed server hosting the custom MBeans in the form

  ```
  service:jmx:t3://<host:t3port>/jndi/weblogic.management.mbeanservers.runtime
  ```

- u - Valid username for the WebLogic domain. Default: "weblogic"
- c - Password associated with the user specified by the -u option. Default: None. If you do not specify a password, you are prompted for the password.
- w - Directory where the metadata and default collection files created by the JMX command-line tool are placed. Default: Current directory.

  When invoking the command-line tool, you must have write permission on this directory to create subdirectories and add files. If the metadata and default collection files already exist within that directory, you have the option of appending to or overwriting the original files.

You can also specify ONE of the following three parameters (-m, -d or -s) to retrieve a subset of MBeans available on the MBeanServer. By default, all MBeans on the MBeanServer are displayed for you to select from if none of these parameters are specified.

- m - MBean ObjectName of the required MBean that needs to be retrieved and examined. If this is an ObjectName pattern-matching multiple MBeans, you are shown a list of all MBeans that match the pattern, and you can select one at a time to work on.
- d - MBean domain of the required application whose MBeans need to be retrieved and examined. For example, you want to browse all MBeans for an application (myApp). MBeans for this application would be available in the JMX domain "myApp".

- s - MBean pattern-matching an existing set of MBeans from which the metrics are to be defined. The -s parameter allows retrieval of JMX Attributes/Statistics from multiple MBeans of a similar type into one Metric.

  If you specify the -s parameter, the resulting metrics created during this `jmxcli` session appear as a table in the Enterprise Manager console with multiple rows - one row representing each MBean that matches the specified pattern (with the MBean ObjectName as a key column if no other key columns are defined). For example, if you specify -s 'com.bea:Type=ServletRuntime,*' the resulting metric will have multiple rows, one for each servlet that matches the ObjectName pattern. Besides the MBean ObjectName key column, other columns would be the attributes or fields from the return object of the operation, selected during the `jmxcli` session.

20.6.1.2 Generating the Files

To start the JMX command-line tool:

1. Go to the $AGENT_HOME/bin directory.
2. Run the following command:

   ./emctl jmxcli -t WebLogic [OPTIONS]

Once invoked, the command-line interface automatically prompts you for the requisite information, as shown in the following example. If you need to abort a JMX command-line tool session, you can press Ctrl+C at any point to exit. Session information will not be saved.

The following example illustrates a sample `jmxcli` session:

Example 20–15 `jmxcli` ***Session***

```
$ ./emctl jmxcli -t WebLogic -l
"service:jmx:t3://stbct14:22048/jndi/weblogic.management.mbeanservers.runtime" -u
weblogic -c welcome1 -s "*:type=soainfra_bpel_requests,*"
```

NOTE 1: The -s option above will result in a metric with as many rows as the number of MBeans which match the ObjectName pattern specified, every time the metric is collected by the agent. If you need to always collect from a specific Mbean then use the -m <ObjectName> option instead of the -s <Mbean pattern> used in above example.

NOTE 2: If you need to use t3s to connect to the weblogic server then the following env variable needs to be set before invoking the `jmxcli`

```
setenv USER_JAVA_PROPS=-Dweblogic.security.TrustKeyStore=CustomTrust
 -Dweblogic.security.CustomTrustKeyStoreFileName=$ORACLE_
HOME/sysman/config/montrust/AgentTrust.jks
 -Dweblogic.security.SSL.enforceConstraints=off
 -Dweblogic.security.SSL.ignoreHostnameVerification=true
 -Djavax.net.ssl.trustStore=$ORACLE_HOME/sysman/config/montrust/AgentTrust.jks
(or set USER_JAVA_PROP= … equivalent on win32)
setenv USER_JARS  <; separated list of jars needed in classpath if custom
authentication modules are involved in SSL connection>
```

A semi-colon is used as a delimiter for the list of jar files.

Example: `setenv USER_JARS "jar1;jar2;jar3"`

In some cases, if MBeans return custom WebLogic objects in their MBeanInfo, the weblogic.jar may need to be set to the above env variable before invoking the `jmxcli`.

Example: `setenv USER_JARS $BEA_HOME/server/lib/weblogic.jar`

```
oracleHome=/ade/sparmesw_egcli/oracle/work/middleware/oms
userJars=
Connecting to server:
 service:jmx:t3://stbct14:22048/jndi/weblogic.management.mbeanservers.runtime
Connecting as user: weblogic
Obtained 3 MBeans matching pattern *:type=soainfra_bpel_requests,*.

Enter the target type for this metric: [myJ2EEApp] myCustomWLApp
Enter the target version: [1.0]
Enter the target metadata file: [./metadata/myCustomWLApp.xml]
Enter the default collections file: [./default_collection/myCustomWLApp.xml]
Enter a label for this target type: [myCustomWLApp]
Enter a description for this target type: [myCustomWLApp]
The available targets are:
0: DMS metric mbean
        (oracle.dms:name=/soainfra/engines/bpel/requests/engine,type=soainfra_
bpel_requests)
1: DMS metric mbean
        (oracle.dms:name=/soainfra/engines/bpel/requests/system,type=soainfra_
bpel_requests)
2: DMS metric mbean
        (oracle.dms:name=/soainfra/engines/bpel/requests/invoke,type=soainfra_
bpel_requests)
Following metric source types are available for selected target(s):
        0: JMX Attributes
Enter the index of your choice or press <Ctrl-C> to quit: 0
Attributes are:
        0: active_count  Return Value: java.lang.Integer
        1: active_maxValue       Return Value: java.lang.Integer
        2: active_minValue       Return Value: java.lang.Integer
        3: active_value  Return Value: java.lang.Integer
        4: Name  Return Value: java.lang.String
        5: Parent        Return Value: java.lang.String
        6: scheduled_count       Return Value: java.lang.Integer
        7: scheduled_maxValue    Return Value: java.lang.Integer
        8: scheduled_minValue    Return Value: java.lang.Integer
        9: scheduled_value       Return Value: java.lang.Integer
        10: threadCount_count    Return Value: java.lang.Integer
        11: threadCount_maxValue         Return Value: java.lang.Integer
        12: threadCount_minValue         Return Value: java.lang.Integer
        13: threadCount_value    Return Value: java.lang.Integer

Select one or more items as comma separated indices: 4,0,1,2
Number of possible columns in the resultant metric are 4.

Enter the name for this metric column at index=0 : [Name]
Is this column a KEY Column <y/n>? [n] y
```

Specifying "y" signifies that the value of this column is unique in case multiple rows are returned.

```
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [Name]
Enter the NLSID for column: [Name]
Enter the UNIT for column "Name": [millisec, kb etc.. ]

Enter the name for this metric column at index=1 : [active_count]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
```

```
Enter the label for column: [active_count]
Enter the NLSID for column: [active_count]
Enter the UNIT for column "active_count": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n] y
Creating threshold!!
Following operators are available for creating thresholds:
 0: GT
 1: EQ
 2: LT
 3: LE
 4: GE
 5: CONTAINS
 6: NE
 7: MATCH

Enter the index of your choice or press <Ctrl-C> to quit: 0
Enter the CRITICAL threshold: [NotDefined] 50
Enter the WARNING threshold: [NotDefined] 45
Enter the number of occurrences that trigger threshold: [6] 3
Enter the message to be used when threshold is triggered: [active_count is %value%
and has crossed warning (%warning_threshold%) or critical (%critical_threshold%)
threshold.]
Enter NLSID for the message used when threshold is triggered: [active_count_cond]
Enter the name for this metric column at index=2 : [active_maxValue]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [active_maxValue]
Enter the NLSID for column: [active_maxValue]
Enter the UNIT for column "active_maxValue": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]

Enter the name for this metric column at index=3 : [active_minValue]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [active_minValue]
Enter the NLSID for column: [active_minValue]
Enter the UNIT for column "active_minValue": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]

Enter the name of this metric: bpel_requests
Enter the label for this metric: [bpel_requests]

Do you want periodic collection for this metric <y/n>? [n] y
Enter the collection interval in seconds: 300
Periodic collection interval is: 300 seconds.

Do you want to create another metric <y/n>? [n]
Written the metadata xml file: ./metadata/myCustomWLApp.xml.
Creating new file: ./default_collection/myCustomWLApp.xml.
Updated the default collection file for myCustomWLApp at location ./default_
collection/myCustomWLApp.xml.
Exiting...
```

Example 20–16 Sample `jmxcli` ***Invocation (using -m and defining multiple metrics from multiple Mbeans in one*** `jmxcli` ***session***

```
$ ./emctl jmxcli -t WebLogic -l
"service:jmx:t3://stbct14:22048/jndi/weblogic.management.mbeanservers.runtime" -u
weblogic -c welcome1 -m
```

```
"com.bea:ApplicationRuntime=soa-infra,WebAppComponentRuntime=soa_server1_/b2b,*"

oracleHome=/ade/sparmesw_egcli/oracle/work/middleware/oms
userJars=
Connecting to server: service:jmx:t3://stbct14:22048/jndi/weblogic.management.mb
eanservers.runtime
Connecting as user: weblogic
Obtained 8 MBeans matching pattern
com.bea:ApplicationRuntime=soa-infra,WebAppComponentRuntime=soa_server1_/b2b,*.

Enter the target type for this metric: [myJ2EEApp] myCustomWLApp

Enter the target version: [1.0]

Enter the target metadata file: [./metadata/myCustomWLApp.xml]

Enter the default collections file: [./default_collection/myCustomWLApp.xml]
The file ./metadata/myCustomWLApp.xml already exists.

Do you want to overwrite the existing file, append to it, or quit <o/a/q>? [a]
```

Note: *Because the file already exists, it will be appended*.

```
Appending to existing file: ./metadata/myCustomWLApp.xml.
The available targets are:
0: (com.bea:ApplicationRuntime=soa-infra,Name=JspServlet,ServerRuntime=soa_
server1,Type=ServletRuntime,WebAppComponentRuntime=soa_server1_/b2b)
1: (com.bea:ApplicationRuntime=soa-infra,Name=transportServlet,ServerRuntime=soa_
server1,Type=ServletRuntime,WebAppComponentRuntime=soa_server1_/b2b)
2: (com.bea:ApplicationRuntime=soa-infra,Name=transportServletV,ServerRuntime=soa_
server1,Type=ServletRuntime,WebAppComponentRuntime=soa_server1_/b2b)
3: (com.bea:ApplicationRuntime=soa-infra,Name=b2b_starter_wls,ServerRuntime=soa_
server1,Type=ServletRuntime,WebAppComponentRuntime=soa_server1_/b2b)
4: (com.bea:ApplicationRuntime=soa-infra,Name=soa_server1_soa_server1_
/b2b,ServerRuntime=soa_server1,Type=PageFlowsRuntime,WebAppComponentRuntime=soa_
server1_/b2b)
5 (com.bea:ApplicationRuntime=soa-infra,Name=WebServiceServlet,ServerRuntime=soa_
server1,Type=ServletRuntime,WebAppComponentRuntime=soa_server1_/b2b)
6: (com.bea:ApplicationRuntime=soa-infra,Name=RedirectUIServlet,ServerRuntime=soa_
server1,Type=ServletRuntime,WebAppComponentRuntime=soa_server1_/b2b)
7: (com.bea:ApplicationRuntime=soa-infra,Name=FileServlet,ServerRuntime=soa_
server1,Type=ServletRuntime,WebAppComponentRuntime=soa_server1_/b2b)
Enter the index of target/MBean you wish to monitor or press <Ctrl-C> to quit: 4
Following metric source types are available for selected target(s):
        0: JMX Attributes
        1: JMX Operations
Enter the index of your choice or press <Ctrl-C> to quit: 0

Attributes are:
        0: AppName       Return Value: java.lang.String
        1: ContextPath   Return Value: java.lang.String
        2: HttpServerName        Return Value: java.lang.String
        3: Name  Return Value: java.lang.String
        4: PageFlows     Return Value: [Ljavax.management.ObjectName;
        5: Parent        Return Value: javax.management.ObjectName
        6: ServerName    Return Value: java.lang.String
        7: Type  Return Value: java.lang.String

Select one or more items as comma separated indices: 3,0,1

Number of possible columns in the resultant metric are 3.
```

```
Enter the name for this metric column at index=0 : [Name]
Is this column a KEY Column <y/n>? [n] y
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [Name]
Enter the NLSID for column: [Name]
Enter the UNIT for column "Name": [millisec, kb etc.. ]

Enter the name for this metric column at index=1 : [AppName]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [AppName]
Enter the NLSID for column: [AppName]
Enter the UNIT for column "AppName": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]

Enter the name for this metric column at index=2 : [ContextPath]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [ContextPath]
Enter the NLSID for column: [ContextPath]
Enter the UNIT for column "ContextPath": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]
Enter the name of this metric: PageFlowsRuntime
Enter the label for this metric: [PageFlowsRuntime]

Do you want periodic collection for this metric <y/n>? [n] y
Enter the collection interval in seconds: 3600
Periodic collection interval is: 3600 seconds.

Do you want to create another metric <y/n>? [n] y
```

This indicates more metrics need to be created in this `jmxcli` session. This process will repeat until you answer "n" to the question.

```
Do you want to create another metric <y/n>? [n]
Written the metadata xml file: ./metadata/myCustomWLApp.xml.
Updated the default collection file for myCustomWLApp at location ./default_coll
ection/myCustomWLApp.xml.
Exiting...
```

After the JMX command-line tool generates the target metadata and collection files, you can create the Management Plug-in archive.

20.6.1.3 Displaying Target Status Information

For the status information of your targets to appear correctly within the Enterprise Manager console, you must define a metric, called Response, that has a column, named Status, with a critical threshold set. The status of target instances of this type appears in the console as "Up" (available) if the metric value is below the critical threshold. When the threshold is exceeded, the target status appears as "Down" in the console.

You can create the Response metric in another `jmxcli` session (append the metric to the metadata and collection files created in an earlier session).

Example 20–17 Adding a Response Metric

```
setenv USER_JARS $T_WORK/middleware/wlserver_10.3/server/lib/weblogic.jar
```

This is required as some MBeans return WebLogic-specific classes which the JMX client (`jmxcli`) needs in its classpath.

```
$ ./emctl jmxcli -t WebLogic -l
"service:jmx:t3://stbct14:22048/jndi/weblogic.management.mbeanservers.runtime" -u
weblogic -c welcome1 -m com.bea:Type=ApplicationRuntime,Name=soa-infra,*"
```

For J2EE applications deployed on WebLogic it may be appropriate to make the ActiveVersionState JMX attribute of the ApplicationRuntime Mbean corresponding to the application deployment as the Status column. However, any other attribute of any other relevant Mbean to the application could also be used.

```
oracleHome=/ade/sparmesw_egcli/oracle/work/middleware/oms
userJars=
Connecting to server:
service:jmx:t3://stbct14:22048/jndi/weblogic.management.mbeanservers.runtime
Connecting as user: weblogic
Obtained 1 MBeans matching pattern
 com.bea:Type=ApplicationRuntime,Name=soa-infra,*.
Enter the target type for this metric: [myJ2EEApp] myCustomWLApp
Enter the target version: [1.0]
Enter the target metadata file: [./metadata/myCustomWLApp.xml]
Enter the default collections file: [./default_collection/myCustomWLApp.xml]
The file ./metadata/myCustomWLApp.xml already exists.

Do you want to overwrite the existing file, append to it, or quit <o/a/q>? [a]
Appending to existing file: ./metadata/myCustomWLApp.xml.
The available targets are:
0: (com.bea:Name=soa-infra,ServerRuntime=soa_server1,Type=ApplicationRuntime)
Enter the index of target/MBean you wish to monitor or press <Ctrl-C> to quit: 0
Following metric source types are available for selected target(s):
        0: JMX Attributes
        1: JMX Operations
Enter the index of your choice or press <Ctrl-C> to quit: 0
Attributes are:
        0: ActiveVersionState    Return Value: java.lang.Integer
        1: ApplicationName       Return Value: java.lang.String
        2: ApplicationVersion    Return Value: java.lang.String
        3: ClassRedefinitionRuntime      Return Value: javax.management.ObjectName
        4: ComponentRuntimes     Return Value: [Ljavax.management.ObjectName;
        5: EAR   Return Value: java.lang.Boolean
        6: HealthState   Return Value: weblogic.health.HealthState
        7: KodoPersistenceUnitRuntimes   Return Value:
[Ljavax.management.ObjectName;
        8: LibraryRuntimes       Return Value: [Ljavax.management.ObjectName;
        9: MaxThreadsConstraintRuntimes  Return Value:
[Ljavax.management.ObjectName;
        10: MinThreadsConstraintRuntimes         Return Value:
[Ljavax.management.ObjectName;
        11: Name         Return Value: java.lang.String
        12: OptionalPackageRuntimes      Return Value:
[Ljavax.management.ObjectName;
        13: Parent       Return Value: javax.management.ObjectName
        14: QueryCacheRuntimes   Return Value: [Ljavax.management.ObjectName;
        15: RequestClassRuntimes         Return Value:
[Ljavax.management.ObjectName;
        16: Type         Return Value: java.lang.String
        17: WorkManagerRuntimes  Return Value: [Ljavax.management.ObjectName;
        18: WseeRuntimes         Return Value: [Ljavax.management.ObjectName;

Select one or more items as comma separated indices: 0
```

```
Number of possible columns in the resultant metric are 1.

Enter the name for this metric column at index=0 : [ActiveVersionState] Status
```

Note: The column name must be "Status".

```
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [Status]
Enter the NLSID for column: [Status]
Enter the UNIT for column "Status": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n] y
Creating threshold!!
Following operators are available for creating thresholds:
 0: GT
 1: EQ
 2: LT
 3: LE
 4: GE
 5: CONTAINS
 6: NE
 7: MATCH

Enter the index of your choice or press <Ctrl-C> to quit: 6
Enter the CRITICAL threshold: [NotDefined] 2
```

Status of target is marked down if a CRITICAL THRESHOLD is triggered on the Status column of the Response Metric.. In this case if value != ACTIVATED (such as: != 2)

```
Enter the WARNING threshold: [NotDefined]
Enter the number of occurrences that trigger threshold: [6] 1
Enter the message to be used when threshold is triggered: [Status is %value% and
has crossed warning (%warning_threshold%) or critical (%critical_threshold%)
threshold.]
Enter NLSID for the message used when threshold is triggered: [Status_cond]

Enter the name of this metric: Response
```

Note: The metric name must be "Response".

```
Enter the label for this metric: [Response]

Do you want periodic collection for this metric <y/n>? [n] y
Enter the collection interval in seconds: 30
Periodic collection interval is: 30 seconds.

Do you want to create another metric <y/n>? [n]
Written the metadata xml file: ./metadata/myCustomWLApp.xml.
Updated the default collection file for myCustomWLApp at location ./default_
collection/myCustomWLApp.xml.
Exiting...
```

20.6.2 Using the Metadata and Default Collection Files

Look at the `<currentDir>/metadata` and `currentDir/default_collection` directories to see the `myTarget.xml` files (for the target type you specified earlier).

You can use these files as follows:

- Convert the files to a Management Plug-in Archive (MPA). See Section 20.7, "Creating a Management Plug-in Archive".
- Move the MPA to the OMS. See Section 20.8, "Importing a Management Plug-in"
- Push the MPA to the Agents. See Section 20.9, "Deploying a Management Plug-in"
- Create custom target instances. See Section 20.10.4, "Adding a Target Instance for a Custom J2EE Application on WebLogic"

If you want the status information of your targets to appear correctly in the Enterprise Manager console, you need to define a Response metric. See Section 20.6.1.3, "Displaying Target Status Information" for more information.

20.7 Creating a Management Plug-in Archive

Note: The Management Plug-in creation process applies to the Web Service, JMX-instrumented, and standalone JVM target types. Section 20.5 through Section 20.9 use Web Services as an illustrative example.

Important: Oracle Enterprise Manager Extensibility Development Kit Version 12.1.0.0.0 or greater must be installed in order to create a Management Plug-in Archive.

To create a Management Plug-in Archive:

1. Create the following directory structure in a convenient location such as /tmp:

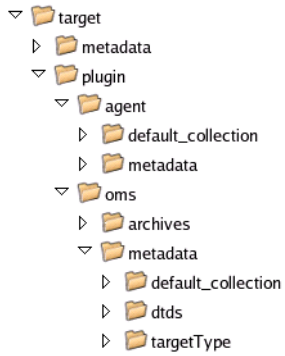

Note: 'targetType' directory under oms/metadata must be literally called 'targetType' and not the name of your custom target type.

---

2. Copy all DTD files from the `$ORACLE_HOME/oracle_common/sysman/admin/dtds` directory to the corresponding DTD directory within the directory structure you created in step one. For example, `/tmp/target/plugin/oms/metadata/dtds`
3. Copy the newly created target metadata and collection files to the appropriate destination folder.

Table 20–1 Copied Metadata and Collection Files

Source File	Destination Folder
/tmp/target/metadata/CalculatorService.xml	/tmp/target/plugin/agent/metadata /tmp/target/plugin/oms/metadata/targetType
/tmp/target/metadata/CalculatorServiceCollection.xml	/tmp/target/plugin/agent/default_collection /tmp/target/plugin/oms/metadata/default_collection
/tmp/target/metadata/TrafficLight.xml	/tmp/target/plugin/agent/metadata /tmp/target/plugin/oms/metadata/targetType
/tmp/target/metadata/ TrafficLightCollection.xml	/tmp/target/plugin/agent/default_collection /tmp/target/plugin/oms/metadata/default_collection

Note: 'If montoring a custom application/target on WebLogic the wlthint3client.jar (client library for weblogic) also needs to be bundled with the plugin.

Specifically, you must copy <WebLogicHome>/wlserver10.x/server/lib/wlthint3client.jar to /tmp/target/plugin/agent/archives/jlib/ (archives/jlib directory needs to be created)

4. Create a new file (plugin.xml) under `/tmp/target/plugin` with the following content:

```
<?xml version = '1.0' encoding = 'UTF-8'?>
<Plugin xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
        xsi:schemaLocation="http://www.oracle.com/EnterpriseGridControl/plugin_
metadata plugin_metadata.xsd"
        xmlns="http://www.oracle.com/EnterpriseGridControl/plugin_metadata">
  <PluginId vendorId="oracle" productId="sysman" pluginTag="demo"/>
  <PluginVersion value="12.1.0.0.0"/>
  <PluginOMSOSAruId value="2000">
    <CertifiedOMSOSAruId value="46"/>
    <CertifiedOMSOSAruId value="912"/>
  </PluginOMSOSAruId>
  <ShortDescription>Demo Plugins</ShortDescription>
  <TargetTypeList>
    <TargetType name="CalculatorService" isIncluded="TRUE">
      <VersionSupport>
        <SupportedVersion supportLevel="Comprehensive" minVersion="1.0.0"/>
      </VersionSupport>
    </TargetType>
    <TargetType name="TrafficLight" isIncluded="TRUE">
      <VersionSupport>
        <SupportedVersion supportLevel="Comprehensive" minVersion="1.0.0"/>
      </VersionSupport>
    </TargetType>
  </TargetTypeList>
  <EMPlatforms>
    <BasePlatform version="11.2.0.1.0"/>
    <CertifiedPlatform version="11.2.0.1.0"/>
  </EMPlatforms>
  <PluginAttributes Type="MP" Category="Others"
                    DisplayName="Demo Plugins" ReleaseDate="2011-01-01"
                    ReleaseStatus="Test" OnlineDocLink="http://otn.oracle.com"
                    LastUpdDate="2011-01-01"/>
```

```
  <AgentSideCompatibility>
    <Version>11.2.0.1.0</Version>
  </AgentSideCompatibility>
</Plugin>
```

5. Create a new file (plugin_registry.xml) under `/tmp/target/plugin/agent` with the following content:

```
<?xml version="1.0"?>
<PlugIn ID="oracle.sysman.demo" Description="Demo Plugins"
        Version="12.1.0.0.0" HotPluggable="false"
        xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
        xsi:schemaLocation="http://www.oracle.com/EnterpriseGridControl/plugin
plugin.xsd">
  <TargetTypes>
    <FileLocation>metadata/CalculatorService.xml</FileLocation>
    <FileLocation>metadata/TrafficLight.xml</FileLocation>
  </TargetTypes>
  <TargetCollections>
    <FileLocation>default_
collection/CalculatorServiceCollection.xml</FileLocation>
    <FileLocation>default_collection/TrafficLightCollection.xml</FileLocation>
  </TargetCollections>
</PlugIn>
```

> **Note:** If you change the plugin ID and/or version in the plugin.xml and plugin_registry.xml files, make sure the values in both files match.

6. Use Oracle Enterprise Manager Extensibility Development Kit to create a Management Plug-in. The following example session illustrates Management Plug-in creation using the EMPDK. Please note the partner PDK is the one that needs to be used.

```
$ empdk create_plugin -stage_dir /tmp/target/plugin/ -out_dir /tmp/target/out
-force
Validating Plugin Xml : Passed
Validating Plugin Structure: Passed
Validating Metadata Syntax: Passed
Validating Metadata Semantic: Passed
Plugin Validation : Passed
Validation Report generated to: /tmp/target/out/plugin_validation_report_
110607.txt
Creating the opar file......
Successfully created the plugin archive . The opar file is
/tmp/target/out/12.1.0.0.0_oracle.sysman.demo_2000_0.opar
```

The Management Plug-in (OPAR) generated is `/tmp/target/out/12.1.0.0.0_oracle.sysman.demo_2000_0.opar`. Once the OPAR file is created, you can import the Management Plug-in to Enterprise Manager.

20.8 Importing a Management Plug-in

The next step is to import the Management Plug-in in the OPAR file into the Enterprise Manager using the Enterprise Manager Command-line Interface (EMCLI). To import the Management Plug-in, run the following EMCLI command:

```
$ emcli import_update -omslocal -file=<full path to OPAR file>
```

Note: (`emcli setup` and `emcli sync` commands must be performed to set up the `emcli` before issuing the `import_update` command. Also, the software library "upload file" location needs to setup on the OMS from the Software Library home page (from the Enterprise menu, choose Provisioning And Patching, and then Software Library) by selecting "Administration" from the **Action** drop down menu and providing a location using the **Add** button)

Example 20–18 Management Plug-in Import EMCLI Session

```
$ emcli import_update -omslocal -file=/tmp/target/out/12.1.0.0.0_
oracle.sysman.demo_2000_0.opar
Processing update: Plug-in - Demo Plugins
Operation completed successfully. Update has been uploaded to Enterprise Manager.
Please use the Self Update Home to manage this update.
```

Note: If the import command generates an error "Software library is not configured, see the Self Update Home for details", do the following:

- Login to the EM console and navigate to Enterprise ' Provisioning and Patching ' Software Library.
- Click on Actions -> Administration.
- Enter a name (like mySFLIB) and a directory location for the library.
- Retry the `emcli` command above.

20.9 Deploying a Management Plug-in

After the Management Plug-in has been imported into Enterprise Manager, you must deploy it to an OMS. You perform the import process directly from the Enterprise Manager console.

1. From the **Setup** menu, choose **Extensibility-->Plug-ins**.
2. Verify the Management Plug-in (Demo Plug-ins) you imported previously is shown on page:

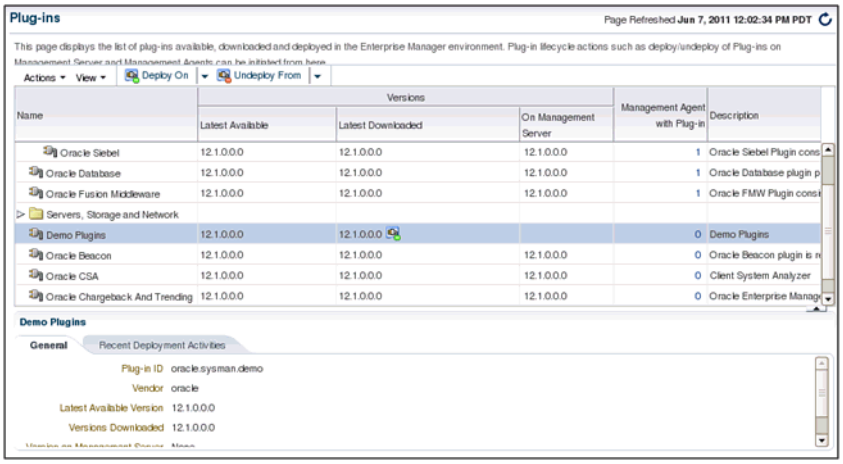

3. Select the plug-in and right-click on the entry to display the popup menu. Choose **Deploy On-->Management Servers**.
4. Enter the Repository SYS Password and click **Next** and wait for the *Pre-requisite Checks* to complete.
5. On the Pre-requisite Checks have completed, click **Next.**

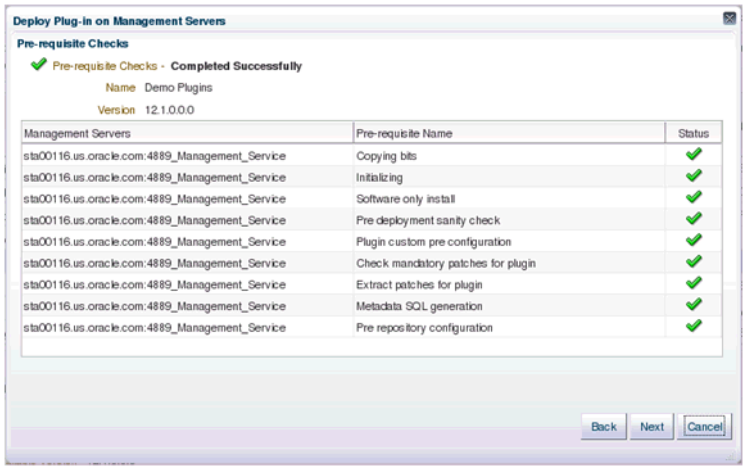

6. From the **Review** window, click **Deploy**.
7. Click **Show Status** to check the deployment status.

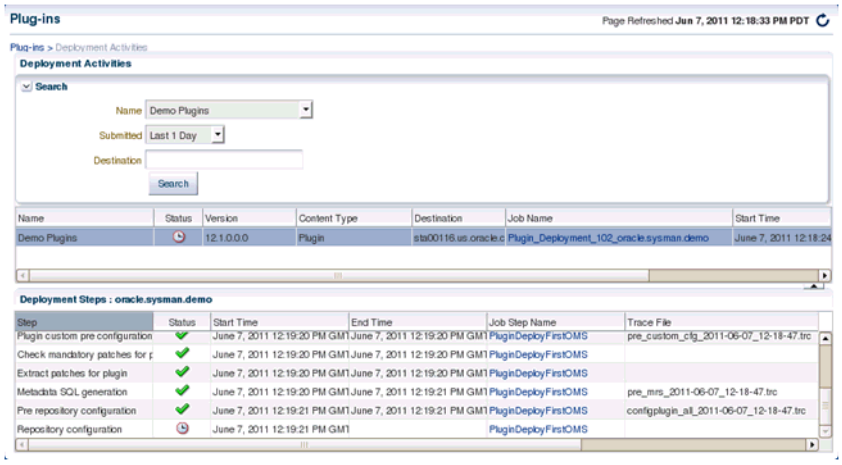

8. Keep checking the deployment status until it is completed.

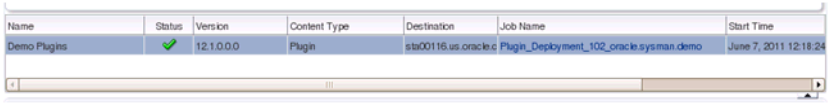

20.10 Adding a Target to Management Agent

Once the Management Plug-in has been deployed to the OMS, you are ready to add targets defined by your Management Plug-in to different monitoring agents.

For illustrative purposes, the following steps show how to add the sample CalculatorService and TrafficLight as targets.

20.10.1 Adding a Web Services Target - CalculatorService

To add the CalculatorService target, perform the following steps:

1. From the **Setup** menu, choose **Add Target** and then **Add Targets Manaully**.
2. Select **Add Non-Host Targets by Specifying Target Monitoring Properties**.
3. Select **CalculatorService** from the **Target Type** drop-down menu.
4. Select the desired agent from the **Monitoring Agent** drop-down menu.

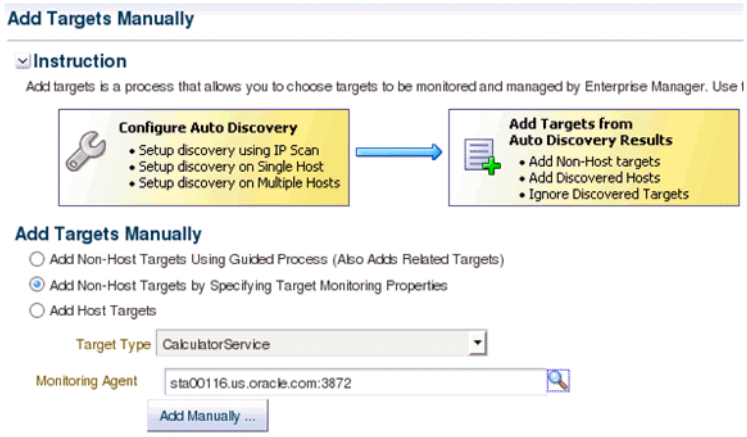

5. Click **Add Manaully** to proceed.
6. Enter the property values of the target to be monitored.

7. Click **OK** to complete the process. The confirmation window displays information on the newly added target.

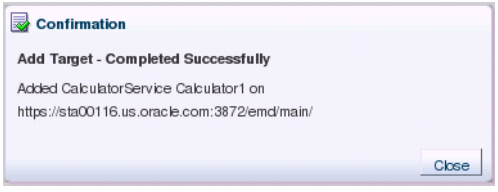

20.10.2 Adding a WS-Management Target - TrafficLight

To add the sample TrafficLight target, perform the following steps:

1. From the **Setup** menu, choose **Add Target** and then **Add Targets Manaully**.
2. Select **Add Non-Host Targets by Specifying Target Monitoring Properties**.
3. Select **TrafficLight** from the **Target Type** drop-down menu.
4. Select the desired agent from the **Monitoring Agent** drop-down menu.

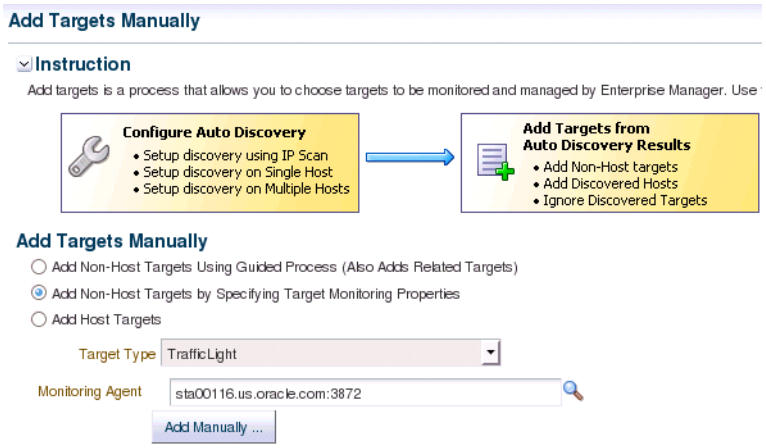

5. Click **Add Manaully** to proceed.
6. Enter the property values of the target to be monitored.

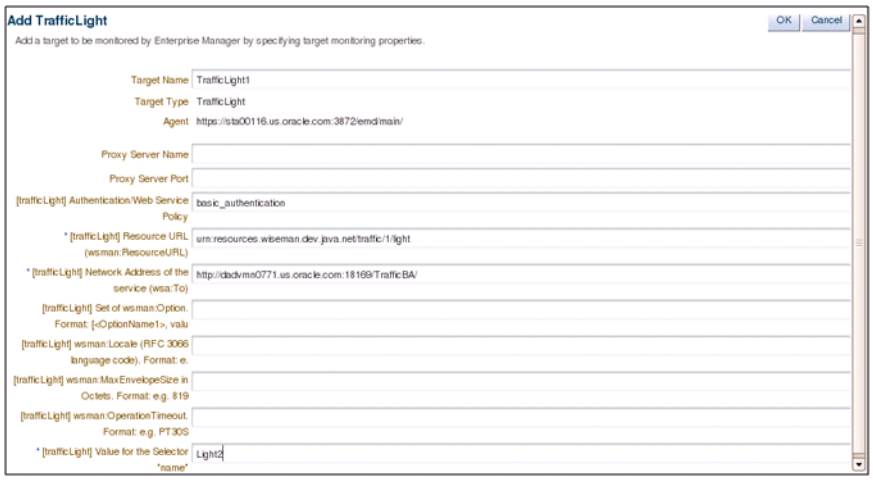

7. Click **OK** to complete the process. The confirmation window displays information on the newly added target..

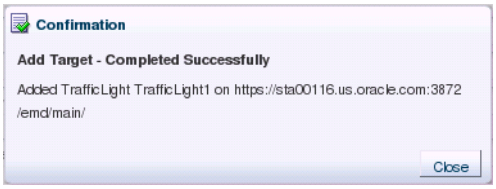

20.10.3 Configuring a Standalone Java Application or JVM Target

If you deployed a Management Plug-in that defines a standalone Java application or you want to use the built-in JVM target type, you can begin configuring your JVM or JMX-enabled Java application targets so that metrics for these targets can be collected in Enterprise Manager Cloud Control.

On the system running the JVM, install an Enterprise Manager Agent version 10.2.0.3 or greater. Although recommended, this is not absolutely necessary for JVM and standalone Java application targets.

To add the JVM target instance, perform the following steps:

1. From the **Setup** menu of Enterprise Manager console (top right) , choose **Add Target** and then **Add Targets Manually.**
2. Select **Add Non-Host Targets by Specifying Target Monitoring Properties**..
3. Select **JVM** from the **Target Type** drop-down menu.
4. Select the desired agent from the **Monitoring Agent** drop-down menu (preferably an Agent local to the JVM being monitored)

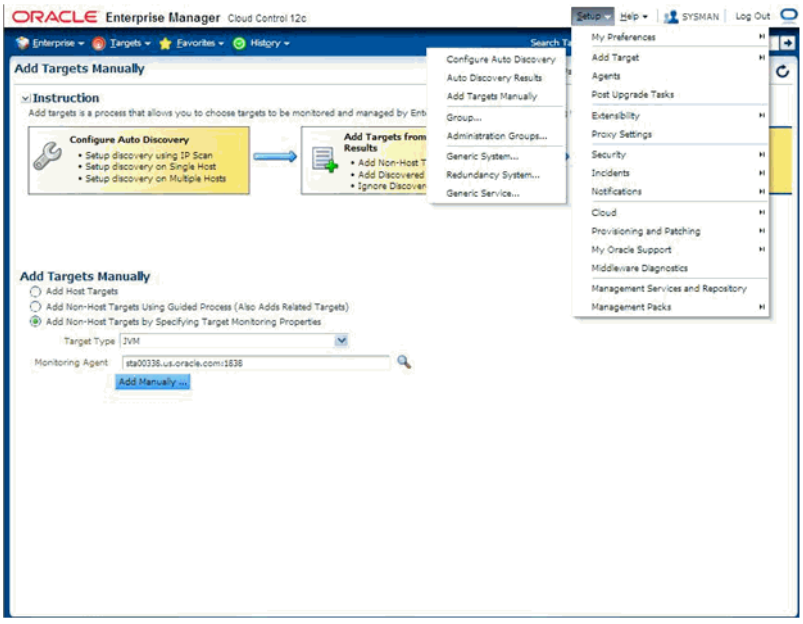

5. Enter the instance properties for this JVM or Java application instance that the Enterprise Manager Agent needs to monitor, then click **OK**.

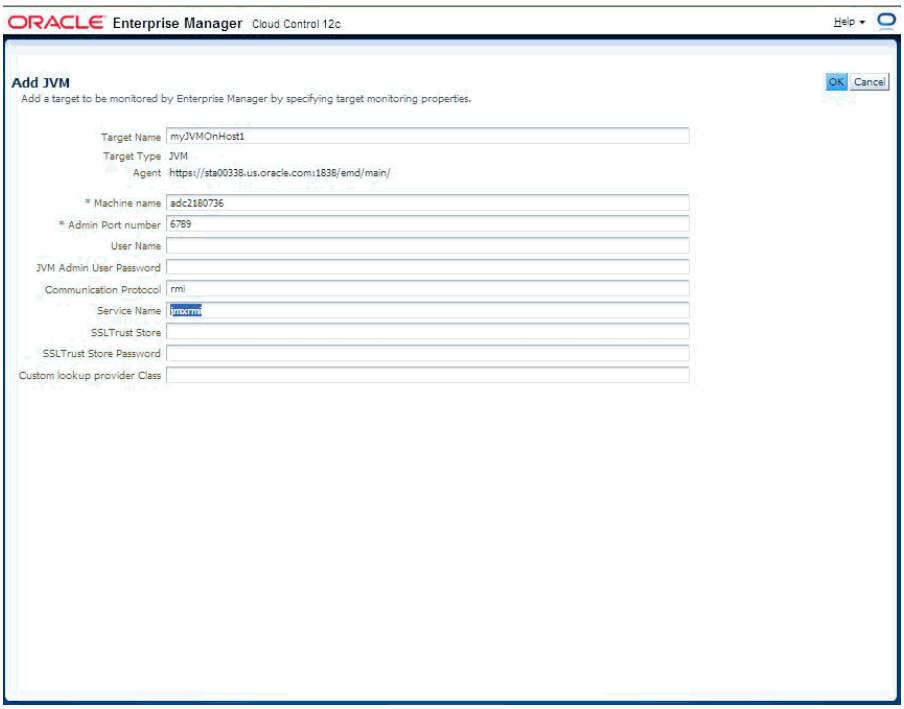

Table 20–2 provides definitions for the instance properties.

Table 20–2 JVM Instance Properties

Property	Definition
Name	Target name for this JVM instance.
MachineName	Host name where this JVM is running.
Admin Port Number	Port number a JSR-160 client can use (such as jconsole when using the "remote" option) to connect to the JVM. (This is the port specified for the `-Dcom.sun.management.jmxremote.port` property when the JVM is started up to enable remote management.)
User Name	Required if JVM started with: `Dcom.sun.management.jmxremote.authenticate=true` with a password and access file.
JVM Admin User Password	See User Name above.
Communication Protocol	Establishes a connection to the MBeanServer on the target JVM. This corresponds to the properties of the JMX ServiceURL needed to establish the JMX connection to the target MBeanServer. The default of rmi should be retained.
Service Name	Establishes a connection to the MBeanServer on the target JVM. This corresponds to the properties of the JMX ServiceURL needed to establish the JMX connection to the target MBeanServer. The default of jmxrmi should be kept.
SSL Trust Store	Location of the SSL Trust Store, which is needed if the target JVM has SSL enabled with `-Dcom.sun.management.jmxremote.ssl=true` on its startup.

Table 20–2 (Cont.) JVM Instance Properties

Property	Definition
SSL Trust Store Password	Password needed to access the SSL Trust Store path.
Custom Lookup Provider Class	Full package name of a user-implemented Java lookup class that can be integrated into the Enterprise Manager client and be used to perform a custom lookup of the MBeanServer through LDAP or other lookup protocols.

6. Navigate to the All Metrics page of the added JVM (Java application) target to see the metrics collected from the JVM (Java application) to Enterprise Manager. These metrics are exposed by the platform MBeans, which is available on JDK1.5 or above, or from application-defined MBeans for your Java application.

   To navigate to JVM target home page from the **Targets** menu, choose **All Targets** and then select your JVM target instance.

   To navigate to the **All Metrics** page, from the **Target** menu, choose **Monitoring** and then choose **All Metrics** from the JVM target's home page menu.

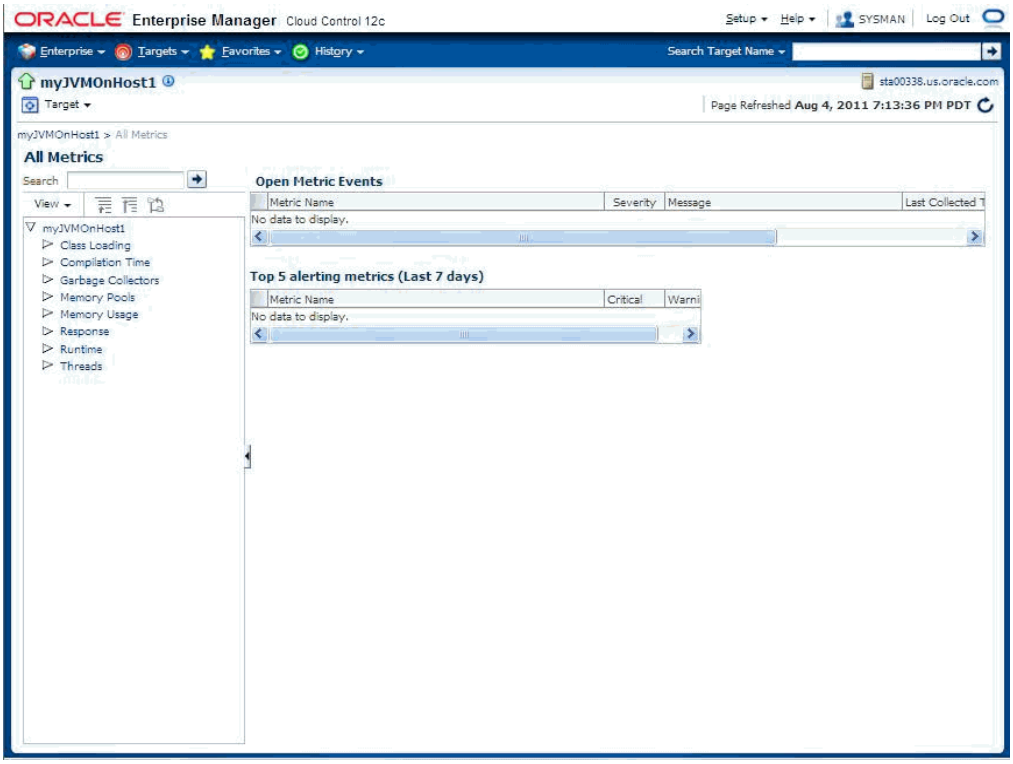

The following graphic shows the collected metric details.

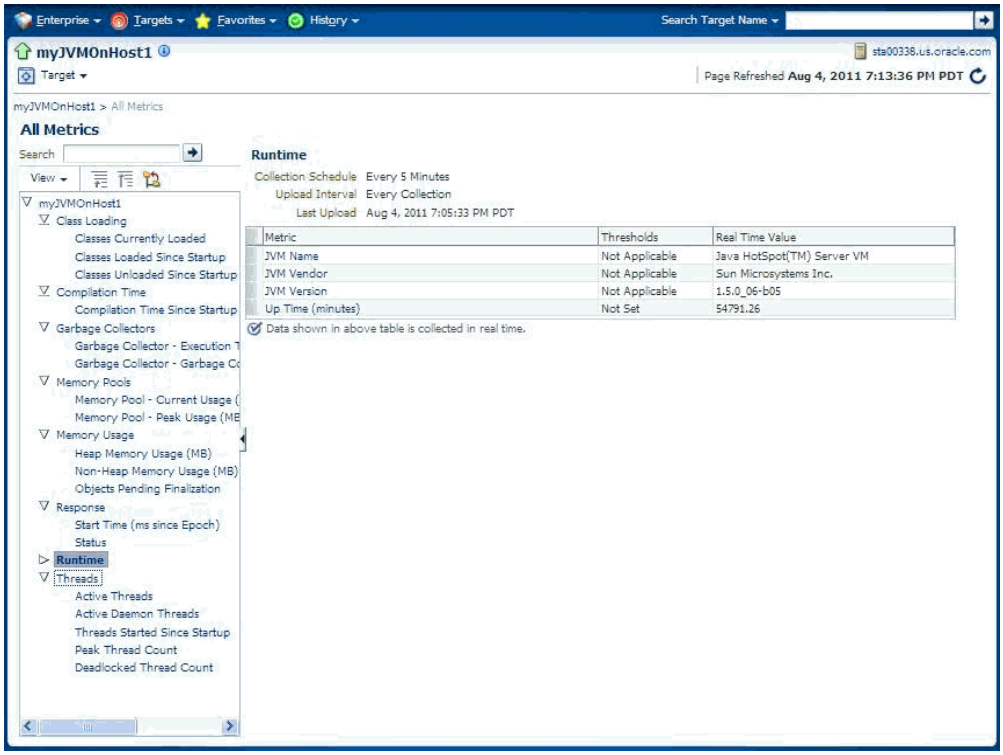

Table 20–3 Properties the Fetchlet Uses

Property	Default	Description
MachineName	localhost	MBean server host machine name.
Port	8888	Port on which the MBean server is listening for connections.
Username	null	User name if required for a connection.
Password	null	Password if required for a connection.
Protocol	rmi	Protocol used for the connection.
Service	jmxrmi	Service used for the connection.
SSLTrustStore	null	Path to the SSLTrustStore.
SSLTrustStorePassword	null	Password needed to access the SSLTrustStore path.

20.10.4 Adding a Target Instance for a Custom J2EE Application on WebLogic

You have a custom J2EE application on WebLogic from which you need to collect custom metrics into Enterprise Manager that are exposed via JMX Mbeans. Once you have defined and deployed a Management Plug-in that defines your custom target type, you can begin configuring your JMX-enabled J2EE application target instances on the various agents where you deployed the Management Plugins to. This is so that metrics for these target instances can be collected in Enterprise Manager.

1. From the **Setup** menu, choose **Add Target** and then **Add Target Manually**. Select the **Add non-Host targets by specifying Target Monitoring Properties** option.
2. Select your custom target type created earlier and deployed to the OMS

3. Select the monitoring agent where you want to create an instance of this target type (this should preferably be an emagent local to the target)

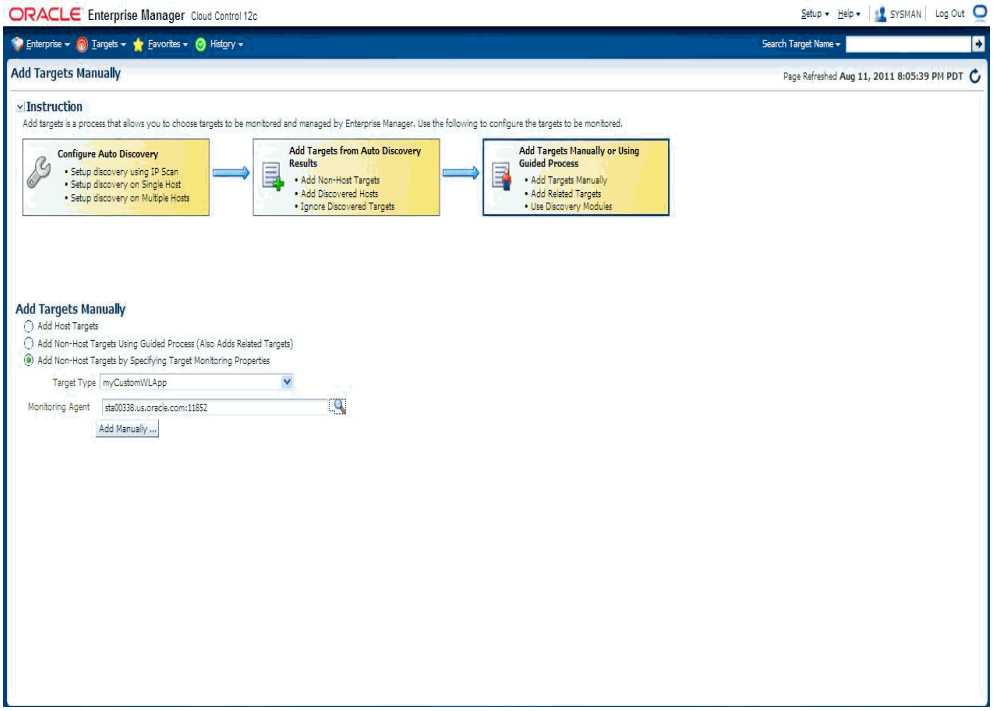

4. Click on **Add Manually** button.

5. Enter the requisite target properties, as shown in the following graphic, then click OK. The newly added target appears in the "All Targets" list.

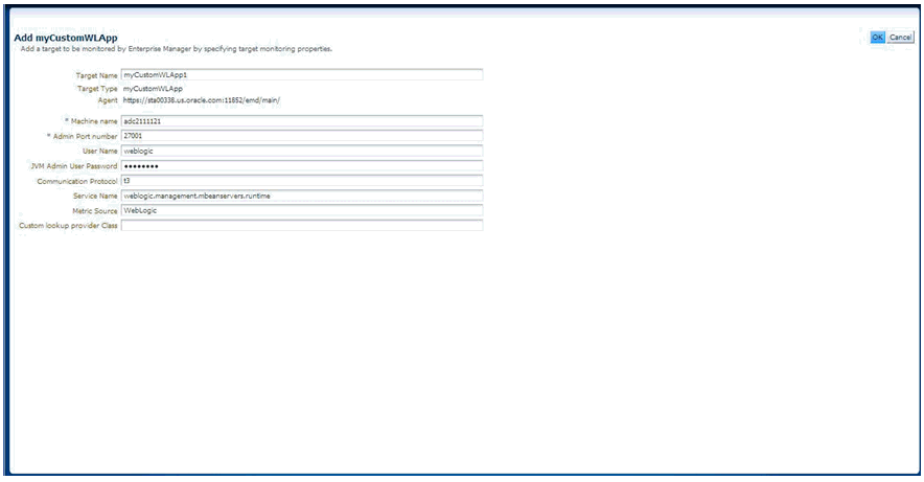

Table 20–4 Target Properties

Property	Definition
Name	Unique name for this target instance.
MachineName	Hostname/IP Address of the machine running the 9.x version or higher of the Oracle WebLogic Application Server..

Table 20–4 (Cont.) Target Properties

Property	Definition
Username	User Name used to establish the JMX connection to the WebLogic server. This could be either an administrator or monitor user.
JVM Admin User Password	Password for above user.
Communication Protocol	t3 (default) or t3s.
Service Name	*weblogic.management.mbeanservers.runtime* (or other MbeanServer where the application registers its Mbeans).
Metric Source	WebLogic

The metrics created can be viewed by navigating to the target instance home page and navigating to the **All Metrics** page (from the **Target** menu, choose **Monitoring** and then **All Metrics**) .

20.11 Monitoring Credential Setup

Some target types require monitoring credentials to be set for target instances. In the demo plug-ins, both CalculatorService and TrafficLight require monitoring credentials. The following steps demonstrate how to set up the credentials:

1. From the **Setup** menu, choose **Security** and then **Monitoring Credentials**.
2. Select **CalculatorService** and then click **Manage Monitoring Credentials**.

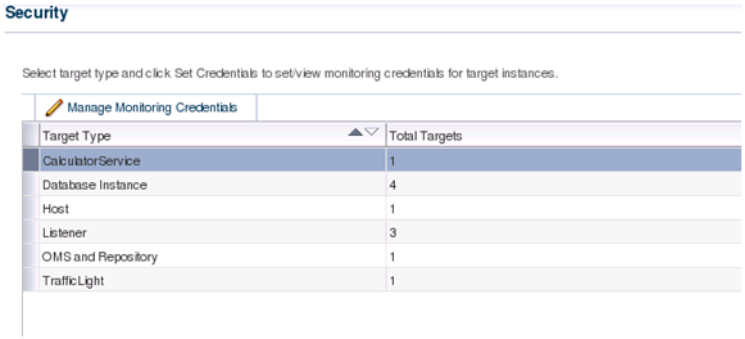

3. Select **Calculator1** and then click **Set Credentials**.

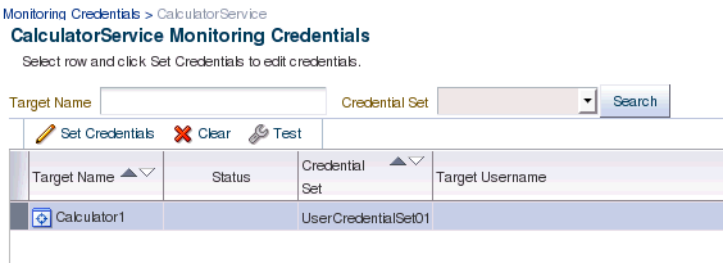

4. Select **AliasCredential** from **Credential Type**. Enter values for **Alias** and **Password**.

5. Click **Save** to finish.
6. Repeat the above steps for the target **TrafficLight1**.

20.12 Viewing Monitored Metrics

With a target instance added to the Agent for monitoring, you can now view metrics defined for your target type. As before, the sample targets are used to illustrate the procedure.

1. From the **All Targets** page, click on the target you added in the previous step. Enterprise Manager takes you to that target's home page.
2. From **Target** menu, select **Monitoring** and then **All Metrics**. The **All Metrics** page appears for the monitored target. An expandable tree list for each metric enables you to drill down to view specific metric parameters, as shown below:

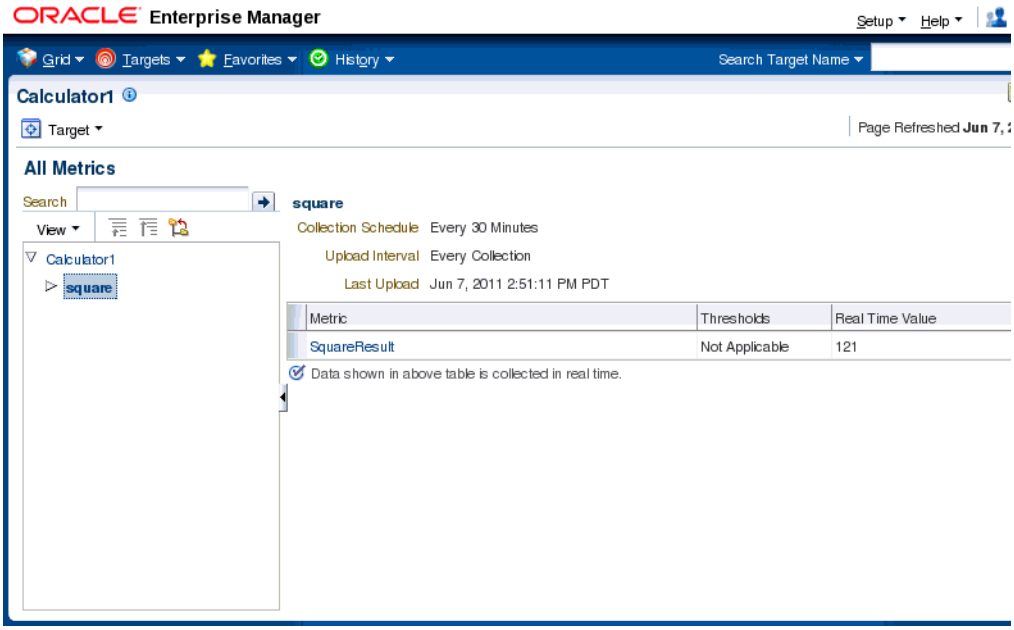

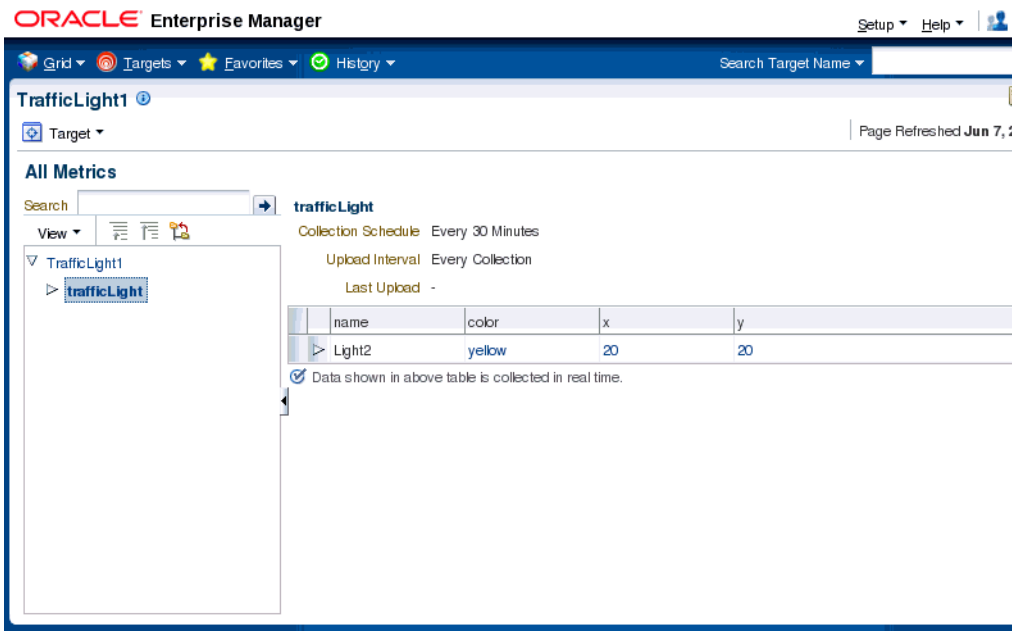

20.13 Creating JMX Metric Extensions

If you wish to collect metrics from your custom J2EE application deployed on Oracle Fusion middleware and exposed via JMX attributes into Enterprise Manager 12c, you can use either the Enterprise Manager console or the `jmxcli` command line tool. The latter also supports defining Metric Extensions from JMX operations and supports the creation of a Metric Extension Archive (MEA) which then needs to be imported into the OMS via the console and then tested and deployed to the desired J2EE application target instances representing your custom application.

Note: While you can select attributes that are not open types using the Mbean browser, the JMX metric extension UI supports only open type attributes. An error will occur if the UI is used to create metrics by selecting attributes which are not open types.

20.13.1 Using the Enterprise Manager Console

1. From the **Enterprise** menu, choose **Monitoring** and the **Metric Extensions**. The Metric Extension page displays.

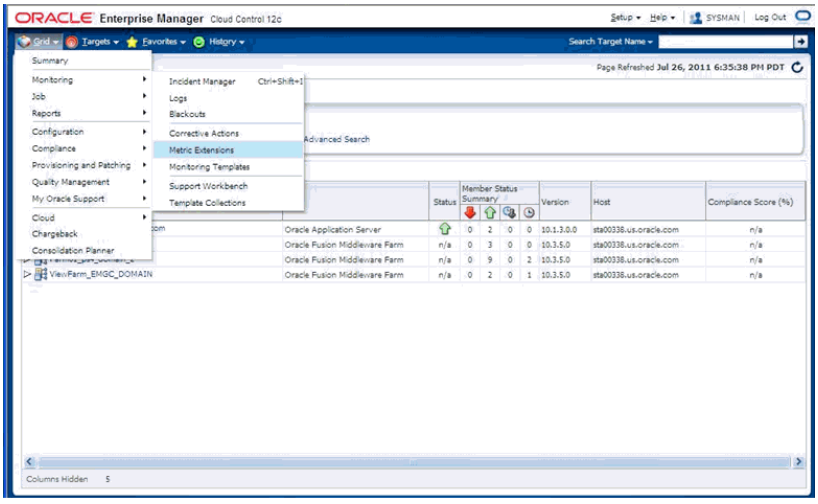

2. Click on the **Create** button to create a new Metric Extension.

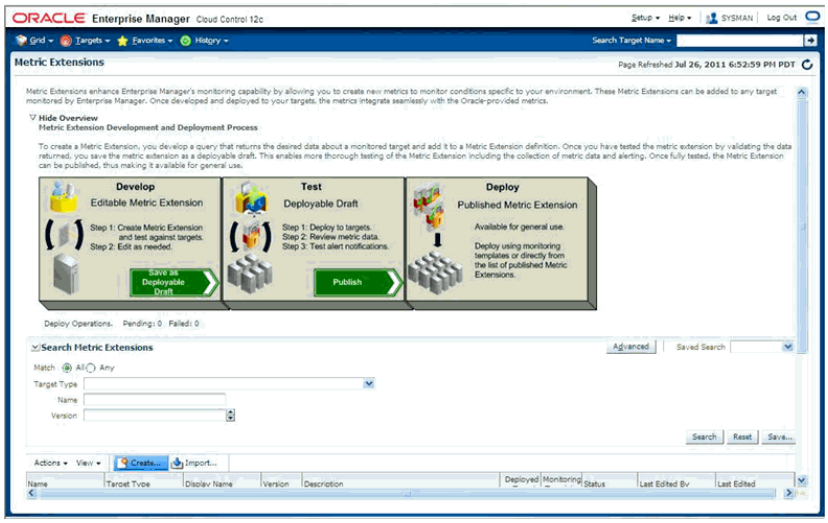

3. Select "Application Deployment" target type (or any other appropriate Enterprise Manager target type for which this metric needs to be defined) and specify a meaningful name for your metric extension. Please keep in mind that you may eventually end up creating additional metric extensions on the "Application Deployment" target type both for this application and for other custom applications so it is desirable to capture both the metric name and the application name in the metric extension name, whenever possible.

   Also select JMX for the Adapter.

   Please note the "Collection Schedule" section below the "General Properties" section. This is where you define how often this metric is to be collected, or if this is realtime-only metric (in which case the **Disabled** button should be selected.).

   If "Alerting and Historical Trending" is selected, you can also select an **Upload Interval**, which indicates which samples (whose frequency is specified in the "Repeat Every" field) are uploaded to the Enterprise Manager repository for historical trending. For example if Collection frequency is specified as 15 minutes and the Upload Interval is 3, then every 3rd sample will be uploaded into the

repository (every 45 minutes) and will be available for historical trending. However "alerts" that are possibly triggered due to threshold violations will be available for every collection (15 minutes)

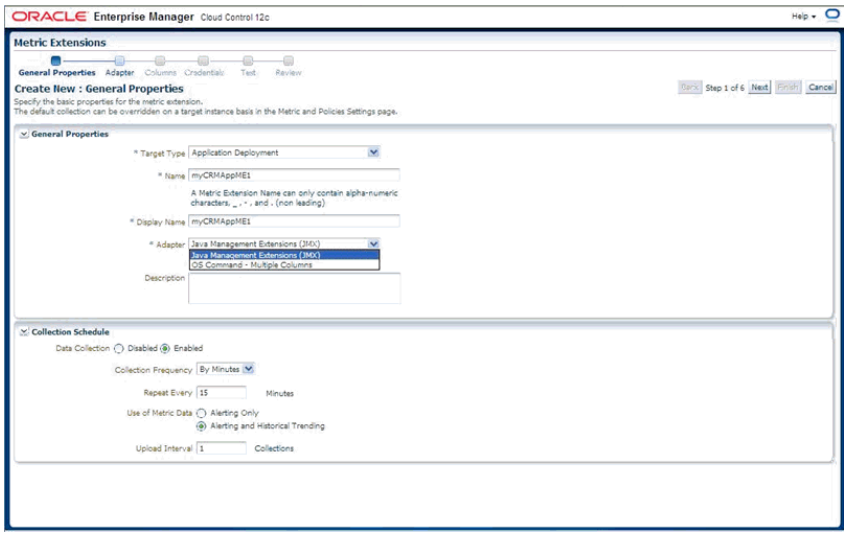

4. Click **Next** and specify the required properties needed for a JMX-based metric. These are defined in the **Basic Properties** section and are:

   - Metric: The Mbean ObjectName or Pattern and
   - Column Order: A semi-colon separated list of JMX attributes for above Mbean (if a metric needs to be defined using a JMX operation, use the `jmxcli` as shown in a following section)

   Please note that the Mbean ObjectName or pattern defined above must not have any server-specific key properties defined. These properties may be replaced with a wildcard ("*").

   For example, if an Mbean object name is *com.bea:Type=foo,Location=Server1,Name=abc* then it may be appropriate to define this as *com.bea:Type=foo,** in the "Metric" property described above. Also, if the Mbean ObjectName is a pattern, please be aware that multiple Mbeans could be returned making this metric a "table" with multiple rows (each row representing the JMX attributes of an Mbean matching the ObjectName pattern). In this case we need to define at least one or more columns as Key columns so that each row is unique in the resultant metric.

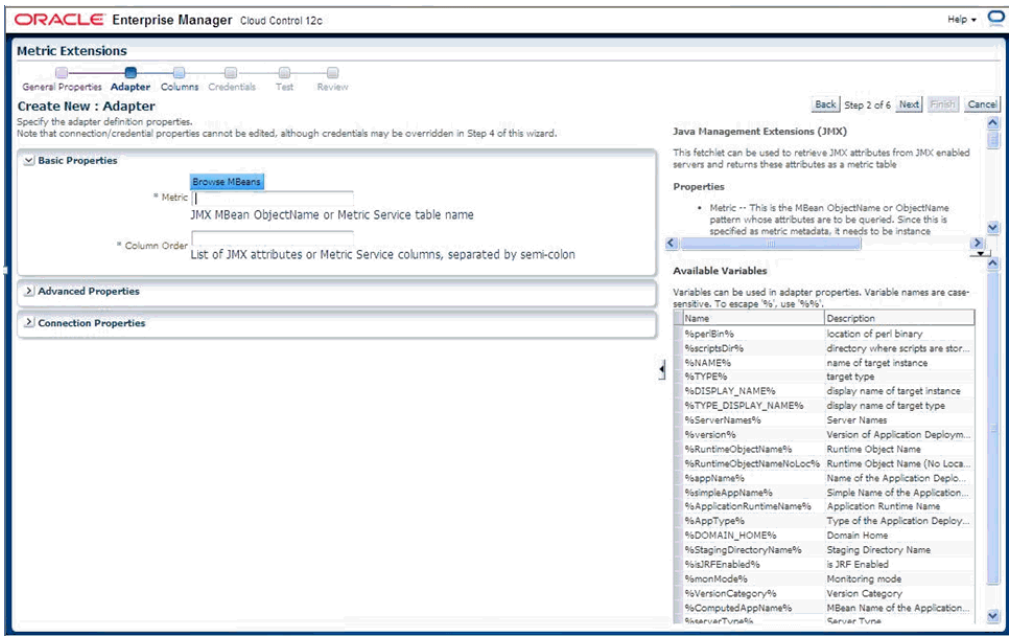

For above step there is a "Browse Mbeans" button that makes it easier to configure these two properties by allowing you to browse an MbeanServer and selecting an Mbean and its JMX attributes that need to be represented by this metric being defined in the metric extension.

If you click on the "Browse Mbeans" button, you must perform he following in sequence.

- Select the Target: Select an instance of the target type that you need to use to define this metric. This target instance is used to help configure the metric and does not have to be the target instance on which the metric is eventually defined.
- Enter the Mbean Pattern: Here, you enter an Mbean Object Name or pattern for the Mbean you are interested in monitoring
- Click on the "List Mbeans" button: This will be displayed in the table under "Select Mbean and Attributes", the Mbeans that match the Mbean pattern or the text "No Mbean listed or none matches Mbean Pattern" if there is no match. You can iteratively update the previous "Enter Mbean pattern" field and click the "List Mbeans" button to refine the list of Mbeans displayed.

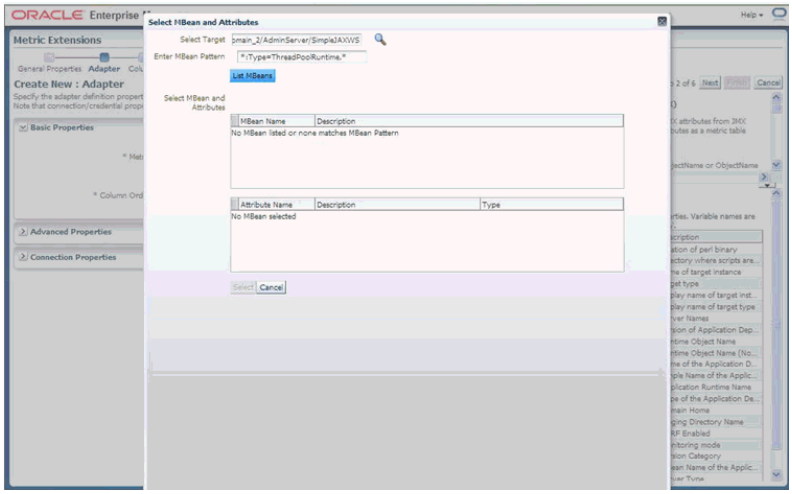

- Select an Mbean of interest: This will automatically populate the table below with the JMX attributes for the selected Mbean.

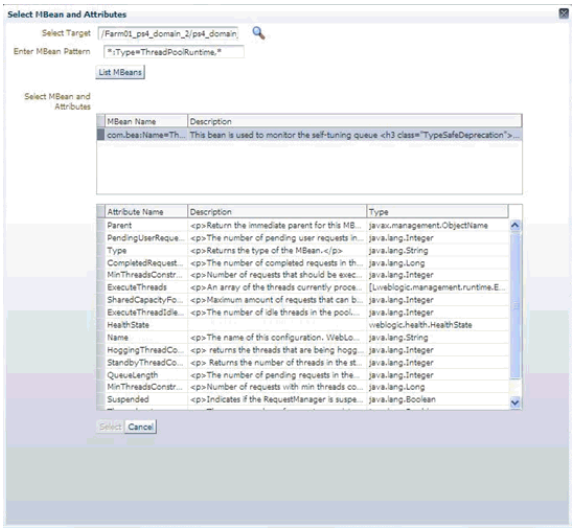

You can multi-select (using <Crtl> click) multiple attributes and then click **Select** to accept the selections.

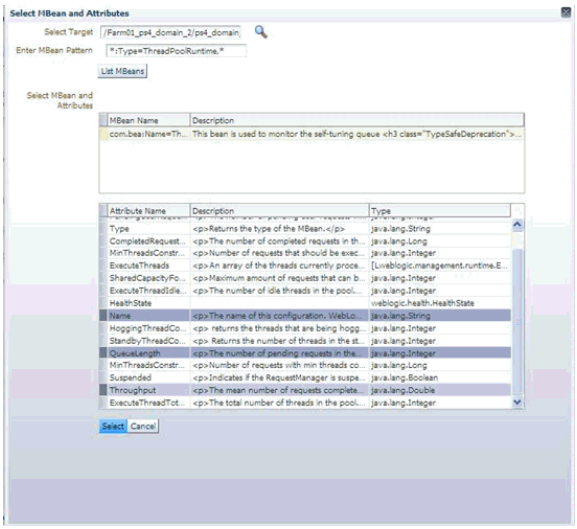

5. You must now specify the required parameters "Metric" and "Column Order" needed to define a JMX based metric extension.

   Please note the Mbean name populated in the "Metric" field should not have any instance specific information in its key properties (like Location=Server1 or ServerName=foo) if this metric extension can be applied to multiple servers besides the one that was selected/used to configure the metric extension using the "Browse Mbean" wizard above. These instance-specific key properties could be replaced with a wildcard "*" as appropriate to make this a valid Mbean ObjectName pattern.

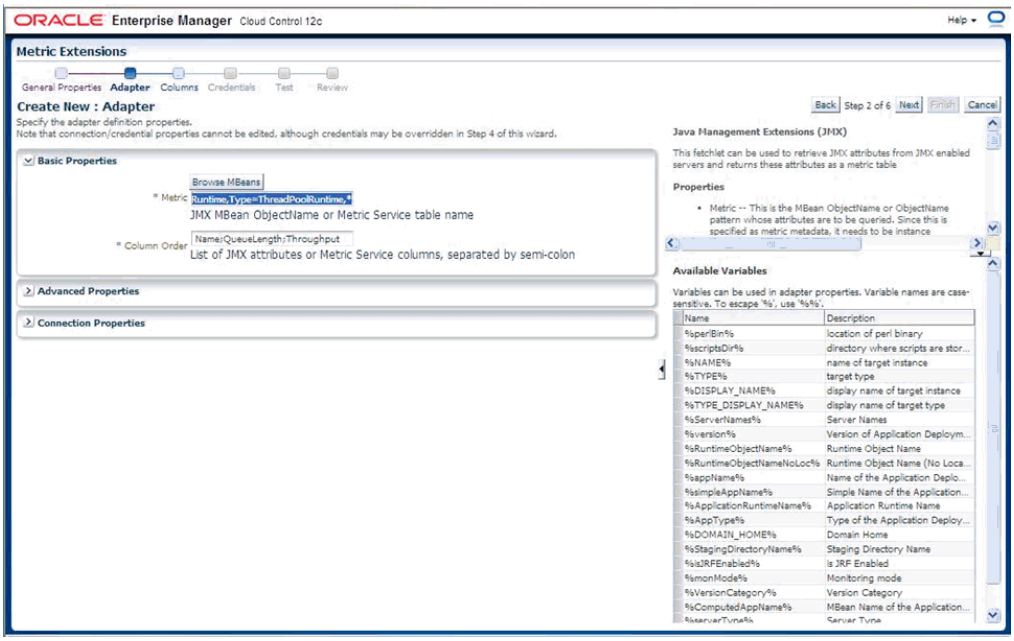

Explanation of Specifiable Properties

Required Properties:

- metric -- This is the MBean ObjectName or ObjectName pattern whose attributes are to be queried. Since this is specified as metric metadata, it needs

to be instance agnostic so instance specific key-properties if any (like servername), on the MBean ObjectName may need to be replaced with wildcards.

- columnOrder -- This is a semi colon separated list of JMX attributes in the order they need to be presented in the metric

Advance Properties:

- identityCol -- This is an MBean key property that needs to be surfaced as a column when it is not available as a JMX attribute. For example: com.myCompany:Name=myName,Dept=deptName, prop1=prop1Val, prop2=prop2Val In above case setting identityCol as Name;Dept (note that separator is a semi-colon) will result in two additional key columns representing Name and Dept besides the columns representing the JMX attributes specified in the columnOrder property above.
- autoRowId -- This is the prefix used for an automatically generated row in case the MBean ObjectName pattern specified in metric property matches multiple MBeans and none of the JMX attributes specified in the columnOrder are unique for each. The autoRowId value specified here will be used as a prefix for the additional key column created. For example, if the metric is defined as com.myCompany:Type=CustomerOrder,* columnOrder is CustomerName;OrderNumber;DateShipped (and assuming customerName;OrderNumber;DateShipped may not be unique if an order is shipped in two parts).

  Setting autoRowId as "ShipItem-" will populate an additional key column for the metric for each row with ShipItem-0, ShipItem-1.
- MetricService -- True/False indicates whether MetricService is enabled on the target WebLogic domain. This would be unchecked or false in most cases for user-defined metrics except when metrics that are exposed via the Oracle DMS MBean needs to be collected. If set to true, then the basic property "metric" above should represent the MetricService table name and the basic property "columnOrder" will represent a semicolon separated list of column names for aforementioned MetricService table.

6. Specify the Columns for this metric (if you have used the "Browse Mbeans" step earlier, then these columns are automatically pre-filled for you). You may need to edit these pre-created columns by the "Browse Mbean" wizard to specify columns that are "Key columns". This done in the event an Mbean pattern is specified in the previous step for the "Metric" property, and multiple Mbeans could match this Mbean pattern for any of the target instances to which this metric extension will be applied to.

   If the order of the columns are changed (using Move Up - Move Down buttons) then the corresponding order of the semi-colon separated columns in the "Column Order" property in the previous step also needs to be updated accordingly (using the Back button).

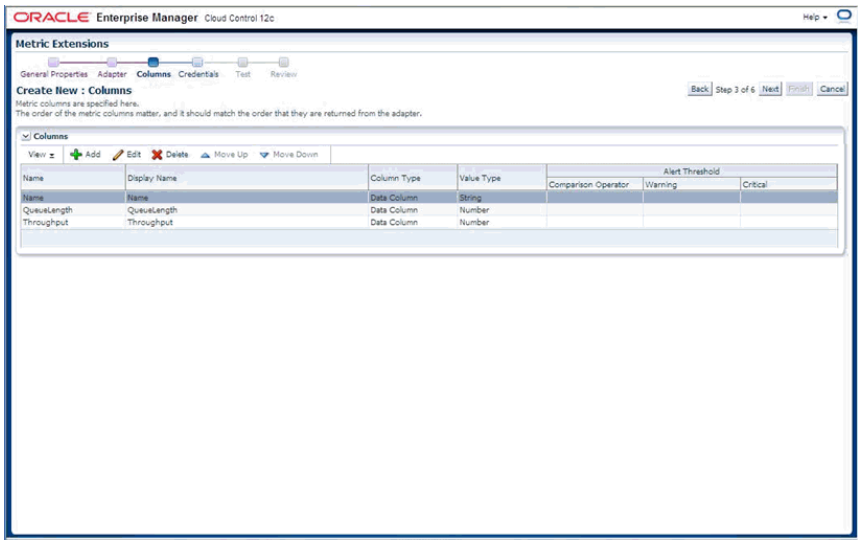

If needed, edit the columns as desired to make them a Key Column as shown in the following graphic.

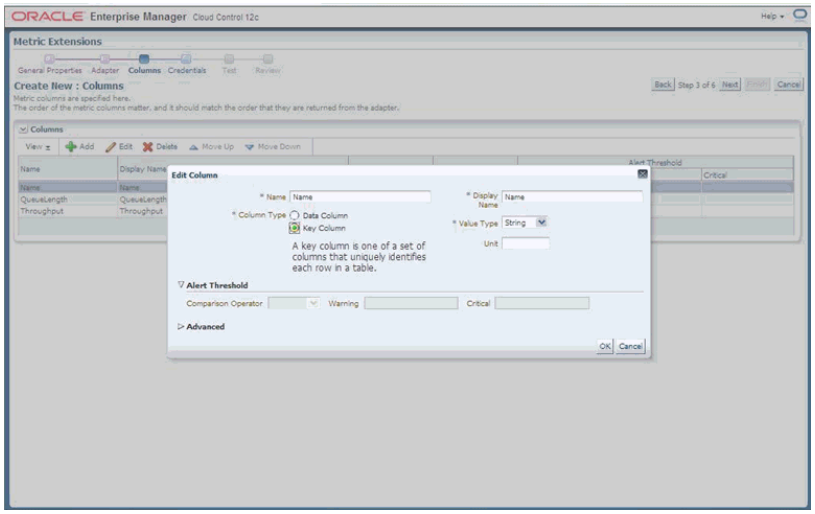

Once columns are labeled and edited, click the "Next" button. We are now ready to test the metric extension

7. Click on the Add button to select a target instance on which to test this metric extension. This could preferably be a different target instance than the one used to define the metric extension (if the Browse Mbeans button was used to help in defining the metric extension earlier).

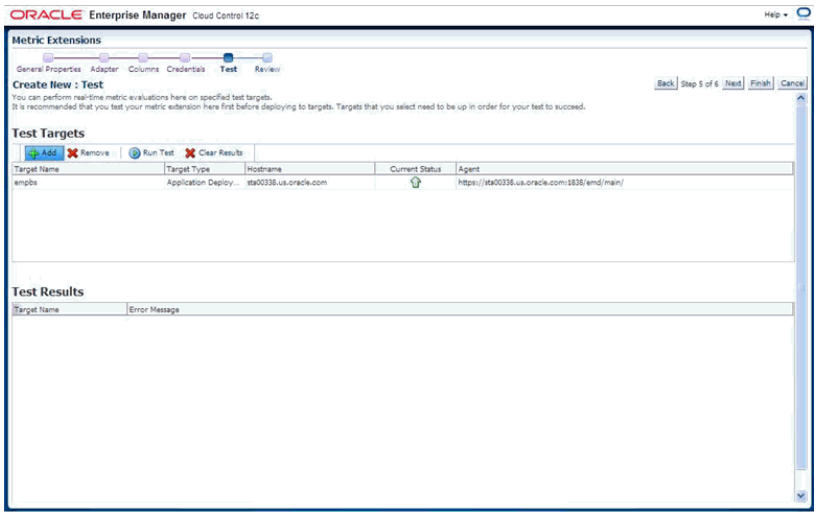

Now select a target instance in the Test Targets table and then click the "Run Test" button above that table.

The metric values are displayed in the "Test Results" table (if there are errors ,then those are also shown).

If errors are present, click the Back button and fix the errors and re-run the test.

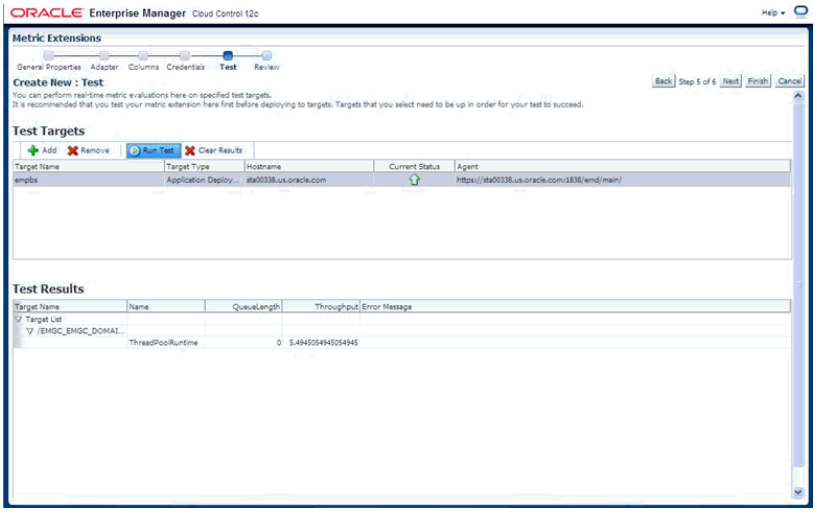

8. Once satisfied with the Test, click the Next button to view a summary of the metric extension and then click the **Finish** button to define the metric extension.

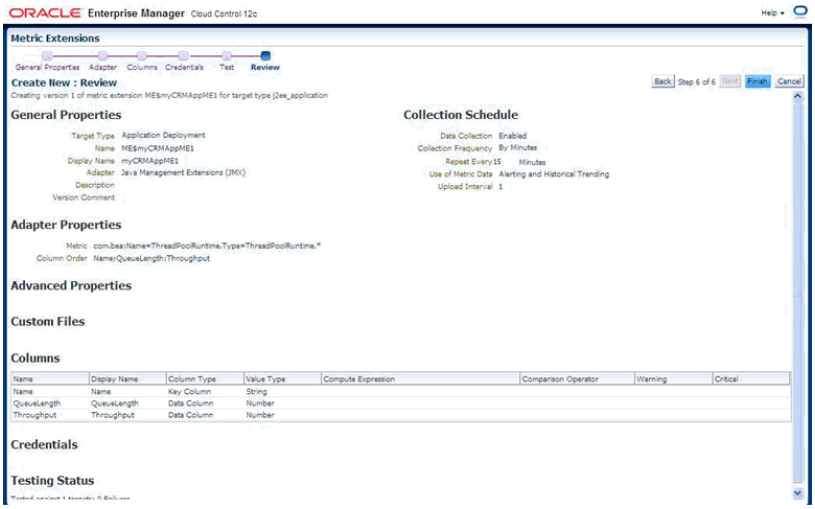

9. Before deploying the metric extension to selected target instances the metric extension needs to be saved as a "Deployable draft". This will let the metric extension designer deploy the metric extension to selected target instances and verify the metric collection but will prevent other administrators from deploying this metric extension until after it has been tested and the designer is satisfied.

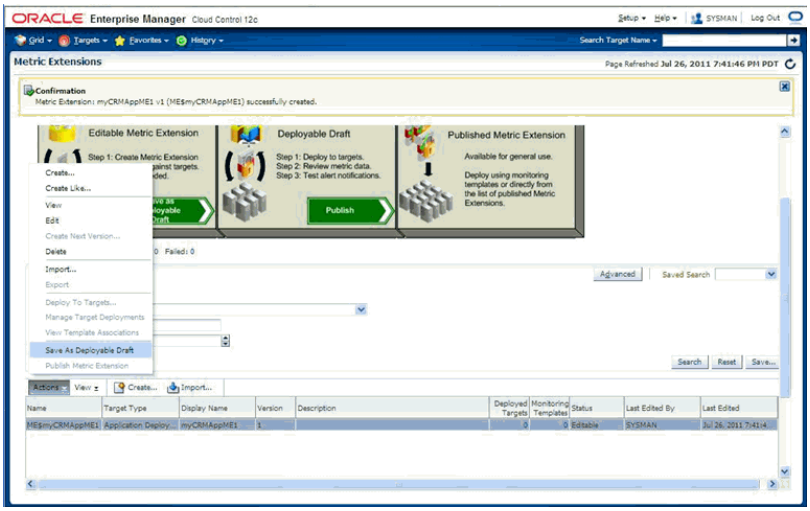

10. Select the metric extension just created and saved as a deployable draft . From the Acitons menu, choose Deploy to Targets.

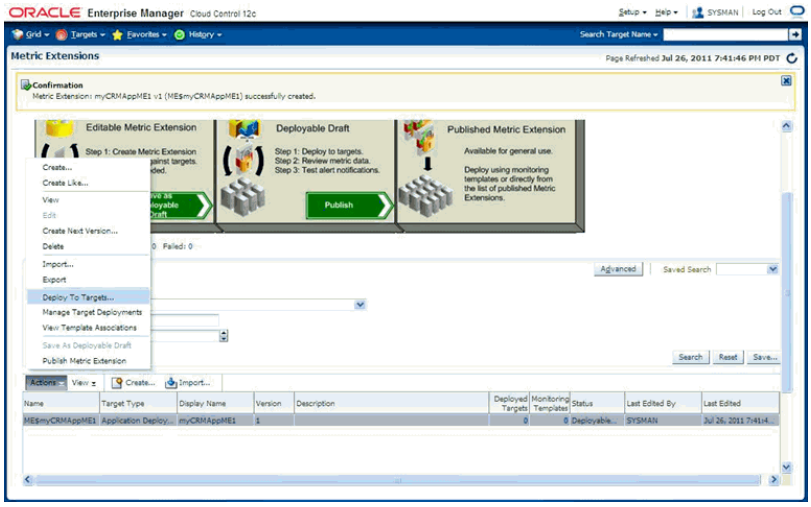

11. Select the target instances that this metric extension needs to be tested on and click **Submit**. For example, if the metric extension is defined on an "Application Deployment" target type and represents a metric from a Custom Mbean registered by a custom JEE application, the instances of that custom application could be selected. This will schedule a job to asynchronously deploy the metric extension to the Agents monitoring the selected targets.

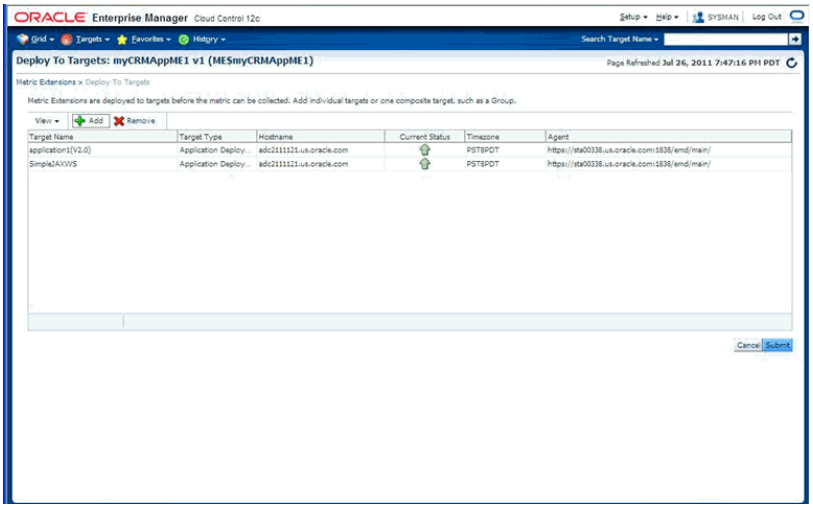

12. Monitor the status of the Pending deploy operation of the metric extension to selected targets by refreshing this page periodically to monitor the Status column and Failed Deploy Operations table for any possible errors during deployment.

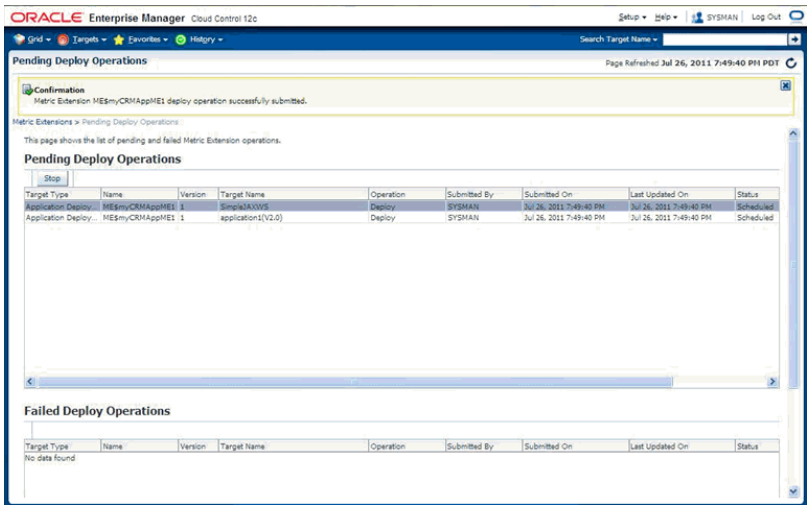

13. From the Enterprise menu, choose Monitoring and then Metric Extension. On the Metric Extension home page, your metric extension appears as a row in the table with a column "Deployed Targets" representing the count of the number of targets this metric extension is deployed to.

    From the **Actions** menu, choose **Manage Target Deployments** from the table after selecting the desired metric extension. This will list the target instances this metric extension is deployed to.

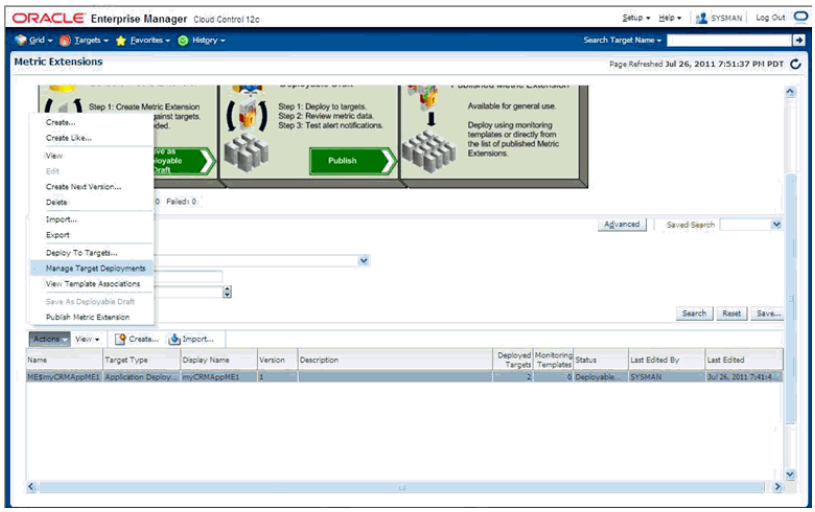

14. Click on the value in the "Target Name" column for the target instance you want to verify the metric extension on. This takes you to the home page of the target.

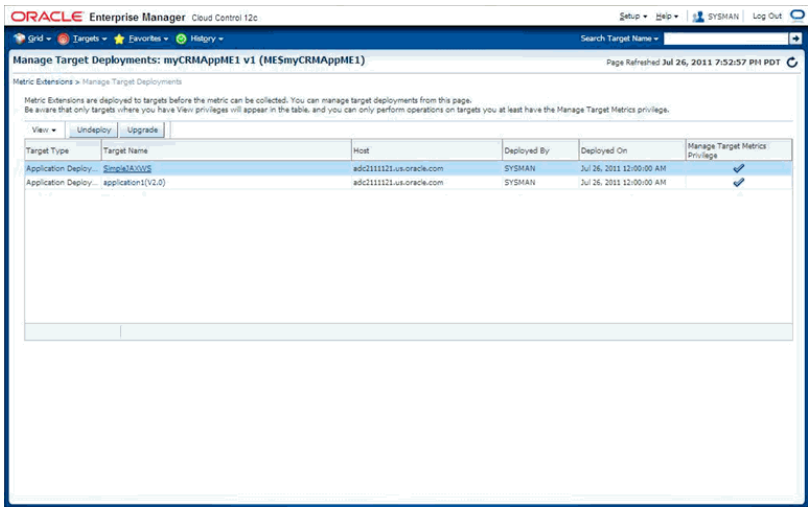

15. For middleware targets, navigate to the <Target Type>/Monitoring/ Performance Summary (or in general to the <Target Type> /Monitoring /All Metrics) page.

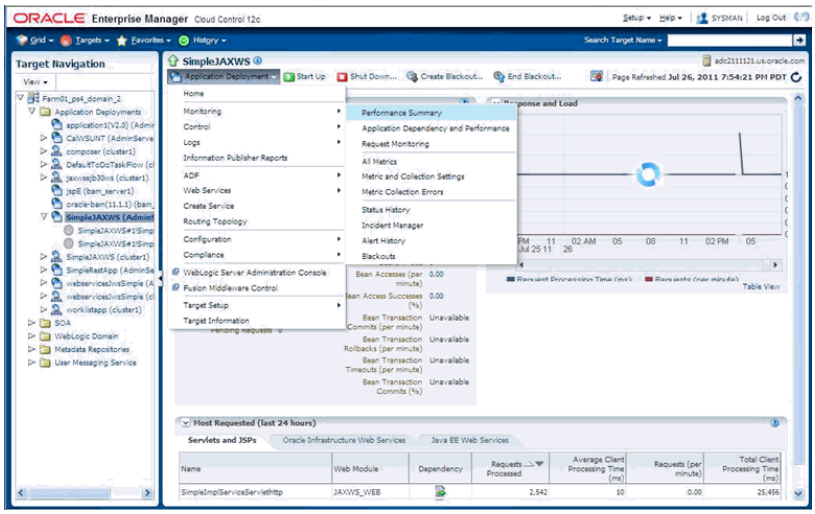

16. From the Performance Summary page, the newly created metric will be visible on the Metric Palette and can be selected and charted on the page.

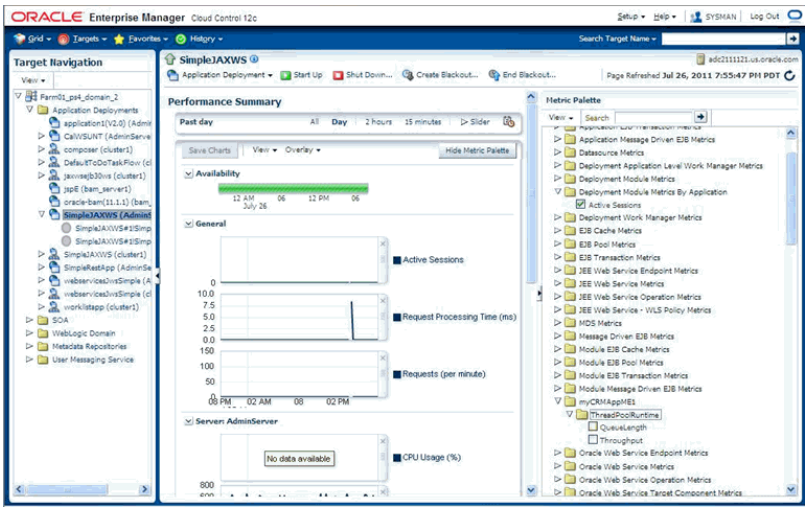

17. Once satisfied with testing the metric extension on one or more target instances, the metric extension can be published from the Metric Extension page (from the Actions menu, choose Publish Metric Extensions) and then deployed to remaining target instances.

20.13.2 Using the JMXCLI to create a Metric Extension Archive

If you do not wish to use the Enterprise Manager console (or do not want to surface an Enterprise Manager metric exposed via a JMX operation), you can use the command line tool JMXCLI to create a Metric Extension Archive. This can then be imported into the OMS, edited, tested, published and then deployed to desired instances of the target type on which it is defined. The following illustrates the use of the jmxcli in creating a Metric Extension archive.

1. cd <Agent Instance Home>/bin
2. setenv USER_JARS $T_WORK/middleware/wlserver_10.3/server/lib/weblogic.jar (this should not be necessary if your Mbeans just return JMX Open Types and not any custom classes).
3. emctl jmxcli -t WebLogic -MEXT -l "service:jmx:t3://sta00338:7018/jndi/weblogic.management.mbeanservers.runtime" -u weblogic -c welcome1 -m "*:Type=ThreadPoolRuntime,*" -w /scratch/TEMP/

   Options:

   -l : JMX serviceURL to connect to the WebLogic server. Replace the host:port above with what is appropriate for your instance

   -u : WebLogic user having access to required MBeans

   -c : Password for the WebLogic user

   -m : Mbean ObjectName or pattern.

   -w : Temporary work directory where the Metric Extension Archive (which can later be imported into the OMS console) is created.

```
Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.0.0
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
Using Plugin Root /ade/sparmesw_
egc802/oracle/emagent/gcagent/plugins/oracle.sysman.emas.agent.plugin_
```

```
12.1.0.0.0
Connecting to server:
service:jmx:t3://sta00338:7018/jndi/weblogic.management.mbeanservers.runtime
Connecting as user: weblogic
Obtained 1 MBeans matching pattern *:Type=ThreadPoolRuntime,*.
Enter an existing target type for this Metric Extension: [j2ee_application]
Enter the name of the Metric Extension: [myMEXT] myAppME_1
Enter the Metric Extension version: [1.0]
Enter the Metric Extension metadata file location: [./metadata/ME#24#myAppME_
1.xml]
Enter the Metric Extension collection file location:
[./collection/ME#24#myAppME_1.xml]
Enter a label for this Metric Extension: [myAppME_1]
Enter a description for this Metric Extension: [myAppME_1]
The available targets are:
0: This bean is used to monitor the self-tuning queue  <h3
class="TypeSafeDeprecation">Deprecation of MBeanHome and Type-Safe
Interfaces</h3>  <p class="TypeSafeDeprecation">This is
a type-safe interface for a WebLogic Server MBean, which you can import into
your client classes and access through
<code>weblogic.management.MBeanHome</code>. As of 9.0, the
<code>MBeanHome</code> interface and all type-safe interfaces for WebLogic
Server MBeans are deprecated. Instead, client classes that interact with
WebLogic Server MBeans should use standard JMX design patterns in which clients
use the <code>javax.management.MBeanServerConnection</code> interface to
discover MBeans, attributes, and attribute types at runtime. For more
information, see "Developing Manageable Applications with JMX" on <a
href="http://www.oracle.com/technology/products/weblogic/index.html"
shape="rect">http://www.oracle.com/technology/products/weblogic/index.html</a>.
</p>
        (com.bea:Name=ThreadPoolRuntime,ServerRuntime=EMGC_
ADMINSERVER,Type=ThreadPoolRuntime)
Enter the index of target/MBean you wish to monitor or press <Ctrl-C> to quit:
0
Following metric source types are available for selected target(s):
        0: JMX Attributes
        1: JMX Operations
Enter the index of your choice or press <Ctrl-C> to quit: 0
Attributes are:
        0: CompletedRequestCount         Return Value: java.lang.Long
        1: ExecuteThreadIdleCount        Return Value: java.lang.Integer
        2: ExecuteThreads        Return Value:
[Lweblogic.management.runtime.ExecuteThread;
        3: ExecuteThreadTotalCount       Return Value: java.lang.Integer
        4: HealthState   Return Value: weblogic.health.HealthState
        5: HoggingThreadCount    Return Value: java.lang.Integer
        6: MinThreadsConstraintsCompleted        Return Value: java.lang.Long
        7: MinThreadsConstraintsPending  Return Value: java.lang.Integer
        8: Name  Return Value: java.lang.String
        9: Parent        Return Value: javax.management.ObjectName
        10: PendingUserRequestCount      Return Value: java.lang.Integer
        11: QueueLength  Return Value: java.lang.Integer
        12: SharedCapacityForWorkManagers        Return Value:
java.lang.Integer
        13: StandbyThreadCount   Return Value: java.lang.Integer
        14: Suspended    Return Value: java.lang.Boolean
        15: Throughput   Return Value: java.lang.Double
        16: Type         Return Value: java.lang.String
Select one or more items as comma separated indices: 5,13
Number of possible columns in the resultant metric are 2.
```

```
Enter the name for this metric column at index=0 : [HoggingThreadCount]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [HoggingThreadCount]
Enter the NLSID for column: [HoggingThreadCount]
Enter the UNIT for column "HoggingThreadCount": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]
Enter the name for this metric column at index=1 : [StandbyThreadCount]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [StandbyThreadCount]
Enter the NLSID for column: [StandbyThreadCount]
Enter the UNIT for column "StandbyThreadCount": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]
Do you want periodic collection for this metric <y/n>? [n] y
Enter the collection interval in seconds: 300
Periodic collection interval is: 300 seconds.
Written the metadata xml file: ./metadata/ME#24#myAppME_1.xml.
Creating new file: ./collection/ME#24#myAppME_1.xml.
Updated the default collection file for j2ee_application at location
./collection/ME#24#myA
ppME_1.xml.
createMextArchive: Adding metadata
createMextArchive: Adding collection file
createMextArchive: Adding mea.xml file

Creating Metric Extension zip archive: ./MEA_ME$myAppME_1.zip
Please import this into Enterprise Manager Cloud Control using the console.
Exiting...
```

The above session creates a zip file MEA_ME$myAppME_1.zip in the directory specified by the -w option when `jmxcli` is invoked (or in current directory if -w is not specified).

Import this into the Enterprise Manager console as shown below. From the Enterprise menu, choose Monitoring and then Metric Extensions to access the Metric Extensions home page.

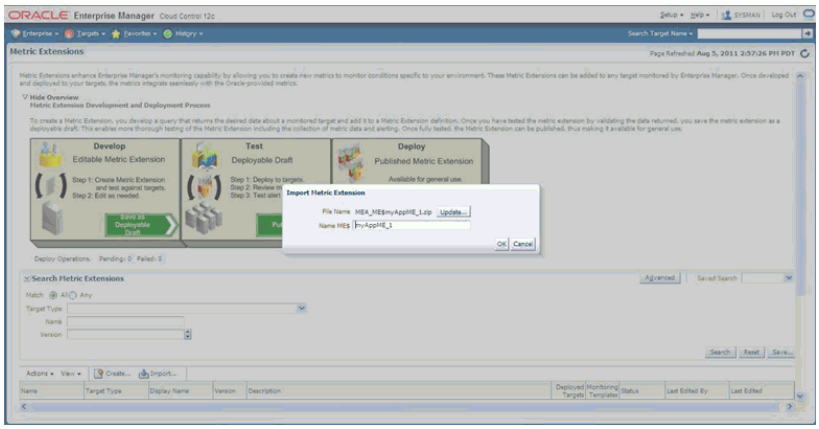

After the Management Extension Archive is imported as shown above, it can be edited (and modified), tested, published and deployed.

20.14 Surfacing Metrics from a Standalone JVM or Oracle Coherence

Users can also use the mechanism outlined in previous section to create additional metrics that are not available out-of- box for Oracle Coherence or JVM targets and the data for which are available via JMX Mbean attributes.

20.14.1 Using the Enterprise Manager Console

The procedure is similar to the ones followed in previous section for extending metrics on j2ee_application target types except that you must select target type "JVM" or "Oracle Coherence xxx" in Step 3 for defining the Metric Extension on JVM or Oracle Coherence target types.

20.14.2 Using JMXCLI

The steps are similar to those for using JMXCLI to define a Metric Extension Archive for custom J2EE applications except that the start-up arguments when `jmxcli` is invoked as follows:

```
emctl jmxcli -t JVM -MEXT  -h adc2180736 -p 6789 -m "*:*"  -w /scratch/TEMP/
```

You must specify target type on which the Metric Extension is defined to be JVM or oracle_coherence as appropriate (instead of the default j2ee_application).

21

Configuring Services

This chapter describes how to configure services in Oracle Enterprise Manager 10*g* Grid Control Console. It contains the following sections:

21.1 Summary of Service Management Tasks

This table provides a summary list of all the service management features and their requirements.

Table 21–1 Summary of Service Management Tasks

Feature	Description	Requirements	Refer to
Test Performance	This feature allows you to proactively monitor services using service tests or synthetic transactions and determine their performance and availability from different user locations using beacons. For Web transactions, you can monitor the transactions at the transaction, step group and step level.	▪ Management Agent for enabling a beacon. ▪ Microsoft Internet Explorer 5.5 or later	Configuring a Service
Root Cause Analysis	The Root Cause Analysis (RCA) feature provides you with the ability to analyze and determine possible causes of service failure. The Topology Viewer provides a graphical representation of the service and its relationship to other services, systems and infrastructure components, with the causes identified by RCA highlighted in the display.	For the Topology Viewer ▪ Microsoft Internet Explorer 5.5 or higher ▪ Adobe SVG Viewer 3.0	Root Cause Analysis Configuration

21.2 Setting up the System

A system is the set of infrastructure components, for example hosts, databases, and application servers that work together to host your applications. Before you create a service, you must specify the system that will be used to host your service. Refer to the Enterprise Manager Online Help for details on defining systems.

After you have selected the system, you must mark one or more components as key components that are critical to running your service. These key components are used to determine the availability of the service and identify possible causes of service failure for root cause analysis.

21.3 Creating a Service

Before you create a service, you must be familiar with the concepts of service management You must also perform the following tasks:

- Install the Management Agent on the hosts on which the components of your service have been installed.
- Discover all the components for your service so that they can be listed as Enterprise Manager targets.
- Define systems on which the service is to be hosted.

To create a service, click the **Targets** tab and **Services** subtab. The Services main page is displayed. Select a service from the Add drop-down list and click **Go**. The following screen is displayed:

Figure 21–1 Create Service - General Page

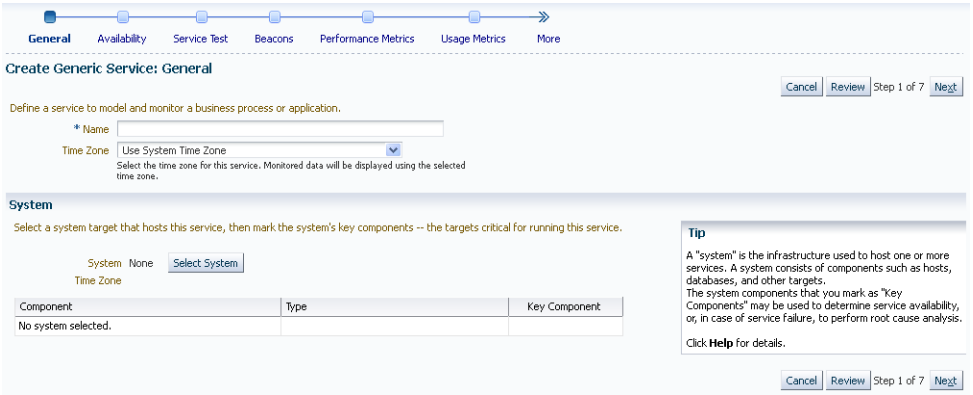

Follow the steps in the wizard to create your service. This involves the following:

- Identifying the type of service to be created. You can define different types of services based on your requirement. Some of the services that you can define are Generic Service, Web Application, Aggregate Service, and Forms Application. A Generic Service is used to monitor a variety of different protocol based services. A Web Application is used to monitor Web transactions. Enterprise Manager provides additional monitoring and diagnostics features for Web applications. A Forms Application is used to monitor Forms transactions. Each Forms transaction can consist of one or more actions that can be monitored. You can also define other services that are specific to an application such as the OCS Service. You can combine one or more services to form an Aggregate Service.
- Specifying the name and time zone for the service.
- Selecting a system target that hosts this service and then marking the system's key components that are critical for running the service. These key components are used to determine the availability of the service and identify possible causes of service failure. For more information on defining systems and monitoring them, refer to the Service Management chapter in *Oracle Enterprise Manager Concepts*.
- Setting up the availability definition for the service. This can be service test-based or system-based. If you select service test, the service's availability is based on the execution of the service test by the one or more key beacons. If availability is based on system, availability is based on the status of one or more key components of the system.
- Adding one or more beacons to monitor service tests. Click **Add** to add one or more beacons for monitoring the service. It is recommended that you use beacons that are strategically located in your key user communities in order for them to proactively test the availability of the service from those locations. If no beacons exist, click **Create** to create a new beacon.

Note: Beacons marked as key beacons will be used to determine the availability of the service. The service is available if one or more service tests can be successfully executed from at least one key beacon.

For Web applications, you can compare the performance of the service test execution from each remote beacon against the local beacon.

- Defining the metrics that will be used to measure the performance of the service. Performance metrics can be based on service tests or system components. After defining the metrics, you can specify the critical and warning thresholds. You can also specify the metric that is to be displayed in a graphical format on the Service Home page.
- Defining the metrics that will be used to measure the user demand for the service. Usage metrics can be based on one or more system components. After defining the metrics, you can specify the critical and warning thresholds. You can also specify the metric that is to be displayed in a graphical format on the Service Home page.

> **Note:** You can define usage metrics for system-based services only.

- After you have completed all the steps in the wizard, click **Finish** to create your service. Refer to the Enterprise Manager Online Help for more details on these pages.

21.4 Configuring a Service

After you have created the service, you can configure it further by selecting an option from the Monitoring Configuration page. To configure a service, select a service from the Services main page and click **Configure** to go to the Monitoring Configuration page. The following screen is displayed.

Figure 21–2 Monitoring Configuration Page

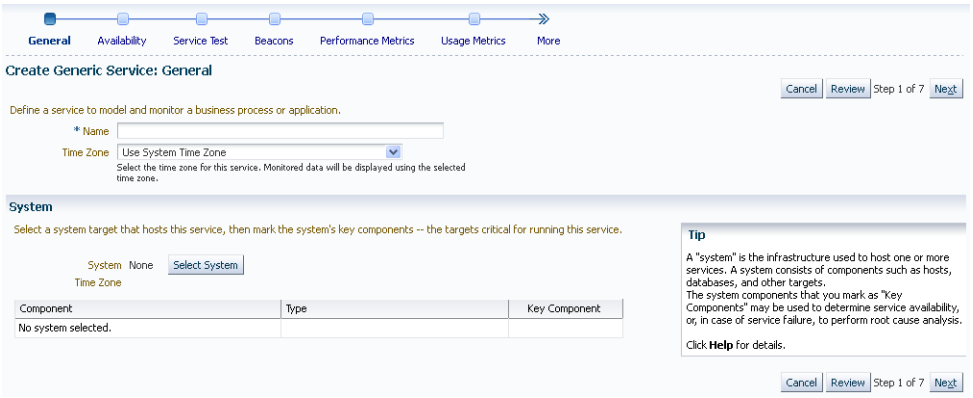

The following options are available:

- Availability Definition
- Performance Metrics
- Usage Metrics
- Business Metrics
- Service Tests and Beacons
- Root Cause Analysis Configuration

21.4.1 Availability Definition

You can modify the availability definition (service test-based or system-based) for the selected service. If availability is based on service tests, you can specify whether the service should be available when:

- All key service tests are successful (Default)
- At least one key service test is successful

> **Note:** A service test is considered available if it can be executed by at least one key beacon. If there are no key beacons, the service test will have an **unknown** status.

If availability is based on the key system components, you can specify whether the service should be available when:

- All key components are up (Default)
- At least one key component is up

You can also mark one or more components as key system components that will be used to compute the availability of the service. Key system components are used to determine the possible root cause of a service failure. For more information, refer to "Root Cause Analysis Configuration" on page 21-12.

You can also indicate whether the service test is a key test by enabling the Key Service Test checkbox. Only key service tests are used to compute the availability of the service. You can then select the beacons that will be used to execute the key tests and determine the availability of the service.

21.4.2 Performance Metrics

Performance metrics are used to measure the performance of the service. If a service test has been defined for this service, then the response time measurements as a result of executing that service test can be used as a basis for the service's performance metrics. Alternatively, performance metrics from the underlying system components can also be used to determine the performance of the service. You can do the following:

- Add a performance metric for a service test. After selecting a metric, you can choose to:
  - Use the metric values from one beacon. Choose this option if you want the performance of the service to be based on the performance of one specific location.
  - Aggregate the metric across multiple beacons. Choose this option if you want to consider the performance from different locations. If you choose this option, you need to select the appropriate aggregation function:

Table 21–2 Beacon Aggregation Functions

Function	Description
Maximum	The maximum value of the metric from data collected across all beacons will be used. Use this function if you want to measure the worst performance across all beacons.

Table 21–2 (Cont.) Beacon Aggregation Functions

Function	Description
Minimum	The minimum value of the metric from data collected across all beacons will be used. Use this function if you want to measure the best performance across all beacons.
Average	The average value of the metric will be used. Use this function if you want to measure the 'average performance' across all beacons.
Sum	The sum of the metric values will be calculated. Use this function if you want to measure the sum of all response times across each beacon.

Note: If you are configuring a Web transaction, you can specify the **Source** which can be transaction, step group or step. Based on this selection, the metric you add will be applicable at the transaction, step group, or step level.

- Add a performance metric for the underlying system components on which the service is hosted. After selecting a metric for a target, you can choose to:
  - Use the metric from a specific component. Choose this option if you want the performance of the service to be based on the performance of one specific system component.
  - Aggregate the metric across multiple components. Choose this option if you want to consider the performance from multiple components. If you choose this option, you need to select the appropriate aggregation function.

Table 21–3 System Aggregation Functions

Function	Description
Maximum	The maximum value of the metric across all components will be used as the value of this performance metric for the service.
Minimum	The minimum value of the metric across all components will be used as the value of this performance metric for the service.
Average	The average value of the metric across all components will be used.
Sum	The sum of values of metrics across all components will be calculated.

Note: When a system is deleted, performance metrics associated with the system will not be collected.

- Edit a performance metric that has been defined. For service test-based performance metrics, you can modify the beacon function that should be used to calculate the metric values. For system-based performance metrics, you can modify the target type, metric, and whether the aggregation function should be used. You can also modify the Critical and Warning thresholds for the metric.
- Delete a performance metric that has been defined.

21.4.3 Usage Metrics

Usage metrics are used to measure the user demand for the service. Usage metrics are collected based on the usage of the underlying system components on which the service is hosted. You can do the following:

- Add a usage metric. After selecting a metric for a target, you can choose to:
  - Use the metric from a specific component. Use this option if you want to monitor the usage of a specific component.
  - Aggregate the metric across multiple components. Use this option if you want to statistically calculate the usage across multiple components. If you choose this option, you need select the appropriate aggregation function.

Table 21–4 Aggregation Functions - Usage Metrics

Function	Description
Maximum	The maximum value of the metric across all components will be used as the value of this usage metric for the service.
Minimum	The minimum value of the metric across all components will be used as the value of this usage metric for the service.
Average	The average value of the metric across all components will be used.
Sum	The sum of the values of metrics across all components will be calculated.

- Edit a usage metric that has been defined.
- Delete a usage metric that has been defined.

21.4.4 Business Metrics

Business metrics are used to measure the performance of business in an organization. These metrics are based on business indicators that can assess the business performance. You can define one or more system based metrics and specify critical and warning thresholds for these metrics. You can define business metrics for Generic Services and Aggregate Services.

> **Note:** This option is available only if one of the system components is a service and has business metrics associated with it.

You can do the following:

- Add a business metric. After selecting a metric for a target, you can choose to:
  - Use the metric from a specific component. Use this option if you want the business metric to be based on the performance of one specific system component
  - Aggregate the metric across multiple components. Use this option if you want to measure the business performance from multiple components. Select the appropriate aggregation function from the drop down list. If you choose this option, you need select the appropriate aggregation function.

Table 21–5 Aggregation Functions - Usage Metrics

Function	Description
Maximum	The maximum value of the metric across all components will be used as the value of this business metric for the service.
Minimum	The minimum value of the metric across all components will be used as the value of this business metric for the service.
Average	The average value of the metric across all components will be used.
Sum	The sum of the values of metrics across all components will be calculated.

- Edit a business metric that has been defined.
- Delete a business metric that has been defined.

You can define system based metrics only. You can configure non-system based metrics by using the Data Exchange feature which facilitates data transfer between Enterprise Manager Cloud Control and other external monitoring systems. For details, refer to the *Oracle Enterprise Manager Integration Guide*.

21.4.5 Service Tests and Beacons

You can add additional service tests and specify one or more beacons that will execute these service tests. To add a service test or modify an existing service test, click the **Service Test and Beacons** link on the Monitoring Configuration page. The Service Tests and Beacons page appears. You can do the following:

- Add one or more service tests for your service. Select the Test Type and click **Add**. Some of the test types that can be defined are ATS, FTP, Web Transaction, DNS, SOAP and others. The Create Service Test page is displayed. Refer to the Enterprise Manager Online Help for details on the various types of service tests.

  Note: While defining a SOAP (Simple Object Access Protocol) service test, if the WSDL URL to be accessed is outside the company's intranet, proxy settings need to be added to the `$OMS_HOME/sysman/config/emoms.properties` file.

  For example, to set up `www-proxy.us.oracle.com` as proxy, specify the values as follows:

  ```
  proxyHost=www-proxy.us.oracle.com
  proxyPort=80
  dontProxyFor=us.oracle.com,oraclecorp.com
  ```

  The `proxyUser`, `proxyPwd`, `proxyRealm`, and `proxyPropsEncrypted` properties are used to configure an authenticated proxy. After you have modified the proxy settings, you must restart the Oracle Management Service for the changes to be effective.

- After you have created the service test, you must enable it. If your service test is not enabled, it will not be executed by any of the beacons. You can define one or more service tests as key tests. These key tests are used to monitor the availability and performance of your service. Only service tests that are enabled can be designated as key tests. To set up a service test as a key test, click the **Availability Definition** link at the bottom of the page.

- Create, add, or remove a beacon. When you identify the beacon locations, select locations on your internal network or on the Internet that are important to your e-business. These are typical locations where your end users are located. For example, if your business is hosted in Canada and you have customers in the United States, use a beacon installed on a host computer in the United States to measure the availability and performance of your applications.
- After you have created the service test, you can verify it by clicking **Verify Service Test**.

Note:

- The Forms Transaction test type has been deprecated in Enterprise Manager 12c. Forms transactions created in earlier releases can still be used but you cannot create new Forms Transaction test types. You must create a Generic Service target and create an ATS Transaction using OATS EBS/Forms Load test scripts. This ATS test type is used to monitor Oracle Forms applications.
- The Web Transaction test type is in maintenance mode only. To monitor Web applications, we recommend that you create an ATS load script and use the ATS Transaction test type to monitor Web applications.

For more details on creating different types of service tests, refer to the Enterprise Manager Online Help.

21.4.5.1 Creating an ATS Service Test Using OATS Load Script

You can use the Oracle Application Test Suite (OATS) to define an Openscript Transaction Service Test. This test is used to enable beacon application transaction monitoring using Openscript load testing scripts. Openscript is a component of OATS and provides advanced capabilities to record and play back various types of Web transactions, such as web/HTTP, Oracle EBS/Forms, Oracle Fusion/ADF, Siebel, Adobe Flex etc. For details on how to create ATS load testing scripts, please refer to *Oracle® Functional Testing OpenScript User's Guide*. For details on creating an ATS test type, refer to the Enterprise Manager Online Help.

21.4.5.2 Configuring the Beacons

This section lists additional beacon related configuration tasks.

- **Configuring SSL Certificates for the Beacon**: To configure SSL certificates for Web transaction and Port Checker service tests, follow the steps given below:
- **Configuring Dedicated Beacons**: Beacon functionality on an agent requires the the use of an internal Java VM. The use of a Java VM can increase the virtual memory size of the agent by several hundred megabytes. Because of memory constraints, it is preferable to create beacons only on agents that run on dedicated hosts. If you are running large numbers of tests (e.g., several hundred per minute) on a given beacon, you may also wish to install that beacon's agent on a dedicated host. To take full advantage of dedicated hardware, edit the agent's `$ORACLE_HOME/sysman/config/emd.properties` file. as follows:
  - Set the property, `ThreadPoolModel=LARGE`. This allows the agent to simultaneously run many threads.

  - Set the property, `useAllCPUs=TRUE`. This allows the agent to run on multiple CPUs simultaneously.
  - Append `-Xms512m -Xmx512m` to the `agentJavaDefines` property. This increases the Java VM heap size to 512 MB.

- **Configuring a Web Proxy for a Beacon**: Depending on your network configuration, the beacon may need to be configured to use a Web proxy. To configure the Web proxy for a beacon, search for the beacon in the All Targets page. Select the beacon you wish to configure and click **Configure**. Enter the properties for the Web proxy. For example, to set up `www-proxy.us.oracle.com` as the beacon's Web proxy, specify the values as the following:

```
Proxy Host: www-proxy.us.oracle.com
Proxy Port: 80
Don't use Proxy for: .us.oracle.com,.oraclecorp.com
```

> **Note:** You cannot play Siebel service tests and Web Transaction (Browser) service tests on the same machine.

21.4.5.3 Configuring Windows Beacons for Web Transaction (Browser) Playback

To run a Web Transaction (Browser) service test, you need beacons that are running on an 10.2.0.4 or later Management Agent on Windows. The beacon drives an Internet Explorer process. This process runs as the same user as the Management Agent service.

Verifying Web Transaction (Browser) test involves the following 3 steps:

1. Navigate to the **Service Tests and Beacons** page and select a Web Transaction (Browser) test from the list.
2. Click **Verify Test**. The Verify Service Test page is displayed.
3. Select a Windows beacon and click **Perform Test**.

One of the common problems that you may encounter is that the **Perform Test** does not respond immediately.

There may be several reasons for this delay. Complicated tests may take longer to run. However, the most probable cause for delayed response is when the Internet Explorer process from the beacon is waiting for manual confirmation, which is invisible when run as a process that does not interact with desktop.

You may need to change the browser settings on the beacon machine. These settings need to be changed for the Local Service account and are account specific. Therefore, any changes to the Internet Explorer process that was opened from the Start menu on the beacon machine, will not affect the Internet Explorer process instantiated from the beacon which runs in an invisible window. To make the Internet Explorer window instantiated from the beacon visible:

1. Login as administrator to the Windows machine on which the Management Agent is running.
2. From the Start menu, click **Run**, type `services.msc` and click **Enter**.
3. Find the Management Agent in the list of Windows services, e.g. `OracleServiceagent1`.
4. Right click the Management Agent and select **Properties**.

5. Click the **Log On** tab.
6. Click the **Select Allow service to interact with desktop** checkbox and click **OK**.
7. Right click the Management Agent and select **Stop**, and then select **Start**.

To check Internet Explorer on the Management Agent machine for any dialog confirmations. For example, SSL Certificates and security warnings.

1. Use the previous procedure to make the Internet Explorer process instantiated from the beacon visible.
2. Launch Enterprise Manager (from any machine), and navigate to the Service Tests and Beacons page of the corresponding service target.
3. Select the Web Transaction (Browser) service test, and click **Verify Service Test**.
4. From the Verify Service Test page, select the beacon running on the Windows, and click **Perform Test**.
5. If it is a SSL Certificates issue, From the Windows machine on which the Management Agent is running, you will see an Internet Explorer window open and a Security Alert with a **View Certificate** option is displayed.
6. Select the **Certificate Path** tab, click the root certificate, which should have a red cross next to the name, and click the **View Certificate** button.
7. Click **Install Certificate** and proceed with the **Certificate Import Wizard**. (Click Next and Yes for any prompts).

> **Note:** Other security warnings may also pause the Internet Explorer automation process. Typically, these security warnings have a check box that allow you disable the display of all future warning messages for all Web sites. These warnings may have already been dismissed on the machine where the transaction was recorded.

8. Once this manual step has been performed, the Internet Explorer process should be in auto-pilot mode until the service test is completed. The warning message will not be displayed when you play back the service test next time.
9. Click **Perform Test** again to make sure the entire service test is completed automatically the second time.

To make the Internet Explorer window instantiated from the beacon invisible, you can repeat the steps 1 to 5, uncheck the **Select Allow service to interact with desktop** checkbox and continue with step 7.

To configure the proxy setting for Web Transaction (Browser) service tests:

1. Make the Internet Explorer process instantiated from the beacon visible.
2. Launch Enterprise Manager (from any machine), and navigate to the Service Tests and Beacons page of the corresponding service target.
3. Select the Web Transaction (Browser) service test, and click **Verify Service Test**.
4. From the Verify Service Test page, select the beacon running on the Windows, and click **Perform Test**.
5. From the Windows machine where the Management Agent is running, you should see two Internet Explorer windows open. From either of the windows, select the **Tools > Internet Options**.

6. Click the **Connections** tab and then click **LAN Settings**, and make all relevant changes there. These changes apply to all service tests running on this beacon.
7. Close both the Internet Explorer windows.
8. Click **Perform Test** again to make sure the entire service test is completed automatically the second time.
9. Make the Internet Explorer process instantiated from the beacon invisible.

> **Note:** At any one time, each test run launches two Internet Explorer windows. One of the windows schedules the steps during playback. The other window actually shows the site being played back.

21.4.6 Root Cause Analysis Configuration

You can use Root Cause Analysis (RCA) to filter a set of events to determine the cause of a higher level system, service, or application problem. RCA can help you to eliminate apparent performance problems that may otherwise appear to be root causes but which are only side effects or symptoms of the actual root cause of the problem, allowing you to more quickly identify problem areas. You can view the RCA results on the Home page or Topology page of any service that is currently down. The Topology page allows you to see a graphical representation of the service, system and component dependencies with the targets highlighted that RCA has implicated as causing the service failure.

Before running RCA, you can choose to:

- Configure the tool to run automatically whenever a service fails.
- Disable RCA by changing the default Analysis Mode to Manual.
- Define component tests for the service and thresholds for individual tests.

To configure Root Cause Analysis, follow these steps:

1. From the Service Home page, click **Monitoring Configuration**.
2. From the Monitoring Configuration page, click **Root Cause Analysis Configuration**.
3. If the current mode is set to Automatic, click **Set Mode Manual** to disable RCA. If you choose to perform the analysis manually, you can perform the analysis from the Service home page at anytime by choosing **Perform Analysis** if the service is down. If the current mode is set for Manual, click **Set Mode Automatic** to enable RCA when the state of the service and its components change
4. Click the link in the **Component Tests** column of the table for the key component you want to manage. You can then manage component tests for the service on the Component Tests page by adding, removing, or editing tests. Refer to the Enterprise Manager Online Help for details on defining component tests.

> **Note:** When you disable RCA and set it back to automatic mode, RCA does not store the previous history results for you, thus providing no history for later reference.

21.4.6.1 Getting the Most From Root Cause Analysis

Root Cause Analysis (RCA) can provide you with great value by filtering through large amounts of data related to your services and identifying the most significant

events that have occurred that are affecting your service's availability. If you are constructing your own services to manage in Enterprise Manager it is important that the services are defined with some thought and planning in order to get the most out of RCA.

The first item to consider in getting the most from RCA is the set of dependencies that your service has on other services or system components. Be sure to identify all of the system components that your service utilizes in order to accomplish its task. If you omit a key component and the service fails, RCA will not be able to identify that component as a possible cause. Conversely, if you include components in the service definition that the service does not actually depend on, RCA may erroneously identify the component as a cause of service failures.

When building service dependencies, keep in mind that you can take advantage of the aggregate service concept that is supported by Enterprise Manager. This allows you to break your service into smaller sub-services, each with its own set of dependencies.

Your services may be easier to manage in the modular fashion, and RCA will consider not only the status of a sub-service (a service that you depend on) but also on any of the system components or service that the sub-service depends on in turn and provides you with the power to encapsulate the services a key component exposes to you in the form of a managed service that your service may then depend on.

The second item to consider in getting the most from RCA is the use of component tests. As you define the system components that your service depends on, consider that there may be aspects of these components that may result in your service failure without the component itself failing. Component tests allow RCA to test the status not only of the target itself but also the status of its key aspects.

The RCA system allows you to create component tests based on any metric that is available for the key component. Remember, this includes any user-defined metrics that you have created for the component, allowing you great flexibility in how RCA tests the aspects of that component should your service fail.

21.5 Recording Web Transactions

You can record a transaction using an intuitive playback recorder that automatically records a series of user actions and navigation paths. You can play back transactions interactively, know whether it is internal or external to the data center, and understand the in-depth break-out of response times across all tiers of the Web application for quick diagnosis.

You must install the transaction recorder in your computer to record transactions. The transaction recorder is also used for playing back and tracing transactions. The transaction recorder is downloaded from the Enterprise Manager Cloud Control server the first time any of these actions is performed. The transaction recorder requires some Microsoft libraries to be installed in your computer. If these libraries are not present during installation, they are automatically downloaded and installed from the Microsoft site. Make sure that your computer has access to the Internet to download these files. After the installation has been completed, you may need to restart your computer to make the changes effective.

21.6 Monitoring Settings

For each service, you can define the frequency (which determines how often the service will be triggered against your application) and the performance thresholds. When a service exceeds its performance thresholds, an alert is generated.

To define metrics and thresholds, click **Monitoring Settings for Tests** link on the Service Tests and Beacons page. The Metric and Policy Settings page is displayed. Click the **Monitoring Settings** link. The Monitoring Settings - Thresholds page appears.

- **View By Metric, Beacon** - In this view, you can click **Add Beacon Overrides** to override the default threshold values for one or more beacons. In this case, the default thresholds will only be used for beacons without any overrides. Any new beacons added to the service will use the default thresholds. Click **Add Metric** to add one or more metrics.
- **View By Beacon, Metric** - In this view, you can click on the **Default** icon to toggle between Edit and View modes for a specific metric. In the Edit mode, you can modify the parameters of the metric. You can also modify the parameters of the metric for a specific beacon. In the View mode, the default parameters of the metric will be used.

  Apart from these procedures, you can also define metrics at the step, and step group level for Web transactions. You can choose either of the following views:

  – **View By Step, Metric, Beacon:** In this view, you can click **Add Beacon Overrides** to override the default threshold values for one or more beacons. In this case, the default thresholds will only be used for beacons without any overrides. Any new beacons added to the Web transaction will use the default thresholds. Click **Add Metric** to define thresholds for one or more metrics. Incidents are generated only if the value of the Data Granularity property is set to 'Transaction' for the service tests. For more information on the Web transaction properties, refer to the Create / Edit Service Test - Web Transaction help page in the Enterprise Manager Online Help.

  – **View By Step, Beacon, Metric:** In this view, you can click on the **Default** icon to toggle between Edit and View modes for a specific metric. In the Edit mode, you can modify the parameters of the metric for a specific beacon. In the View mode, the default parameters of the metric will be used. Incidents are generated only if the value of the Data Granularity property is set to ' Step'.

To define the default collection frequency and collection properties, click the **Collection Settings** tab on the Monitoring Settings page. You can do the following:

- Specify the default collection frequency for all the beacons. To override the collection frequency for a specific beacon, click **Add Beacon Overrides**.
- Specify the collection properties and their corresponding values for one or more beacons.

Refer to the Enterprise Manager Online Help for more details on the defining the collection intervals and performance thresholds.

21.7 Configuring Aggregate Services

Aggregate services consist of one or more services, called subservices. A subservice is any service created in Enterprise Manager. The availability, performance, business criteria, and usage for the aggregate service depend on the availability, performance, business criteria, and usage for the individual subservices comprising the service. To create an aggregate service, navigate to the Services main page, select Aggregate Service from the Add drop-down list and click **Go**. The Add Aggregate Service page appears. Creating an Aggregate Service involves the following:

- Specifying the name and time zone for the service.

- Adding the services that make up this aggregate service.
- Establishing the availability definition for the aggregate service. Availability of an aggregate service depends on the availability of its constituent subservices. The availability for a subservice may depend on the successful execution of a service test or on the availability of the system components on which the subservice runs, depending how the subservice was defined.
- Defining the metrics used to measure the performance of your aggregate service. You can add performance metrics from single subservices, or based on statistical aggregations of more than one metric. Once you have selected the performance metrics, you can set the thresholds used to trigger incidents, or remove metrics you no longer want.
- Defining the metrics used to measure the usage of your aggregate service. Usage metrics can be based on the metrics of one or more system components. You can add usage metrics from single subservices, or based on statistical aggregations of more than one metric. Once you have selected the usage metrics, you can set the thresholds used to trigger incidents, or remove metrics you no longer want.
- Defining the metrics that are used to measure of the performance of business in the organization. These metrics are based on business indicators that can assess the business performance.You can add business metrics from single subservices, or based on statistical aggregations of more than one metric. Once you have selected the business metrics, you can set the thresholds to trigger incidents, or remove metrics you no longer want.

After you have created an aggregate service, you can add or remove its constituent subservices, modify the availability definition and add or delete performance or usage metrics. Refer to the Enterprise Manager Online Help for details on these operations.

WARNING: If you delete or remove a subservice from an aggregate service, the aggregate service performance, usage, and business metrics may be affected if they are based on a deleted subservice's metrics.

21.8 Setting Up Monitoring Templates

A monitoring template for a service contains definitions of one or more service tests, as well as a list of monitoring beacons. A monitoring template can be used to create service tests on any number of service targets, and specify a list of monitoring beacons.

A monitoring template must be created from a service target. Once the template is created, the user can edit the template, create copies, or delete it. Finally, the user can apply the template to other targets, which creates the service tests on the other targets and adds the monitoring beacons.

To create a Monitoring Template, follow the steps given below:

1. Click **Setup** to navigate to the main Setup page in Enterprise Manager.
2. Click the **Monitoring Templates** link in the left panel.
3. Click **Create** to create a monitoring template.
4. In the target selection box, enter or select a service target and click **Continue**.
5. In the Monitoring Template General Page, enter the name of the template that you wish to create.

6. Click **Tests** to add / remove or configure service tests associated with the selected service target. Make the required changes to this page and click **OK** to save the template to the repository.

After you have created the Monitoring Template, use the **Apply** option to apply this template to a service test. You can click **Edit** to modify the template. For more details on these operations, refer to the Online Help.

21.8.1 Configuring Service Tests and Beacons

You can configure the service tests and beacons associated with the template by using the options in the **Tests** page. A service test-based template contains the following elements:

- **Variables**: A variable may occur at multiple locations in the service tests. The Variables table allows you to specify default values for all the variables. These default values will be stored in the template along with the variables. You can specify values other than the default while applying the template to a target. You can perform the following operations:
  - **Add** a variable. The variable can consist of letters, numbers and underscores only.
  - **Rename** a variable. When you rename a variable, all variable references in the service tests will be replaced with the new name.
  - **Remove** variables for properties within service tests. If you remove a non-password variable, all references to the variable in test properties will be replaced with the variable's default value
  - **Replace Text** in test properties with a variable definition.
- **Service Tests:** You can edit the test definition and define variables for various properties. You can select the tests from the original target that are to be part of the template by clicking the **Add / Remove** button. You can specify whether the service test is a key test and if it should be enabled. You can also click **Monitoring Settings** to drill down to this page and define metrics and thresholds for the service tests.
- **Beacons:** Use the **Add / Remove** button to specify which beacons are to be included in the template. You can also specify whether each beacon is a key beacon.

Refer to the Enterprise Manager Online Help for detailed instructions on these operations.

21.9 Configuring Service Levels

A service level rule is defined as an assessment criteria used to determine service quality. It allows you to specify availability and performance criteria that your service must meet during business hours as defined in your Service Level Agreement. For example, e-mail service must be 99.99% available between 8am and 8pm, Monday through Friday.

A service level rule 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 rule is based on the following:

- **Business Hours:** Time range during which the service level should be calculated as specified in your Service Level Agreement.
- **Availability:** Allows you to specify when the service should be considered available. This will only affect the service level calculations and not the actual availability state displayed in the console. You can choose a service to be considered up when it is one or more of the following states:
  - Up: By default the service is considered to be Up or available.
  - Under Blackout: This option allows you to specify service blackout time (planned activity that renders the service as technically unavailable) as available service time.
  - Unknown: This option allows you to specify time that a service is unmonitored because the Management Agent is unavailable be counted as available service time.
- **Performance Criteria**: You can optionally designate poor performance of a service as a Service Level violation. For example, if your Website is up, but it takes 10 seconds to load a single page, your service may be considered unavailable.
- **Business Criteria:** Business criteria are useful in determining in the health of the business processes for a particular service. You can optionally define business metrics that can affect the Service Level. A Service Level violation occurs when a critical alert is generated for a specified business metric.

> **Note:** The **Business Criteria** column is displayed only if one or more key business indicators are associated with the service. Refer to *Oracle Enterprise Manager Integration Guide*.

- **Actual Service Level:** This is calculated as percentage of time during business hours that your service meets the specified availability, performance, and business criteria.
- **Expected Service Level:** Denotes a minimum acceptable service level that your service must meet over any relevant evaluation period.

You can define only one service level rule for each service. The service level rule will be used to evaluate the **Actual Service Level** over a time period and compare it against the **Expected Service Level**.

21.9.1 Defining Service Level Rules

When you create a service, the default service rule is applied to the service. However, you must edit the service level rule for each service to accurately define the assessment criteria that is appropriate for your service. To define a service level rule:

1. Click the **Targets** tab and **Services** subtab. The Services main page is displayed.
2. Click the service name link to go to the Service Home page.
3. In the Related Links section, click **Edit Service Level Rule**.
4. On the Edit Service Level Rule page, specify the expected service level and the actual service level and click **OK**. The expected service level specifies the percentage of time a service meets the performance, usage, availability, and business criteria defined in the Service Level Rule. The actual service level defines the baseline criteria used to define service quality and includes business hours, availability, performance criteria, usage criteria, and business criteria.

Note: Any Super Administrator, owner of the service, or Enterprise Manager administrator with OPERATOR_TARGET target privileges can define or update the Service Level Rule.

21.9.2 Viewing Service Level Details

You can view service level information directly from the either of the following:

- **Enterprise Manager Cloud Control Console** -From any Service Home page, you can click on the Actual Service Level to drill down to the Service Level Details page. This page displays what Actual Service Level is achieved by the service over the last 24 hours/ 7 days / 31 days, compared to the Expected Service Level. In addition, details on service violation and time of each violation are presented in both graphical and textual formats.
- **Information Publisher** - Information Publisher provides an out-of-box report definition called the Services Dashboard that provides a comprehensive view of any service. From the Report Definition page, click on the **Services Monitoring Dashboard** report definition to generate a comprehensive view of an existing service. By default, the availability, performance, status, usage, business, and Service Level of the service are displayed. The Information Publisher also provides service-specific report elements that allow you to create your own custom report definitions. The following report elements are available:
  - **Service Level Details**: Displays **Actual Service Level** achieved over a time-period and violations that affected it.
  - **Service Level Summary**: Displays service level violations that occurred over selected time-period for a set of services.
  - **Services Monitoring Dashboard**: Displays status, performance, usage, business, and service level information for a set of services.
  - **Services Status Summary**: Information on one or more services' current status, performance, usage, business, and component statuses.

    Refer to the Online Help for more details on the report elements.

21.10 Configuring a Service Using the Command Line Interface

Using the Command Line Interface, you can define service targets, templates and set up incidents. EM CLI is intended for use by enterprise or system administrators writing scripts (shell/batch file, perl, tcl, php, etc.) that provide workflow in the customer's business process. EM CLI can also be used by administrators interactively, and directly from an operating system console. Refer to *Enterprise Manager Command Line Interface Guide* for details.

21.11 Troubleshooting Service Tests

This section lists some of the common errors you may encounter while using the Web Transaction test type. Some of the common errors you may encounter while recording and playing back Web transactions are listed below:

1. **Scenario:** Verify Service Test displays: `Connection establishment timed out -- http://..../`

   Possible Cause: The beacon can only access that URL via a proxy server and it has not been configured.

Solution: From the All Targets page, select the beacon, click **Configure** and set the beacon proxy setting.

2. Scenario: Verify Service Test displays: `Authorization Required -- https://...../`

   Possible Cause: The Basic Authentication information is not recorded automatically.

   Solution: To resolve this error, follow these steps:

   1. From the Service Tests and Beacons page, select the service test, click Edit.
   2. Make sure you enter all the Basic Authentication information: Username, Password, and Realm.

   > **Note:** Realm usually appears above the Username label in the Browser's authorization dialog box.

3. Scenario: Verify Service Test displays `sun.security.validator.ValidatorException: No trusted certificate found -- https://....../.`

   Possible Cause: The beacon does not know about this SSL Certificate.

   Possible Solution: From the Service Tests and Beacons page, select the service test, and click **Edit**. Under **Advanced Properties**, and set **Authenticate SSL Certificates** to **No**.

4. **Scenario**: Verify Service Test displays: `Timeout of 300000 exceeded for https://...../ Response time = 3000000`

   Possible Cause: The test may be too complex to complete within the allotted time. Or, this may be an actual performance issue with the server.

   Possible Solution: From the Service Tests and Beacons page, select the service test, and click **Edit**. If this is not a server performance issue, under **Advanced Properties**, increase the **Timeout Value**.

5. Scenario: The Verify Service Test option reports that the service as down, but the Web application is up and you can successfully play back the Web transaction.

   Possible Cause: The Web application is only compatible with Internet Explorer or Mozilla-based browsers.

   Possible Solution: From the Service Tests and Beacons page, select the service test, and click **Edit**. Under **Advanced Properties**, set the **User Agent Header** as `Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1) OracleEMAgentURLTiming/3.0.`

   > **Note:** For Grid Control 10.2.0.4 and beyond, this User Agent Header is set automatically during Web transaction recording.

6. **Scenario**: Test Performance Page does not show any step metrics.

   Possible Cause: By default, only transaction-level metrics are collected.

   Possible Solution: From the Service Tests and Beacons page, select the service test, click **Edit**, and set **Data Granularity** to Step.

22

Identity Management

This chapter describes how you can use Cloud Control to manage your Identity Management targets.

This chapter contains the following sections:

22.1 About Oracle Identity Management

Oracle Identity Management provides a unified, integrated security platform designed to manage user identities, provision resources to users, secure access to corporate resources, enable trusted online business partnerships, and support compliance (identity analytics) across the enterprise.

Enterprise Manager supports monitoring of the following Oracle Identity Management components:

- Oracle Access Manager 10g and 11g
- Oracle Authorization Policy Manager 11g
- Oracle Adaptive Access Manager 11g
- Oracle Directory Server Enterprise Edition 6.x, 7.x, and 11g
- Oracle Directory Integration Platform 11g
- Oracle Identity Manager 9.x and 11g
- Oracle Identity Federation 10g and 11g
- Oracle Identity Management Suite 10g (including Oracle Internet Directory, Oracle Directory Integration Platform, Oracle Delegated Administration Services, and Oracle Single Sign-On)
- Oracle Internet Directory 11g
- Oracle Virtual Directory 11g

22.2 Using Cloud Control for Monitoring Identity Management Targets

Enterprise Manager helps you monitor the availability and diagnose the health of Identity Management components within your enterprise configuration. By deploying a Management Agent on each host, you can use Enterprise Manager to discover the Identity Management components on these hosts, and automatically begin monitoring them using default monitoring levels, notification rules, and so on.

22.2.1 Identity and Access Dashboard

In Cloud Control 12c, a new Identity and Access dashboard provides a centralized view of all deployed Oracle Identity Management components - including both Identity Management 10g and Identity Management 11g components. This dashboard enables you to monitor the health of complex Identity Management deployment by providing an integrated interface for Component Type Overview, Member Summary, Resource Usage, Systems, Services and individual component types (Oracle Internet Directory, Oracle Access Manager, etc.) regions. Based on the deployment criteria, you can select the regions that best fit your deployment and display those in the dashboard. Following are the Oracle Identity Management components for which these regions are displayed:

- Oracle Access Manager Server
- Oracle Adaptive Access Manager Server
- Oracle Directory Server Enterprise Edition Server
- Oracle Identity Manager Cluster
- Oracle Internet Directory Server
- Oracle Virtual Directory Server

Each individual component type region displays the most critical metrics for the discovered target members of the specified Identity Management component type. Besides showing current values of these critical metrics, the region displays performance trends of these critical metrics for the last 24 hours so that you can visualize the performance of all target members in a single region.

You can access Identity and Access dashboard (shown in Figure 22–1) from the **Middleware Features** menu when you click on **Targets->Middleware** from the Cloud Control home page.

Figure 22–1 Identity and Access Dashboard

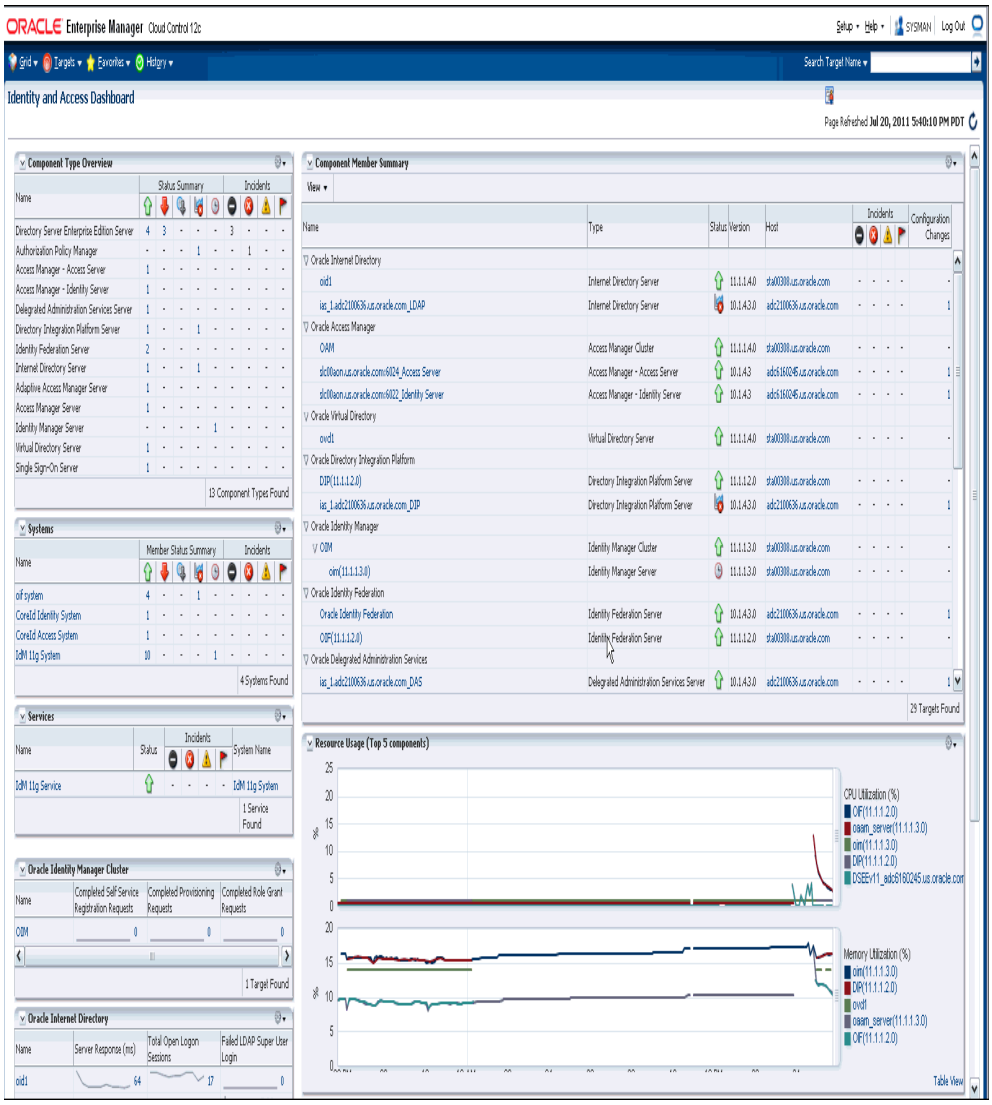

22.2.2 Identity Management Component Server Home Page

All Identity Management targets, whether Access, Identity, Identity Federation, and Identity Manager have their own server home pages that provide easy access to key information required by the administrators. Each Identity Management Component Server home page provides the following information:

- Server availability, responsiveness, and performance data. This includes a wide range of out-of-box performance metrics such as server up/down status, average response time, CPU utilization, memory utilization, provisioning metrics, failed logins, and total connections.
- Customizable performance summaries with a **Metric Palette** that allows users to drag and drop performance charts to drill down into usage and performance statistics.
- Resource usage for the host or WebLogic Server
- Functionality to start, stop, and restart components

Figure 22–2 shows the Oracle Directory Server Enterprise Edition server home page.

Figure 22–2 Oracle Directory Server Enterprise Edition Server Home Page

.

****

22.2.3 Configuration Management

You can perform key configuration management tasks such as keeping track of configuration changes, taking snapshots to store configurations, and comparing component configurations. To ensure that the configurations of all critical Oracle Identity Management components in your production environment are consistent with your staging or test environments, you can use Configuration Snapshots to save working configurations into the Management Repository or into an external XML file and then use the Configuration Comparison tool to compare the configuration in the production environment against the test or staging environments. Configuration Comparison helps you ensure the consistency of configurations in your environment, thus reducing "configuration drift." To diagnose performance problems that may be related to system configuration changes, you can use the Configuration History tool

(Figure 22–3) to keep track of all configuration changes to locate the root cause of performance problems.

Figure 22–3 Compare Results Page

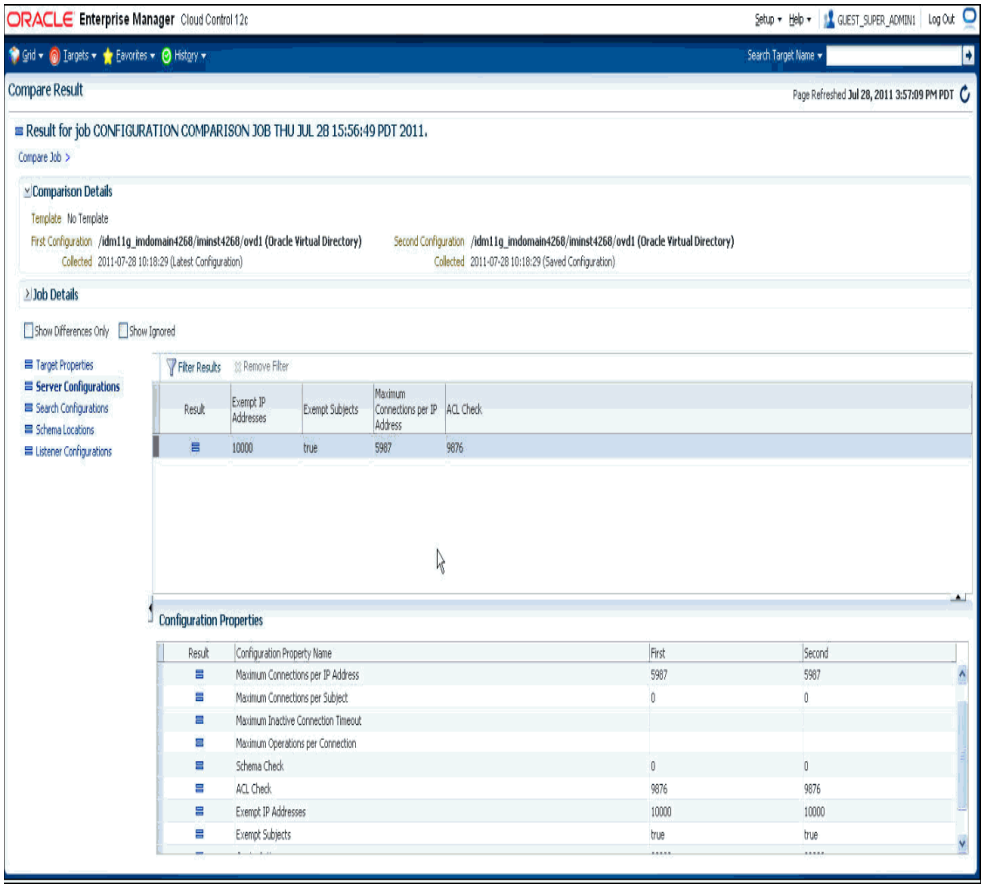

22.2.4 Identity Management Systems

Identity Management services run on Identity Management systems defined in Cloud Control. The system includes the software infrastructure components that the Identity services rely on. This system can be created using the Identity and Access System Create wizard, that can be accessed from Systems page.

The system is a collection of server targets that are grouped together in Cloud Control to give you a view of the "data-center" components that comprise your Identity Management deployment. Identity Management Systems are created when Identity suite components are discovered using Cloud Control. Cloud Control also monitors the performance and availability of these components and provides a System Dashboard to view the health of the Identity Management system in a single window.

Figure 22–4 shows an Identity and Access System home page:

Figure 22–4 Identity and Access System Home Page

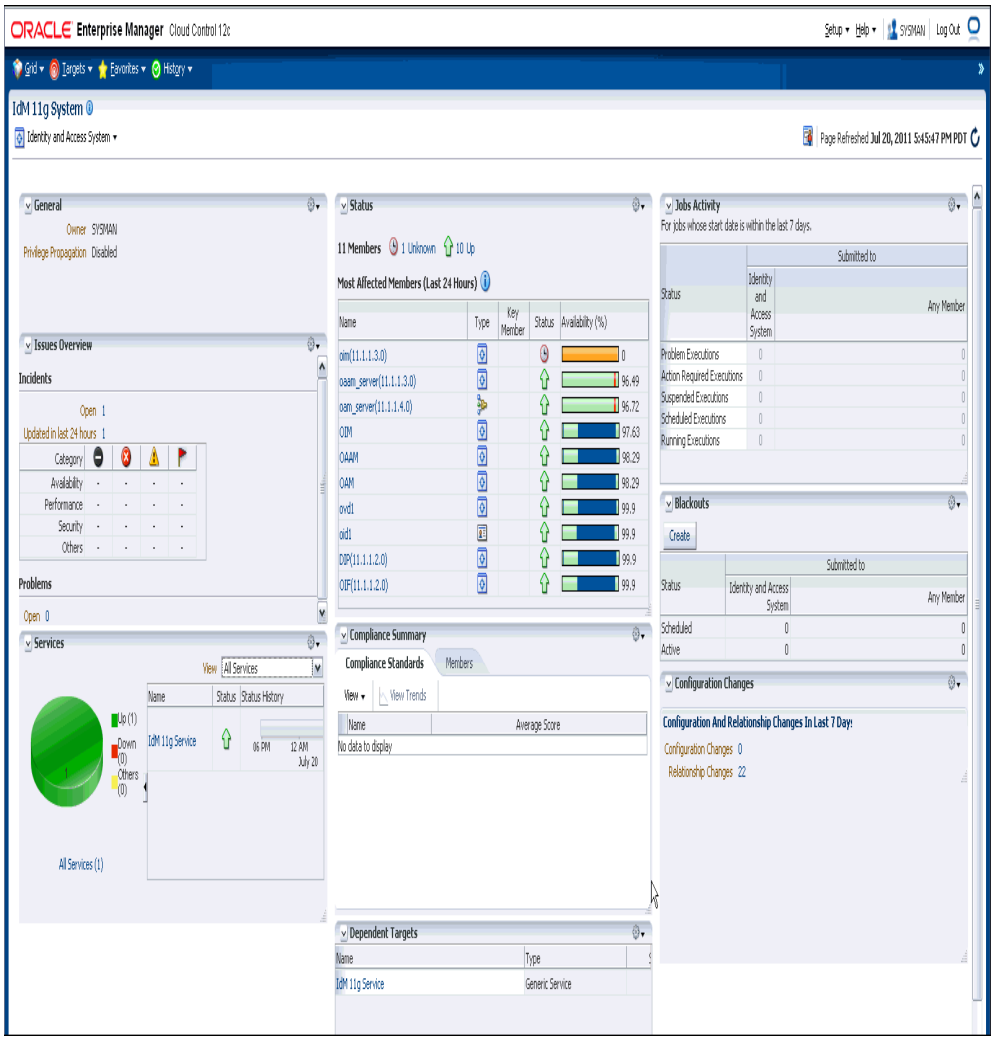

22.2.4.1 Identity Management Services

An Identity Management service is a logical target configured by Cloud Control. You use Cloud Control to step you through the process of configuring a web application service for your Identity component instances. After you configure a service, that service is displayed on the Services page.

Critical application functions are defined and monitored as services in Cloud Control. Each service is monitored by Cloud Control beacons, which run service tests that simulate real user access to the service. Service availability and performance are monitored automatically, and problems are immediately reported to the administrator. By monitoring availability and performance of Identity Management services, you can identify and resolve user-visible problems more quickly and thus minimize the impact on users.

22.2.4.2 Monitoring Services

Cloud Control enables you to monitor all of your Identity Management services. Each service is monitored for performance, usage, and availability.

Each service has its own home page. The Service Home pages in Cloud Control provide:

- Status, responsiveness, and performance data
- Resource usage data for the service
- Summary information such as status, performance alerts, usage alerts, and policy violations for the service's subcomponents, including other services and associated systems
- Links to home pages for the service's subcomponents
- Alerts and diagnostic drill-downs so that you can identify and resolve problems quickly
- Services Dashboard

  The Services Dashboard provides a high-level view of the status, performance, and usage of each Identity Management target. Service-level compliance for various time periods are also included for each service on the dashboard. You can launch the dashboard directly from Identity system target home page. You can also publish the Services Dashboard so that it can be viewed by non-Enterprise Manager users. This allows you to provide a self-service status web page to your end users.

- Related Links to do the following:
  - View metrics for the service
  - View client configurations
  - Edit the service
  - View the service target's properties
  - Manage blackouts
  - View and manage metric thresholds and policies

See Also: Chapter 21, "Configuring Services"

22.3 Identity Management Root Cause Analysis

Individual services in Identity Management are associated with critical system components. This allows Enterprise Manager to perform Root Cause Analysis down to the system level whenever a service outage is detected. When you are configuring an Identity Management service in Cloud Control, as mentioned in Identity Management Services, you also mention the critical system components of this service. When an Identity Management service goes down, Enterprise Manager automatically performs a root cause analysis to determine which critical system component is responsible for this.

22.4 Automated Identity Management Monitoring and Alerts

Enterprise Manager automatically gathers and evaluates diagnostic information from Identity Management targets distributed across the enterprise. As with all targets managed by Enterprise Manager, an extensive number of Identity Management performance metrics are automatically monitored against predefined thresholds. Alerts are generated in Cloud Control when metrics exceed these thresholds.

22.5 Diagnosing Identity Management Performance and Availability Problems

You can use Cloud Control to diagnose performance and availability problems with your Identity Management services. For example, if a service outage occurs, Root Cause Analysis will determine if the primary cause is an outage of a critical service or system component. If a service performance issue is found, an administrator can examine detailed metrics over time related to that service and any of the service or system components used by that service. When you suspect there is a problem with one or more server components in the Identity Management system, the system home pages provide metrics and charts for diagnosing the issue.

Administrators can monitor the health of all critical Identity Management components, including both Identity Management 10g and Identity Management 11g components. Thresholds may be defined against server and component statistics such as CPU utilization, the number of failed and successful authentications or authorizations, average response time, provisioning metrics (e.g. number of newly provisioned, created, deleted, disabled, locked users), Identity Provider and Service Provider metrics, and up/down status of servers and components.

In addition to relying on system performance metrics, you may use Management Pack for Identity Management Service Tests to record synthetic web transactions that include a combination of one or more navigation paths within the application to be used as the criteria for determining the availability of the service. For example, Oracle Access Manager requires that a user be successfully authenticated and authorized against a certain WebGate for the service to be considered available. Enterprise Manager uses these logical tasks or transactions to define the availability of the Identity Management environment. In addition to synthetic web transactions, Enterprise Manager also supports LDAP tests that allow you to record LDAP operations against a specific LDAP server (including Oracle Virtual Directory). With the LDAP tests, you can specify the username or password, Search Filter, Search Base, and Compare Attribute Name or Value. These synthetic web transactions are recorded, and the stored transaction or service test can be launched at a user-defined interval from strategic locations across the user-base."

22.6 Leveraging the Cloud Control Management Framework

Cloud Control includes many general features that are useful to an Identity Management administrator, including:

- **Job Automation:** You can use the Cloud Control job system to schedule tasks you want to automate.
- **Policies:** You can utilize the policy framework to ensure your Identity Management infrastructure adheres to your site-specific standards.
- **Database and Application Server Management:** Using the single Cloud Control console, you can also manage the specific databases and application servers in your Identity Management deployment if needed.
- **Extensions:** Cloud Control also includes monitoring of key network components that may be part of your Identity Management deployment. You can also extend Cloud Control to monitor other components that are not recognized out-of-box by Enterprise Manager.

23

Lifecycle Management

This chapter provides an overview of the server and software provisioning and patching features offered by Enterprise Manager Cloud Control. This chapter contains the following:

- Overview of Lifecycle Management Solutions
- Provisioning Operating System
- Provisioning Database and Middleware
- Upgrading Databases and Software
- Patching Database and Middleware Targets
- Patching Linux Hosts
- Customizing Deployment Procedures

The provisioning and patching features together make up the Lifecycle Management solution area of Enterprise Manager Cloud Control. To read more about this solution area, access the following URL:

`http://www.oracle.com/technetwork/oem/automation/index.html`

To learn how to access the provisioning and patching features within the Enterprise Manager Cloud Control console, see the *Oracle® Enterprise Manager Administrator's Guide for Software and Server Provisioning and Patching*. This guide is available in the Enterprise Manager documentation library available at:

`http://www.oracle.com/technetwork/indexes/documentation/index.html`

Overview of Lifecycle Management Solutions

Enterprise Manager Cloud Control offers provisioning deployment procedures that automate the deployment of software and applications, and patch plans that automate the patching of systems across your network. These features make critical data center operations easy, efficient, and scalable resulting in lower operational risk and cost of ownership. The ability to provision and patch the entire software stack that includes the operating system, the middleware, database, third party software, and applications supplemented by comprehensive reporting tools make these features extremely significant entities in the overall System Management space.

As shown in Figure 23–1, with the help of provisioning deployment procedures and patch plans, Enterprise Manager Cloud Control covers the entire lifecycle management of software, applications, and servers. The deployment procedures

orchestrate the initial reference sandbox deployment and then the mass unattended deployment of gold images created from these reference deployments.

Going forward, as the computation demand for the resources decline, Enterprise Manager Cloud Control allows you to deactivate and de-provision the resources making them available for a different purpose.

Figure 23–1 Lifecycle Management Overview

Advantages of Using Provisioning and Patching Features

The following are the advantages of using the provisioning deployment procedures and patch plans in Enterprise Manager Cloud Control:

- Provides a repeatable, reliable, and automated solution for performing mass, unattended, and schedulable deployment of
  - Software and servers based on Gold Images created using reference deployment or installation media
  - Software and operating system updates

    - Complex and multi-tier software like Oracle Real Application Clusters (RAC), Oracle Cloud Infrastructure, and Fusion Middleware Clusters
- Allows new resources to be provisioned at short notice based on compliant and tested gold images.
- Allows multiple operations to be accommodated in a single change window.
- Supports SUDO, PAM, and Privilege Delegation authentication.
- Offers a single interface for multiple players. For example, component designers responsible for creating Gold Images based on corporate standards and the operators all use the same Enterprise Manager Cloud Control console.
- Provides automation of repeatable installation and patching operations across the stack leads to substantial cost savings in terms of costs and man-hours.
- Automates the patching operation across your network using patch plans. A patch plan is a collection of patches which you might want to consider applying as a group to one or more targets. Patch plans can be created using the *My Oracle Support Patches and Updates* functionality that is tightly integrated with Enterprise Manager Cloud Control.

Enterprise Manager Cloud Control also provides command-line interface support to all out-of-box provisioning and patching deployment procedures. These features can hence be invoked by custom scripts.

Enterprise Manager Cloud Control also allows you to customize these default deployment procedures to suit your requirements.

Wide Coverage Across the Stack

Enterprise Manager Cloud Control provides provisioning and patching capabilities across the stack for:

- Operating Systems, with Bare Metal Provisioning on Linux and operating system patching
- Databases, with Real Application Clusters (RAC) provisioning, extension, and deletion; Cloud Infrastructure provisioning for standalone servers and clustered environments; and flexible patching for Oracle Database and Oracle Real Application Clusters
- Middleware, with Oracle Fusion Middleware provisioning, Oracle SOA Suite provisioning, and SOA Artifacts provisioning, BPEL provisioning, Oracle Service Bus provisioning

Note that these features require Oracle Management Agents to be present on the destination hosts where the software has to be provisioned.

Provisioning Operating System

Bare metal or operating system provisioning application provides server lifecycle management to build, manage, and optimize server infrastructure. The application:

- Automates deployment of consistent, certified Linux operating system images along with larger number of servers on physical and virtual servers.
- Automates deployment of hypervisors and virtual machines.

- Provides a template-based approach for provisioning a variety of Linux configurations servers (RedHat 3.0/4.0, SuSE/SLES9). This also ensures compliance to standards and consistency across all deployments.
- Reduces errors with standardized gold image-based server provisioning.
- Supports heterogeneous hardware and network configuration.
- Automatically discovers bare metal and live target servers for provisioning.
- Especially for Oracle software, the application encodes best practices out-of-the-box for patching.
- Results in considerable reduction in manual labor that leads to substantial cost savings.

Provisioning Database and Middleware

Software Provisioning automates the deployment of Databases and Fusion Middleware. It makes critical data center operations easy, efficient, and scalable resulting in lower operational risk and cost of ownership.

The features are as follows:

- Enables mass deployment of Oracle software (Database, Oracle Real Application Clusters stack, and Fusion Middleware)
- Supports all versions up to 11.2 and Cloud Infrastructure Architecture
- Enables standardized software deployment via Provisioning Profiles
- Allows segregation of duties using Designer and Operator roles
- Provides lock down access for controlled and error free deployments
- Provides pre-requisite checks and fix-ups
- Enables Group-based operations
- Supports user-defined deployment procedures

Upgrading Databases and Software

Enterprise Manager Cloud Control offers the ability to mass upgrade Oracle Databases in an easy and efficient way. Upgrading to the latest version of Oracle Database enables you to access the latest technology, thereby increasing efficiency and providing secure data management for your applications.

The features are as follows:

- Supports mass upgrade of Oracle Databases
- Supports single instance database upgrade in first release
- Upgrades from versions 10.2.0.x and 11.1.0.x to 11.2.0.x
- Upgrades software and database instances combined or separately

Patching Database and Middleware Targets

Enterprise Manager Cloud Control offers patch plans that simplify the patching of targets such as Oracle Database, Oracle RAC, Oracle ASM, Oracle Clusterware, and Oracle Fusion Middleware.

A patch plan is a collection of patches which you might want to consider applying as a group to one or more targets. Each target can have a separate group of patches. A patch plan can include a description and a deployment date for the plan and one or more patches.

A patch can be added to a target in a plan only if the patch has the same release and platform as the target to which it is being added. You will receive a warning if the product for the patch being added is different from the product associated with the target to which the patch is being added. The warning does not prevent you from adding the patch to the plan.

Patch plans can be created using the *My Oracle Support Patches and Updates* feature that is tightly integrated with Enterprise Manager Cloud Control. This integration greatly simplifies the patching operations by offrering timely and easy access to invaluable patch and support information during the patch planning phase, thus enabling you to utilize the wealth of information from My Oracle Support to implement the best possible patch rollout for your organization. Not only does this help you make important decisions regarding your patch and deployment process, it also helps you resolve conflicts quickly.

In addition to creating new patch plans, you can also create patch plan templates using these patch plans.

A patch plan template is a predesigned plan based on an existing successfully analyzed or deployable patch plan, however without any targets selected.

A patch plan template enables you to create new patch plans using a predetermined set of patches and deployment options saved from the source patch plan, and by selecting a completely new set of targets. Doing this reduces the time and effort required to create new patch plans and enables patch designers to expose only approved plans to patch operators.

Patching Linux Hosts

The Patch Linux Hosts application facilitates the automated management of Linux hosts in an enterprise. You can use this feature to keep the Linux hosts in your enterprise up to date with vital software updates from your Linux vendor.

Patch Linux Hosts uses a reference-based grouped patching model, where you can create one or more reference package repositories containing up-to-date versions of various packages, and associate a group of Linux hosts with these package repositories.

The Patch Linux Hosts tool uses package repositories to patch the hosts as well as to monitor the deviation of the packages installed on the hosts. You can create different groups suited to your administrative needs and even associate different package repositories with different priorities for each group. You can independently control when and how often to update the hosts in the group, and how to determine their compliance with respect to the package repositories.

Note: To use this feature, make sure you have the following:

- Licenses for the Provisioning and Patch Automation Pack
- Linux Management Pack
- "Operator" privileges on the host that you want to patch
- Ability to do sudo to the root user

The Linux patching feature provides the following functionalities:

- Setting up and managing RPM Repositories by subscribing to Unbreakable Linux network (ULN) channels
- Setting up and managing custom RPM Repositories and channels (cloning channels, copying packages from one channel into another, and deleting channels)
- Setting up Linux Patching Group to update a group of Linux hosts and compliance reporting from the Linux Patching group
- Scheduling Patching for non-compliant groups
- Managing Configuration file channels (creating/deleting channels, uploading files, and copying files from one channel into another)
- Patching through deployment procedures and emergency patching
- Undo Patching feature

Enhanced Linux Patching for ULN

Enhanced Linux Patching feature of Enterprise Manager supports the Unbreakable Linux Network (ULN) subscribers through EM. ULN provides access to Linux software patches, updates and fixes for its customers. Oracle provides three levels of Unbreakable Linux support:

- Network Support - access to patches and updates via ULN
- Basic Support - access to patches and updates via ULN, 24x7 support, complete Linux server lifecycle management
- Premier Support - access to patches and updates via ULN, 24x7 support, Linux server lifecycle management, backporting, lifetime support

The Linux RPM Repository Server Setup page in Enterprise Manager allows you to set up a RPM repository server for Linux patching. You can select the Host to setup the RPM repository server and register the host to the Unbreakable Linux Network (ULN).

Patching Features

- **Linux Host Patching Groups:** You can group a set of Linux hosts together to update all at once. Each group is associated with one or more package repositories that contain all the certified and appropriate versions of the software packages for the hosts of that group. Each group is configured with an update schedule for a recurring job to run to update the hosts with the associated package repositories.
- **RPM Repository:** RPM repository is a directory that contains RPM packages. The RPM repository is accessible via http or ftp. A RPM repository can be organized to contain packages from multiple channels.
- **Custom Channel:** A custom channel is a channel created by the user to store a set of custom RPM packages. Custom channels can be added to the RPM repository.
- **Configuration Channel:** A channel that is created by the user to store a set of Linux configuration files. Configuration channels can be used in the Linux patching application user interface to update configuration files on Linux hosts.
- **Compliance and automatic updates:** The compliance page contains information on the number of hosts in a group that are in compliance, as well as the number of "rogue" packages on a particular host. You can see metrics and charts to measure

compliance for all Linux Host Patching Groups, as well as historical compliance data.

- **Emergency Patching:** This feature gives you the option of performing "forced" updates, outside of the established schedule, to immediately respond to critical bugs or security alerts for all configured Linux hosts.
- **Undo Patching:** This feature adds flexibility by allowing you to roll back the software to its previous stable version, or even de-install the unstable version completely if that software version was found to be unsuitable or to have a bug or security vulnerability.
- **Patching through Deployment Procedures:** You can use deployment procedures to set up RPM repository, patch linux hosts, and perform other custom patching procedures.

Customizing Deployment Procedures

The provisioning and patching deployment procedures offered by Enterprise Manager Cloud Control are default procedures that have been created considering all the best practices in the industry. The steps embedded within a deployment procedure ensure that they meet all your provisioning and patching requirements. You can, of course, use them with the default settings to provision or patch your targets in the environment, however, you also have the choice of customizing them to include additional custom steps, disable unwanted steps, and use authentication tools to run some steps as another user.

You can also customize the deployment procedures to run them as another user, ignore the steps that require special privileges, add new steps, run custom scripts as part of the procedure, implement different error handling methods, and so on. You can run the above-mentioned deployment procedures using EMCLI commands.

24

Starting and Stopping Enterprise Manager Components

This chapter explains how to use the Enterprise Manager command line utility (emctl) to start and stop the Management Service, the Management Agent, and Cloud Control.

This chapter also explains the various emctl commands and how to use log information to troubleshoot emctl.

Following are the sections in this chapter:

24.1 Controlling the Oracle Management Agent

The following sections describe how to use the Enterprise Manager command line utility (`emctl`) to control the Oracle Management Agent:

24.1.1 Starting, Stopping, and Checking the Status of the Management Agent on UNIX

To start, stop, or check the status of the Management Agent on UNIX systems:

1. Change directory to the *AGENT_HOME*`/bin` directory.
2. Use the appropriate command described in Table 24–1.

   For example, to stop the Management Agent, enter the following commands:

```
$PROMPT> cd AGENT_HOME/bin
$PROMPT> ./emctl stop agent
```

Table 24–1 Starting, Stopping, and Checking the Status of the Management Agent

Command	Purpose
emctl start agent	Starts the Management Agent
emctl stop agent	Stops the Management Agent
emctl status agent	If the Management Agent is running, this command displays status information about the Management Agent, including the Agent Home, the process ID, and the time and date of the last successful upload to the Management Repository (Example 24–1).

Example 24–1 Checking the Status of the Management Agent

```
$ emctl status agent
Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
---------------------------------------------------------------
Agent Version     : 12.1.0.1.0
OMS Version       : null
Protocol Version  : 12.1.0.1.0
Agent Home        : /ade/jexample_username/oracle/work/agentStateDir
Agent Binaries    : /scratch/jklein/view_storage/example_
username/emagent/gcagent/agent
Agent Process ID  : 22299
Parent Process ID : 22226
Agent URL         : https://example.us.oracle.com:11852/emd/main/
Repository URL    : https://example.us.oracle.com:14487/empbs/upload
Started at        : 2011-08-09 06:19:12
Started by user   : user
Last Reload       : (none)
Last successful upload                       : (none)
Last attempted upload                        : (none)
Total Megabytes of XML files uploaded so far : 0
Number of XML files pending upload           : 787
Size of XML files pending upload(MB)         : 2.21
Available disk space on upload filesystem    : 47.91%
Collection Status                            : Collections enabled
Last attempted heartbeat to OMS              : 2011-08-09 06:20:51
Last successful heartbeat to OMS             : (none)

---------------------------------------------------------------
```

On IBM AIX environment with a large memory configuration where the Management Agent is monitoring a large number of targets, the Agent may not start. To prevent this issue, prior to starting the Management Agent, set the following variables in the shell:

```
LDR_CNTRL="MAXDATA=0x80000000"@NOKRTL
AIX_THREADSCOPE=S
```

The LDR_CNTRL variable sets the data segment size and disables loading of run time libraries in kernel space. The AIX_THREADSCOPE parameter changes AIX Threadscope context from the default Processwide 'P' to Systemwide 'S'. This causes less mutex contention.

24.1.2 Starting and Stopping the Management Agent on Windows

When you install the Oracle Management Agent on a Windows system, the installation procedure creates one new service in the Services control panel.

The procedure for accessing the Services control panel varies, depending upon the version of Microsoft Windows you are using. For example, on Windows 2000, locate the Services Control panel by selecting **Settings** and then **Administrative Tools** from the **Start** menu.

> **Note:** The `emctl` utility described in Section 24.2.1 is available in the `bin` subdirectory of the Oracle home where you installed the Management Agent; however, Oracle recommends that you use the Services control panel to start and stop the Management Agent on Windows systems.

Table 24–2 describes the Windows service that you use to control the Management Agent.

Table 24–2 Summary of Service Installed and Configured When You Install the Management Agent on Windows

Component	Service Name Format	Description

Oracle Management Agent	`Oracle<agent_home>Agent`
	For example:
	`OracleOraHome1Agent`
		Use this to start and stop the Management Agent.

> **Note:** If you are having trouble starting or stopping the Management Agent on a Windows NT system, try stopping the Management Agent using the following emctl command:
>
> ```
> $PROMPT> <AGENT_HOME>\bin\emctl istop agent
> ```
>
> After stopping the Management Agent using the emctl istop agent command, start the Management Agent using the Services control panel.
>
> This problem and solution applies only to the Windows NT platform, not to other Windows platforms, such as Windows 2000 or Windows XP systems.

24.1.3 Checking the Status of the Management Agent on Windows

To check the status of the Management Agent on Windows systems:

1. Change directory to the following location in the AGENT_HOME directory:

   AGENT_HOME\bin

2. Enter the following emctl command to check status of the Management Agent:

   ```
   $PROMPT> .\emctl status agent
   ```

   If the Management Agent is running, this command displays status information about the Management Agent, including the Agent Home, the process ID, and the

time and date of the last successful upload to the Management Repository (Example 24–1).

24.2 Controlling the Oracle Management Service

The following sections describe how to control the Oracle Management Service:

- Controlling the Management Service on UNIX
- Controlling the Management Service on Windows

24.2.1 Controlling the Management Service on UNIX

To start and stop the Oracle Management Service on UNIX systems, use a set of `emctl` commands.

- Using emctl to Start, Stop, and Check the Status of the Oracle Management Service

24.2.1.1 Using emctl to Start, Stop, and Check the Status of the Oracle Management Service

To start, stop, or check the status of the Management Service with the Enterprise Manager command-line utility:

1. Change directory to the `ORACLE_HOME/bin` directory in the Management Service home.
2. Use the appropriate command described in Table 24–3.

   For example, to stop the Management Service, enter the following commands:

```
$PROMPT> cd bin
$PROMPT> ./emctl stop oms
```

Table 24–3 Starting, Stopping, and Checking the Status of the Management Service

Command	Purpose
emctl start oms	Starts the Fusion Middleware components required to run the Management Service. Specifically, this command starts HTTP Server, the Node Manager, OPMN process, and the managed server on which the Management Service is deployed. In addition if this command is run on the host that has the Administration Server, the Administration Server is started too.
emctl stop oms	Stops the OMS managed server and HTTP server but leaves Node Manager and Administration Server running. **Note:** The `emctl stop oms` command does not stop Fusion Middleware. Use `emctl stop oms -all` to stop all processes including Administration Server, HTTP Server, Node Manager, and management server.
emctl status oms	Displays a message indicating whether or not the Management Service is running. Run `emctl status oms -details` to view information about the configuration of the management service such as ports being used and the URLs for console and upload.

24.2.2 Controlling the Management Service on Windows

When you install the Oracle Management Service on a Windows system, the installation procedure creates three new services in the Services control panel.

The procedure for accessing the Services control panel varies, depending upon the version of Microsoft Windows you are using. For example, on Windows 2000, locate the Services control panel by selecting **Settings** and then **Administrative Tools** from the **Start** menu.

> **Note:** The `emctl` utility described in Section 24.2.1 is available in the `bin` subdirectory of the Oracle home where you installed the Management Service; however, Oracle recommends that you use the Services control panel to start and stop the Management Service on Windows systems.

Table 24–4 describes the Windows services that you use to control the Oracle Management Service.

Table 24–4 Summary of Services Installed and Configured When Installing the Oracle Management Service on Windows

Component	Service Name Format	Description
WebLogic Server	`OracleWeblogicNodeManager_EMGC_OMS1_1`	Use this service to start and stop the node manager of the WebLogic Server that was installed and configured to deploy the Management Service J2EE application.
Oracle Management Server	`OracleManagementServer_EMGC_OMS1_1`	Use this service to start and stop all components that were installed and configured as part of the Management Service J2EE application.

24.3 Controlling Fusion Middleware Control

Fusion Middleware Control is a component of Oracle Fusion Middleware 11*g* that is installed as part of any WebLogic Server installation. For information about starting and stopping Fusion Middleware Control, see the chapter on Starting and Stopping Oracle Fusion Middleware in the *Oracle® Fusion Middleware Administrator's Guide* available on OTN.

24.4 Guidelines for Starting Multiple Enterprise Manager Components on a Single Host

Oracle Enterprise Manager 12c components are used to manage a variety of Oracle software products. In most cases, in a production environment, you will want to distribute your database and WebLogic Server instances among multiple hosts to improve performance and availability of your software resources. However, in cases where you must install multiple WebLogic Servers or databases on the same host, consider the following guidelines.

When you start Fusion Middleware Control, the Management Agent, or the Database Control, Enterprise Manager immediately begins gathering important monitoring data about the host and its managed targets. Keep this in mind when you develop a process for starting the components on the host.

Specifically, consider staggering the startup process so that each Enterprise Manager process has a chance to start before the next process begins its startup procedure.

When you start up all the components (for example, after a restart of the system), use a process such as the following:

1. Use the `emctl start` command to start Oracle Management Service.
2. Wait 15 seconds.
3. Use the `emctl start agent` command to start the Management Agent for the host.

Using a staggered startup procedure such as the preceding example will ensure that the processes are not in contention for resources during the CPU-intensive startup phase for each component.

24.5 Starting and Stopping Oracle Enterprise Manager 12c Cloud Control

As described in the previous sections, you use separate commands to control the Oracle Management Service and Oracle Management Agent.

The following sections describe how to stop and start all the Cloud Control components that are installed by the Oracle Enterprise Manager 12c Cloud Control Console installation procedure.

You can use this procedure to start all the framework components after a system reboot or to shutdown all the components before bringing the system down for system maintenance.

24.5.1 Starting Cloud Control and All Its Components

The following procedure summarizes the steps required to start all the components of the Cloud Control. For example, use this procedure if you have restarted the host computer and all the components of the Cloud Control have been installed on that host.

To start all the Cloud Control components on a host, use the following procedure:

1. If your Oracle Management Repository resides on the host, change directory to the Oracle Home for the database where you installed the Management Repository and start the database and the Net Listener for the database:
   a. Set the ORACLE_HOME environment variable to the Management Repository database home directory.
   b. Set the ORACLE_SID environment variable to the Management Repository database SID (default is asdb).
   c. Start the Net Listener:

      ```
      $PROMPT> $ORACLE_HOME/bin/lsnrctl start
      ```

   d. Start the Management Repository database instance:

      ```
      ORACLE_HOME/bin/sqlplus /nolog
      SQL> connect SYS as SYSDBA
      SQL> startup
      SQL> quit
      ```

      See Also: *Oracle Database Administrator's Guide* for information about starting and stopping an Oracle Database.

2. Start the Oracle Management Service:

```
$PROMPT> OMS_HOME/bin/emctl start oms
```

> **See Also:** "Controlling the Oracle Management Service" on page 24-4

3. Change directory to the home directory for the Oracle Management Agent and start the Management Agent:

```
$PROMPT> AGENT_HOME/bin/emctl start agent
```

> **See Also:** "Controlling the Oracle Management Agent" on page 24-1

Note: Be sure to run the `emctl start agent` command in the Oracle Management Agent home directory and not in the Management Service home directory.

24.5.2 Stopping Cloud Control and All Its Components

The following procedure summarizes the steps required to stop all the components of the Cloud Control. For example, use this procedure if you have installed all the components of the Cloud Control on the same host you want to shut down or restart the host computer.

To stop all the Cloud Control components on a host, use the following procedure:

1. Stop the Oracle Management Service:

```
$PROMPT> $ORACLE_HOME/bin/emctl stop oms -all
```

> **See Also:** "Controlling the Oracle Management Service" on page 24-4

2. Change directory to the home directory for the Oracle Management Agent and stop the Management Agent:

```
$PROMPT> AGENT_HOME/bin/emctl stop agent
```

> **See Also:** "Controlling the Oracle Management Agent" on page 24-1

Note: Be sure to run the `emctl stop agent` command in the Oracle Management Agent home directory and not in the Oracle Management Service home directory.

---

3. If your Oracle Management Repository resides on the same host, change directory to the Oracle Home for the database where you installed the Management Repository and stop the database and the Net Listener for the database:
   a. Set the ORACLE_HOME environment variable to the Management Repository database home directory.
   b. Set the ORACLE_SID environment variable to the Management Repository database SID (default is asdb).
   c. Stop the database instance:

```
$PROMPT> ORACLE_HOME/bin/sqlplus /nolog
SQL> connect SYS as SYSDBA
SQL> shutdown
SQL> quit
```

See Also: *Oracle Database Administrator's Guide* for information about starting and stopping an Oracle Database.

d. Stop the Net Listener:

```
$PROMPT> $ORACLE_HOME/bin/lsnrctl stop
```

24.6 Additional Management Agent Commands

The following sections describe additional `emctl` commands you can use to control the Management Agent:

- Uploading and Reloading Data to the Management Repository
- Specifying New Target Monitoring Credentials
- Listing the Targets on a Managed Host
- Controlling Blackouts

24.6.1 Uploading and Reloading Data to the Management Repository

Under normal circumstances, the Management Agent uploads information about your managed targets to the Management Service at regular intervals.

To use these commands, change directory to the `AGENT_HOME/bin` directory (UNIX) or the `AGENT_HOME\bin` directory (Windows) and enter the appropriate command as described in Table 24–5.

Table 24–5 Manually Reloading and Uploading Management Data

Command	Description
`emctl upload (agent)`	Use this command to force an immediate upload of the current management data from the managed host to the Management Service. Use this command instead of waiting until the next scheduled upload of the data.
`emctl reload` (agent)	This command can be used to apply the changes after you have manually modified the `emd.properties` file. For example, to change the upload interval, emd.properties can be modified, and `emctl reload` can then be run. **Note:** Oracle does not support manual editing of the `targets.xml` files unless the procedure is explicitly documented or you are instructed to do so by Oracle Support.

24.6.2 Specifying New Target Monitoring Credentials

To monitor the performance of your database targets, Enterprise Manager connects to your database using a database user name and password. This user name and password combination is referred to as the database monitoring credentials.

Note: The instructions in this section are specific to the monitoring credentials for a database target, but you can use this procedure for any other target type that requires monitoring credentials. For example, you can use this procedure to specify new monitoring credentials for your Oracle Management Service and Management Repository.

When you first add an Oracle9i Database target, or when it is added for you during the installation of the Management Agent, Enterprise Manager uses the DBSNMP database user account and the default password for the DBSNMP account as the monitoring credentials.

When you install Oracle Database 11*g*, you specify the DBSNMP monitoring password during the database installation procedure.

As a result, if the password for the DBSNMP database user account is changed, you must modify the properties of the database target so that Enterprise Manager can continue to connect to the database and gather configuration and performance data.

Similarly, immediately after you add a new Oracle Database 11*g* target to the Cloud Control, you may need to configure the target so it recognizes the DBSNMP password that you defined during the database installation. Otherwise, the Database Home page may display no monitoring data and the status of the database may indicate that there is a metric collection error.

Note: You can modify the Enterprise Manager monitoring credentials by using the Oracle Enterprise Manager 12*c* Cloud Control Console.

24.6.3 Listing the Targets on a Managed Host

There are times when you need to provide the name and type of a particular target you are managing. For example, you must know the target name and type when you are setting the monitoring credentials for a target.

To list the name and type of each target currently being monitored by a particular Management Agent:

1. Change directory to the `AGENT_HOME/bin` directory (UNIX) or the `AGENT_HOME\bin` directory (Windows).
2. Enter the following command to specify new monitoring credentials:

   ```
   $PROMPT>./emctl config agent listtargets
   ```

   Example 24–2 shows the typical output of the command.

Example 24–2 Listing the Targets on a Managed Host

```
ade:[ exampname_1208_qc_ag ] [example_username@example emagent]$ emctl config
agent listtargets
Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
[example.us.oracle.com:11852, oracle_emd]
[example.us.oracle.com, host]
[chronos_test, oracle_webcache]
[chronos_apache_test, oracle_apache]
[mytestBeacon, oracle_beacon]
```

```
[CSAcollector, oracle_csa_collector]
[database, oracle_database]
[database2, oracle_database]
[database3, oracle_database]
[listener, oracle_listener]
[listener2, oracle_listener]
[listener3, oracle_listener]
[Management Services and Repository, oracle_emrep]
ade:[ example_username_1208_qc_ag ] [example_username@example emagent]$
```

24.6.4 Controlling Blackouts

Blackouts allow Enterprise Manager users to suspend management data collection activity on one or more managed targets. For example, administrators use blackouts to prevent data collection during scheduled maintenance or emergency operations.

You can control blackouts from the Oracle Enterprise Manager 12c Cloud Control Console or from the Enterprise Manager command line utility (`emctl`). However, if you are controlling target blackouts from the command line, you should not attempt to control the same blackouts from the Cloud Control Console. Similarly, if you are controlling target blackouts from the Cloud Control Console, do not attempt to control those blackouts from the command line.

> **See Also:** "Creating, Editing, and Viewing Blackouts" in the Enterprise Manager online help for information about controlling blackouts from the Cloud Control Console

From the command line, you can perform the following blackout functions:

- Starting Immediate Blackouts
- Stopping Immediate Blackouts
- Checking the Status of Immediate Blackouts

> **Note:** When you start a blackout from the command line, any Enterprise Manager jobs scheduled to run against the blacked out targets will still run. If you use the Cloud Control Console to control blackouts, you can optionally prevent jobs from running against blacked out targets.

To use the Enterprise Manager command-line utility to control blackouts:

1. Change directory to the `AGENT_HOME/bin` directory (UNIX) or the `AGENT_HOME\bin` directory (Windows).
2. Enter the appropriate command as described in Table 24–6.

> **Note:** When you start a blackout, you must identify the target or targets affected by the blackout. To obtain the correct target name and target type for a target, see "Listing the Targets on a Managed Host".

Table 24–6 Summary of Blackout Commands

Blackout Action	Command
Set an immediate blackout on a particular target or list of targets	`emctl start blackout <Blackoutname> [<Target_name>[:<Target_Type>]]....[-d <Duration>]` Be sure to use a unique name for the blackout so you can refer to it later when you want to stop or check the status of the blackout. The `-d` option is used to specify the duration of the blackout. Duration is specified in `[days] hh:mm where:` ■ days indicates number of days, which is optional ■ hh indicates number of hours ■ mm indicates number of minutes If you do not specify a target or list of targets, Enterprise Manager will blackout the local host target. All monitored targets on the host are not blacked out unless a list is specified or you use the `-nodelevel` argument. If two targets of different target types share the same name, you must identify the target with its target type.
Stop an immediate blackout	`emctl stop blackout <Blackoutname>`
Set an immediate blackout for all targets on a host	`emctl start blackout <Blackoutname> [-nodeLevel] [-d <Duration>]` The `-nodeLevel` option is used to specify a blackout for all the targets on the host; in other words, all the targets that the Management Agent is monitoring, including the Management Agent host itself. The `-nodeLevel` option must follow the blackout name. If you specify any targets after the `-nodeLevel` option, the list is ignored.
Check the status of a blackout	emctl status blackout `[<Target_name>[:<Target_Type>]]....`

Use the following examples to learn more about controlling blackouts from the Enterprise Manager command line:

- To start a blackout called "bk1" for databases "db1" and "db2," and for Oracle Listener "ldb2," enter the following command:

  ```
  $PROMPT> emctl start blackout bk1 db1 db2 ldb2:oracle_listener -d 5 02:30
  ```

  The blackout starts immediately and will last for 5 days 2 hours and 30 minutes.

- To check the status of all the blackouts on a managed host:

  ```
  $PROMPT> emctl status blackout
  ```

- To stop blackout "bk2" immediately:

  ```
  $PROMPT> emctl stop blackout bk2
  ```

- To start an immediate blackout called "bk3" for all targets on the host:

  ```
  $PROMPT> emctl start blackout bk3 -nodeLevel
  ```

- To start an immediate blackout called "bk3" for database "db1" for 30 minutes:

  ```
  $PROMPT> emctl start blackout bk3 db1 -d 30
  ```

- To start an immediate blackout called "bk3" for database "db2" for five hours:

```
$PROMPT> emctl start blackout bk db2 -d 5:00
```

24.6.5 Changing the Management Agent Time Zone

The Management Agent may fail to start after the upgrade if it realizes that it is no longer in the same time zone that it was originally configured with.

You can reset the time zone used by the Management Agent using the following command:

```
emctl resetTZ agent
```

This command will correct the Management Agent side time zone and specify an additional command to be run against the Management Repository to correct the value there.

> **IMPORTANT: Before you change the Management Agent time zone, first check to see if there are any blackouts that are currently running or scheduled to run on any target managed by that Management Agent.**

To check for blackouts:

1. From the Cloud Control home page, click **Targets** and then **All Targets**. In the All Targets page, locate the Management Agent in the list of targets. Click on the Management Agent's name. This brings you to the Management Agent's home page.
2. The list of targets monitored by the Management Agent are listed in the Monitored Targets section.
3. For each of target in the list:
   a. Click the target name. This brings you to the target's home page.
   b. From the <Target> menu, select **Monitoring** and then click **Blackouts**. This allows you to check any currently running blackouts or blackouts that are scheduled in the future for this target.

If such blackouts exist, then:

1. From the Cloud Control Console, stop all currently running blackouts on all targets monitored by that Management Agent.
2. From the Cloud Control Console, stop all scheduled blackouts on all targets monitored by that Management Agent.

Once you have stopped all currently running and scheduled blackouts, you can run the `emctl resetTZ agent` command to change the Management Agent's time zone.

Once you have changed the Management Agent's time zone, create new blackouts on the targets as needed.

24.6.6 Reevaluating Metric Collections

Use the following command to perform an immediate reevaluation of a metric collection:

```
emctl control agent runCollection <targetName>:<targetType> <colletionItemName>
```

where <collectionItemName> is the name of the Collection Item that collects the metric.

Performing this command causes the reevaluated value of the metric to be uploaded into the Management Repository, and possibly trigger alerts if the metric crosses its threshold.

Related metrics are typically collected together; collectively a set of metrics collected together is called a Metric Collection. Each Metric Collection has its own name. If you want to reevaluate a metric, you first need to determine the name of the Metric Collection to which it belongs, then the CollectionItem for that Metric Collection.

When you run the previous command to reevaluate the metric, all other metrics that are part of the same Metric Collection and Collection Item will also be reevaluated.

Perform the following steps to determine the Metric Collection name and Collection Item name for a metric:

1. Go to $INSTALL_BASE/ngagent/plugins directory, where $INSTALL_BASE is the Oracle Home of the Management Agent.
2. Locate the XML file for the target type. For example, if you are interested in the host metric 'Filesystem Space Available(%)' metric, look for the host.xml file.
3. In the xml file, look for the metric in which you are interested. The metric that you are familiar with is actually the display name of the metric. The metric name would be preceded by a tag that started with:

   ```
   <Label NLSID=
   ```

   For example, in the host.xml file, the metric 'Filesystem Space Available(%)" would have an entry that looks like this:

   ```
   <Label NLSID="host_filesys_pctAvailable">Filesystem Space Available (%)
   </Label>
   ```

4. Once you have located the metric in the xml file, you will notice that its entry is part of a bigger entry that starts with:

   ```
   <Metric NAME=
   ```

   Take note of the value defined for "Metric NAME". This is the Metric Collection name. For example, for the 'Filesystem Space Available(%)' metric, the entry would look like this:

   ```
   <Metric NAME="Filesystems"
   ```

   So for the 'Filesystem Space Available(%)' metric, the Metric Collection name is 'Filesystems'.
5. The Collection Item name for this Metric Collection needs to be determined next. Go to the $INSTALL_BASE/ngagent/plugins/default_collection directory, where $INSTALL_BASE is the Oracle Home of the Management Agent.
6. In this directory, look for the collection file for the target type. In our example, this would be host.xml.
7. In cases where a Metric Collection is collected by itself, there would be a single Collection Item of the same name in the collection file. To determine if this is the case for your Metric Collection, look for an entry in the collection file that starts with:

   ```
   <CollectionItem NAME=
   ```

where the value assigned to the CollectionItem NAME matches the Metric NAME in step (4).

For the 'Filesystem Space Available(%)' metric, the entry in the collection file would look like:

```
<CollectionItem NAME = "Filesystems"
```

8. If you find such an entry, then the value assigned to "CollectionItem NAME" is the collection item name that you can use in the emctl command.
9. Otherwise, this means the Metric Collection is collected with other Metric Collections under a single Collection Item. To find the Collection Item for your Metric Collection, first search for your Metric Collection. It should be preceded by the tag:

   `<MetricColl NAME=`

   Once you have located it, look in the file above it for: `<CollectionItem NAME=`

   The value associated with the CollectionItem NAME is the name of the collection item that you should use in the emctl command.

   For example if the you want to reevaluate the host metric "Open Ports", using the previous steps, you would do the following:

   a. Go to the `$INSTALL_BASE/ngagent/plugins` directory where $INSTALL_BASE is the Oracle Home of the Management Agent. Look for the host.xml file and in that file locate: `<Metric NAME="openPorts"`.
   b. Then go to the `$INSTALL_BASE/ngagent/plugins/default_collection` directory. Look for the host.xml file and in that file look for `<CollectionItem NAME="openPorts"`.

      Failing this, look for `<MetricColl NAME="openPorts"`.
   c. Look above this entry in the file to find the `<CollectionItem NAME=` string and find `<CollectionItem NAME="oracle_security"`.

   The CollectionItem NAME oracle_security is what you would use in the emctl command to reevaluate the Open Ports metric.

24.7 emctl Commands

This section lists the emctl commands for the Enterprise Manager Agent and Management Service.

Table 24–7 explains the emctl commands for OMS.

Table 24–7 emctl Commands for OMS

emctl Command	Description

emctl [getversion] oms	Gets the version of the Management Service. Sample output is as follows:
	`./emctl getversion oms`
	`Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`Enterprise Manager 12c OMS Version 12.1.0.1.0`

emctl [start] oms	Starts the Management Service.

Table 24–7 (Cont.) emctl Commands for OMS

emctl Command	Description
emctl stop oms -all	Stops the Management Service including Administration Server, HTTP Server, Node Manager, and management server.
emctl stop oms -all -force and emctl stop oms -force	-force can be used with both emctl stop oms -all and emctl stop oms. If the emctl stop oms commands do not shutdown the relevant processes, using -force option will forcefully stop the relevant processes.
emctl status oms	Lists the status of the Management Service
emctl status oms -details	Lists Management Service details such as port numbers, lock status, domain information, and so on.
emctl config oms -list_ repos_details	Lists the Management Service repository details.
emctl config oms -store_ repos_details [-repos_host <host> -repos_port <port> -repos_sid <sid> | -repos_ conndesc <connect descriptor>] -repos_user <username> [-repos_pwd <pwd>]	Configures the settings used by Management Service to connect to the Management Repository.
emctl config oms -change_ repos_pwd [-old_pwd <old_pwd>] [-new_pwd <new_pwd>] [-use_sys_ pwd [-sys_pwd <sys_ pwd>]]	Configures the password used by Management Service to connect to the Management schema in the Management Repository.
emctl config oms -change_ view_user_pwd [-sysman_ pwd <sysman_pwd>] [-user_pwd <user_pwd>] [-auto_generate]	Configures the password used by Management Service for MGMT_VIEW user that is used for report generation.
emctl upload	Uploads xml files that are pending to upload to the OMS under the upload directory.

Table 24–8 explains emctl commands for Management Agent.

Table 24–8 emctl Commands for Management Agent

emctl Command	Description
emctl start | stop agent	Starts or stops agent.
emctl status agent	Lists the status of agent.

Table 24–8 (Cont.) emctl Commands for Management Agent

emctl Command	Description

emctl status agent -secure	Lists the secure status of the agent and the port on which the agent is running in secure mode and also the OMS security status of the agent it points to. This command also gives the OMS secure port. Below is an example output:
	`bash-3.00$ emctl status agent -secure`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`Checking the security status of the Agent at location set in /ade/example_username_cpap4_ag/oracle/sysman/config/emd.properties... Done.`
	`Agent is secure at HTTPS Port 1838.`
	`Checking the security status of the OMS at http://example.us.oracle.com:7654/em/upload/... Done.`
	`OMS is secure on HTTPS Port 4473`
	`bash-3.00$`

emctl status agent scheduler	Lists all Running, Ready, and Scheduled Collection threads.

emctl status agent jobs	Lists the status of the jobs that are running at present on the agent. The following is an example output:
	`bash-3.00$ emctl status agent jobs`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`----------------------------------------------------`
	`step id typ pid stat command line`
	`------- --- --- ---- ------------`
	`----------------------------------------------------`
	`Agent is Running and Ready`

Table 24–8 (Cont.) emctl Commands for Management Agent

emctl Command	Description
emctl status agent target <target name>,<target type>,<metric>	Lists the detailed status of the specified targets in the order of target name, target type. The following is an example of an oracle_database target. You can also provide a particular metric name in the emctl command to get the status of a particular metric of a target.

```
bash-3.00$ emctl status agent target database,oracle_
database
Oracle Enterprise Manager 12c Release 1 Cloud Control
12.1.0.1.0.
Copyright (c) 1996, 2011 Oracle Corporation.  All
rights reserved.
---------------------------------------------------
Target Name : database
Target Type : oracle_database
Current severity state
----------------------
Metric      Column name    Key   State  Timestamp
---------------------------------------------------
DeferredTrans errortrans_count  n/a CLEAR 2011-07-09
02:38:07
DeferredTrans deftrans_count    n/a   CLEAR 2011-07-09
02:38:07
ha_recovery missing_media_files n/a   CLEAR 2011-07-09
02:28:57
ha_recovery corrupt_data_blocks n/a  CLEAR 2011-07-09
02:28:57
ha_recovery datafiles_need_recovery n/a CLEAR
2011-07-09 02:28:57
Response     Status        n/a     CLEAR  2011-07-09
02:38:04
Response userLogon         n/a CLEAR  2011-07-09
02:38:04
Response    State          n/a  CLEAR  2011-07-09
02:38:04
OCMInstrumentation  NeedToInstrument n/a CLEAR
2011-07-09 02:31:55
health_check  Status  n/a  CLEAR   2011-07-09 02:40:00
health_check  Unmounted  n/a  CLEAR  2011-07-09
02:40:00
health_check   Mounted  n/a CLEAR  2011-07-09 02:40:00
health_check Unavailable  n/a  CLEAR  2011-07-09
02:40:00
health_check Maintenance n/a  CLEAR  2011-07-09
02:40:00
sql_response       time  n/a  CLEAR  2011-07-09
02:38:50
sga_pool_wastage java_free_pct n/a  CLEAR  2011-07-09
02:28:
58
UserAudit username DBSNMP_example  CLEAR 2011-07-09
02:32:48
-----------------------------------------------------
Agent is Running and Ready
```

Table 24–8 (Cont.) emctl Commands for Management Agent

emctl Command	Description

emctl status agent mcache <target name>,<target type>,<metric>	Lists the names of the metrics for which the values are present in the metric cache. See the following example for a simple host target:
	`bash-3.00$ emctl status agent mcache example.us.oracle.com,host`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`---------- Metric cache contains value for following metrics at 2011-07-09 02:54:47`
	`CPUUsage`
	`DiskActivity`
	`FileMonitoring`
	`LPAR Performance on AIX`
	`Load`
	`Network`
	`PagingActivity`
	`---------- Agent is Running and Ready`
	The metrics listed above are the ones whose values are present in the metric cache.

emctl reload agent dynamicproperties [<Target_name>:<Target_Type>]...	Recomputes the dynamic properties of a target and generates the dynamic properties for the target.
	Sample output for oracle_database is as follows:
	`bash-3.00$ emctl reload agent dynamicproperties database:oracle_database`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`---------- EMD recompute dynprops completed successfully`

emctl pingOMS [agent]	Pings the OMS to check if the agent is able to connect to the OMS. Agent will wait for the reverse ping from the OMS so that agent can say the pingOMS is successful.
emctl config agent getTZ	Gets the current timezone set in the environment.
emctl config agent getSupportedTZ	Prints the supported timezone based on the setting in the environment.
emctl config console <fileloc> [<EM loc>]	Allows you to configure the console based on the configuration entries that you have mentioned in the file <fileloc>. <EM loc> is optional and can be used to operate on a different Oracle Home.
emctl config [agent] listtargets [<EM loc>]	Lists all targets present in targets.xml. <EM loc> is optional and can be used to operate on a different Oracle Home.

Table 24–8 (Cont.) emctl Commands for Management Agent

emctl Command	Description

emctl control agent runCollection <target_name>:<target_type> <metric_name>	Allows to manually run the collections for a particular metric of a target. Sample output is as follows:
	`emctl control agent runCollection example.us.oracle.com:host CPUUsage`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`----------`
	`EMD runCollection completed successfully`

emctl getcurdir agent	Prints the current working directory you are in (pwd).
emctl resetTZ agent	Resets the timezone of the agent. Stop the agent first and then run this command to change the current timezone to a different timezone. Then start the agent.

emctl getversion agent	Prints the version of the agent. Sample output is as follows:
	`./emctl getversion agent`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control Agent 12.1.0.1.0`

emctl dumpstate agent <component> . . .	Generates the dumps for the agent. This command allow you to analyze the memory/cpu issues of the agent. Sample output is as follows:
	`./emctl dumpstate agent`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`Dumpstate succeeded`

emctl gensudoprops	Generates the sudo properties of the agent.
emctl clearsudoprops	Clears the sudo properties.
emctl clearstate	Clears the state directory contents. The files that are located under $ORACLE_HOME/sysman/emd/state will be deleted if this command is run. The state files are the files which are ready for the agent to convert them into corresponding xml files.

emctl getemhome	Prints the agent home directory. The sample output is as follows:
	`bash-3.00$ emctl getemhome`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`EMHOME=/scratch/aime/gcagent/ngagent/agent_inst`

Table 24–8 (Cont.) emctl Commands for Management Agent

emctl Command	Description
emctl start blackout <Blackoutname> [-nodeLevel] [<Target_ name>[:<Target_Type>]].... [-d <Duration>]	Starts blackout on a target. <Target_name:Target_type> defaults to local node target if not specified. If -nodeLevel is specified after <Blackoutname>,the blackout will be applied to all targets and any target list that follows will be ignored. Duration is specified in [days] hh:mm
emctl stop blackout <Blackoutname>	Stops the blackout that was started on a particular target. Only those blackouts that are started by the emctl tool can be stopped using emctl. This command cannot stop the blackouts that are started using the Console or emcli.
emctl status blackout [<Target_name>[:<Target_ Type>]]....	Provides the status of the blackout of the target. The status includes the type of blackout, whether one time, repeating, or a scheduled blackout. This command also specifies whether the blackout has started or stopped.
emctl secure agent [registration password]	Secures the agent against an OMS. The registration password must be provided.
emctl unsecure agent	Unsecures the agent. This will make the agent unsecure and the agent's port will be changed to http port.
emctl verifykey	Verifies the communication between the OMS and agent by sending pingOMS.
emctl deploy agent [-s <install-password>] [-o <omshostname:consoleSrvP ort>] [-S] <deploy-dir> <deploy-hostname>:<port> <source-hostname>	'agent' creates and deploys only the agent. [-s <password>]: Install password for securing agent. [-S]: Password will be provided in STDIN. [-o <omshostname:consoleSrvPort>]: The OMS Hostname and console servlet port. Choose the unsecured port. <deploy-dir> : Directory to create the shared (state-only) installation port. <deploy-hostname:port> : Host name and port of the shared (state-only) installation. Choose unused port. <source-hostname>: The host name of the source install. Typically the machine where EM is installed. This is searched and replaced in targets.xml by the host name provided in argument <deploy-hostname:port>. <sid>: The instance of the remote database. Only specified when deploying "dbconsole".

24.8 Using emctl.log File

The `emctl.log` file is a file that captures the results of all emctl commands you run. For Management Agent, the log file resides in the `$AGENT_INSTANCE_HOME/sysman/log` directory of the Management Agent, and for Management Service, the log file resides in the `$OMS_INSTANCE_HOME/sysman/log` directory. The file is updated every time you run an emctl command. If your emctl command fails for some reason, access this log file to diagnose the issue.

For example, run the following command from the Oracle home directory of the Management Agent to check its status:

```
<agent_home>emctl status agent
```

After running the command, navigate to the log directory to view the following information in the `emctl.log` file:

```
1114306 :: Wed Jun 10 02:29:36 2011::AgentLifeCycle.pm: Processing status agent
1114306 :: Wed Jun 10 02:29:36 2011::AgentStatus.pm:Processing status agent
1114306 :: Wed Jun 10 02:29:37 2011::AgentStatus.pm:emdctl status returned 3
```

Here, the first column, that is, 1114306, is the PID that was used to check the status. The second column shows the date and time when the command was run. The third column mentions the Perl script that was run for the command. The last column describes the result of the command, where it shows the progress made by the command and the exit code returned for the command. In this case, the exit code is 3, which means that the Management Agent is up and running.

Similarly, for the Management Service, you can run the following command from the Oracle home directory of the Management Service to check its status:

```
<agent_home>emctl status oms
```

In another example, run the following command from the Oracle home directory of the Management Agent to upload data:

```
<agent_home>emctl upload agent
```

After running the command, navigate to the log directory to view the following information in the `emctl.log` file:

```
1286220 :: Tue Jun  9 07:13:09 2011::AgentStatus.pm:Processing upload
1286220 :: Tue Jun  9 07:13:10 2011::AgentStatus.pm:emdctl status agent returned 3
1286220 :: Tue Jun  9 07:13:41 2011::AgentStatus.pm: emdctl upload returned with
exit code 6
```

Here, the entries are similar to the entries in the first example, but the exit code returned is 6, which means the upload operation is failing for some reason.

The exit codes returned depend on the emctl command executed. In general, exit code of zero means success and any exit code other than zero means failure. For details about the cause of failure, view the error message.

25

Locating and Configuring Enterprise Manager Log Files

When you install the Oracle Management Agent (Management Agent) or the Oracle Management Service (OMS), Enterprise Manager automatically configures the system to save certain informational, warning, and error information to a set of log files.

Log files can help you troubleshoot potential problems with an Enterprise Manager installation. They provide detailed information about the actions performed by Enterprise Manager and whether or not any warnings or errors occurred.

This chapter not only helps you locate and review the contents of Enterprise Manager log files, but also includes instructions for configuring the log files to provide more detailed information to help in troubleshooting or to provide less detailed information to save disk space.

This chapter contains the following sections:

25.1 Managing Log Files

Many Enterprise Manager components generate log files containing messages that record errors, notifications, warnings, and traces.

Table 25–1 describes the columns in the Log Message table. For any given component, the optional column may not be populated in the message.

Table 25–1 Message Columns

Column Name	Description
Time	The date and time when the message was generated. This reflects the local time zone.
Message Type	The type of message. Possible values are: Incident Error Warning, Notification, and Trace. In addition, the value Unknown may be used when the type is not known.
Message ID	The ID that uniquely identifies the message within the component. The ID consists of a prefix that represents the component, followed by a dash, then a 5-digit number. For example: `OHS-51009`
Message	The text of the error message.

Table 25–1 (Cont.) Message Columns

Column Name	Description
Target (Expanded)	Expanded target name.
Target	Target name
Target Type	Target type
Execution Context	The Execution Context ID (ECID), which is a global unique identifier of the execution of a particular request in which the originating component participates. You can use the ECID to correlate error messages from different components. The Relationship ID, which distinguishes the work done in one thread on one process, from work done by any other threads on this and other processes, on behalf of the same request.
Component	The component that originated the message.
Module	The identifier of the module that originated the message.
Incident ID	The identifier of the incident to which this message corresponds.
Instance	The name of the Oracle instance to which the component that originated the message belongs.
Message Group	The name of the group to which this message belongs.
Message Level	The message level, represented by an integer value that qualifies the message type. Possible values are from 1 (highest severity) through 32 (lowest severity).
Hosting Client	The identifier for the client or security group to which this message relates.
Organization	The organization ID for the originating component. The ID is `oracle` for all Oracle components.
Host	The name of the host where the message originated.
Host IP Address	The network address of the host where the message originated.
User	The name of the user whose execution context generated the message.
Process ID	The ID for the process or execution unit that generated the message.
Thread ID	The ID of the thread that generated the message.
Upstream Component	The component that the originating component is working with on the client (upstream) side.
Downstream Component	The component that the originating component is working with on the server (downstream) side.
Detail Location	A URL linking to additional information regarding the message.
Supplemental Detail	Supplemental information about the event, including more detailed information than the message text.
Target Log Files	Link to the log files page for this target.
Log File	Log file that this message contains.

Using Log Viewer, you can do the following:

- Viewing Log Files and Their Messages
- Searching Log Files
- Downloading Log Files

25.1.1 Viewing Log Files and Their Messages

You can view the messages for all of the entities in a domain, an Oracle WebLogic Server, a component, or an application.

For example, to view the log files and their messages:

1. From the **Enterprise** menu, select **Monitoring**, then select **Logs**.

   or

   From the **Targets** menu, select **Middleware**, click a domain. From the **Farm** menu, select **Logs**, then select **View Log Messages**.

   The Log Messages page is displayed.

2. Expand **Selected Targets** and in the row for a particular component or application, click the **Target Log Files** icon.

   The Log Files page is displayed. On this page, you can see a list of log files related to the Managed Server.

3. Select a file and click **View Log File.**

   The View Log File page is displayed. On this page, you can view the list of messages.

4. To view the details of a message, select the message.

   By default, the messages are sorted by time, in ascending order. You can sort the messages by the any of the columns, such as Message Type, by clicking the column name.

5. To view messages that are related by time or ECID, click **View Related Messages** and select **by Time** or **by ECID (Execution Context ID).**

   The Related Messages page is displayed.

25.1.2 Searching Log Files

You can search for diagnostic messages using the Log Messages page. By default, this page shows a summary of the logged issues for the last hour.

You can modify the search criteria to identify messages of relevance. You can view the search results in different modes, allowing ease of navigation through large amounts of data.

The following sections describe how to search log files:

- Searching Log Files: Basic Searches
- Searching Log Files: Advanced Searches

25.1.2.1 Searching Log Files: Basic Searches

You can search for all of the messages for all of the entities in a domain, an Oracle WebLogic Server, a component, or an application.

For example, to search for messages for a domain:

1. From the **Targets** menu, select **Middleware**, click a domain. From the **Farm** menu, select **Logs**, then select **View Log Messages**.

   or

   From the **Enterprise** menu, select **Monitoring**, then select **Logs**.

The Log Messages page displays a Search section and a table that shows a summary of the messages for the last hour.

2. In the Date Range section, you can select either:
   - **Most Recent:** If you select this option, select a time, such as 3 hours. The default is 1 hour.
   - **Time Interval:** If you select this option, select the calendar icon for **Start Date.** Select a date and time. Then, select the calendar icon for **End Date.** Select a date and time.
3. In the Message Types section, select one or more of the message types.
4. You can specify more search criteria, as described in Searching Log Files: Advanced Searches.
5. Click **Search.**
6. To help identify messages of relevance, in the table, for **Show,** select one of the following modes:
   - **Messages:** Shows the matching messages.

     To see the details of a particular message, click the message. The details are displayed below the table of messages.

     To view related messages, select a message, then click **View Related Messages** and select **by Time** or **by ECID (Execution Context ID).**
   - **Group by Message Type:** Summarizes the matching messages by grouping them based on message type at the target level.

     To see the messages, click the count in one of the message type columns. The Messages by Message Type page is displayed. To see the details of a particular message, click the message. The details are displayed below the table of messages.
   - **Group by Message ID:** Summarizes the matching messages by grouping them based on message ID, message type, and module IDs at the target level.

     To see the associated messages, click the count in the **Occurrences** column. The Messages by Message ID page is displayed. To see the details of a particular message, click the message. The details are displayed below the table of messages.

25.1.2.2 Searching Log Files: Advanced Searches

You can refine your search criteria using the following controls in the Log Messages page:

- **Message:** You can select an operator, such as **contains** and then enter a value to be matched.
- **Add Fields:** Click this to specify additional criteria, such as Host, which lets you narrow the search to particular hosts. Then click **Add.**

  For each field you add, select an operator, such as **contains** and then enter a value to be matched.
- **Selected Targets:** Expand this to see the targets that are participating in the search. To add targets, click **Add** and provide information in the dialog box. To remove targets, select the target and click **Remove.**

25.1.3 Downloading Log Files

You can download the log messages to a file. You can download either the matching messages from a search or the messages in a particular log file.

To download the matching messages from a search to a file:

1. From the **Enterprise** menu, select **Monitoring**, then select **Logs**.

   or

   From the **Targets** menu, select **Middleware**, click a domain. From the **Farm** menu, select **Logs**, then select **View Log Messages**.

   The Log Messages page is displayed.

2. Search for particular types of messages as described in Searching Log Files: Basic Searches.

3. Select a file type by clicking **Export Messages to File** and select one of the following:

   - **As Oracle Diagnostic Log Text (.txt)**
   - **As Oracle Diagnostic Log Text (.xml)**
   - **As Comma-Separated List (.csv)**

   An Opening dialog box is displayed.

4. Select either **Open With** or **Save to Disk.** Click **OK.**

To export specific types of messages or messages with a particular Message ID to a file:

1. From the **Enterprise** menu, select **Monitoring**, then select **Logs**.

   or

   From the **Targets** menu, select **Middleware**, click a domain. From the **Farm** menu, select **Logs**, then select **View Log Messages**.

   The Log Messages page is displayed.

2. Search for particular types of messages as described in Searching Log Files: Basic Searches.

3. For **Show,** select **Group by Message Type** or **Group by Message ID.**

4. To download the messages into a file, if you selected Group by Message Type, select the link in one of the columns that lists the number of messages, such as the Errors column. If you selected Group by Message ID, select one of the links in the Occurrences column.

   The Messages by Message Type page or Message by Message ID is displayed.

5. Select a file type by clicking the arrow near **Export Messages to File.**

   You can select one of the following:

   - **As Oracle Diagnostic Log Text (.txt)**
   - **As Oracle Diagnostic Log Text (.xml)**
   - **As Comma-Separated List (.csv)**

   An Opening dialog box is displayed.

6. Select either **Open With** or **Save to Disk.** Click **OK.**

To download the log files for a specific component:

1. From the **Enterprise** menu, select **Monitoring**, then select **Logs**.

   or

   From the **Targets** menu, select **Middleware**, click a domain. From the **Farm** menu, select **Logs**, then select **View Log Messages**.

   The Log Messages page is displayed.

2. Click **Target Log Files.**

   The Log Files page is displayed. On this page, you can see a list of log files related to the component or application.

3. Select a log file and click **Download.**

4. An Opening dialog box is displayed.

5. Select either **Open With** or **Save to Disk.** Click **OK.**

25.2 Locating and Configuring Management Agent Log and Trace Files

The following sections provide information on the log and trace files for the Oracle Management Agent:

- About the Management Agent Log and Trace Files
- Locating the Management Agent Log and Trace Files

25.2.1 About the Management Agent Log and Trace Files

Oracle Management Agent log and trace files store important information that support personnel can later use to troubleshoot problems. The agent main log is located in `$EMSTATE/sysman/log`. The log is segmented by default to 11 segments, 5MB each. The segments are named gcagent.log and gcagent.log.# where # is a number in the range of 1-10. These settings are controlled by properties in emd.properties as explained in the following sections. The latest segment is always gcagent.log and the oldest is the gcagent.log.X where X is the highest number.

The Management Agent uses these log files:

- Oracle Management Agent log file (`gcagent.log`)

  This log file contains trace, debug, information, error, or warning messages from the agent.

- Oracle Management Agent errors log file (`gcagent_errors.log`)

  This error log file contains information about errors and alerts.

- Oracle Management Agent log file (`gcagent_mdu.log`)

  This log tracks the metadata updates to the agent.

- Enterprise Manager Control log file (`emctl.log`)

  The information is saved to `emctl.log` file, when you run the Enterprise Manager Control commands. For more information about `emctl.log` file, see chapter *Starting and Stopping Enterprise Manager Components*.

25.2.1.1 Structure of Agent Log Files

The log contain individual log messages with the following format:

```
YYYY-MM-DD HH:MM:SS,### [<tid>:<thread code or code:name>] <level> -<the message>
```

Where:

- YYYY-MM-DD HH:MM:SS,### is a timestamp (in 24 hours format and ### is the fraction in msec).
- <tid> is the thread id (as a decimal number)
- <thread name or code> is the thread full name or an abbreviated hexadecimal code (see the following example).
- <level> is the logging level that can be one of (in ascending order of importance): DEBUG, INFO, WARN, ERROR, FATAL.
- <the message> is the free text message that is being logged. The message can contain new lines and spawn multiple lines.

For example:

```
2011-06-07 15:00:00,016 [1:3305B9:main] DEBUG - ADR_BASE='/ade/example_
user/oracle/example/agentStateDir' 2011-06-07 15:00:01,883 [1:3305B9] INFO - Agent
is starting up
```

25.2.2 Locating the Management Agent Log and Trace Files

The log and trace files for the Agent are written in the Agent runtime directory. You can find the runtime directory by using this command:

```
$ emctl getemhome
```

The log and trace files will be located at `<EMHOME>/sysman/log`.

25.2.3 Setting Oracle Management Agent Log Levels

Every log message is logged using a specific log level. The log levels are ordered in priority order: DEBUG, INFO, WARN, ERROR, and FATAL. The log setting determines the minimum level that will be included in the log. For example, if the log level is set to INFO (the default), only log messages of level INFO and above (INFO, WARN, ERROR and FATAL) are going to be included in the log.

Agent logging is based on log4j so the logging configuration is also log4j-based. The logging configuration properties are located in emd.properties and they all start with the prefix "Logger".

The "rootCategory" property controls the default log level. It is by default set to INFO. For example:

```
Logger.log4j.rootCategory=INFO, Rolling, Errors, Test
```

The agent uses an embedded HTTP server (jetty) to service client requests made on HTTP. There is a dedicated class which enables jetty logging as shown below:

```
Logger.log4j.category.oracle.sysman.gcagent.comm.agent.http.HTTPListener=INFO
```

This can be configured to DEBUG to try to troubleshoot certain communication problems, but the number of entries logged for DEBUG setting cannot be controlled. It is recommended that you configure the DEBUG option only under the direction of Oracle Support as in doing so may cause other valuable logging to be sacrificed.

25.2.3.1 Setting gcagent.log

The gcagent.log is configured using properties that starts with "Logger.log4j.appender.Rolling". The following is a sample gcagent.log:

```
Logger.log4j.appender.Rolling=org.apache.log4j.RollingFileAppender
Logger.log4j.appender.Rolling.File=/ade/user_
name/oracle/agentStateDir/sysman/log/gcagent.log
Logger.log4j.appender.Rolling.Append=true
Logger.log4j.appender.Rolling.MaxFileSize=5000000
Logger.log4j.appender.Rolling.MaxBackupIndex=10
Logger.log4j.appender.Rolling.layout=oracle.sysman.gcagent.util.logging.GCPattern
```

Where:

- "File" property controls the location and name of the main log. It is recommended that you do not change the setting for it.
- "Append" determines if the log file should be appended (true) or overwritten (false) on agent startup.
- " MaxFileSize" determines the size of each of the log segments.
- "MaxBackupIndex" determines the number of "backup" segments for the log (the total number of segments are this number plus one).

25.2.3.2 Setting gcagent_error.log

The gcagent_errors.log is a subset of the gcagent.log. The following is a sample gcagent_error.log file:

```
Logger.log4j.appender.Errors=org.apache.log4j.FileAppender
Logger.log4j.appender.Errors.File=/ade/user_
name/oracle/agentStateDir/sysman/log/gcagent_errors.log
Logger.log4j.appender.Errors.Append=true
Logger.log4j.appender.Errors.Threshold=ERROR
```

The above gcagent_errors.log contains log messages of ERROR and FATAL levels. The log is not segmented and it has no size limit.

25.2.3.3 Setting the Log Level for Individual Classes and Packages

The logging level for individual class and/or packages can also be set. The following are examples that are currently configured by default:

```
# Set the class loaders to level INFO
Logger.log4j.category.oracle.sysman.gcagent.metadata.impl.ChainedClassLoader=INFO

Logger.log4j.category.oracle.sysman.gcagent.metadata.impl.ReverseDelegationClassLo
ader=INFO
Logger.log4j.category.oracle.sysman.gcagent.metadata.impl.PluginLibraryClassLoader
=INFO
Logger.log4j.category.oracle.sysman.gcagent.metadata.impl.PluginClassLoader=INFO
```

The above configuration changed the default level of logging for the four classes to be INFO. When the default level of logging is INFO it does not make any difference but if the default log level is set to DEBUG (when debugging the code) it will prevent those four classes from logging at DEBUG level (as they are normally too verbose).

The reverse is also true, for example if the following configuration is added (not set by default):

```
Logger.log4j.category.oracle.sysman.gcagent.metadata.impl.collection=DEBUG
```

It will cause all classes in the "oracle.sysman.gcagent.metadata.impl.collection" package to log at DEBUG level even if the default log level is INFO.

25.2.3.4 Setting gcagent_mdu.log

A set of entries are created in the gcagent_mdu.log file for each client command that modifies target instances, target instance collections, or blackouts. Entries are as follows:

```
2011-08-18 22:56:40,467 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] - SAVE
TARGET(S)
<Target IDENTIFIER="TARGET_GUID=6A3A159D0BB320C50B7926E0671A1A98"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_NAME="EM Management
Beacon" NAME="EM Management Beacon" TYPE="oracle_beacon"/>
<Target IDENTIFIER="TARGET_GUID=51F9BBC6F5B833058F4278B51E496000"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_NAME="mytestBeacon"
NAME="mytestBeacon" TYPE="oracle_beacon"><Property VALUE="***"
NAME="proxyHost"/><Property VALUE="***" NAME="proxyPort"/><Property VALUE="***"
NAME="dontProxyFor"/></Target>
<Target IDENTIFIER="TARGET_GUID=7C4336B536C9F241DBCAC4D1D082AD22"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_NAME="CSAcollector"
NAME="CSAcollector" TYPE="oracle_csa_collector"><Property VALUE="***"
NAME="recvFileDir"/></Target>
<Target IDENTIFIER="TARGET_GUID=207B57A3FE300C86F81FE7D409F5DD1C"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_NAME="Oemrep_Database"
NAME="Oemrep_Database" TYPE="oracle_database"><Property VALUE="***"
NAME="MachineName"/><Property VALUE="***" NAME="Port"/><Property VALUE="***"
NAME="SID"/><Property VALUE="***" NAME="OracleHome"/><Property ENCRYPTED="FALSE"
VALUE="***" NAME="UserName"/><Property ENCRYPTED="FALSE" VALUE="***"
NAME="Role"/><Property ENCRYPTED="FALSE" VALUE="***" NAME="password"/></Target>
<Target IDENTIFIER="TARGET_GUID=0C48C5AE0FAFB42ED91F897FF398FC84"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_NAME="Management Services
and Repository" NAME="Management Services and Repository" TYPE="oracle_
emrep"><Property VALUE="***" NAME="ConnectDescriptor"/><Property ENCRYPTED="FALSE"
VALUE="***" NAME="UserName"/><Property ENCRYPTED="FALSE" VALUE="***"
NAME="password"/><AssocTargetInstance ASSOC_TARGET_TYPE="oracle_oms" ASSOC_TARGET_
NAME="adc2190447.us.oracle.com:41034_Management_Service" ASSOCIATION_NAME="app_
composite_contains"/><AssocTargetInstance ASSOC_TARGET_TYPE="oracle_oms" ASSOC_
TARGET_NAME="adc2190447.us.oracle.com:41034_Management_Service" ASSOCIATION_
NAME="internal_contains"/><CompositeMembership><Member ASSOCIATION=""
NAME="adc2190447.us.oracle.com:41034_Management_Service_CONSOLE" TYPE="oracle_oms_
console"/><Member ASSOCIATION="" NAME="adc2190447.us.oracle.com:41034_Management_
Service_PBS" TYPE="oracle_oms_pbs"/><Member ASSOCIATION=""
NAME="adc2190447.us.oracle.com:41034_Management_Service" TYPE="oracle_
oms"/></CompositeMembership></Target>
<Target IDENTIFIER="TARGET_GUID=DF64B4A7C0F2EEBA7894EA3AD4CAF61E"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_
NAME="adc2190447.us.oracle.com:41034_Management_Service"
NAME="adc2190447.us.oracle.com:41034_Management_Service" TYPE="oracle_
oms"><Property VALUE="***" NAME="InstanceHome"/><Property VALUE="***"
NAME="OracleHome"/><AssocTargetInstance ASSOC_TARGET_TYPE="oracle_oms_console"
ASSOC_TARGET_NAME="adc2190447.us.oracle.com:41034_Management_Service_CONSOLE"
ASSOCIATION_NAME="app_composite_contains"/><AssocTargetInstance ASSOC_TARGET_
TYPE="oracle_oms_pbs" ASSOC_TARGET_NAME="adc2190447.us.oracle.com:41034_
Management_Service_PBS" ASSOCIATION_NAME="app_composite_
contains"/><AssocTargetInstance ASSOC_TARGET_TYPE="oracle_oms_console" ASSOC_
TARGET_NAME="adc2190447.us.oracle.com:41034_Management_Service_CONSOLE"
ASSOCIATION_NAME="internal_contains"/><AssocTargetInstance ASSOC_TARGET_
TYPE="oracle_oms_pbs" ASSOC_TARGET_NAME="adc2190447.us.oracle.com:41034_
```

```
Management_Service_PBS" ASSOCIATION_NAME="internal_
contains"/><CompositeMembership><MemberOf ASSOCIATION="" NAME="Management Services
and Repository" TYPE="oracle_emrep"/><Member ASSOCIATION=""
NAME="adc2190447.us.oracle.com:41034_Management_Service_CONSOLE" TYPE="oracle_oms_
console"/><Member ASSOCIATION="" NAME="adc2190447.us.oracle.com:41034_Management_
Service_PBS" TYPE="oracle_oms_pbs"/></CompositeMembership></Target>
<Target IDENTIFIER="TARGET_GUID=4D290260F13596502EFD8F3E22752404"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_
NAME="adc2190447.us.oracle.com:41034_Management_Service_CONSOLE"
NAME="adc2190447.us.oracle.com:41034_Management_Service_CONSOLE" TYPE="oracle_oms_
console"><Property VALUE="***" NAME="InstanceHome"/><Property VALUE="***"
NAME="OracleHome"/><CompositeMembership><MemberOf ASSOCIATION="" NAME="Management
Services and Repository" TYPE="oracle_emrep"/><MemberOf ASSOCIATION=""
NAME="adc2190447.us.oracle.com:41034_Management_Service" TYPE="oracle_
oms"/></CompositeMembership></Target>
<Target IDENTIFIER="TARGET_GUID=D0A23AE06A9E678221B075A216364541"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_
NAME="adc2190447.us.oracle.com:41034_Management_Service_PBS"
NAME="adc2190447.us.oracle.com:41034_Management_Service_PBS" TYPE="oracle_oms_
pbs"><Property VALUE="***" NAME="InstanceHome"/><Property VALUE="***"
NAME="OracleHome"/><CompositeMembership><MemberOf ASSOCIATION="" NAME="Management
Services and Repository" TYPE="oracle_emrep"/><MemberOf ASSOCIATION=""
NAME="adc2190447.us.oracle.com:41034_Management_Service" TYPE="oracle_
oms"/></CompositeMembership></Target>
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
```

For the batch of saved targets in the above example, the original request came in at 22:56:40 and the list of targets saved are found in the line(s) following the SAVE TARGET(S) message. In this case, there were 8 targets. The result of saving the targets is available in the next 8 lines (for the same thread) and in this case all were saved successfully by 22:57:10.

The pattern is the same for saved collection items (or collections) and blackouts.

The logging configuration for the gcagent_mdu log is specified in `emd.properties` but you must not be modify this log. For example, these entries are logged at INFO level, which means that if you decided to save space and change this to WARN only by editing the mdu log entries in the `emd.properties` file, you will have effectively disabled this log.

25.2.3.5 Setting the TRACE Level

The following `_enableTrace` property when set to "true" will enable the TRACE logging level that shows as DEBUG messages.

```
Logger._enableTrace=true
```

The default log level for the agent log must be set to DEBUG for the tracing level to work.

25.3 Locating and Configuring Oracle Management Service Log and Trace Files

The following sections describe how to locate and configure the OMS log files:

- Locating Oracle Management Service Log and Trace Files
- Controlling the Size and Number of Oracle Management Service Log and Trace Files
- Controlling the Contents of the Oracle Management Service Trace File
- Controlling the Oracle WebLogic Server and Oracle HTTP Server Log Files

25.3.1 About the Oracle Management Service Log and Trace Files

OMS log and trace files store important information that Support personnel can later use to troubleshoot problems. OMS uses the following six types of log files:

- Oracle Management Service log file (`emoms.log`)

  The Management Service saves information to the log file when it performs an action (such a starting or stopping) or when it generates an error. This is a log file for console application.

- Oracle Management Service trace file (`emoms.trc`)

  OMS trace file provides an advanced method of troubleshooting that can provide support personnel with even more information about what actions the OMS was performing when a particular problem occurred. This is a trace file for Console application.

- Oracle Management Service log file (`emoms_pbs.log`)

  The Management Service saves information to this log file for background modules such as the loader, job system, event system, notification system, and so on. This file contains messages logged at ERROR or WARN levels.

- Oracle Management Service trace file (`emoms_pbs.trc`)

  This trace file provides additional logging for the background modules such as the loader, job system, event system, notification system, and so on when DEBUG or INFO level logging is enabled for these modules. This file can provide Support personnel with even more information about actions these modules were performing when a particular problem occurred.

- Enterprise Manager Control log file (`emctl.log`)

  The information is saved to `emctl.log` file, when you run the Enterprise Manager Control commands. For more information about emctl.log file, see chapter *Starting and Stopping Enterprise Manager Components*.

25.3.2 Locating Oracle Management Service Log and Trace Files

OMS log and trace files are stored in the following location:

```
<EM_INSTANCE_BASE>/em/<OMS_NAME>/sysman/log/
```

Where, <EM_INSTANCE_BASE> is the OMS Instance Base directory. By default, the OMS Instance Base directory is gc_inst, which is present under the parent directory of the Oracle Middleware Home.

For example, if the Oracle Middleware Home is /u01/app/Oracle/Middleware, then the instance base directory is /u01/app/Oracle/gc_inst, and the log and trace files are available in /u01/app/Oracle/gc_inst/em/EMGC_OMS1/sysman/log/ directory path.

25.3.3 Controlling the Size and Number of Oracle Management Service Log and Trace Files

OMS log and trace files increases in size over time as information is written to the files. However, the files are designed to reach a maximum size. When the files reach the predefined maximum size, the OMS renames (or rolls) the logging information to a new file name and starts a new log or trace file. This process keeps the log and trace files from growing too large.

As a result, you will often see multiple log and trace files in the OMS log directory. The following example shows one archived log file and the current log file in the /u01/app/Oracle/gc_inst/em/EMGC_OMS1/sysman/log/ directory:

```
emoms.log
emoms.log.1
```

To control the maximum size of the OMS log and OMS trace files, as well as the number of rollover files, run the following command, and specify details as described in Table 25–2:

```
emctl set property -name <property> -value <property value> -module logging
```

Note: For Oracle Enterprise Manager 12*c*, you do not have to restart OMS for the changes to take effect.

Table 25–2 Oracle Management Service Log File Properties in the emomslogging.properties File

Property	Purpose	Example
log4j.appender.emlogAppender.MaxFileSize	When OMS log file reaches this size, then OMS copies the logging data to a new rollover file and creates a new emoms.log log file. The size of the log is specified in units of bytes.	log4j.appender.emlogAppender.MaxFileSize=20000000

Table 25–2 (Cont.) Oracle Management Service Log File Properties in the emomslogging.properties File

Property	Purpose	Example
log4j.appender.emlogAppender. MaxBackupIndex	This optional property indicates how many times OMS will rollover the log file to a new file name before deleting logging data. **Note:** Because the log file does not contain as much data as the trace file, it is usually not necessary to create more than one rollover file.	log4j.appender.emlogAppender. MaxBackupIndex=1
log4j.appender.emtrcAppender. MaxFileSize	When the OMS trace file reaches this size, then OMS copies the logging data to a new rollover file and creates a new `emoms.trc` log file.	log4j.appender.emtrcAppender. MaxFileSize=5000000
log4j.appender.emtrcAppender. MaxBackupIndex	This property indicates how many times the OMS will rollover the trace file to a new file name before deleting tracing data.	log4j.appender.emtrcAppender. MaxBackupIndex=10

25.3.4 Controlling the Contents of the Oracle Management Service Trace File

By default, the OMS will save all critical and warning messages to the `emoms.trc` file. However, you can adjust the amount of logging information that the OMS generates.

To change the amount of logging information generated by the OMS, run the following command:

```
emctl set property -name "log4j.rootCategory" -value "<LEVEL>, emlogAppender,
emtrcAppender" -module logging
```

Note: If you change the `root` logging level for the `emoms.trc` file, then a lot of messages are written to the trace file filling up the space quickly, and potentially slowing down the system. Run the following command to enable debug selectively for specific modules that need to be assessed:

```
emctl set property -name <logging module> -value DEBUG -module
logging
```

Where, `<logging module>` represents the logging module from a specific subsystem.

For example, `oracle.sysman.emdrep.dbjava.loader.`

The logging level can be changed for specific modules by running the following command:

```
emctl set property -name "<CATEGORY>" -value "<LEVEL>" -module
logging
```

where LEVEL can be DEBUG, INFO, WARN, or ERROR, and CATEGORY is specific to the module for which level has to be changed. To change the logging module, contact Oracle Support.

25.3.5 Controlling the Oracle WebLogic Server and Oracle HTTP Server Log Files

Oracle Management Service is a J2EE application deployed on an Oracle WebLogic Server. Different components of the Oracle WebLogic Server generate their own log files. These files contain important information that can be used later by support personnel to troubleshoot problems.

Table 25–3 lists the location of the log files for some components.

Table 25–3 Component Log File Location

Component	Location

Oracle HTTP Server (OHS)	`<EM_INSTANCE_BASE>/<webtier_instance_name>/diagnostics/logs/OHS/<ohs_name>`
	For example,
	`/u01/app/Oracle/gc_inst/WebTierIH1/diagnostics/logs/OHS/ohs1`

OPMN	`<EM_INSTANCE_BASE>/<webtier_instance_name>/diagnostics/logs/OPMN/<opmn_name>`
	For example,
	`/u01/app/Oracle/gc_inst/WebTierIH1/diagnostics/logs/OPMN/opmn`

Oracle WebLogic	The log data from WebLogic will be at:
	`<EM_INSTANCE_BASE>/user_projects/domains/<domain_name>/servers/<SERVER_NAME>/logs/<SERVER_NAME>.log`
	This log can be restricted, rotated by size, time, and other conditions from the WebLogic Console. The default settings are:
	■ In production mode, they are rotated at a default of 5MB.
	■ The log level is WARNING.
	■ The number files are restricted to 10.
	For example,
	`/u01/app/Oracle/gc_inst/user_projects/domains/GCDomain/servers/EMGC_OMS1/logs/EMGC_OMS1.log`
	The messages written to sysout and syserr will be available in the .out files. They cannot be rotated by size or time. They are rotated only when the server starts. They are located at:
	`<EM_INSTANCE_BASE>/user_projects/domains/<domain_name>/servers/<SERVER_NAME>/logs/<SERVER_NAME>.out`
	For example,
	`/u01/app/Oracle/gc_inst/user_projects/domains/GCDomain/servers/EMGC_OMS1/logs/EMGC_OMS1.out`

By default, the Enterprise Manager Cloud Control configures OHS logs to rollover periodically to a new file, so that each file does not grow too large in size. You must also ensure that you delete the old rollover files periodically to free up the disk space. You can use an operating system scheduler, like cron on UNIX, to periodically delete the rollover files.

For instructions on controlling the size and rotation of these log files, refer to chapter "Managing Log Files and Diagnostic Data" in *Oracle Fusion Middleware Administrator's Guide*.

For information about configuring Enterprise Manager to view Fusion Applications PL/SQL and C diagnostic log files, see chapter "Managing Oracle Fusion Applications Log Files and Diagnostic Tests" in *Oracle Fusion Applications Administrator's Guide*.

26

Maintaining and Troubleshooting the Management Repository

This chapter describes maintenance and troubleshooting techniques for maintaining a well-performing Management Repository.

Specifically, this chapter contains the following sections:

26.1 Management Repository Deployment Guidelines

To be sure that your management data is secure, reliable, and always available, consider the following settings and configuration guidelines when you are deploying the Management Repository:

- Install a RAID-capable Logical Volume Manager (LVM) or hardware RAID on the system where the Management Repository resides. At a minimum the operating system must support disk mirroring and stripping. Configure all the Management Repository data files with some redundant configuration.
- Use Real Application Clusters to provide the highest levels of availability for the Management Repository.
- If you use Enterprise Manager to alert administrators of errors or availability issues in a production environment, be sure that the Cloud Control components are configured with the same level of availability. At a minimum, consider using Oracle Data Guard to mirror the Management Repository database. Configure the Data Guard environment for no data loss.

> **See Also:** *Oracle Database High Availability Architecture and Best Practices*
>
> *Oracle Data Guard Concepts and Administration*

- Oracle strongly recommends that archive logging be turned on and that a comprehensive backup strategy be in place prior to an Enterprise Manager implementation going live in a production environment. The backup strategy should include both incremental and full backups as required.

  See Also: *Oracle Enterprise Manager Cloud Control Installation and Basic Configuration* for information about the database initialization parameters required for the Management Repository

26.2 Management Repository Data Retention Policies

When the various components of Enterprise Manager are configured and running efficiently, the Oracle Management Service gathers large amounts of raw data from the Management Agents running on your managed hosts and loads that data into the Management Repository. This data is the raw information that is later aggregated, organized, and presented to you in the Cloud Control Console.

After the Oracle Management Service loads information into the Management Repository, Enterprise Manager aggregates and purges the data over time.

The following sections describe:

- The default aggregation and purging policies used to maintain data in the Management Repository.
- How you can modify the length of time the data is retained before it is aggregated and then purged from the Management Repository.

26.2.1 Management Repository Default Aggregation and Purging Policies

Enterprise Manager aggregates collected metric data by hour and by day to enhance query performance and help minimize the size of the Management Repository. Before the data is aggregated, each data point is stored in a raw metric data table. Once a day, the previous day's raw metric data is rolled up, or aggregated, into a one-hour and a one-day table. These hourly and daily records will have hourly and daily metric data averages, minimums, maximums and standard deviations respectively.

After Enterprise Manager aggregates the data, the data is then considered eligible for purging. A certain period of time must pass for data to actually be purged. This period of time is called the retention time.

The raw data, with the highest insert volume, has the shortest default retention time, which is set to 7 days. As a result, 7 days after it is aggregated into a one-hour record, a raw data point is eligible for purging.

Note: This data retention policy varies for JVMD and ADP data.

Hourly aggregate metric data records are purged after 31 days. The highest level of aggregation, one day, is kept for 12 months (roughly 365 days).

The default data retention policies are summarized in Table 26–1.

Table 26–1 Default Repository Purging Policies

Aggregate Level	Retention Time
Raw metric data	7 days
Hourly aggregated metric data	31 days

Table 26–1 (Cont.) Default Repository Purging Policies

Aggregate Level	Retention Time
Daily aggregated metric data	12 months (~365 days)

If you have configured and enabled Application Performance Management, Enterprise Manager also gathers, saves, aggregates, and purges response time data. The response time data is purged using policies similar to those used for metric data. The Application Performance Management purging policies are shown in Table 26–2.

Table 26–2 Default Repository Purging Policies for Application Performance Management Data

Aggregate Level	Retention Time
Raw response time data	24 hours
One-hour aggregated response time data	7 days
One-hour distribution response time data	24 hours
One-day aggregated response time data	31 days
One-day distribution aggregated response time data	31 days

26.2.2 Management Repository Default Aggregation and Purging Policies for Other Management Data

Besides the metric data and Application Performance Monitoring data, other types of Enterprise Manager data accumulates over time in the Management Repository.

For example, the last availability record for a target will also remain in the Management Repository indefinitely, so the last known state of a target is preserved.

26.2.3 Modifying the Default Aggregation and Purging Policies

The Enterprise Manager default aggregation and purging policies were designed to provide the most available data for analysis while still providing the best performance and least disk-space requirements for the Management Repository. As a result, you should not modify these policies to improve performance or increase your available disk space.

However, if you plan to extract or review the raw or aggregated data using data analysis tools other than Enterprise Manager, you may want to increase the amount of raw or aggregated data available in the Management Repository. You can accomplish this by increasing the retention times for the raw or aggregated data.

A PL/SQL API has been provided to modify the default retention time for the core metric data tables in the Enterprise Manager repository. Table 26–3 shows the default number of partitions retained for each of the three tables and the size of the partitions for each table. The API will allow you to change the number of partitions retained only.

Table 26–3 Core EM Metric Data Tables and Default Data Retention in the Management Repository

Table Name	Partitions Retained	Partition Size
EM_METRIC _VALUES	7	DAY
EM_METRIC_VALUES_HOURLY	32	DAY
EM_METRIC_VALUES_DAILY	12	MONTH

The PL/SQL API is:

```
gc_interval_partition_mgr.set_retention(<repository schema name>,
                                        <table name>,
                                        <number of partitions to retain>);
```

An example of using the PL/SQL API to change the number of partitions retained in EM_METRIC_VALUES (raw data) from the default of 7 to 10 follows:

```
BEGIN
  gc_interval_partition_mgr.set_retention('SYSMAN', 'EM_METRIC_VALUES', 10);
END;
/
```

26.2.4 Modifying Data Retention Policies When Targets Are Deleted

By default, when you delete a target from the Grid Control console, Enterprise Manager automatically deletes all target data from the Management Repository.

However, deleting raw and aggregated metric data for database and other data-rich targets is a resource consuming operation. Targets can have hundreds of thousands of rows of data and the act of deleting this data can degrade performance of Enterprise Manager for the duration of the deletion, especially when several targets are deleted at once.

To avoid this resource-consuming operation, you can prevent Enterprise Manager from performing this task each time you delete a target. When you prevent Enterprise Manager from performing this task, the metric data for deleted targets is not purged as part of target deletion task; instead, it is purged as part of the regular purge mechanism, which is more efficient.

In addition, Oracle strongly recommends that you do not add new targets with the same name and type as the deleted targets within 24 hours of target deletion. Adding a new target with the same name and type will result in the Grid Control console showing data belonging to the deleted target for the first 24 hours.

To disable raw metric data deletion:

1. Use SQL*Plus to connect to the Management Repository as the Management Repository user.

   The default Management Repository user is SYSMAN. For example:

   ```
   SQL> connect sysman/sysman_password;
   ```

2. To disable metric deletion, run the following SQL command.

   ```
   SQL> EXEC MGMT_ADMIN.DISABLE_METRIC_DELETION();
   SQL> COMMIT;
   ```

To enable metric deletion at a later point, run the following SQL command:

1. Use SQL*Plus to connect to the Management Repository as the Management Repository user.

   The default Management Repository user is SYSMAN. For example:

   ```
   SQL> connect sysman/oldpassword;
   ```

2. To enable metric deletion, run the following SQL command.

   ```
   SQL> EXEC MGMT_ADMIN.ENABLE_METRIC_DELETION();
   SQL> COMMIT;
   ```

26.2.5 How to Modify the Retention Period of Job History

Enterprise Manager Cloud Control has a default purge policy which removes all finished job details which are older than 30 days. This section provides details for modifying this default purge policy.

The actual purging of completed job history is implemented via a DBMS_SCHEDULER job that runs once a day in the repository database. When the job runs, it looks for finished jobs that are 'n' number of days older than the current time (value of sysdate in the repository database) and deletes these jobs. The value of 'n' is, by default, set to 30 days.

The default purge policy cannot be modified via the Enterprise Manager console, but it can be changed using SQL*Plus.

To modify this purge policy, follow these steps:

1. Log in to the repository database as the SYSMAN user, via SQL*Plus.
2. Check the current values for the purge policies using the following command:

   ```
   SQL> select * from mgmt_job_purge_policies;

   POLICY_NAME                    TIME_FRAME
   ------------------------------ ----------
   SYSPURGE_POLICY                        30
   REFRESHFROMMETALINKPURGEPOLICY          7
   FIXINVENTORYPURGEPOLICY                 7
   OPATCHPATCHUPDATE_PAPURGEPOLICY         7
   ```

   The purge policy responsible for the job deletion is called SYSPURGE_POLICY. As seen above, the default value is set to 30 days.
3. To change the time period, you must drop and re-create the policy with a different time frame:

   ```
   SQL> execute MGMT_JOBS.drop_purge_policy('SYSPURGE_POLICY');

   PL/SQL procedure successfully completed.

   SQL> execute MGMT_JOBS.register_purge_policy('SYSPURGE_
   POLICY', 60, null);

   PL/SQL procedure successfully completed.

   SQL> COMMIT;

   Commit complete.

   SQL> select * from mgmt_job_purge_policies;

   POLICY_NAME                    TIME_FRAME
   ------------------------------ ----------
   SYSPURGE_POLICY                        60
   ```

....

The above commands increase the retention period to 60 days. The time frame can also be reduced below 30 days, depending on the requirement.

You can check when the purge job will be executed next. The actual time that the purge runs is set to 5 AM repository time and can be verified using these steps:

1. Login to the Repository database using the SYSMAN account.
2. Execute the following command:

```
SQL> select job_name,
            to_char(last_start_date, 'DD-MON-YY HH24:MI:SS') last_run,
            to_char(next_run_date,   'DD-MON-YY HH24:MI:SS') next_run
from all_scheduler_jobs
where job_name ='EM_JOB_PURGE_POLICIES';

JOB_NAME              LAST_RUN           NEXT_RUN
--------------------- ------------------ ------------------
EM_JOB_PURGE_POLICIES                    07-SEP-11 05:00:00
```

   The schedule can also be verified from the Enterprise Manager console by following these steps:

   a. Select Setup > Management Service and Repository
   b. Click on the **Repository Operations** tab
   c. Find the Next Scheduled Run and Last Scheduled Run information for Job Purge in the list.

   Please note that the time of the next scheduled execution of the Job Purge does not represent the cutoff time for the retention period; the cutoff time is determined by the purge policy at the time the Job Purge runs.

26.2.6 DBMS_SCHEDULER Troubleshooting

Enterprise Manager uses the database scheduler (dbms_scheduler) to run various processes in the repository. When the dbms_scheduler is stopped or has insufficient resources to operate, the Enterprise Manager processes do not run or are delayed. The following is a list of common causes that may prohibit the dbms_scheduler from running normally.

Job Queue Processes

The dbms_scheduler uses a separate job-queue process for each job it runs. The maximum number of these processes is controlled by the database parameter, *job_queue_processes*. If all processes are in use, no new jobs will be started.

The following query returns the number of currently running jobs.

```
SELECT count(*)
FROM dba_scheduler_running_jobs;
```

If the count is close to the setting of *job_queue_processes*, it could mean that Enterprise Manager dbms_scheduler jobs cannot be started (on time). Determine if any of the running dbms_scheduler jobs are stuck and consider increasing the setting for job_queue_processes.

Job Slave Processes

The dbms_scheduler also depends on the setting of the dbms_scheduler property MAX_JOB_SLAVE_PROCESSES. If the number of running dbms_scheduler jobs exceeds this setting, no new jobs will be started. This attribute can be checked using this query.

```
SELECT value
FROM dba_scheduler_global_attribute
WHERE attribute_name='MAX_JOB_SLAVE_PROCESSES';
```

If the count equals the number of running dbms_scheduler jobs, then determine if any of the running dbms_scheduler jobs are stuck and consult the dbms_scheduler documentation about how to adjust this attribute.

DBMS_SCHEDULER Program Disabled

The dbms_scheduler has an attribute that can be set to disable this feature in the database. When set, the Enterprise Manager dbms_scheduler jobs will not run. To check if this attribute has been set (inadvertently), run this query.

```
SELECT *
FROM dba_scheduler_global_attribute
WHERE attribute_name = 'SCHEDULER_DISABLED';
```

When a row is returned, the dbms_scheduler is disabled. Execute `dbms_scheduler.set_scheduler_attribute('SCHEDULER_DISABLED', 'FALSE');`

Consult the dbms_scheduler documentation about how to remove this attribute.

Too Many Database Sessions

Each dbms_scheduler job requires two database sessions. When no more sessions are available, Enterprise Manager dbms_scheduler jobs will not run. The following two queries give the maximum number of allowed sessions and the current number of active sessions:

```
SELECT value
FROM v$parameter
WHERE name='sessions';

SELECT count(*)
FROM v$session;
```

When the current number of sessions approaches the maximum, then you should determine if any of the sessions are stuck and consult the Oracle Database documentation about how to increase the maximum number of sessions.

Also the high water mark of the number of sessions may indicate that this issue has played a role in the past:

```
select *
from v$resource_limit
where resource_name = 'sessions' ;
```

If the MAX_UTILIZATION column indicates a value that is close the maximum number of sessions, it could explain why some of the Enterprise Manager dbms_scheduler jobs may not have run (on time) in the past.

Insufficient Memory

The database may not be able to spawn a new job queue process when there is insufficient memory available. The following message in the database alert file, *Unable to spawn jobq slave processes*, in combination with, *(free memory = 0.00M)*, would be indicative of this problem. Please consult the Oracle Database documentation about how to diagnose this memory problem further.

26.3 Repository and Sizing Requirements for Fusion Middleware Monitoring in Enterprise Manager Release 12g

A Fusion Middleware target is like any other Enterprise Manager target. Therefore any repository or sizing guideline that is applicable for an Enterprise Manager target would be applicable on a Fusion Middleware target.

One major concern in the case of Fusion Middleware discovery is that too many targets may be discovered, created and monitored. This adds additional load on the OMS instance, repository and agent. In the case of very large number of targets, after target discovery Oracle recommends that users should review all the targets and their respective metrics.

Based on requirements, users should finalize which targets and metrics should be monitored and the required frequency those targets should be monitored.

After discovery, Oracle recommends you allow Fusion Middleware/ADP/JVMD monitoring to run for some duration (a few days to possibly a few weeks) and continuously monitor the database size and Operating System file system growth (in the case of ADP; ADP Manager requires a minimum of 10GB of disk space) until it becomes constant. You can then fine tune various parameters associated with these different features.

In version 12g of Enterprise Manager, both ADP and JVMD use Enterprise Manager repository as their repository. Their data are stored in the MGMT_AD4J_TS tablespace.

26.3.1 ADP Monitoring

Use the following information when utilizing ADP Monitoring.

- ADP Manager Resources Requirement

  While managing 70K managed entities, if the number of managed entities is high you must allocate resources accordingly.

Resource	Amount
Physical Memory	2 GB
Minimum Disk Space	10 GB

- ADP Data requirement

  To monitor each entity per JVM, the MGMT_AD4J_TS tablespace must have 8 MB available.

- ADP Data Retention Policy

  ADP maintains sophisticated multi-tiered logic for aggregation (or compression) of performance data. This helps to optimize performance of interaction with the internal data repository both when querying data for presentation or inserting new performance metrics.

Users who want to store longer term data should look for this section in *Acsera.properties*:

Example 26–1

```
########################
# Production setting
# NOTE: use Model.GlobalSamplingRateSecs to configure Metric.Grain.0
########################
Metric.Grain.0 0s
Metric.TableInterval.0 = 4h
Metric.DataLife.0 = 2d

Metric.Grain.1 = 3m
Metric.TableInterval.1 =1d
Metric.DataLife.1 = 8d

#Metric.Grain.2 = 30m
#Metric.TableInterval.2 = 7d
#Metric.DataLife.2 = 420d
```

Uncomment the last 3 lines for the *Metric.*.2* properties.

26.3.2 JVMD Monitoring

Use the following information when employing JVMD Monitoring.

- JVMD Manager Resources Requirement

  To manage 200-300 jvms, JVMD manager requires physical memory of 1 GB. JVMD manager caches monitoring data in the TEMP space for each pool and flushes to the database frequently. Usually, depending on the number of pools the manager is monitoring and the amount of data being gathered from each pool, the size requirement of these temporary cache files varies, but it is rare to see more than a few MBs for each pool. If this is a concern, the TEMP space should be allocated accordingly.

- JVMD Data requirement

  To monitor every JVM with OOB settings, the MGMT_AD4J_TS tablespace must have 50-100MB available.

- JVM Diagnostics Historical Data and its Retention policy

  Historical data is available at three summary levels 0, 1 and 2.

  – Summary level 0 - is raw sample data taken at the specified pool polling interval (default 2 seconds). If you look at data within one hour on the Performance Diagnostics page, it shows summary level 0 data. Level 0 data is retained for 24 hours and subsequently purged. It can be changed via the Console Setup page, but before increasing the value, you should ensure that the repository is tuned properly to handle such large amounts of data.

  – Summary level 1 - is aggregated data. If you view data after more than one hour but less than 5 hours, it is summary level 1 data. The default aggregation interval is 90 seconds. This value can be changed via the Console Setup page. Level 1 data is retained for 20 days and subsequently purged.

  – Summary level 2 - is further aggregated data. If you view data more than five hours old, it is summary level 2 data. This data is aggregated every 60 minutes. Level 2 data is retained for 400 days and subsequently purged.

There are two JVMD features that can drastically affect MGMT_AD4J_TS tablespace usage:

- JVMD Heap Dumps

  Analyzing heap requires massive tablespace resources. Oracle recommends having 5 times the size of the heap dump file you are loading free in your tablespace. Since you will have the heap dump file and know its size before you run the load script, you should ensure that you have adequate space to accommodate the dump before you load it into your database.

- Thread Traces

  While these are smaller than heaps by an order of magnitude, these are loaded into the database automatically by default when you initiate a trace at the console. The size of these traces 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. They should generally be under 100MB each, but a user utilizing a large number of these could manually fill up the database quickly. Again, since these are created only by manual intervention, you should ensure that there is adequate space to accommodate traces before initiating them.

26.4 Changing the SYSMAN Password

The SYSMAN account is the default super user account used to set up and administer Enterprise Manager. It is also the database account that owns the objects stored in the Oracle Management Repository. From this account, you can set up additional administrator accounts and set up Enterprise Manager for use in your organization.

The SYSMAN account is created automatically in the Management Repository database during the Enterprise Manager installation. You also provide a password for the SYSMAN account during the installation.

If you later need to change the SYSMAN database account password, use the following procedure:

1. Shut down all the Oracle Management Service instances that are associated with the Management Repository.
2. Stop the agent that is monitoring the target OMS and Repository.

   Failure to do this will result in the agent attempting to connect to the target with a wrong password once it is changed with SQL*Plus. This may also result in the SYSMAN account being locked which can subsequently prevent logins to the Cloud Control console to change the password of the target OMS and Repository.
3. Change the password of the SYSMAN database account using the following SQL*Plus commands:

   ```
   SQL>connect sysman/oldpassword;

   SQL>alter user sysman identified by newpassword;
   ```
4. To change the password of the SYSMAN user, enter the following command:

   ```
   emctl config oms -change_repos_pwd [-old_pwd <old_pwd>]
   [-new_pwd <new_pwd>] [-use_sys_pwd [-sys_pwd <sys_pwd>]]
   ```

   You must run this command on each Management Service in your environment.

Parameter	Description
-old_pwd	This is the current SYSMAN password.
-new_pwd	This is the new password.
-use_sys_pwd	This parameter is optional and is used to connect to the database as a SYS user.
-sys_pwd	This is the password for the SYS user.

5. In the Cloud Control console, click the **Targets** tab and then click **All Targets** on the sub tab.
6. Select the **Management Services and Repository** target and click **Configure**. Enterprise Manager displays the Monitoring Configurations page.
7. Enter the new password in the Repository password field and click **OK**.

To change the password of the SYSMAN user, enter the following command:

```
emctl config oms -change_repos_pwd [-old_pwd <old_pwd>] [-new_
pwd <new_pwd>] [-use_sys_pwd [-sys_pwd <sys_pwd>]]
```

You must run this command on each Management Service in your environment.

26.5 Dropping and Recreating the Management Repository

This section provides information about dropping the Management Repository from your existing database and recreating the Management Repository after you install Enterprise Manager.

26.5.1 Dropping the Management Repository

To recreate the Management Repository, you first remove the Enterprise Manager schema from your Management Repository database. You accomplish this task using the `-action drop` argument to the `RepManager` script, which is described in the following procedure.

To remove the Management Repository from your database:

1. Locate the `RepManager` script in the following directory of the Oracle Application Server Home where you have installed and deployed the Management Service:

   ```
   ORACLE_HOME/sysman/admin/emdrep/bin
   ```

2. At the command prompt, enter the following command:

   ```
   $PROMPT> RepManager repository_host repository_port repository_SID
   -sys_password password_for_sys_account -action drop
   ```

   In this syntax example:

   - *repository_host* is the machine name where the Management Repository database is located
   - *repository_port* is the Management Repository database listener port address, usually 1521 or 1526
   - *repository_SID* is the Management Repository database system identifier
   - *password_for_sys_account* is the password of the SYS user for the database. For example, `change_on_install`

- `-action drop` indicates that you want to drop the Management Repository

Alternatively, you can use a connect descriptor to identify the database on the RepManager command line. The connect descriptor identifies the host, port, and name of the database using a standard Oracle database syntax.

For example, you can use the connect descriptor as follows to create the Management Repository:

```
$PROMPT> ./RepManager -connect "(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)
(HOST=host1)(PORT=1521)) (CONNECT_DATE=(SERVICE_NAME=servicename)))"
-sys_password efkl34lmn -action drop
```

> **See Also:** "Establishing a Connection and Testing the Network" in the *Oracle Database Net Services Administrator's Guide* for more information about connecting to a database using connect descriptors.

26.5.2 Recreating the Management Repository

The preferred method for creating the Management Repository is to create the Management Repository during the Enterprise Manager installation procedure, which is performed using Oracle Universal Installer.

> **See Also:** *Oracle Enterprise Manager Cloud Control Installation and Basic Configuration* for information about installing Enterprise Manager.

However, if you need to recreate the Management Repository in an existing database, you can use the `RepManager` script, which is installed when you install the Management Service. Refer to the following sections for more information:

- Using the RepManager Script to Create the Management Repository
- Using a Connect Descriptor to Identify the Management Repository Database

26.5.2.1 Using the RepManager Script to Create the Management Repository

To create a Management Repository in an existing database:

1. Review the hardware and software requirements for the Management Repository as described in *Oracle Enterprise Manager Cloud Control Installation and Basic Configuration*. and review the section "Management Repository Deployment Guidelines" on page 26-1.
2. Locate the `RepManager` script in the following directory of the Oracle Management Service home directory:

   ```
   ORACLE_HOME/sysman/admin/emdrep/bin
   ```

3. At the command prompt, enter the following command:

   ```
   $PROMPT> ./RepManager repository_host repository_port repository_SID
   -sys_password password_for_sys_account -action create
   ```

In this syntax example:

- *`repository_host`* is the machine name where the Management Repository database is located
- *`repository_port`* is the Management Repository database listener port address, usually 1521 or 1526

- *repository_SID* is the Management Repository database system identifier
- *password_for_sys_account* is the password of the SYS user for the database. For example, change_on_install

Enterprise Manager creates the Management Repository in the database you specified in the command line.

26.5.2.2 Using a Connect Descriptor to Identify the Management Repository Database

Alternatively, you can use a connect descriptor to identify the database on the RepManager command line. The connect descriptor identifies the host, port, and name of the database using a standard Oracle database syntax.

For example, you can use the connect descriptor as follows to create the Management Repository:

```
$PROMPT> ./RepManager -connect "(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)
(HOST=host1)(PORT=1521)) (CONNECT_DATA=(SERVICE_NAME=servicename)))"
-sys_password efkl34lmn -action create
```

> **See Also:** "Establishing a Connection and Testing the Network" in the *Oracle Database Net Services Administrator's Guide* for more information about connecting to a database using a connect descriptor

The ability to use a connect string allows you to provide an address list as part of the connection string. The following example shows how you can provide an address list consisting of two listeners as part of the RepManager command line. If a listener on one host becomes unavailable, the second listener can still accept incoming requests:

```
$PROMPT> ./RepManager -connect "(DESCRIPTION=
(ADDRESS_LIST=
(ADDRESS=(PROTOCOL=TCP)(HOST=host1)(PORT=1521)
(ADDRESS=(PROTOCOL=TCP)(HOST=host2)(PORT=1521)
(CONNECT_DATE=(SERVICE_NAME=servicename)))"
-sys_password efkl34lmn -action create
```

> **See Also:** *Oracle Database High Availability Architecture and Best Practices*
>
> *Oracle Enterprise Manager Cloud Control Installation and Basic Configuration*

26.6 Troubleshooting Management Repository Creation Errors

Oracle Universal Installer creates the Management Repository using a configuration step at the end of the installation process. If the repository configuration tool fails, note the exact error messages displayed in the configuration tools window, wait until the other configuration tools have finished, exit from Universal Installer, and then use the following sections to troubleshoot the problem.

26.6.1 Package Body Does Not Exist Error While Creating the Management Repository

If the creation of your Management Repository is interrupted, you may receive the following when you attempt to create or drop the Management Repository at a later time:

```
SQL> ERROR:
```

```
ORA-00604: error occurred at recursive SQL level 1
ORA-04068: existing state of packages has been discarded
ORA-04067: not executed, package body "SYSMAN.MGMT_USER" does not exist
ORA-06508: PL/SQL: could not find program unit being called
ORA-06512: at "SYSMAN.SETEMUSERCONTEXT", line 5
ORA-06512: at "SYSMAN.CLEAR_EMCONTEXT_ON_LOGOFF", line 4
ORA-06512: at line 4
```

To fix this problem, see "General Troubleshooting Techniques for Creating the Management Repository" on page 26-14.

26.6.2 Server Connection Hung Error While Creating the Management Repository

If you receive an error such as the following when you try to connect to the Management Repository database, you are likely using an unsupported version of the Oracle Database:

```
Server Connection Hung
```

To remedy the problem, upgrade your database to the supported version as described in *Oracle Enterprise Manager Cloud Control Installation and Basic Configuration*.

26.6.3 General Troubleshooting Techniques for Creating the Management Repository

If you encounter an error while creating the Management Repository, drop the repository by running the `-drop` argument to the `RepManager` script.

See Also: "Dropping the Management Repository" on page 26-11

If the `RepManager` script drops the repository successfully, try creating the Management Repository again.

If you encounter errors while dropping the Management Repository, do the following:

1. Connect to the database as SYSDBA using SQL*Plus.
2. Check to see if the SYSMAN database user exists in the Management Repository database.

   For example, use the following command to see if the SYSMAN user exists:

   ```
   prompt> SELECT username FROM DBA_USERS WHERE username='SYSMAN';
   ```

3. If the SYSMAN user exists, drop the user by entering the following SQL*Plus command:

   ```
   prompt> DROP USER SYSMAN CASCADE;
   ```

4. Check to see if the following triggers exist:

   ```
   SYSMAN.EMD_USER_LOGOFF
   SYSMAN.EMD_USER_LOGON
   ```

   For example, use the following command to see if the EMD_USER_LOGOFF trigger exists in the database:

   ```
   prompt> SELECT trigger_name FROM ALL_TRIGGERS
           WHERE trigger_name='EMD_USER_LOGOFF';
   ```

5. If the triggers exist, drop them from the database using the following commands:

   ```
   prompt> DROP TRIGGER SYSMAN.EMD_USER_LOGOFF;
   ```

```
prompt> DROP TRIGGER SYSMAN.EMD_USER_LOGON;
```

26.7 Cross Platform Enterprise Manager Repository Migration

There are user requirements for migrating an Enterprise Manager repository across servers - same and cross platforms.

The Enterprise Manager repository migration process is not exactly the same as database migration. In case of Enterprise Manager Repository migration you must take care of Enterprise Manager specific data, options, and pre-requisites for the repository move. You should make sure data integrity is maintained from both the Enterprise Manager and Oracle database perspective.

This brings up need for defining the process that can be followed by end users for successful and reliable migration of repository in minimum time and with maximum efficiency.

The overall strategy for migration depends on:

- The source and target database version
- The amount of data/size of repository
- Actual data to migrate [selective/full migration]

Cross platform transportable tablespace along with data pump (for metadata) is the fastest and best approach for moving large Enterprise Manager Cloud Control repository from one platform to another. Other option that can be considered for migration is to use Data Pump for both data and metadata moves but this would require more time than the cross platform transportable tablespace approach for the same amount of data. The advantage to using the data pump approach is that it provides granular control over options and the overall process, as in the case of selective data being migrated and not the whole of source data. If the source and target is not on version 11g then export/import is the only way to get the data migrated cross platform.

More details on cross platform transportable tablespace, data pump, and export/import options can be found at the *Oracle Technology Network* (OTN) or in the *Oracle Database Administrator's Guide*.

26.7.1 Common Prerequisites

The following lists the common prerequisites for a repository migration:

- Source and target database must use the same character set and should be at same version.
- Source and target database should meet all the pre-requisites mentioned for Enterprise Manager Repository software requirements mentioned in Enterprise Manager install guide.
- If source and target database are NOT on 11g - only Export/Import can be used for cross platform migration.
- If Source and target database are on 11g - either of three options Cross platform transportable tablespaces migration, Data Pump or Export/Import can be used for cross platform repository migration.
- You cannot transport a tablespace to a target database in which a tablespace with the same name already exists. However, you can rename either the tablespace to be transported or the destination tablespace before the transport operation.

- To plug a transportable tablespace set into an Oracle Database on a different platform, both databases must have compatibility set to at least 10.0.
- Most of the platforms (but not all) are supported for cross-platform tablespace transport. You can query the V$TRANSPORTABLE_PLATFORM view to see the platforms that are supported, and to determine their platform IDs and their endian format (byte ordering).
- Source and Destination host should have Enterprise Manager agent running and configured to the instance which is to be migrated.
- If target database has Enterprise Manager repository installed, it should be first dropped using RepManager before target database related steps are carried out.

26.7.2 Methodologies

The following sections discuss the methodologies of a repository migration.

26.7.2.1 Cross Platform Transportable Tablespaces

Oracle's transportable tablespace feature allows users to quickly move a user tablespace across Oracle databases. It is the most efficient way to move bulk data between databases. Prior to Oracle Database Release 11g, if you want to transport a tablespace, both source and target databases need to be on the same platform. Oracle Database Release 11g adds the cross platform support for transportable tablespaces. With the cross platform transportable tablespace, you can transport tablespaces across platforms.

Cross platform transportable tablespaces allows a database to be migrated from one platform to another (use with Data Pump or Import/Export).

26.7.2.1.1 Preparation for Transportable Tablespaces

Use these steps to prepare for transportable tablespaces:

1. Prepare set of user tablespaces and Check for containment violation.

   ```
   execute DBMS_TTS.TRANSPORT_SET_CHECK('MGMT_TABLESPACE,MGMT_
   ECM_DEPOT_TS', TRUE);

   select * FROM transport_set_violations;
   ```

2. Shutdown OMS instances and prepare for migration.

   Shutdown OMS, set job queue_processes to 0 and run:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   remove_dbms_jobs.sql
   ```

3. Make the tablespaces to be transported read only.

   ```
   alter tablespace MGMT_TABLESPACE read only;

   alter tablespace MGMT_ECM_DEPOT_TS read only;
   ```

26.7.2.1.2 Extract metadata

Extract Metadata for transportable tablespaces using Data Pump Utility:

1. Create the data pump directory.

   ```
   create directory data_pump_dir as
   '/scratch/gachawla/EM102/ttsdata';
   ```

2. Extract the metadata using data pump (or export).

```
expdp DUMPFILE=ttsem102.dmp TRANSPORT_TABLESPACES=MGMT_
TABLESPACE,MGMT_ECM_DEPOT_TS TRANSPORT_FULL_CHECK=Y
```

3. Extract other objects (packages, procedures, functions, temporary tables. and so on -- Not contained in user tablespaces).

```
expdp SCHEMAS=SYSMAN CONTENT=METADATA_ONLY
EXCLUDE=INDEX,CONSTRAINT DUMPFILE=data_pump_dir:postexp.dmp
LOGFILE=data_pump_dir:postexp.log JOB_NAME=expmet
```

26.7.2.1.3 Endian check and conversion

Run Endian check and convert the datafiles if endian is different between source and destination:

1. For Endian check, run this on both source and destination database:

```
SELECT endian_format

FROM v$transportable_platform tp, v$database d

WHERE tp.platform_name = d.platform_name;
```

If the source platform and the target platform are of different endianness, then an additional step must be done on either the source or target platform to convert the tablespace being transported to the target format. If they are of the same endianness, then no conversion is necessary and tablespaces can be transported as if they were on the same platform.

Example:

```
Source Endian
Linux IA (32-bit) - Little

Destination Endian
Solaris[tm] OE (32-bit) - Big
```

2. Ship datafiles, metadata dump to target and Convert datafiles using RMAN:

Ship the datafiles and the metadata dump to target and On target convert all datafiles to destination endian:

```
CONVERT DATAFILE
'/d14/em10g/oradata/em102/mgmt.dbf',
'/d14/em10g/oradata/em102/mgmt_ecm_depot1.dbf'
FROM PLATFORM 'Linux IA (32-bit)';
```

Conversion via RMAN can be done either on source or target (For more details refer RMAN doc). Parallelism can be used to speed up the process if the user tablespaces contains multiple datafiles.

26.7.2.1.4 Import metadata and plugin tablespaces

Use the following steps to import metadata and plugin tablespaces:

1. Run RepManager to drop the target repository (if the target database has the Enterprise Manager repository installed):

```
RepManager repository_host repository_port repository_SID
-sys_password password_for_sys_account -action drop
```

2. Run the pre-import steps to create the sysman user and grant privileges on the target database:

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
create_repos_user.sql

@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
pre_import.sql
```

3. Invoke the Data Pump utility to plug the set of tablespaces into the target database:

```
impdp DUMPFILE=ttsem102.dmp DIRECTORY=data_pump_dir

TRANSPORT_
DATAFILES=/d14/em10g/oradata/em102/mgmt.dbf,/d14/em10g/oradat
a/em102/mgmt_ecm_depot1.dbf
```

4. Import other objects (packages, procedures, functions, and so on):

```
impdp CONTENT=METADATA_ONLY EXCLUDE=INDEX,CONSTRAINT
DUMPFILE=data_pump_dir:postexp.dmp LOGFILE=data_pump_
dir:postexp.log
```

26.7.2.1.5 Post Plug In Steps

Follow these post plug-in steps:

1. Run post plug-in steps to recompile any invalids, create public synonyms, create other users, enable VPD policy, repin packages:

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
create_synonyms.sql

@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
post_import.sql
```

   Check for invalid objects -- compare source and destination schemas for any discrepancy in the counts and invalids.

2. Bring user tablespaces back to read write mode:

```
alter tablespace MGMT_TABLESPACE read write;

alter tablespace MGMT_ECM_DEPOT_TS read write;
```

3. Submit Enterprise Manager dbms jobs.

   Reset back job_queue_processes to original value and run:

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
submit_dbms_jobs.sql
```

4. Update OMS properties and startup the OMS.

   Update emoms.properties to reflect the migrated repository. Update host name - *oracle.sysman.eml.mntr.emdRepServer* and port with the correct value and start the OMS.

5. Relocate Management Services and Repository target.

   If the Management Services and repository target need to be migrated to the destination host, run *em_assoc. handle_relocated_target* to relocate the target or recreate the target on the target host.

6. Discover/relocate Database and database Listener targets.

   Discover the target database and listener in Enterprise Manager or relocate the targets from source agent to destination agent.

26.7.2.2 Data Pump

Oracle Data Pump technology enables high-speed, parallel movement of bulk data and metadata from one database to another. Data Pump uses APIs to load and unload data instead of usual SQL commands. Data pump operations can be run via Enterprise Manager interface and is very useful for cross platform database migration.

The migration of the database using the Data Pump export and Data Pump import tools comprises these steps: export the data into a dump file on the source server with the *expdp* command; copy or move the dump file to the target server; and import the dump file into Oracle on the target server by using the *impdp* command; and run post import Enterprise Manager specific steps.

Tuning parameters that were used in original Export and Import, such as BUFFER and RECORDLENGTH, are neither required nor supported by Data Pump Export and Import

26.7.2.2.1 Prepare for Data Pump

Use the following steps to prepare for data pump:

1. Pre-requisite for using Data pump for Enterprise Manager repository:

   Impdp fails for Enterprise Manager repository because of data pump bug - Bug 4386766 - IMPDP WITH COMPRESSED INDEXES FAILS WITH ORA-14071 AND ORA-39083. This bug is fixed in 10.2. Backport is available for 10.1.0.4. This RDBMS patch has to be applied to use expdp/impdp for the Enterprise Manager repository migration or workaround is to use exp/imp for extract and import.

2. Create the data pump directory:

   ```
   create directory data_pump_dir as
   '/scratch/gachawla/EM102/ttsdata';
   ```

3. Shutdown OMS instances and prepare for migration.

   Shutdown the OMS, set job *queue_processes* to 0 and run `@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_remove_dbms_jobs.sql`

   To improve throughput of a job, the PARALLEL parameter should be used to set a degree of parallelism that takes maximum advantage of the current conditions. In general, the degree of parallelism should be set to more than twice the number of CPUs on an instance.

   All data pump actions are performed by multiple jobs (server processes not DBMS_JOB jobs). These jobs are controlled by a master control process which uses Advanced Queuing. At runtime an advanced queue table, named after the job name, is created and used by the master control process. The table is dropped on completion of the data pump job. The job and the advanced queue can be named using the JOB_NAME parameter.

   DBMS_DATAPUMP APIs can also be used to do data pump export/import. Please refer to Data pump section in 10g administration manual for all the options.

26.7.2.2.2 Data Pump Export

Use these steps to run data pump export:

1. Run data pump export:

   ```
   expdp FULL=y DUMPFILE=data_pump_dir:dpfull1%U.dmp, data_pump_dir:dpfull2%U.dmp
   PARALLEL=4 LOGFILE=data_pump_dir:dpexpfull.log JOB_NAME=dpexpfull
   Verify the logs for any errors during export
   ```

Data pump direct path export sometimes fails for mgmt_metrics_raw and raises ORA 600. This is due to Bug 4221775 (4233303). This bug is fixed in release 10.2. The workaround: if using expdp data pump for mgmt_metrics_raw, run expdp with ACCESS_METHOD+EXTERNAL_TABLE parameter.

```
expdp directory=db_export dumpfile=exp_st2.dmp logfile=exp_
st2.log tables=sysman.mgmt_metrics_raw access_
method=external_table
```

26.7.2.2.3 Data Pump Import

Use these steps to run data pump import:

1. Run RepManager to drop target repository (if target database has Enterprise Manager repository installed):

   ```
   RepManager repository_host repository_port repository_SID
   -sys_password password_for_sys_account -action drop
   ```

2. Prepare the target database:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   create_tablespaces.sql
   ```

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   create_repos_user.sql
   ```

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   pre_import.sql
   ```

3. Run data pump import:

   ```
   Impdp FULL=y DUMPFILE=data_pump_dir:dpfull1%U.dmp, data_pump_
   dir:dpfull2%U.dmp PARALLEL=4 LOGFILE=data_pump_
   dir:dpimpfull.log JOB_NAME=dpimpfull
   ```

   Verify the logs for any issues with the import.

26.7.2.2.4 Post Import Enterprise Manager Steps

Use the following steps for post import Enterprise Manager steps:

1. Run post plugin steps to recompile any invalids, create public synonyms, create other users, enable VPD policy, repin packages:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   create_synonyms.sql
   ```

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   post_import.sql
   ```

   Check for invalid objects - compare source and destination schemas for any discrepancy in counts and invalids.

2. Submit Enterprise Manager dbms jobs.

   Reset back *job_queue_processes* to original value and run:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   submit_dbms_jobs.sql
   ```

3. Update OMS properties and startup the OMS.

Update emoms.properties to reflect the migrated repository. Update host name - *oracle.sysman.eml.mntr.emdRepServer* and port with the correct value and start the OMS.

4. Relocate Management Services and Repository target.

   If Management Services and the repository target needs to be migrated to the destination host, run *em_assoc. handle_relocated_target* to relocate the target or recreate the target on the target host.

5. Discover/relocate Database and database Listener targets.

   Discover the target database and listener in Enterprise Manager or relocate the targets from source agent to destination agent.

26.7.2.3 Export/Import

If the source and destination database is non-10g, then export/import is the only option for cross platform database migration.

For performance improvement of export/import, set higher value for BUFFER and RECORDLENGTH. Do not export to NFS as it will slow down the process considerably. Direct path can be used to increase performance. Note - As Enterprise Manager uses VPD, conventional mode will only be used by Oracle on tables where policy is defined.

Also User running export should have EXEMPT ACCESS POLICY privilege to export all rows as that user is then exempt from VPD policy enforcement. SYS is always exempted from VPD or Oracle Label Security policy enforcement, regardless of the export mode, application, or utility that is used to extract data from the database.

26.7.2.3.1 Prepare for Export/Import

Use the following steps to prepare for Export/Import:

1. Mgmt_metrics_raw partitions check:

   ```
   select table_name,partitioning_type type,
   partition_count count, subpartitioning_type subtype from
   dba_part_tables where table_name = 'MGMT_METRICS_RAW'
   ```

   If MGMT_METRICS_RAW has more than 3276 partitions please see Bug 4376351 - This is fixed in release 10.2. The work around is to export mgmt_metrics_raw in conventional mode.

2. Shutdown OMS instances and prepare for migration

   Shutdown OMS, set job *queue_processes* to 0 and run `@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_remove_dbms_jobs.sql`

26.7.2.3.2 Export

Follow these steps for export:

1. Export data:

   ```
   exp full=y constraints=n indexes=n compress=y file=fullem102_
   1.dmp log=fullem102exp_1.log
   ```

2. Export without data and with constraints:

   ```
   exp full=y constraints=y indexes=y rows=n ignore=y
   file=fullem102_2.dmp log=fullem102exp_2.log
   ```

26.7.2.3.3 Import

Follow these steps to import:

1. Run RepManager to drop the target repository (if the target database has the Enterprise Manager repository installed):

   ```
   RepManager repository_host repository_port repository_SID
   -sys_password password_for_sys_account -action drop
   ```

2. Pre-create the tablespaces and the users in target database:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   create_tablespaces.sql
   ```

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   create_repos_user.sql
   ```

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   pre_import.sql
   ```

3. Import data:

   ```
   imp full=y constraints=n indexes=n file=fullem102_1.dmp
   log=fullem102imp_1.log
   ```

4. Import without data and with constraints:

   ```
   imp full=y constraints=y indexes=y rows=n ignore=y
   file=fullem102_2.dmp log=fullem102imp_2.log
   ```

26.7.2.3.4 Post-Import Enterprise Manager Steps

Follow these steps for post-import Enterprise Manager steps:

1. Run post plug-in steps to recompile any invalids, create public synonyms, create other users, enable VPD policy, repin packages:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   create_synonyms.sql
   ```

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   post_import.sql
   ```

   Check for invalid objects -- compare source and destination schemas for any discrepancy in counts and invalids.

2. Submit the Enterprise Manager dbms jobs.

   Reset back *job_queue_processes* to its original value and run:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   submit_dbms_jobs.sql
   ```

3. Update the OMS properties and startup the OMS:

   Update emoms.properties to reflect the migrated repository. Update host name *oracle.sysman.eml.mntr.emdRepServer* and port with the correct value and start the OMS.

4. Relocate Management Services and the Repository target.

   If Management Services and the repository target need to be migrated to the destination host, run *em_assoc. handle_relocated_target* to relocate the target or recreate the target on the target host.

5. Discover/relocate Database and database Listener targets.

Discover the target database and listener in Enterprise Manager or relocate the targets from the source agent to the destination agent.

26.7.3 Post Migration Verification

These verification steps should be carried out post migration to ensure that the migration was completely successful:

- Verify any discrepancy in objects by comparing source and target databases through Enterprise Manager.
- Verify the migrated database through Enterprise Manager to determine whether the database is running without any issues.
- Verify the repository operations, dbms jobs and whether any management system errors are reported.
- Verify that all Enterprise Manager functionalities are working correctly after the migration.
- Make sure Management Services and the Repository target is properly relocated by verifying it through Enterprise Manager.

27

Maintaining Enterprise Manager

Enterprise Manager provides extensive monitoring and management capabilities for various Oracle and non-Oracle products. Used to manage your heterogeneous IT infrastructure, Enterprise Manager plays an integral role in monitoring and maintaining the health of your IT resources. It is therefore essential to make sure Enterprise Manager itself is operating at peak efficiency.

To help you maintain your Enterprise Manager installation, a variety of enhanced self-monitoring and diagnostic functionality is available from the Enterprise Manager console. These functions are designed to help you understand and monitor various components of Enterprise Manger, monitor/measure the quality of services Enterprise Manager provides, diagnose failures quickly, and manage Agents more easily.

This chapter covers the following topics:

27.1 Overview: Managing the Manager

Although Enterprise Manager functions as a single entity to manage your IT infrastructure, in reality it is composed of multiple components working in concert to provide a complete management framework from a functional standpoint. Beginning with Enterprise Manager 12c, the functions used to ensure reliability and performance for your monitored targets can be used to maintain Enterprise Manager itself. All major components of Enterprise Manager have been grouped into a single system. A special set of services has been created (based on the system) to model Enterprise Manager functions.

Management Features

- Topology view that allows you to see all major components of Enterprise Manager and their current status.
- Enterprise Manager dashboard displaying the overall health of Enterprise Manager.
- Full control of the Agent directly from the Enterprise Manager console. Functions include:
  - View/edit Agent configuration properties.

- View Agent(s) configuration history and compare the results against other Agents.
- Perform Agent control operations (start/stop/secure).

27.2 Management Services and Repository

The Management Services and Repository home page provides a detailed overview of the OMS and Repository.

Accessing the Management Services and Repository Home Page

From the **Setup** menu, click **Management Services and Repository**.

Figure 27–1 Management Services and Repository Home Page

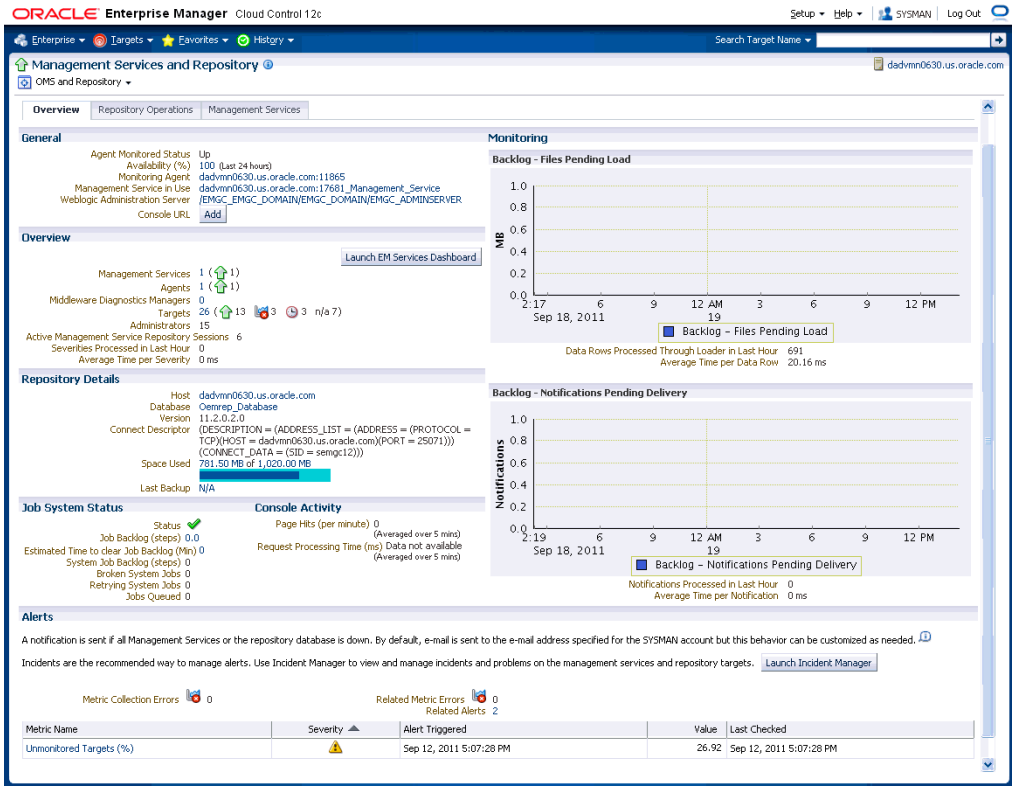

Overview Page

The Overview page displays a summary of the current status of the OMS and Repository. The Overview page also provides details on the status of the Job system and the console activity and two monitoring charts that indicate the backlog in terms of the files pending and the backlog on notifications pending delivery.

Each region provides specific information on the various operational areas of the OMS and Repository.

- **General**: provides the Agent monitored status, availability and the host details of the Management Service in use, the monitoring Agent and the WebLogic administration server.

- **Overview**: Displays the number and breakdown in term of status for the OMS, Agents, targets, active OMS repository sessions and severities detected within the past hour.
- **Repository Details**: Provides physical information about the Management Repository and the host on which the database is located.
- **Alerts**: Provides details on the metric errors recorded and when an alert was triggered. In-context links to Incident Manager are also provided.

Repository Operations Page

The Repository Operations page provides you with an overview of the status and performance of the Repository DBMS Jobs that handle part of Enterprise Manager's maintenance and monitoring functionality. These DBMS jobs run within the Management Repository and require no user input. Charts showing the **Data Upload** to **Repository** and the **Backlog** in Repository collection are provided The **Scheduler Status** region provides the status of the scheduler and the number of Job Queue Processes.

Figure 27–2 Repository Operations Page

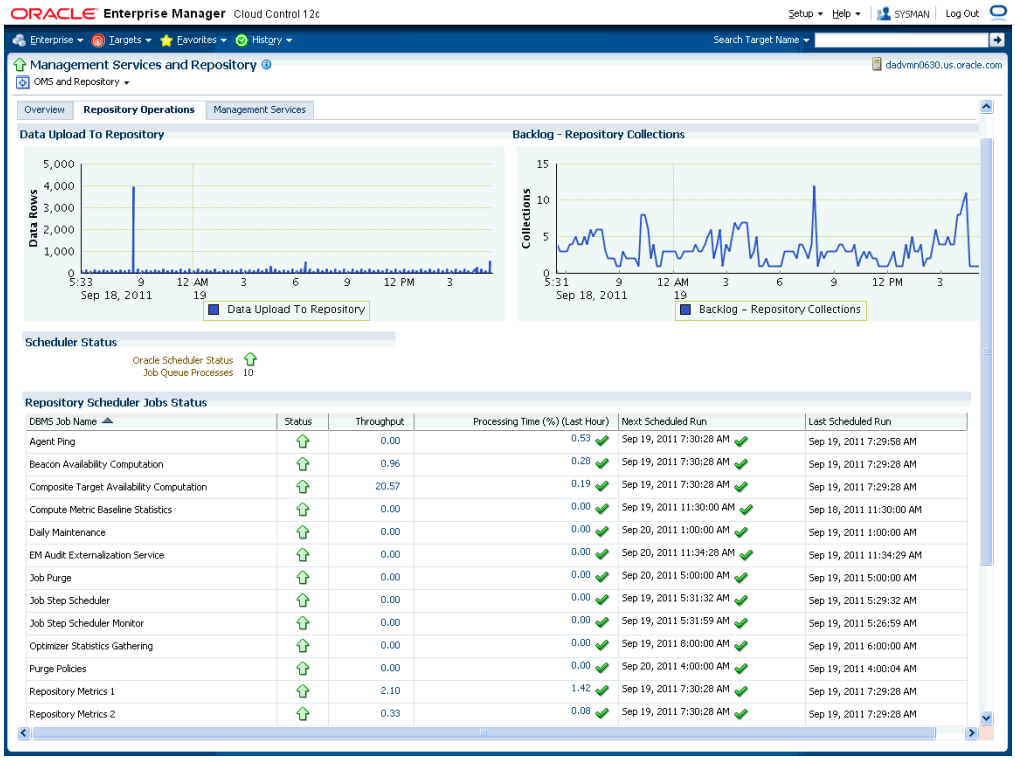

The **Repository Scheduler Jobs Status** region provides details of the DBMS Jobs regarding their status, throughput, processing time, the next scheduled run and the last scheduled run.

Use the Repository Operations page to view the performance of the Repository DBMS jobs. If you want more information, clicking a link brings you to a metrics detail page.

To determine how well the Management Repository is handling its share of the Enterprise Manager functionality, view the Throughput per second and Processing

Time Percent (Last Hour) columns. If the Processing Time Percent (Last Hour) is large and the Throughput is low, there may be problems in that area of management.

Note: Processing Time (%) Last Hour may exceed 100% if a job runs continuously for more than an hour. For example, if you see 125.00 in this column, it means that the job ran for 75 minutes (125 % of one hour).

Management Services Page

The **Management Services** page lists the names of the Management Services and their status. When an OMS is decommissioned, status will be shown as down and delete button will be displayed, allowing you to remove the decommissioned OMS from EM.

Figure 27–3 Management Services Page

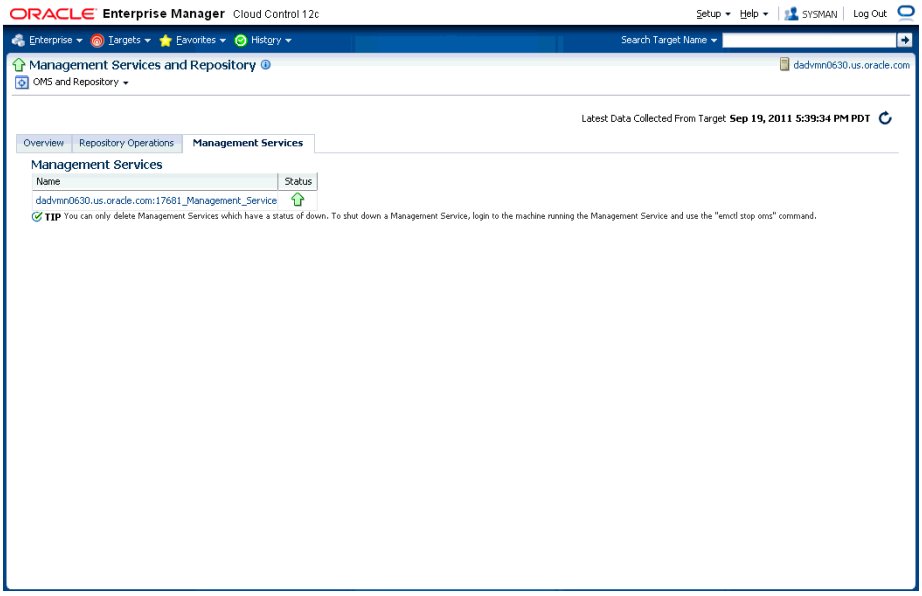

By clicking on the individual Management Service, you can drill down to that OMS' target home page for explicit information and status.

27.3 Viewing Enterprise Manager Topology/Charts

The Enterprise Manager Topology page provides a graphical representation of the Enterprise Manager infrastructure components and their association. Each node in the hierarchy displays key information about the member type, the host on which it resides, and the number of incidents, if any. The incident icons on each of the nodes expand to display a global view of current status for each node in the hierarchy.

Accessing the Enterprise Manager Topology

1. From the Setup menu, choose Management Services and Repository.
2. On the Management Services and Repository page, click on the **OMS and Repository** drop-down menu.
3. Choose **Members** and then **Topology**.

Figure 27–4 Enterprise Manager Topology

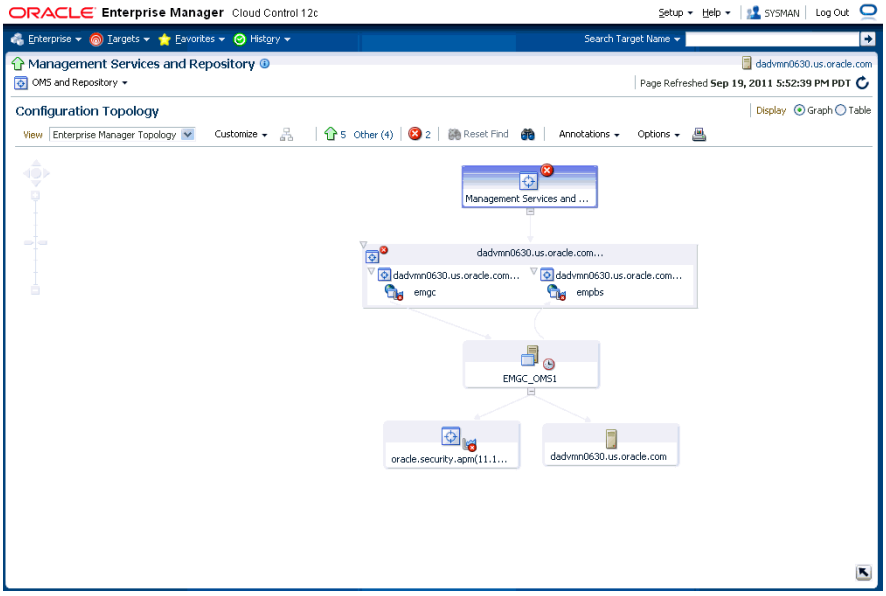

Enterprise Manager Charts

The Enterprise Manager Charts page displays eight charts representing key areas that together indicate the overall health of Enterprise Manager. These are Overall Files Pending Load -Agent, Job Step Backlog, Job Step Throughput (per second), Request Processing Time (ms), Database Time (centiseconds per second), CPU Utilization (%), Pages Paged-in (per second), Pages Paged-out (per second). Data can be viewed for the Last 24 hours, last 7 days or last 31 days.

Accessing the Enterprise Manager Charts

1. From the Setup menu, choose Management Services and Repository.
2. On the Management Services and Repository page, click on the **OMS and Repository** drop-down menu.
3. Choose **Monitoring** and then **Charts**.

Figure 27–5 Enterprise Manager Charts

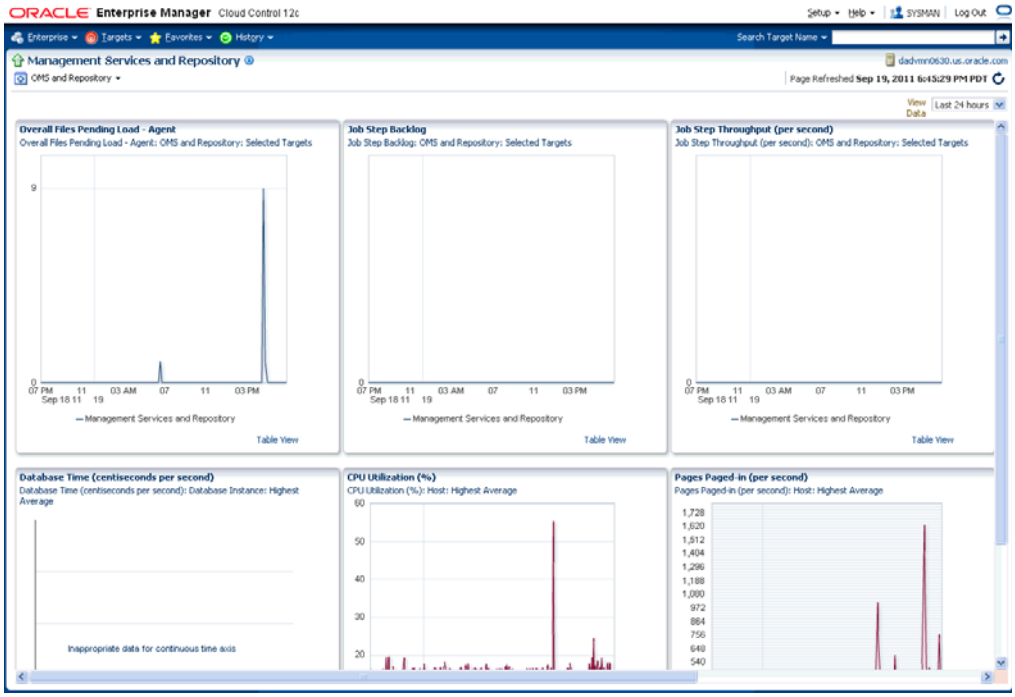

27.4 Viewing Enterprise Manager Services

The Enterprise Manager Services dashboard provides details about the two primary services in-context: Enterprise Manager Jobs Service and the Enterprise Manager Console Service. The **Status**, **Performance** (gauged by means of specific metrics), **Usage and Business Indicators** (again through specific metrics), **Component Status** and the **Service Level** (%) for the two primary services for the Last 24 hours, last 7 days and last 31 days are detailed.

Accessing the Enterprise Manager Topology

1. From the Setup menu, choose Management Services and Repository.
2. On the Management Services and Repository page, click **Launch EM Services Dashboard** in the Overview region.

Figure 27–6 Enterprise Manager Services Dashboard

27.5 Controlling and Configuring Management Agents

Beginning with Enterprise Manager Cloud Control 12c, Agents management can be performed directly from the Enterprise Manager console. This provides a central point where all Agents for your managed targets can be compared, configured and controlled.

27.5.1 Agent Home Page

The Agent home page provides details for a single Agent. This page also lets you drill down for more detailed information. You can access an Agent home page by selecting it from the All Targets page.

Figure 27–7 Agent Home Page

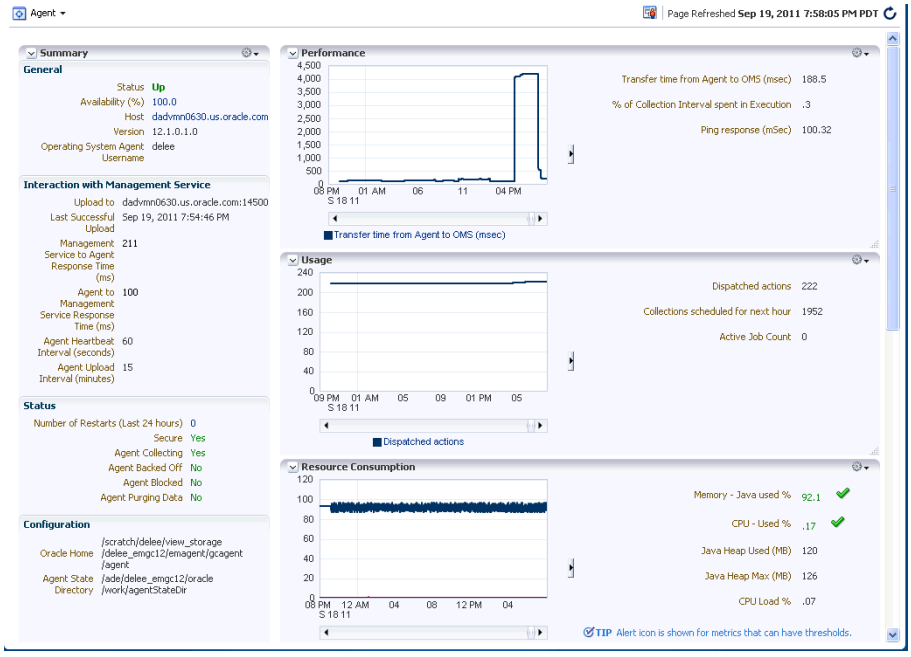

The **Summary** region provides primary details of the Agent such as its status and availability. The **Interaction with Management Service** region provides details on the

communication between the OMS and the Agent. The **Status** region provides further details on the Agent status such as the number of restarts, the action that the Agent is performing currently. The **Performance, Usage and Resource Consumption** charts provide further details on the Agent in graphical format. The **Incidents** region lists the incidents recorded for the Agent. The **Monitoring** section provides details on the targets that are being monitored by the Agent, metric extensions and management plug-ins deployed in the Agent.

27.5.2 Controlling a Single Agent

Control operations for a single Agent can be performed on the Agent home page for that Agent.

1. Navigate to the desired Agent home page.
2. From the **Agent** drop-down menu, choose **Control** and then one of the control operations (Start Up/Shut Down, or Restart)

> **Note:** You must have at least operator privileges in order to perform Agent control operations.

Figure 27–8 Control Operations from the Agent Home Page

Upon choosing any of the above control menu options, a pop-up dialog requesting the credentials of the user displays. These operations require the credentials of the OS user who owns the Agent. At this point, you can either choose from a previously stored username/password, preferred or named credential. You also have the option of choosing a new set of credentials which could also be saved as the preferred credential or as a named credential for future use. When storing a named credential or creating a new set of credentials, a **Test** button is shown, thus allowing you to verify if the credential information you just entered is valid. Once you are authenticated, the chosen control operation begins and continues even if the pop-up dialog is closed. Any message of failure/success of the task is displayed in the pop-up dialog.

When choosing the Secure/Resecure/Unsecure options, you must provide the requisite Registration Password.

27.5.3 Configuring Single Agents

Configuration operations for a single Agent can be performed from the Agent home page. To access the Agent properties page:

1. Navigate to the desired Agent home page.
2. From the **Agent** drop-down menu, choose **Properties**.

> **Note:** You must have at least Configure privileges in order to perform Agent configuration operations.

Figure 27–9 Agent Properties Page

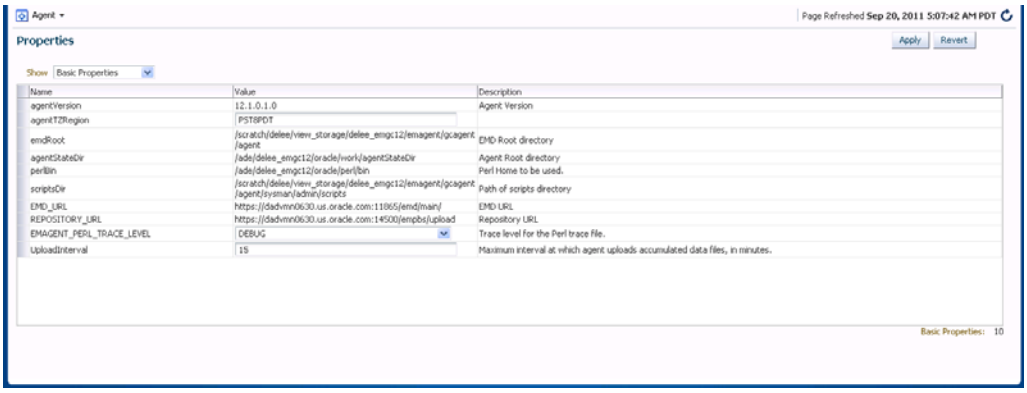

The properties on this page can be filtered to show **All Properties**, **Basic Properties**, or **Advanced Properties**. The **Basic Properties** are a simple name, value combination of a property and its value. **Advanced Properties** are also a combination of name and value but can also be grouped into categories. You must have at least *configure* privileges in order to modify the existing properties and set custom properties.

27.5.4 Controlling Multiple Agents

In order to perform control operations on multiple Agents, Enterprise Manager makes use of the Job system to automate repetitive tasks. Therefore, you must have Job privileges for controlling multiple Agents through a single action. To access

1. From the **Setup** menu, choose **Agents**. The Management Agent setup page displays.
2. Select multiple Agents from the list.
3. Click one of the control operation buttons (**Start Up**/**Shut Down**/**Restart**/**Secure**/**Resecure**/**Unsecure**).

When you click on any of the control operations, you are taken to the Job creation wizard where you schedule a new job to perform the action on the selected Agents.

Figure 27–10 Multiple Agent Control Operation: Job Creation

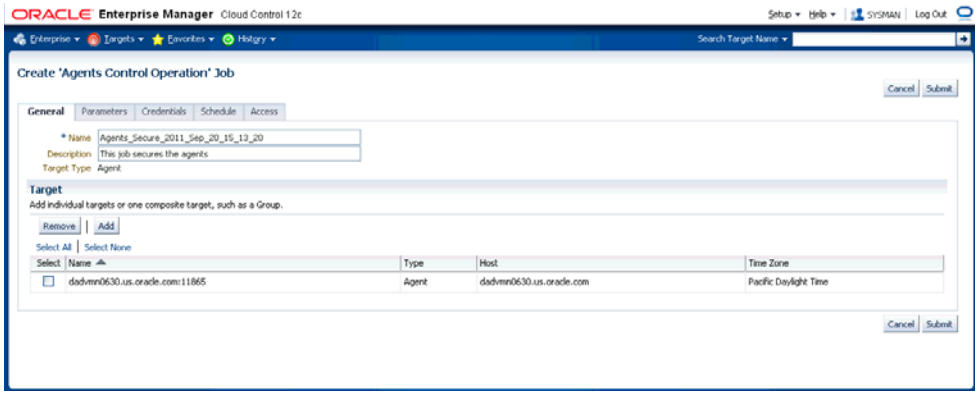

In the Jobs page, you can view the chosen Agents in Target section in the General tab. You can add more Agents by clicking the **Add** button. You then provide the parameters for the operation in the **Parameters** tab, if needed. The credentials must be specified in the **Credentials** tab where you can either choose from a previously stored username/password, preferred, or named credential. You also have the option of choosing a new set of credentials which could also be saved as the preferred credential or as a named credential for future use.

You are given the option to start the job immediately or schedule the job for a later time. At this point, you can also create a repeating job by specifying the job start time, the frequency, and the end time.

The Access tab displays the Administrator details and the access levels they have to the job. You can then add a new administrator or modify the access level to **View** or **Full**, if you have the requisite privileges.

Figure 27–11 Job Creation: Access Tab

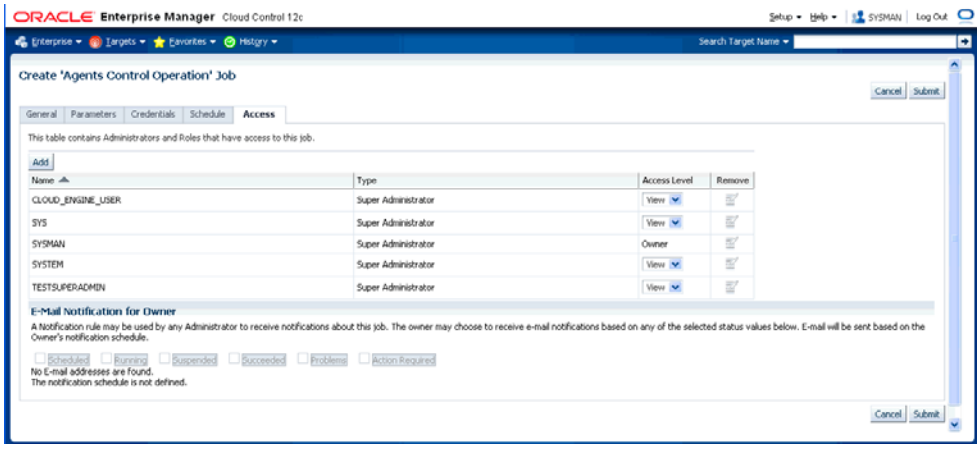

Note: Administrators with insufficient privileges can also schedule jobs for these control operations, but in this situation, the jobs will not complete successfully.

27.5.5 Configuring Multiple Agents

As with multi-Agent control operations, you can also perform Agent configuration on multiple Agents in the same way. This greatly simplifies standardizing Agent configurations across your enterprise. To access Agent properties:

1. From the **Setup** menu, choose **Agents**. The Management Agent setup page displays.
2. Select multiple Agents from the list.
3. Click **Properties**. As with any multi-Agent operation, configuration is implemented using the Job system.

Figure 27–12 Agent

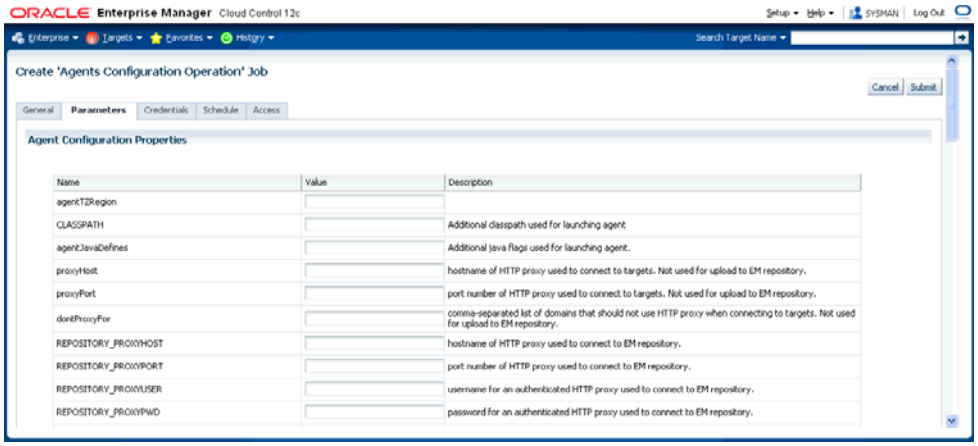

In the Jobs page, you can view the chosen agents in Target section in the General tab. You can add more Agents by clicking the **Add** button if necessary. In the **Parameters** tab, you provide the modified value for a particular set of properties that you want to change. You can also set a custom property for the chosen agents. No credentials are required for modifying Agent properties.

The **Access** tab displays the administrator details and the access levels they have to the job. You can then add a new administrator or modify the access level to View or Full if you have the requisite privileges.

Figure 27–13 Multi-Agent Configuration: Job Access

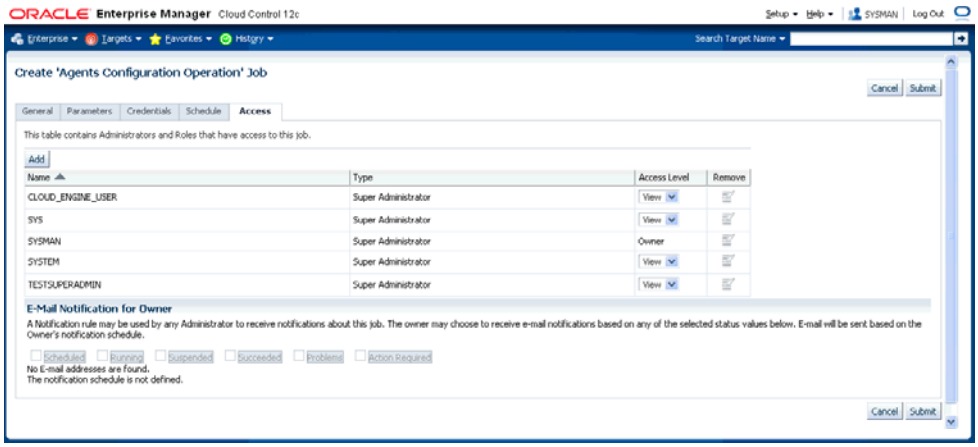

28

Discovering and Managing Exadata Targets and Systems

Oracle Database Machine is an integrated data warehousing solution that eases data warehousing by integrating the whole hardware and software stack into one solution. DB Machine management simplifies monitoring and managing the DB Machine by integrating all hardware and software components into one entity. You do not need to monitor each target individually but instead you can view the whole DB Machine as a single target. You can view all critical issues in the system, monitor performance, and drill down to individual targets from the DB Machine target homepage.

Enterprise Manager automatically or manually discovers the components of the DB Machine and adds these components as managed targets. You can modify the member targets of the DB Machine by adding and removing targets.

The DB Machine homepage includes the following components:

- The hardware schematic allows you to view hardware components of the DB Machine (for example, compute nodes, Exadata cells, and Infiniband switches), monitor critical hardware metrics and view aggregated alerts and faults from all components.
- Alerts sourced from hardware components.
- Easily accessible links and flows to other key feature such as viewing the topology or modifying the schematic

You can also monitor all components of the DB Machine. DB Machine monitors all subcomponent targets, whether hardware or software. This includes the database, ASM, CRS, hosts, Exadata and the Infiniband network. The two targets available to facilitate DB Machine monitoring are:

- Exadata Storage Server
- Infiniband Network Fabric

For more information about discovering and managing Exadata targets, see the following topics:

28.1 Automatically Discovering an Oracle Database Machine

To automatically discover a DB Machine target, follow these steps:

1. From the Enterprise Manager Home page, click **Setup**.
2. Choose **Add Target** and then select **Add Targets Manually**.

   Enterprise Manager displays the Add Targets Manually page displayed in Figure 28–1.

Figure 28–1 Add Targets Manually Page

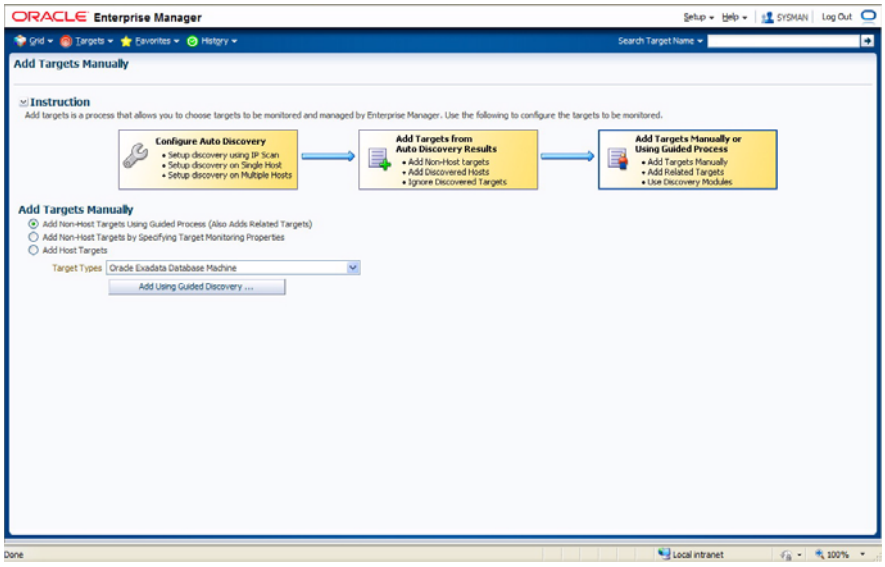

3. Choose **Add Non-Host Targets Using Guided Process (Also Adds Related Targets)**
4. From the Target Types drop-down, choose **Oracle Exadata Database Machine** and click **Add Using Guided Discovery**.

   Enterprise Manager displays the Oracle Database Machine Discovery page. From here you can add the hardware components such as Exadata Storage Servers and Infiniband Switches in the Oracle Database Machine as managed targets. You can choose to discover a new DB Machine and its hardware components as targets or instead discover newly added hardware components in an existing DB Machine as targets.

5. Choose **Discover a new Database Machine and its hardware components as targets** and then click **Discover Targets**.

   Enterprise Manger displays the Database Machine Discovery wizard that steps you through the process of discovering the Database Machine. The first page of the wizard, the Discovery Inputs page, appears.

6. In the Discovery Agent section of the Discovery Inputs page, choose or enter the **Agent URL** that exists on one of the compute nodes to perform the discovery and then enter the **Database Oracle Home** of the database installation (version 11.2 or later) on the agent host. In the Infiniband Discovery section, specify the host name of one of the Infiniband switches in the Database Machine and specify the password of the on the Infiniband switch. Click **Next**.

   The Components page displays as seen in Figure 28–2.

Figure 28–2 Components Page

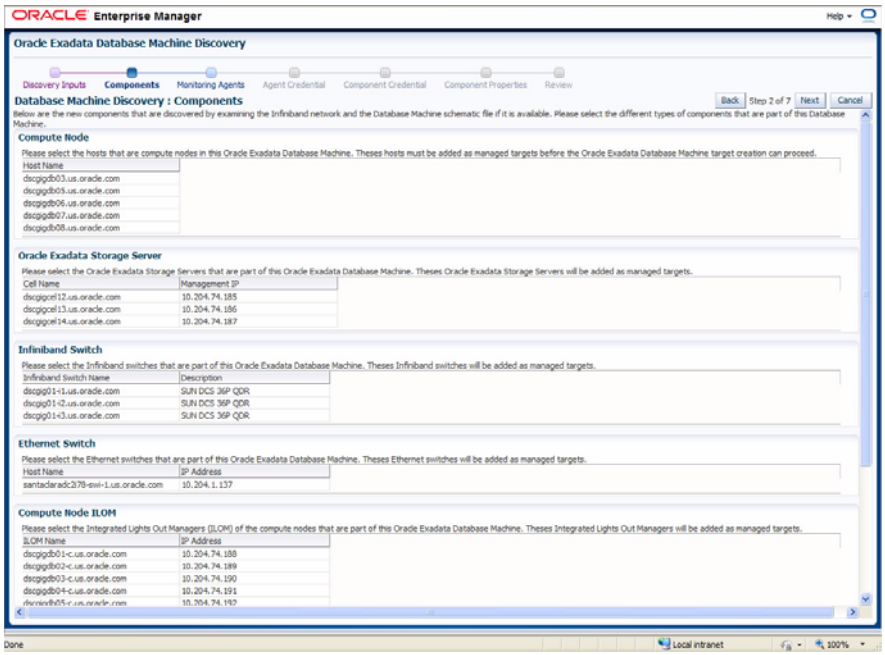

7. On the Components page, the new components discovered in the Oracle Database Machine display.

   In the Compute Node section, select the hosts that are compute nodes in the Oracle Database Machine. The hosts must be added as managed targets before the Oracle Database Machine target promotion can proceed.

   In the Infiniband Switch section, select the Infiniband Switches that are part of the Oracle Database Machine. These also will be added as managed targets.

   In the Ethernet Switch section, select the Ethernet switches that are part of the Oracle Database Machine. The Ethernet switches will be added as managed targets.

   In the Compute Node ILOM section, select the Integrated Lights Out Managers (ILOM) of the compute nodes that are part of this Oracle Database Machine. Theses Integrated Lights Out Managers will be added as managed targets.

   In the KVM section, select the KVM switches that are part of the Oracle Database Machine. The KVM switches will be added as managed targets.

   Click **Next** to continue.

   Enterprise Manager displays the Monitoring Agents page.

8. On the Monitoring Agents page, choose whether you want the monitoring agents to be selected automatically or manually. If you choose **Manually select the agents**, you must add the **Monitoring Agent** and optionally a **Backup Monitoring Agent** for each of the Exadata Cells and Infiniband Switches.

   Click **Next** to continue.

   Enterprise Manager displays the Agent Credential page.

9. On the Agent Credential page, specify whether the agent host users and passwords are the same for all agents. The agent users and passwords are required

to set up password-less SSH between the agents and the cells monitored by the agents.

If the users and passwords are the same, choose **Same for all agents** and enter the user and password combination. If they are not the same for each agent, choose **Different for all agents** and enter each combination for each agent.

Click **Next**.

Enterprise Manager displays the Component Credential page as seen in Figure 1-3.

Figure 28–3 Component Credential Page

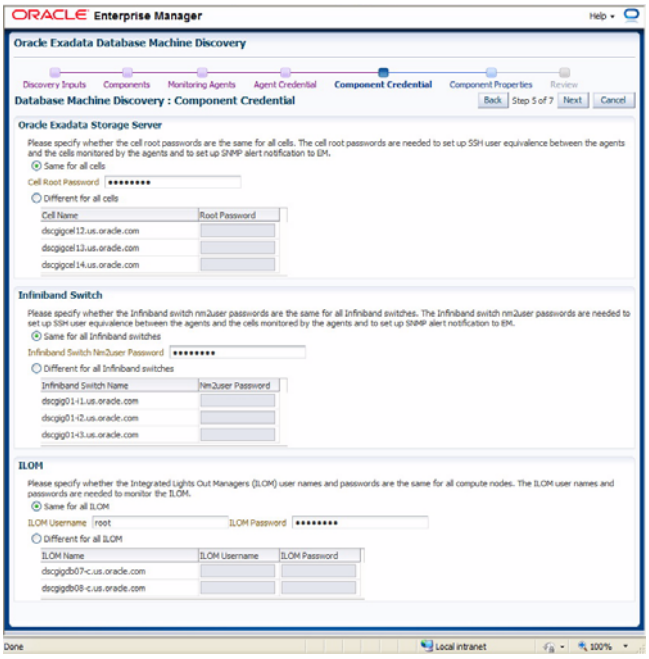

10. On the Component Credential page, enter the credentials for the components. In all cases you can choose to enter the same user/password combinations for all components or you can enter the credentials separately for each occurrence.

    Click **Next** to continue.

    Enterprise Manager displays the Component Properties page.

11. On the Component Properties page, specify the target properties of the different components.

    For each component, enter the required information.

    Click **Next** to continue.

    Enterprise Manager displays the Review page as seen in Figure 28–4.

Figure 28–4 Review Page

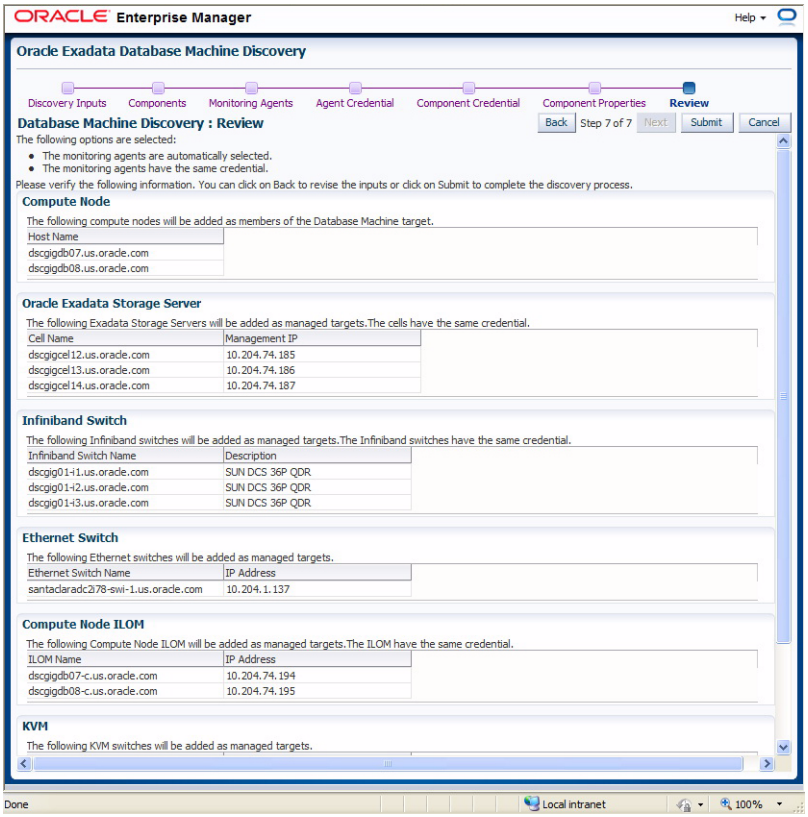

12. Use the Review page to view the selections you have made in the previous pages of the wizard. When you have verified your selections, click **Submit**.

    Enterprise Manager displays the Target Promotion Summary page that displays the targets that are now managed targets.

28.2 Viewing the Topology of an Existing DB Machine Target

You can view the DB Machine topology of the DB Machine target.

DB Machine management simplifies monitoring and managing the DB Machine by integrating all hardware and software components into one entity. You do not need to monitor each target individually but instead you can view the whole DB Machine as a single target. You can view all critical issues in the system, monitor performance, and drill down to individual targets from the DB Machine target homepage.

Use the Topology page of DB Machine to view the topology of the system by Cluster or by Database. Clusters are a complete software system starting with a RAC database, the underlying ASM, and CRS. Clusters define one logical entity that is interconnected. The DB Machine could include several clusters, one cluster, or could just be a number of individual databases. While cabinets define the hardware topology of the DB Machine, clusters define the logical or system topology of the DB Machine.

You can view the Topology by Cluster or Database. Click on an element in the Topology and view alert data associated with the element.

You can monitor all components of the DB Machine. DB Machine monitors all subcomponent targets, whether hardware or software. This includes the database, ASM, CRS, hosts, Exadata and the Infiniband network.

To view the topology of an existing DB Machine target, follow these steps:

1. From the Enterprise Manager Home page, click **Targets**, and then select **Exadata**.

   Enterprise Manager displays the Oracle Exadata Database Machines page showing all the available DB Machine targets. From this page you can add hardware components (such as Oracle Exadata Storage Servers, Infiniband switches, Ethernet Switches, KVM switches, PDU, and compute node ILOM) in the Oracle Database Machine as managed targets.

2. From the Oracle Exadata Database Machines page, select the Oracle Database Machine target whose topology you want to view.

3. From the Oracle Database Machine Home page, click **Target** and then select **Members Topology** from the drop-down menu.

   Cloud Control displays the Configuration Topology page.

28.3 Drilling Down to Individual Targets

You can drill down immediately to a subcomponent target of the DB Machine (such as RAC, a database instance, or an Exadata cell).

To drill down to individual targets, follow these tasks:

1. From the Enterprise Manager Home page, click **Targets**, and then select **Exadata**.

   Enterprise Manager displays the Oracle Exadata Database Machines page showing all the available DB Machine targets. From this page you can add hardware components (such as Oracle Exadata Storage Servers, Infiniband switches, Ethernet Switches, KVM switches, PDU, and compute node ILOM) in the Oracle Database Machine as managed targets.

2. From the Oracle Exadata Database Machines page, select the Oracle Database Machine target whose components you want to view.

   Enterprise Manager displays the Oracle Database Machine Home page showing an Overview, Schematic, and Incident section for the selected DB Machine.

3. From the Oracle Database Machine Home page, use the left navigation panel to expand the list of available targets that comprise the Database Machine.

4. Click on the target to which you want to drill down.

28.4 Viewing Critical Hardware Information for the DB Machine

You can view critical metrics for all the hardware subcomponents of the DB Machine such as DB hosts, Exadata cells, Infiniband switches and so on. These metrics vary for different component targets. For example, databases server nodes and Exadata servers include the CPU, I/O, and Storage metrics.

To view critical hardware-centric information for the entire DB machine, follow these steps:

To drill down to individual targets, follow these tasks:

1. From the Enterprise Manager Home page, click **Targets**, and then select **Exadata**.

   Enterprise Manager displays the Oracle Exadata Database Machines page showing all the available DB Machine targets. From this page you can add hardware components (such as Oracle Exadata Storage Servers, Infiniband switches, Ethernet Switches, KVM switches, PDU, and compute node ILOM) in the Oracle Database Machine as managed targets.

2. From the Oracle Exadata Database Machines page, select the Oracle Database Machine target whose hardware information you want to view.
3. From the Oracle Database Machine Home page, view the hardware schematic of the Database Machine.

28.5 Viewing DB Machine Alerts

You can view alerts on the DB Machine and drill down to details about each alert. These alerts may be performance/configuration metrics or hardware faults.

To view DB Machine alerts, follow these steps:

1. From the Enterprise Manager Home page, click **Targets**, and then select **Exadata**.

   Enterprise Manager displays the Oracle Exadata Database Machines page showing all the available DB Machine targets. From this page you can add hardware components (such as Oracle Exadata Storage Servers, Infiniband switches, Ethernet Switches, KVM switches, PDU, and compute node ILOM) in the Oracle Database Machine as managed targets.
2. From the Oracle Exadata Database Machines page, select the Oracle Database Machine target whose machine configuration information you want to view.

   Enterprise Manager displays the Oracle Database Machine home page on which you can see all alerts associated with the current DB Machine.

28.6 Adding Exadata Components Manually

You can add Exadata components manually using the following steps:

1. Click **Setup** and then choose **Add Target**, then select **Add Targets Manually** from the menu.

   Enterprise Manager displays the Add Targets Manually page where you can choose the type of target you want to add.
2. From the Add Targets Manually section, choose **Add Non-Host Targets by Specifying Target Monitoring Properties**.
3. In the **Target Type** combo box, choose the appropriate target type, for example, KVM for kvm, PDU for pdu, Cisco switch for Cisco, and Oracle ILOM Server for ilom plug-in).
4. Choose the monitoring agent using the search option.
5. Click **Add Manually** and provide the required properties

28.7 About Oracle Exadata Storage Server

An Exadata cell is a network-accessible storage array with Exadata software installed on it. Use the Exadata Home page to manage and monitor the HP Oracle Exadata Storage Server (also known as Exadata cell) by managing the Exadata cells as Enterprise Manager Cloud Control targets. You can discover and consolidate management, monitoring and administration of a single or a group of Oracle Exadata Storage Servers in a datacenter using Enterprise Manager.

Exadata cells can be discovered automatically or manually. Once discovered, you can add the Exadata cells as Enterprise Manager targets. The individual Exadata cell is monitored and managed as an Enterprise Manager target and provides the exception, configuration and performance information.

Grouping of Exadata cells is used for easy management and monitoring of the set of Exadata cells. You can group Exadata cells both manually and automatically. The grouping function provides an aggregation of exceptions, configuration and performance information of the group of cells.

You can view performance analysis by linking Exadata performance both at a cell level and group level to ASM and database performance. You can drill down to Exadata configuration and performance issues from both the database and ASM targets.

Storage Grid (for example, multiple database/ASM instances sharing the same Exadata cell) is supported to the same extent as dedicated storage.

For more information about managing Exadata servers, see the following topics:

- Using Exadata As a Cloud Control Target
- Performing Administration Tasks on Exadata Cells
- Performing Administration Tasks on Infiniband Networks
- Launching the IORM Performance Page
- Viewing an Exadata Cell Configuration
- Managing a Single I/O Resource Management Allocation
- Accessing Oracle Support Workbench for Exadata Cell
- Changing the IORM Mode and Updating the IORM Objective

28.7.1 Using Exadata As a Cloud Control Target

Use Oracle Exadata to manage and monitor the HP Oracle Exadata Storage Server (also known as Exadata cell) by managing the Exadata cells as Enterprise Manager Cloud Control targets. You can discover and consolidate management, monitoring and administration of a single or a group of Oracle Exadata Storage Servers in a datacenter using Enterprise Manager.

Exadata cells can be discovered automatically or manually. Once discovered, you can add the Exadata cells as Enterprise Manager targets.

The individual Exadata cell is monitored and managed as an Enterprise Manager target and provides the exception, configuration and performance information.

Grouping of Exadata cells is used for easy management and monitoring of the set of Exadata cells. You can group Exadata cells both manually and automatically. The grouping function provides an aggregation of exceptions, configuration and performance information of the group of cells.

You can view performance analysis by linking Exadata performance both at a cell level and group level to ASM and database performance. You can drill down to Exadata configuration and performance issues from both the database and ASM targets.

28.7.2 Performing Administration Tasks on Exadata Cells

To perform an administration operation on an Exadata cell, such as executing a cell command, follow these steps:

1. Navigate to the Exadata Cell home page by choosing the Exadata target for which you want to perform an administrative task from the All Targets page of Enterprise Manager.

   Enterprise Manager displays the Exadata Cell Home page for the target you selected.

2. Click on Target and then choose **Administration**.

   From this menu you can choose either **Execute Cell Command, Support Workbench,** or **Manage I/O Resources**.

3. Click **Execute Cell Command**.

   The Command page of the Exadata Cell Administration wizard appears. Enter a CELLCLI command as the administrative command to be executed on the cell. You must read the Command Instructions before you enter a command. Only a single CELLCLI command is allowed to execute. You must enter the command without the 'cellcli -e' prefix, which is automatically appended when you submit the command. Finally, you cannot use the following characters: ; / ' < > / |.

4. Click **Next** to continue.

   Enterprise Manager displays the Admin Credentials page. Select or enter the Administration credentials to execute the command. The credentials you enter are used when submitting the operation. You can choose between Preferred Credentials, Named Credentials, and New Credentials. You can also click **More Details** to view information about Credential Type, Last modified, Credential Name, Credential Owner, Last Modified Date, Last Modified By, and Preferred Credentials Set At.

5. Click **Next**.

   Enterprise Manager displays the Schedule page. Use the Schedule page to schedule the administration task. Enter the Job Name and the Job Description, then provide the job information in the Schedule the Administration Job section. You can choose to begin the job immediately or enter the time you want the job to begin.

6. Click **Next** to continue.

   The Summary page displays. Use the Summary page to ensure you have entered the correct values and then submit the command. The Summary page lists the Job Name, Description, Command to Execute, when the job is Scheduled, and the Selected Cell.

7. Click **Submit Command** to submit the job.

   The Job Status page displays. Use the Job Status page to link to the Job Detail page of the administration task.

28.7.3 Performing Administration Tasks on Infiniband Networks

To perform an administration operation on an Infiniband Network, follow these steps:

1. Navigate to the DB Machine home page of the Infiniband Network by choosing the DB Machine for which you want to perform an administrative task from the All Targets page of Enterprise Manager.

   Enterprise Manager displays the DB Machine Home page for the target you selected.

2. Select the IB Network for which you want to perform an administrative task.

3. From the Target menu item, choose **Administration**.

   The Target & Command page of the Infiniband Network Administration wizard appears.

4. Choose the **Target Type** and then select the target on which you want to perform the administrative task from the Target drop-down list. Enter the administrative

command you want to execute. The available operations from which you can select are dependent on the target type and target you selected. Once you choose the operation, you may need to select a value that will appear after choosing the operation.

5. Click **Next** to continue.

   Enterprise Manager displays the Credentials & Schedule page. Select or enter the credentials to execute the command. The credentials you enter are used when submitting the operation. You can choose between Preferred Credentials, Named Credentials, and New Credentials. Schedule the administration task. Provide the job information in the **Administration Job Schedule** section. You can choose to begin the job immediately or enter the time you want the job to begin

6. Click **Next** to continue.

   The Review page appears. Use the Review page to ensure you have entered the correct values and then submit the command. The Review page lists the Job Name, Description, Command to Execute, when the job is Scheduled, the Target Type, and the Selected Target.

7. Click **Submit Command** to submit the job.

   When you click on Submit command, a popup is shown if the job is successful. You can go to the Job Detail Page or back to the page from where this wizard was launched.

28.7.4 Launching the IORM Performance Page

To launch the IORM Performance page, follow these steps:

1. Navigate to the Exadata Cell home page by choosing the Exadata target for which you want to view the IORM Performance page from the All Targets page of Enterprise Manager.

   Enterprise Manager displays the Exadata Cell Home page for the target you selected.

2. From the Administration menu item, choose **Administration**. and then **Perform IORM Management**.

   The IORM Performance page appears where you can view the status of the current IORM configuration.

28.7.5 Viewing an Exadata Cell Configuration

You can view the configuration of an Oracle Exadata target by following the steps below:

1. Navigate to the Exadata Cell home page by choosing the Exadata target for which you want to view the IORM Performance page from the All Targets page of Enterprise Manager.

   Enterprise Manager displays the Exadata Cell Home page for the target you selected.

2. From the Target menu, choose Configuration and then Topology.

   Enterprise Manager displays the Configuration Topology page for the selected Exadata Cell. The topology page provides a visual layout of a the target's relationships with other targets. From this page you can:

   - Do a target search filtered by target status/events/target type)

- Select from a set of relationships to represent in the graph
- Select annotations to display in the graph, such as alerts and link labels
- Select from a set of options: view navigator, expand or collapse all, toggle graph layout, reload topology
- Print
- Zoom via the slide control
- Pan via the navigator control
- Toggle the presentation from graph to table

When you hover over a node or group member, a popup displays detailed information about the entity. A link can appear in the popup to more detailed information such as customer documentation.

28.7.6 Managing a Single I/O Resource Management Allocation

You can manage the I/O resource allocation for each cell in your Exadata environment. The Manage I/O Resource page displays the current mode (active or inactive) and allows you to change the mode by using the Update IORM Mode function. You cans also change the IORM Objective using the Update IORM Objective function. The page also shows statistics and metrics for the current Exadata cell.

To manage the I/O Resource Management for a single cell, follow these steps:

1. From the Target menu of the Exadata cell for which you want to manage I/O resources, choose **Administration**.
2. From the Administration menu, choose **Manage I/O Resource**.

   The I/O Resource Management page displays where you can change the IORM Mode between Active and Inactive or you can change the IORM Objective.
3. Use the IORM Objective chart to display the historical value of objectives.
4. Use the IORM Wait chart to display the IORM Wait times for all databases or a selected database. You can select the database (or all databases) to display in the IORM Wait chart by choosing the database from the **Select Database** drop-down list.

For more information about changing the IORM mode or updating the IORM objective, see Changing the IORM Mode and Updating the IORM Objective.

28.7.7 Accessing Oracle Support Workbench for Exadata Cell

You can access the Oracle Support Workbench for the current Exadata cell to access diagnostic data for problems and incidents related to the cell.

To access the Support Workbench for a single Exadata cell, follow these steps:

1. From the Target menu of the Exadata cell for which you want to access the Oracle Support Workbench, choose **Administration**.
2. From the Administration menu, choose **Support Workbench**.

   The Login to Support Workbench page displays where you can enter the credentials required to access the Workbench.
3. Choose the type of cell administration credential you want to use from the Credential section.

You can choose **Preferred Credential**, Named **Credential**, or **New Credentials**.

4. Depending on what you chose in Step 3, enter the Credential Details (**Username** and **Password**) in the appropriate fields. You can turn on the More Details Option to see more details about the credentials.
5. Click **Continue** to display the Oracle Support Workbench.

28.7.8 Changing the IORM Mode and Updating the IORM Objective

Follow these steps to change the IORM Mode or update the IORM Objective.

Changing the IORM Mode

To change the IORM mode from the I/O Resource Management page, follow these steps:

1. From the Target menu of the Exadata cell for which you want to manage I/O resources, choose **Administration**.
2. From the Administration menu, choose **Manage I/O Resource**.

   The I/O Resource Management page displays where you can change the IORM Mode between Active and Inactive. The current IORM Mode is displayed on the page.
3. In the IORM Mode and Objective section, click on the mode (Active or Inactive) to which you want to change and then click on the **Change IORM Mode** button.
4. Enterprise Manager displays the

Updating the IORM Objective

To update the IORM Objective from the I/O Resource Management page, follow these steps:

1. From the Target menu of the Exadata cell for which you want to manage I/O resources, choose **Administration**.
2. From the Administration menu, choose **Manage I/O Resource**.

   The I/O Resource Management page displays where you can update the IORM Objective. The current IORM Objective is displayed on the page.
3. From the Change IORM Objective box within the IORM Mode and Objective section, use the Select drop-down list to update the IORM Objective and then click the **Update IORM Objective** button.

   You can choose one of the IORM Objectives from the list: Low Latency, Balanced, High Throughput, or Auto.
4. Enterprise Manager displays the

29

Using Active Reports

Active Reports provide a flash based interactive report that allows you to save data in an HTML file which can then be used for offline viewing or e-mailing to your peers. Active Reports is not related to either the existing Enterprise Manager Information Publisher reports or BI publisher. The Flash player is a plug-in from Adobe that is required to view flash swf files. You must install the plug-in to view the Active Reports.

The following features are available as part of Active Reports:

- Save or Mail reports

  This feature is common to all active reports embedded in Enterprise Manager. You can click the Save button so that it can be viewed later. Clicking the Save button shows the platform-specific Save dialog box, and allows you to save the HTML file. Clicking the Mail button displays a dialog box where you can fill in the recipients or any message which will be included in the email. The report will be sent as an attachment of the mail.

  The reports that display when you initiate Save look similar to those that are generated by Enterprise Manager but differ in the following ways:

  – Active Reports provide language drop-down menus that allow you to change to a different language

  – Active Reports display no hyperlinks to other Enterprise Manager pages

  – Active Reports display no menu items to other online Enterprise Manager pages

- Show Active Report for SQL Details

  The SQL detail Active Report displays the data in the SQL details report. The Active Report shows similar content as the existing SQL Details report in Enterprise Manager.

- Report for SQL Performance Analyzer (SPA) Report

  The Flex SPA report page is similar to the existing SPA Report but provides a clearer and more interactive interface. The overall report page includes three tabs, one on global statistics, one on SQL statement count, and one using a scatter plot chart or a tree map to display the impact on SQL statements in detail.

For more information about using active reports, see the following sections:

- Using Active Reports to Mail a Report
- Using Active Reports to Save a Report
- Generating an Active Report from the SQL Details Page

29.1 Using Active Reports to Mail a Report

To mail an Active Report, follow these steps:

1. Navigate to the Enterprise Manager page displaying an Active Report Mail button and click **Mail**.

   The Active Reports Mail dialog box appears.

2. In the Active Reports Mail dialog box, enter your email address in the **From** field and then enter the addresses of the person or persons to whom you want to send the report in the **To** field. For multiple addresses, separate each address with a semi-colon (;). You can also enter a message in the **Message** field that describes the report and any other information you want to convey.

3. Click **Send** to send the email with the attached report.

To successfully send mail, you must set up the email server information. If the server information is not set, the dialog box will prompt you to configure the server and provide instructions on how to complete the task. To complete this information manually, you can navigate to the Enterprise Manager Setup link and click **Notification Methods** to fill in the mail server information.

29.2 Using Active Reports to Save a Report

To save an Active Report, follow these steps:

1. Navigate to the Enterprise Manager page displaying an Active Report Save button and click **Save**.

   The Save dialog box displays.

2. In the Active Reports Save dialog box, enter the location on your system where you would like to save the HTML report.

3. Click **Save** to save the report.

Note: The Save feature in Active Reports is different from the Save As function on your browser. The browser saves all the files required to render the page to your local file system. In most cases this will not allow you to see the report in the original form. The Active Reports save feature will save the file in HTML only. All the required files, such as the javascript and swf files used to render the report are hosted remotely on the Oracle Technology Network (OTN).

Note that the report that displays when you initiate Save appears similar to those that are generated by Enterprise Manager but differ in the following ways:

- Active Reports provide a language drop-down menu that allows you to switch to a different language
- Active Reports display no hyperlinks to other Enterprise Manager pages
- Active Reports display no menu items to other online Enterprise Manager pages

29.3 Generating an Active Report from the SQL Details Page

To generate an Active Report from the SQL Details page, follow these steps:

1. Navigate to the SQL Details page and click the **Show Active Report** button.

   The Save dialog box displays.

2. In the Active Reports Save dialog box, enter the location on your system where you would like to save the HTML report.

3. Click **Save** to save the report.

   The generated Active Report displays similar contents as the existing SQL Details page in Enterprise Manager.

   Note: The Save feature in Active Reports is different from the Save As function on your browser. The browser saves all the files required to render the page to your local file system. In most cases this will not allow you to see the report in the original form. The Active Reports save feature will save the file in HTML only. All the required files, such as the javascript and swf files used to render the report are hosted remotely on the Oracle Technology Network (OTN).

   Note that the report that displays when you initiate Save appears similar to those that are generated by Enterprise Manager but differ in the following ways:

   – The Active Reports report provides a language drop-down menu that allows you to switch to a different language

   – Active Reports display no hyperlinks to other Enterprise Manager pages

   – Active Reports display no menu items to other online Enterprise Manager pages

29.4 Generating an Active Report from the SQL Performance Analyzer (SPA) Page

To generate an Active Report from the SQL Performance Analyzer (SPA) page, follow these steps:

1. Navigate to the SQL Performance Analyzer page and click the **Show Active Report** button.

   The Save dialog box displays.

2. In the Active Reports Save dialog box, enter the location on your system where you would like to save the HTML report.

3. Click **Save** to save the report.

   The Flex SPA report page is similar to the existing SPA page in Enterprise Manager, however it uses features in Flex to provide a clear and more interactive display. The overall report page includes three tabs; one on global statistics, one on SQL statement count, and one using a scatter plot chart or a tree map to display the impact on SQL statements in detail.

Note: The Save feature in Active Reports is different from the Save As function on your browser. The browser saves all the files required to render the page to your local file system. In most cases this will not allow you to see the report in the original form. The Active Reports save feature will save the file in HTML only. All the required files, such as the Javascript and SWF files used to render the report are hosted remotely on the Oracle Technology Network (OTN).

Note that the report that displays when you initiate Save appears similar to those that are generated by Enterprise Manager but differ in the following ways:

- The Active Reports report provides a language drop-down menu that allows you to switch to a different language
- Active Reports display no hyperlinks to other Enterprise Manager pages
- Active Reports display no menu items to other online Enterprise Manager pages

30

Using Oracle Exalogic Elastic Cloud

Use Oracle Exalogic Elastic Cloud to discover and monitor instances of Exalogic Elastic Cloud (Middleware Machine). You can then display status information, including alerts and key performance metrics, of the following target types in an Exalogic Elastic Cloud:

- Application Deployments
- Weblogic Domains
- IB Switch - An IB Switch system has IB Switches as members. This is a one to many association. Each IB switch can have many IB Switch Ports as target components.
- Coherence Clusters

Exalogic Elastic Cloud (Middleware Machine) is modeled as a System target rather than a group target.

For more information about using Oracle Exalogic Elastic Cloud, see the following sections:

30.1 Using the Exalogic Elastic Cloud Discovery Wizard

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

One important prerequisite for discovery is that all the components running on the Exalogic Elastic Cloud must already exist in Enterprise Manager as targets. Exalogic Elastic Cloud discovery does not add or discover any of these targets. It simply identifies the targets present in the Exalogic Elastic Cloud and maps them to Enterprise Manager targets and then adds these Enterprise Manager targets as Exalogic Elastic Cloud system members.

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

1. In Enterprise Manager, navigate to the Systems page.

2. From the Add drop-down, menu choose **Exalogic Elastic Cloud** and click **Go**.

   Enterprise Manager displays the Discover Exalogic Elastic Cloud page, seen in Figure 30–1, which allows you to enter the parameters and values required to discover an Oracle Exalogic target.

Figure 30–1 Discover Exalogic Elastic Cloud Page

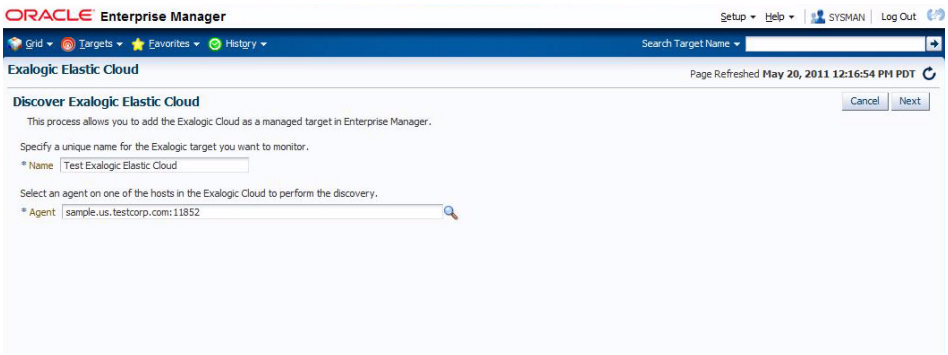

3. Specify a unique name for the Oracle Exalogic target you want to monitor in the Name field.
4. Select an Agent on one of the hosts in the Exalogic System to perform the discovery.

   If you choose an agent that is not part of the Exalogic System, an error message appears stating that no Exalogic Property File can be found and indicating that you must e choose an Agent which is on a Exalogic System Host.
5. Click **Next**. Enterprise Manager displays the Discover Oracle Exalogic Targets: Discovered Targets page. The page displays the Hosts that are discovered as part of the system.
6. Click **Finish** to complete discovery. You can also choose Back to return to the Discover Exalogic Elastic Cloud page or Cancel to terminate the discovery process.

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

30.2 Displaying and Using the Exalogic Elastic Cloud Home Page and Dashboard

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

- Application Deployments
- Weblogic Domains
- Coherence Clusters
- Hosts

You can also use the Home page to access the Hardware tab where you can view information about the hardware and infrastructure of the Exalogic Elastic Cloud.

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

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

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

2. You can view detailed information about each component by choosing the component name from the Exalogic Elastic Cloud drop-down menu.

   Enterprise Manager displays the component page you selected. For example, select Weblogic Domains Summary page to see the charts showing the status of weblogic servers, Request Processing Time metric information, CPU Usage, Requests per minute, and Heap Usage data. For more information about each of those component pages, see the topics in the Related Topics section below.

3. You can return to this page at any time by choosing by choosing Home from the Exalogic Elastic Cloud drop-down menu.

4. You can display General Information about the Exalogic target by choosing General Information from the Exalogic Elastic Cloud drop-down menu.

30.3 Viewing Application Deployments in Exalogic Elastic Cloud Targets

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

To display and use the Application Deployments page, follow these steps:

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

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

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

   Enterprise Manager displays the Application Deployments page.

3. You can choose to show All Domains or filter by specific domains by choosing the domain from the Show menu.

4. You can drill down to specific applications, targets, domains, or dependencies by clicking on its related value in each row.

5. You can filter the list of applications by choosing a value from the Status drop-down. You can select from Up, Down, Unknown, Blackout, and All.

6. You can change the column appearance of the table by clicking **View** and choosing which Columns to display, expanding or collapsing rows, or scrolling to the first or last row. You can also reorder columns.

30.4 Viewing Weblogic Domains in Exalogic Elastic Cloud Targets

Use the Weblogic Domains page in the Exalogic Elastic Cloud target area to view details about the domains hosted on the virtual machines running on Exalogic Elastic Cloud target.

To display and use the Weblogic Domain page, follow these steps:

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

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

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

   Enterprise Manager displays the related Weblogic Domain page.

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

4. On the Members page, you can view the Status information along with alerts and policy violation and metric data for each Weblogic Server or Domain. Use the Performance Summary section to view metrics for each, such as Host and Cluster information and metrics such as Heap Usage and Request Processing Time.

30.5 Viewing Coherence Clusters in Exalogic Elastic Cloud Targets

Use the Coherence Clusters page in the Exalogic Elastic Cloud target area to view details about the Coherence targets hosted on the virtual machines running on the Exalogic Elastic Cloud target.

To display and use the Coherence Clusters page, follow these steps:

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

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

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

   Enterprise Manager displays the Coherence Clusters page.

3. You can view a chart that shows the status of the Coherence Clusters and displays the percentage of clusters that are up and down.

4. You can drill down to specific values for each cluster such as Alerts and Policy Violations along with Node information.

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

30.6 Viewing Hosts in Exalogic Elastic Cloud Targets

Use the Hosts page in the Exalogic Elastic Cloud target area to view details about the host targets hosted on the virtual machines running on the Exalogic Elastic Cloud target.

To display and use the Hosts page, follow these steps:

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

   Enterprise Manager displays the Exalogic Elastic Cloud Home page where you can monitor the status of the Exalogic target and its components.
2. Choose Hosts from the Exalogic Elastic Cloud drop-down menu.

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

31

Installation to Support Real-time Configuration Change Monitoring

The Enterprise Manager Compliance features include the ability to monitor certain elements of your targets in real time to watch for configuration changes or actions that may result in configuration changes.

These features include Operating System level file change monitoring, process starts and stops, Operating System user logins and logouts, Oracle database changes and more.

The real-time monitoring for these features takes place from the Enterprise Manager agent. Some of these monitoring capabilities require specific setup steps depending on the type of monitoring you will do and what Operating System is being monitored.

This section outlines the specific requirements and pre-requisites that exist to use the Compliance Real-time Monitoring features.

31.1 Real-Time Monitoring

Real-time monitoring is configured through the Enterprise Manager Server. Users with the EM_COMPLIANCE_DESIGNER role create Compliance Standard Rules that are of type "Real-time Monitoring Rule." These rules are then associated with Compliance Standards and these standards are subsequently associated with one or more targets.

After the Compliance Standard to target association is complete, the set of monitoring rules are sent to the agent to enable real-time monitoring.

31.2 Resource Consumption Considerations

The Real-time monitoring features are built into the Enterprise Manager agent. There are some specific resource considerations if you use the Real-time monitoring features. The following sections describe issues you should consider when using Real-time monitoring features.

31.2.1 OS File Monitoring Archiving

An optional setting when monitoring for file changes in real time is to make an archive copy of the file on the agent. When monitoring first begins, a copy of the file at that time is made and stored into a private directory in the ORACLE_HOME directory of the agent. Then, any subsequent changes to that file will result in additional copies of the file being archived in that same directory. This feature allows you to later perform a file diff from the user interface or to issue a job to roll back a file to a previous version.

This feature however will use disk space to make copies of the file. Care should be taken to ensure that this feature is only enabled for files that are critical.

31.2.2 OS File Read Monitoring

The Operating System File level monitoring can monitor many types of changes to files, but can also monitor reads to files. If you have a Rule to monitor a facet that has file patterns that are read frequently, this may result in a very large number of observations. You can reduce the number of observations by ensuring that your Rule includes a filter on either time or a user that you want to ensure does not read the file.

For instance, monitoring the */etc/passwd* file for reads for All users will result in many observations being created. However, if you only monitor the */etc/passwd* file by a specific user, then you will only receive an observation when that specific user attempts to read the file.

31.2.3 Creating Facets That Have Very Broad Coverage

It is possible to create a facet for any entity type (Operating System File for instance) that monitors every file on a host. This will result in increased overhead on the agent as well as significant numbers of observations coming into the Enterprise Manager Server. It is important to remember that facets are created to specify files that are very important to monitor for security/compliance purposes. For instance, monitoring all modifies to a log file that change every few seconds will add a great deal of resource usage to the agent and server. Instead, in this case, it may be appropriate to create a rule to monitor the log file for all changes, but filter only when the log change is made by a non-application user. This would only capture the log file change if a regular user attempted to change or tamper with the log rather than when the log is simply being updated by an application.

31.2.4 Enterprise Manager Repository Sizing

Database sizing considerations for Real-time monitoring depend on several factors. The most important factor is the number of observations expected in a month. The second factor is the number of months data will be retained in the repository. Repository retention rates are explained later in this chapter.

In general, each observation consumes roughly 1.5KB of space in the database. This is a guideline and this number can vary depending on many factors for each installation.

If a customer expected a total of 10 million Real-time observations per month across all targets and wanted to retain the data for 12 months, then the database size required for this would be roughly 180GB.

10,000,000 Observations x 12 Months x 1500 Bytes = 180,000,000,000 Bytes

This size represents Real-time monitoring data only and does not include database storage needs for other areas of Enterprise Manager.

The number of observations to expect per month can vary from environment to environment and can also depend on what types of monitoring are configured. You may be required to tune the expected size over time after Rules and Facets have been enabled for some time and configured to fit the organizations requirements. You can easily find your observations usage over a month by using the **Real-time Observations** > **Browse By Systems** screen to select your systems and see the related counts of observations for each system over a period of time.

31.3 Configuring Monitoring Credentials

Many of the real-time monitoring capabilities require monitoring credentials that maintain the ability to launch monitoring programs with root privileges. These processes begin with the prefix *nmxc*. Low-level monitoring uses operating system APIs that are not available to regular users.

Before starting to use the Real-time monitoring features on a target host for the first time, the following settings must be configured from the Enterprise Manager Console.

1. Ensure that *root.sh* script is run after agent installation

   After installing the agent, the `root.sh` script must be run as the root user. This script must be run before configuring the rest of these credential steps.

2. Configuring Privilege Delegation

   Privilege Delegation settings are found from the Setup Menu, **Setup** > **Security** > **Privilege Delegation**. On this page, you can either set privilege delegation for each host manually or you can create a Privilege Delegation Setting Template.

   Privilege delegation for each host that will have real-time monitoring must have SUDO setting enabled with the appropriate SUDO command filled in (for example, `/usr/local/bin/sudo`).

3. Configuring Monitoring Credentials

   Monitoring Credential settings are found from the Setup Menu, **Setup** > **Security** > **Monitoring Credentials**. From this page, select the Host target type and click on the **Manage Monitoring Credentials** button.

   For each entry with the credential "Host Credentials For Real-time Configuration Change Monitoring" set, select the entry and click on the **Set Credentials** button. You will be asked for a credential set to use. Ensure you also add "root" to the Run As entry. If "Run As" is not visible, then the privilege delegation was not set properly in the previous step.

   To set monitoring credentials in bulk on multiple hosts at once, you can use EMCLI. For more information on using EMCLI to set monitoring credentials, see the section, *Managing Credentials Using EMCLI* in the Security chapter of *Oracle Enterprise Manager Administration*. Likewise, for more information about configuring monitoring credentials in Oracle Enterprise Manager, the Security chapter of *Oracle Enterprise Manager Administration*.

31.4 Preparing To Monitor Linux Hosts

The following sections describe how to prepare Linux hosts for monitoring.

31.4.1 OS File Monitoring

Before using Real-time file monitoring for Linux, a loadable kernel module must be installed on the host. This loadable kernel module provides you with the most efficient way of monitoring the host. This loadable kernel module is referred to as the File Audit Module, or Audit Module for short.

Acquiring the Kernel Module

The kernel audit module is available from `http://oss.oracle.com/projects/fileauditmodule`. There are two ways to get the file audit kernel module:

1. **Prebuilt .ko files**. for which Oracle has already prebuilt, you can use this in your environment. You can look for the Prebuilt kernel modules under the **Downloads** link. To find the matching prebuilt version, run the `uname -r` command on the host being monitored and compare that version to the version of the prebuilt modules. The complete version string must match perfectly.
2. **Build your own kernel module.** To build your own kernel module, you can download the following RPM from the **Downloads** link:

   Fileauditmodule-version-noarch.rpm

   You should always retrieve the latest version available at the time you are installing this module.

   Install this RPM on the host you want to monitor as root. The installation of this RPM depends on the kernel-devel package matching your running kernel also existing on the host. This kernel-devel package comes with the same media as the Linux installers.

   In addition to installing this package, you must ensure that the version of gcc available on your host matches the version with which the kernel was built. To do this, view the */proc/version* file to see what gcc version the kernel was built with and then run the command `gcc –version` to see what version of gcc is being used. These two versions should match.

   Also check that the file */boot/System.map-{version}* exists where {version} must match the kernel version you see when you run the `uname -r` command. This file contains system symbols that are required to decode the kernel symbols we are monitoring for real-time changes. Without this file, real-time file monitoring will not function. This file is standard on all default Linux installations.

   After installing this package and checking prerequisites successfully, go to the directory where the package contents were installed (defaults to */opt/fileauditmodule*) and run the *compmod.sh* script. This will build the kernel module file (.ko, .k64, or .o extension depending on the OS version) and place it in the */opt/fileauditmodule* directory.

   If the audit module file is not created, check the *make.log* and *build.log* files for any errors in building the module.

If all of your hosts have the exact same kernel version as shown using the command `uname –r`, then you only need to compile the module on one machine. You can then copy the .ko, .k64, or .o file to the other servers without having to build on that specific host.

Deploying the Kernel Module

Once you have either the prebuilt .ko file or a .ko file that exists from building it from the source RPM, the .ko file must be located in the proper directory. The default location for this file is in the bin folder under the agent home directory. You can also place the file in any location on the host and change the *nmxc.properties* file under the *AGENT_INST/sysman/config* directory of the agent home. The property *nmxcf.kernel_module_dir* specifies the absolute path to the .ko directory.

Install Kernel Module Job

In addition to manually placing the .KO file on the agent, there is an Enterprise Manager job named *CCC Kernel Module Installation*. This job is configured with a list of Linux hosts on which you can install the kernel module. It will search in a directory locally on the Enterprise Manager server disk for prebuilt .ko files or the source RPM

file. If it finds a matching prebuilt .ko file, it will send this to the matching agents; otherwise it will send the RPM to the agent and install it resulting in a new .KO file.

Prior to using this job, files from OSS.ORACLE.COM must be manually retrieved by the user and placed into the *%ORACLE_HOME%/gccompliance/fileauditmodule/resources/linux* directory. This directory already exists on the server with a README file indicating this is the location to place these files. The files that must be placed here are either prebuilt .KO files or the source RPM file. If you have built your own .KO files in your environment, you can also place those .KO files into this directory on the server and deploy it to other hosts in your environment.

Special Considerations for Enterprise Linux 5 and Greater

For Enterprise Linux 5 and greater, the kernel audit module is not required. The monitoring will use the built-in audit subsystem. However, the functionality of the audit subsystem is not as robust as the capability that the kernel audit module can provide.

You will lose the functionality that provides the granularity of what type of change there has been to a file, whether it was a create action or a modify action. Without the kernel module, all changes to a file will appear as a modify action.

It is recommended that you use the kernel audit module even with the newer versions of Linux, if possible.

31.4.2 Debugging Kernel Module Issues

You may detect a problem with the kernel module in a few different ways:

1. You may have noticed that you do not receive real-time file changes on the Enterprise Manager console for file changes that you know should occur.
2. In the Compliance Standard Target Associations page on the user interface, you may see an agent warning indicating a kernel module problem.
3. When examining the *nmxcf.log* file under *AGENT_INST/sysman/logs*, you may see errors indicating that the kernel module could not be loaded or used for various reasons.

If you encounter any of these issues, most likely there was a problem with compiling or inserting the Linux kernel module at run time.

You can confirm whether the auditmodule was loaded properly by running the following command.

```
grep -i auditmodule /proc/modules
```

If you do not get any output, then the auditmodule is not loaded and the agent will not be able to do real time file monitoring.

If the audit module file was generated properly and it does not show up in the module list above, you can attempt to manually load the module to see if there are any errors. Use the following command where you replace {audit module file name} with the entire name of the .ko file that was created from *compmod.sh*:

```
insmod {audit module file name}
```

If you experience no errors during this command, you can check the module list again by using the grep command above. If the audit module now appears, then the file monitoring capability should work. An agent restart is not necessary; however there

still may be a problem with the file monitoring process finding the .ko file which you will experience again next time your host is rebooted.

If the module still is not able to load and if you need to contact Oracle support about the issue, please be sure to include the following information with your support ticket:

- Output of the command: `uname -a`
- Output of the command: `grep -i auditmodule /proc/modules`
- Output of the command: `rpm -q -a |grep -i kernel-devel`
- The make.log and build.log files from the */opt/fileauditmodule* directory where you ran `compmod.sh` if you built your own .ko file
- The files *AGENT_INST/sysman/logs/nmxc*.log*

This information will help Oracle Support to determine if the real time file monitoring audit module of the agent can be built on your environment.

31.5 Preparing To Monitor Windows Hosts

The Real-time monitoring modules for Windows rely on various capabilities of the operating system to collect all of the information on actions. One part of this is to capture the user that made changes from the Windows Event Log. If you do not configure Windows to capture users that make changes, the agent will not capture this information. However it will still capture that a change occurred and when it occurred.

To configure the event log to work with real time monitoring, perform the following steps:

1. From Windows Explorer, select the directory that is being monitored by a Real-time Monitoring Rule, right-click and select **Properties**
2. Go to the Security tab
3. Click the **Advanced** button
4. Select the Auditing tab
5. Click the **Add** button. (In Microsoft XP, double click the **Auditing Entries** window)
6. Select the Name **Everyone** and click **OK**. You can also choose specific users if you are only monitoring for changes by specific users in Configuration Change Console rules. The rules filter the results by user as well, so even if you enable audit for everyone, only users that you want to monitor changes of in your rules will be captured
7. Select the following options (Successful and/or Failed) from the Access window:
   - Create Files/Write Data
   - Create Folders/Append Data
   - Delete Files Subfolders and Files
   - Delete
8. Click **OK** to exit out of the screen
9. Repeat steps 1 through 7 for all other monitored directories and/or files
10. Go **Start** --> **Settings** --> **Control Panel** --> **Administrative Tools** --> **Local Security Policy** --> **Local Policies** --> **Audit Policy**. Double-click, and turn on the following policies (Success and/or Failure):

- Audit account logon events
- Audit logon events
- Audit object access

11. Close the Local Security Settings screen
12. Go to **Start** --> **Settings** --> **Control Panel** -->**Administrative Tools** --> **Event Viewer**
13. Select **System Log**, and click on **Action** from the menu bar and select **Properties**
14. From the System Log Properties panel, on the General tab, set the Maximum log size to at least 5120 KB (5 megabytes) and select **Overwrite Events as Needed**. Note that the log size depends on the number of events generated in the system during a two-minute reporting interval. The log size must be large enough to accommodate those events. If you extend the monitoring time for file events because you expect the change rate to be lower, you need to ensure that the audit log in Windows is large enough to capture the events.
15. Click **Apply** and **OK** to exit.

If Windows auditing is not configured properly, you will see warnings on the Compliance Standard Target Association page on the Enterprise Manager user interface. This is the same page where you associated your Real-time Monitoring compliance Standards to your targets.

31.5.1 Verifying Auditing Is Configured Properly

To verify that the host records login and logout events, follow these steps:

1. Log out of the host and then log back into the host.
2. Go to **Start** --> **Settings** --> **Control Panel** --> **Administrative Tools** --> **Event Viewer**
3. Select **Security Log** and go to **View** --> **Filter**. Select **Security for the Event Source** and **Logon/Logoff** for the Category fields
4. Click **OK**

The Event Viewer should have the activity recorded as Event 528.

31.5.2 Subinacl External Requirements

As mentioned earlier, the agent will send warnings to the server when audit settings are not set properly. It, however, can only do this if the windows feature SUBINACL is installed. If this feature is not installed on the host, a warning will be sent to the server saying that the agent cannot detect whether audit settings are correct. This warning will be visible from the Compliance Standard Associate Targets page.

You can specify the absolute path to the directory that contain subinacl by setting the following property in the *AGENT_INST/sysman/config/nmxc.properties* file:

```
nmxcf.subinacl_dir=
```

SubInACL is available for download from `Microsoft.com`

31.6 Preparing To Monitor Solaris Hosts

Real-time monitoring on Solaris systems utilizes the Solaris audit system which is part of the Solaris Basic Security Model (BSM). BSM auditing allows system administrators to monitor events and to detect user account logins and logouts as well as file changes.

Verify that BSM auditing is enabled by running the following command:

```
/usr/sbin/auditconfig -getcond
```

You should see the following output:

```
audit condition = auditing
```

If you need to enable BSM auditing, you can use the following command:

```
/etc/security/bsmconv
```

See the *Solaris BSM Auditing* manuals for additional details on setting up BSM auditing.

If auditing is already enabled on the server, simply verify that the audit system configuration matches the configurations detailed below.

The audit file can be configured to include specific events. The */etc/security/audit_control* file controls which events will be included in the audit file. This section summarizes the configuration; for further details, refer to the Sun Product Online Documentation site.

For monitoring entity types OS FILE (file changes) and OS USER (user logins/logouts), the flags line in the file */etc/security/audit_control* should be set as follows:

flags: +fw,+fc,+fd,+fm,+fr,+lo

This configuration enables success/fail auditing for file writes (fw), file creates (fc), file deletes (fd), file attribute modifies (fm), file reads (fr) and login/logout events (lo); where '+' means to only log successful events.

If you are interested in logging the failed events as well, remove the "+" sign before each event in the flag.

Note: Installing BSM on an existing host has the requirement that the host is rebooted.

31.6.1 Auditing Users

The *audit_user* file controls which users are being audited. The settings in this file are for specific users and override the settings in the *audit_control* file, which applies to all users.

31.6.2 Audit Logs and Disk Space

The *audit_control* file also has entries to control where the audit logs are stored, and the maximum amount of disk space used by the audit system. The minimum requirement for file monitoring is approximately 10 minutes worth of data stored on the hard drive or the configured reporting interval time.

31.6.3 Managing Audit Files

Enterprise Manager Real-time Monitoring only reads the audit logs; it does not delete the logs. This might flood the system with log files and prevent it from logging additional events. To manage and delete old audit events while maintaining minimum monitoring requirements, do the following:

1. The auditing policy can be set to automatically drop new events (keeping only a count of the dropped events) rather than suspending all processes by running the following command:

   ```
   # auditconfig -setpolicy cnt
   ```

2. Run the following command to force the audit daemon to close the current audit log file and use a new log file.

   ```
   /usr/sbin/audit -s
   ```

3. Run the following command to merge all existing closed auditing log files into a single file with an extension of .trash and then delete the files.

   ```
   /usr/sbin/auditreduce -D trash
   ```

4. Create a cron job to periodically run the commands in step 2 and 3 above. The frequency at which these two commands are run can be adjusted based on the anticipated event volume and the amount of disk space allocated to auditing. The only requirement is that the time between the audit -s command and the auditreduce - D trash command is at least 15 minutes or twice the reporting interval if that is changed.

31.7 Preparing To Monitor the Oracle Database

This section describes the steps involved in setting up auditing within an Oracle database. If you are going to monitor an Oracle database with any of the Oracle entity types, you will need to perform these steps before events will be captured.

Before configuring auditing it is suggested you review the Auditing Database Use section of the *Oracle Database Administrator's Guide*. This document provides an overview of Oracle's auditing functionality, as well as basic concepts and guidelines for auditing configurations. Note that this document does not cover all details of configuring and fine tuning the Oracle audit system. Instead, this document serves as an example of the basic steps involved to configure the Oracle audit system to enable Real-time monitoring through Real-time monitoring rules.

31.7.1 Setting Auditing User Privileges

When you create a Real-time Monitoring Compliance Standard Rule to monitor an Oracle instance target, the agent will read the audit trail to perform its monitoring.

Real-time monitoring for Oracle entity types requires the audit trail to be stored in the database as opposed to a file. To verify if a setting is correct, follow these steps:

1. In Enterprise Manager, go to the target home page for the Oracle Database target for which you want to enable Real-time Monitoring.
2. Click on the **Administration** menu and then choose **Initialization Parameters**.
3. Log in to the database as a sys user, connecting as SYSDBA.
4. Find the parameter *audit_trail* and ensure it is set to DB. If not, this parameter needs to be changed in the Oracle Database.

5. This change will require a restart of the database.

31.7.2 Specifying Audit Options

Through SQL plus, an Oracle DBA can use audit and noaudit statements to configure audit options for the database. The audit statement allows you to set audit options at three levels:

Table 31–1 Audit Options Table

Level	Effect
Statement	Audits specific SQL statements or groups of statements that affect a particular type of database object. For example, AUDIT TABLE audits the CREATE TABLE, TRUNCATE TABLE, COMMENT ON TABLE, and DELETE [FROM] TABLE statements.
Privilege	Audits SQL statements that are executed under the umbrella of a specified system privilege. For Example, AUDIT CREATE ANY TRIGGER audits statements issued using the CREATE ANY TRIGGER system privilege.
Object	Audits specific statements on specific objects, such as ALTER TABLE on the employee table

To use the audit statement to set statement and privilege auditing options a DBA must be assigned AUDIT SYSTEM privileges. To use the audit statement to set object audit options, the DBA must own the object to be audited or be assigned the AUDIT ANY privilege within Oracle. Privilege assignments are covered in the following section.

Audit statements that set statement and privilege audit options can also include a BY clause to supply a list of specific users or application proxies to audit, and thus limit the scope of the statement and privilege audit options.

Some examples of audit statements can be seen below. Feel free to use these as a basis for the audit settings you specify within your database. Once all audit settings are in place you can create application policies, using the Oracle (SQL Trace) agent module with which to monitor the Oracle database instance.

Statement Audit Options (User sessions)

The following statement audits user sessions of users Bill and Lori.

```
AUDIT SESSION BY scott, lori;
```

Privilege Audit Options

The following statement audits all successful and unsuccessful uses of the DELETE ANY TABLE system privilege:

```
AUDIT DELETE ANY TABLE BY ACCESS WHENEVER NOT SUCCESSFUL;
```

Object Audit Options

The following statement audits all successful SELECT, INSERT, and DELETE statements on the dept table owned by user jward:

```
AUDIT SELECT, INSERT, DELETE ON jward.dept BY ACCESS WHENEVER
SUCCESSFUL;
```

Example Oracle Audit Monitor Configurations

The following command audits all basic statements. Extra statements are not audited.

```
Audit all by access;
```

The following statement audits all extra statements:

```
audit ALTER SEQUENCE, ALTER TABLE, DELETE TABLE, EXECUTE
PROCEDURE, GRANT DIRECTORY, GRANT PROCEDURE, GRANT SEQUENCE,
GRANT TABLE, GRANT TYPE, INSERT TABLE, LOCK TABLE, UPDATE TABLE
by access;
```

The following command displays audit settings for statements:

```
SELECT * FROM DBA_STMT_AUDIT_OPTS;
```

Once you have specified your audit configuration you can then create real-time monitoring rules from the Enterprise Manager Server that uses the Oracle Database entity types.

31.8 Setting Up Change Request Management Integration

This section explains how to install and configure integration with a Change Management Server and to be able to determine whether changes that occur are authorized automatically.

31.8.1 BMC Remedy Action Request System 7.1 Integration

Remedy ARS 7.1 is a supported Change Management system for automatic reconciliation of observations. The following steps outline how to setup Remedy and also configure the integration with Enterprise Manager.

31.8.1.1 Remedy Installation and Customization

Follow these steps to install and customize Remedy ARS 7.1.

1. Install Remedy ARS 7.1. Ensure the following components are all installed and properly licensed:

   ARS 7.1.00 Patch 011

   Midtier 7.1.00 Patch 011

   Flashboard Server 7.0.03

   Assignment Engine 7.1

   Asset Management 7.0.03*

   CMDB 2.1.00 Patch 4

   CMDB Extension Loader

   Approval Server 7.1

   Change Management Server 7.0.03 Patch 008*

   Problem Management Server 7.0.03*

   Incident Management Server 7.0.0*3

   User Client

   Administrator Client

   These packages all come with the IT Service Management Pack. Oracle provides example customizations for the Remedy under ITSM 7.0.03 Patch 008 environment. For different versions, the customizations may need to be adjusted to account for changes in the version of Remedy.

2. Install the Enterprise Manager EMCLI_Client on the same host that Remedy is installed on. This will need to be able to communicate to your Enterprise Manager Server.
   a. Login to the Enterprise Manager console
   b. Navigate to **Setup** > **My Preferences** > **Command Line Interface**
   c. Click **Download the EM CLI kit to your workstation** link and download the jar to your Remedy server.
   d. Follow the steps given in the page to install the EMCLI client on the Remedy server.
3. Get the latest version of the Change Request Management connector self-update package. Also acquire the latest version of the example Remedy ARS customizations for Enterprise Manager version 12c.

   These definition files provide a guideline of customizations that must be made in your environment for the integration. These customization files assume a fresh install of Remedy ARS. When integrating with a production instance of Remedy, care should be taken to make sure these customizations are compatible with any previous customizations that have been made to the Remedy instance.
   - ActiveLinks _Customization.def
   - Forms_Customization.def
   - Menus_Customization.def
   - Webservices_Customization.def
4. Install the four definition files (.DEF) files in the running Remedy environment by completing these steps:
   a. Log into the Remedy Administrator tool
   b. Select the **Remedy** instance from the hierarchy on the left
   c. Click **Tools** menu > **Import Definitions** > **From Definition File...**
   d. Select the definition file to import from the list above
   e. Check the box labeled **Replace Objects on the Destination Server**
   f. Choose the drop down option **Replace With New Type**
   g. Click **Import**
   h. You should not encounter any errors during this process. At the end of import there should be an Import Complete message.
   i. When done, repeat for the rest of the customization files.
5. Customize Web Services
   a. Log into Remedy Administrator tool
   b. Select **Webservices** and then select the webservice **EMCCC_GetCR**. Right click and select **Open**
   c. Select the **WSDL** tab
   d. In the input on top, modify the midtier_server and servername values in the **WSDL Handler URL**.
   e. If midtier is on localhost, you can enter localhost right after http://

f. If the midtier uses port 80, you can omit the port, otherwise include the port after the server name

g. For the servername after "public/", enter the name of the Remedy server.

h. Click the **View** button

i. You should see an XML representing the webservice WSDL file for the webservice

j. If you see an error, check the midtier_server name, port, or servername. Also, you can try adding/removing the domain part of the servername.

k. If you see the XML content after clicking View, then close this window and save the changes

l. Repeat all above steps with the webservices EMCCC_PublishCSData and EMCCC_UpdateChangeRequest

6. Customize Active Links

   a. Log into Remedy Administrator tool

   b. Select active links and then select the active link **EMCCCC_ApprovedCR**. Right click and select **Open**

   c. Click on the tab **If Action**

   d. Click on the Current Action **Run Process** at the end of the list of actions

   e. In the Command Line field, change the path to *emcli.bat* to match that of where you installed the emcli on the local host.

7. Create a user in Remedy that will be used for creating requests that will be used for automatic observation reconciliation:

   a. Log into BMC Remedy User Client as an administrative user

   b. Click on **Application Administration Console** on the User Client Home Page

   c. Click on **Create** for each step 1 through 4 in this wizard

   d. When adding the person, add the support group under the Support Groups tab.

   e. Under the Support Groups Tab, select sub tab **Support Group Functional Roles**

   f. Add Support groups with functional role of *Infrastructure Change Management*. Without this, you will not be able to create change requests as the Infrastructure Change Manager fields support group will not have values.

   g. Go to AR System Administrator Console

   h. On the left side bar, click **Application** > **Users/Groups/Roles** > **Select Users**.

   i. This will load the user search page. Click the **Search** button at the top right

   j. Double click the newly created user above to bring up the user form

   k. Click the down arrow next to "Group List" field and select I**nfrastructure Change Master**

   l. Repeat the previous step and add the following Groups to this user as well

      Infrastructure Change Submit

      Infrastructure Change User

Infrastructure Change Viewer

m. Save the changes to this user by clicking the **Save** button in the upper right hand corner of the window.

31.8.1.1.1 Adding the Connector to Enterprise Manager

Follow these steps to add the connector to Enterprise Manager.

1. Add the Change Management Connector to Enterprise Manager
   a. Log into Enterprise Manager as an Administrative user that has privileges to create a connector.
   b. Go to **Enterprise** > **Change Management** > **Software Library**
   c. Click **Actions** > **Administration**
   d. Click **Add**
   e. Provide a name, such as "self update swlib"
   f. Provide a location where the swlib files will be located on the Enterprise Manager server. This can be anywhere, but must be a path that the Enterprise Manager user can access. You must put the full absolute path in this input.
   g. This process will take several minutes to complete
   h. Locate the connector self-update package file.

      The connectors jar can be downloaded from the Enterprise Manager store to EM@Customer using the Self Update console, and can be exported to any local directory using the export functionality of Self Update.

   i. Run: `emcli import_update -file=<full path>/connector.zip -omslocal` (where *connector.zip* is an example name of the self update package)
   j. If you have errors with the previous step, make sure the user you run emcli as has permissions to access this directory and file. Also, be sure you are using absolute path for the *-file* switch
   k. When successful, you will receive the following message:

      Operation completed successfully. Update has been uploaded to Enterprise Manager. Please use the Self Update Home to manage this update.

   l. Log into the Enterprise Manager console
   m. Navigate to **Setup** > **Extensibility** > **Self Update**
   n. Find the type "Management Connector" and click on the link "1" under "Downloaded Updates" for this entry.
   o. Select the Connector from the table and click the **Apply** button.
2. Create a Change Management Connector instance
   a. Login into Enterprise Manager console
   b. Navigate to **Setup** > **Extensibility** > **Management Connectors**
   c. Select "Remedy Change Management Connector" from the dropdown after "Create Connector" and click **Go**
   d. Provide a name and description for the connector. This name is used to choose the connector when creating a Real-time Monitoring Compliance Standard Rule.

e. After returning to the management connector listing page, select the newly added row and click **Configure**.

f. Under the Web Service End Points label, change the [servername] and [port] to match that of your Remedy instance Web Services. The values you put here will be similar to what you configured in the Web Services step earlier in these instructions.

g. Enter the Remedy username and password you are using for the connector integration

h. Enter the locale ('en', for example)

i. Enter the time zone offset of the remedy server from UTC, ('-08:00', for example)

j. Enter the Change ID to use as a test. This should be a valid Change Request ID currently existing in Remedy that is used to test the connectivity between Enterprise Manager and Remedy.

31.8.1.1.2 Using Automatic Reconciliation Rules

Once Remedy is customized and the Enterprise Manager connector is configured, to utilize the automatic reconciliation features, you need to create Real-time Monitoring Rules that are configured to use automatic reconciliation. Use the following steps:

1. Create a Real-time monitoring Rule:

   a. Follow normal steps to create a Real-time monitoring Rule

   b. On the Settings page, choose **Authorized Observations Automatically** using Change Request Management System

   c. Select the connector from the drop down

   d. Click to annotate change requests with authorized observations check box

   e. Continue to save the rule after this. The Real-time Monitoring Rule can be used like any other Real-time Monitoring rule. Create a Compliance Standard, add this rule to the Compliance Standard, and associate this compliance standard to one or more targets.

31.8.1.1.3 Creating Change Requests for Upcoming Changes

Now that integration is set up and Real-time monitoring rules have been created, Change Requests can be created by Remedy users in the Remedy interface. These Change Requests will be compared to observations that occur to automatically determine if these observations are from actions that were authorized by change requests or not.

To make this correlation, some new fields that have been added to the Change Request form must be filled out by the change request filer. Not all fields are required; correlation only occurs on the fields that are present in the Change Request.

For instance, the following fields have been added to the Change Request form under the Oracle Enterprise Manager Integration tab:

- **Connector**: Choose the Enterprise Manager connector this Change Request will use to integrate with Enterprise Manager
- **Hostname**: the hostname(s) this change request is for. These are the hosts that this change request is specifying someone needs to make changes to. An empty value in this field indicates that all hosts will be correlated to this change request.

- **Target User List**: the user name(s) this change request is for based on target users. These are the target users you expect to log in to the target to make a change. An empty value in this field means that all users on the target will be correlated to this change request.
- **Target Type**: the target type this change request is against. An empty value in this field means that any target type will be correlated to this change request.
- **Target**: The target this change request is specifically for. An empty value in this field indicates that any target will be correlated to this change request.
- **Facet**: The facet this change request is specifically for. An empty value in this field indicates that all facets on the above target type and target will be correlated for this change request.

When creating a change request that you want to use to authorize changes detected by Real-time monitoring rules, follow these steps in addition to whatever requirements your organization implements for creation of Change Requests:

1. Under the Dates tab of the Change Request form, fill out the Scheduled Start date and Scheduled End Date. These are the date ranges the request is valid for reconciliation. If an action occurs outside this time, it is marked as unauthorized by the Real-time Monitoring feature.
2. Click on the Oracle Enterprise Manager **Integration** tab
3. Select the Enterprise Manager connector from the drop-down list
4. Optionally select values for the five reconciliation criteria as described above: Hostname, Target User List, Target type, Target and Facet. The last three -- Target Type, Target, and Facet -- will be Choice lists based on content in Real-time Monitoring Rules that have been created in Enterprise Manager that belong to Compliance Standards which are associated to targets. You can add multiple values separated by commas.

> **Note:** This form can be customized in Remedy to look differently. The example form elements from the customizations loaded earlier are only examples.

5. Change auditable status to True. This configures Remedy to use this change request by Enterprise Manager for reconciliation of Real-time Observations that are detected.
6. Save the change request.
7. A popup displays notifying you that active links will send the content to Enterprise Manager. You will see a DOS command window open and then close.

31.8.1.1.4 Overview of Reconciliation Functionality

After creating a change request that references a target and/or facet that is being monitored by Real-time Monitoring rules, any observations that happen against that rule will be correlated to all open and matching change requests.

When the observation arrives at the Enterprise Manager server, all open change requests that were active (based on Scheduled Start/Stop time) and have matching correlation criteria from the Enterprise Manager Integration tab will be evaluated. If any change request exists that matches the criteria of the observation, this observation will be marked with an “authorized” audit status. If the annotation check box was

checked in the Rule configuration, details of these authorized observations will be put into a table in the Enterprise Manager Integration tab of the Remedy Change Request.

If no open change requests can be correlated to the observation and the rule was configured to use automatic reconciliation, then this observation is set to an Unauthorized audit status. The Observation bundle to which this observation belonged will be in violation and results in an Enterprise Manager event being created. This event can further be used through creation of an Enterprise Manager Event Rule.

An observations audit status can be seen whenever looking at observation details either in the Compliance > Real-time Observations > Observation Search or either of the Browse By screens. A user with the proper role can also override the audit status for individual observations from these pages.

Any bundles that are in violation because they contain unauthorized observations will be reflected as violations in the Compliance Results page. These violations cause the compliance score skew lower. If these violations are cleared, the score becomes higher; however, the history of these audit status changes will be retained for the given observation.

31.9 Repository Views Related to Real-time Monitoring Features

The following views exist to allow access to Real-time Monitoring data.

View: mgmt$ccc_all_observations

Description: This view returns all observations that have occurred. Any query against this view should ensure that filtering is done on appropriate fields with *action_time* being the first to take advantage of partitions.

Fields:

Field	Description
OBSERVATION_ID	Unique ID given to the observation when detected by the agent
BUNDLE_ID	Bundle to which this observation belongs based on rule bundle settings
TARGET	Target this observation was found against
TARGET_TYPE	Type of the target
ENTITY_TYPE	Entity type of the entity that had an action against it
ACTION	Action that was observed
ACTION_TIME	Time the action occurred
USER_TYPE	Type of user that performed the action (for example, OS user versus DB user)
USER_PERFORMING_ACTION	Name of the user that performed the action
ORIGINAL_USER_NAME	Previous user name in the case of a SU/SUDO action (only applicable to some entity types)
AFFECTED_ENTITY_NAME	Name of the entity that was affected by this action (file name, and so on)
SOURCE_HOST_IP	Source IP of a connection when an action comes from another host (only applicable to some entity types)

Field	Description
ACTION_PROCESS_ID	PID of the process that performed the action (only applicable to some entity types)
ACTION_PROCESS_NAME	Name of the process that performed the action (only applicable to some entity types)
ACTION_PARENT_PROCESS_ID	PID of the parent process of the process that performed the action (only applicable to some entity types)
ACTION_PARENT_PROCESS_NAME	Name of the parent process of the process that performed the action (only applicable to some entity types)
ENTITY_PREVIOUS_VALUE	Previous value of the entity (only applicable to some entity types)
ENTITY_NEW_VALUE	New value of the entity (only applicable to some entity types)
FILE_ENTITY_PREVIOUS_MDS_HASH	Previous MD5 hash value of the entity (only applicable to some entity types)
FILE_ENTITY_NEW_MDS_HASH	New MD5 hash value of the entity (only applicable to some entity types
AUDIT_STATUS	Current audit status of the observation (unaudited, authorized, unauthorized, and so on)
AUDIT_STATUS_SET_DATE	Date the most recent audit status was set
AUDIT_STATUS_SET_BY_USER	User who set the most recent audit status

View: mgmt$ccc_all_obs_bundles

Description: This view returns a summary of all observations bundles. Any query against this view should ensure that filtering is done on appropriate fields with *bundle_start_time* being the first to take advantage of partitions.

Fields:

Field	Description
BUNDLE_ID	Bundle to which this observation belongs based on rule bundle settings
TARGET	Target this observation was found against
TARGET_TYPE	Type of the target
RULE_NAME	Name of the Real-time Monitoring Compliance Standard Rule
ENTITY_TYPE	Entity type of the entity that had an action against it
USER_PERFORMING_ACTION	Name of the user that performed the action
BUNDLE_IN_VIOLATION	Boolean value if the bundle currently is in violation. This means at least one observation in the bundle is unauthorized. True indicates the bundle is in violation.
BUNDLE_START_TIME	Date of the first observation in this bundle
BUNDLE_CLOSE_TIME	Date when this bundle was closed

Field	Description
BUNDLE_CLOSE_REASON	Explanation of why this bundle was closed
DISTINCT_OBS_COUNT	Total number of observations in this bundle
AUTHORIZED_OBS_COUNT	Number of observations in this bundle that are currently authorized
UNAUTHORIZED_OBS_COUNT	Number of observations in this bundle that are currently unauthorized
UNAUTH_CLEARED_OBS_COUNT	Number of observations in this bundle that are currently cleared (that were at one point unauthorized)
UNAUDITED_OBS_COUNT	Number of observations in this bundle that are currently unaudited. They have not been evaluated manually or with Change Management integration to determine audit status.

31.10 Modifying Data Retention Periods

Real-time Monitoring features use partitioning and data retention configuration.

The following are the tables along with their default retention periods. When changing any retention periods, all tables related to Real-time monitoring must be changed to the same value to ensure that data is consistent across various features.

Note: For more information about modifying data retention values, see the chapter "Maintaining and Troubleshooting the Management Repository" in the book *Oracle Enterprise Manager Administration*.

Table Name	Default Retention Period
EM_CCC_WATCHDOG_ALERTS	366 Days
EM_CCC_HISTORY_JOBEXEC	366 Days
EM_CCC_OBSERVATION	366 Days
EM_CCC_OBSGROUP	366 Days
EM_CCC_OBS_GROUP_MAP	366 Days
EM_CCC_HISTORY_OBS_STATUS	366 Days
EM_CCC_HA_OBS	366 Days
BUNDLE_START_TIME	366 Days
BUNDLE_CLOSE_TIME	366 Days
BUNDLE_CLOSE_REASON	366 Days
EM_CCC_HA_OBSGROUP	366 Days
EM_CCC_FILEOBS_DIFF	366 Days
EM_CCC_AUTHOBS_CR_MAP	366 Days

Part II

Enterprise Manager High Availability

This section covers Enterprise Manager high availability best practices and strategies that allow you to safeguard your Oracle Enterprise Manager installation.

32

Setting Up Enterprise Manager High Availability

This chapter covers the following topics:

32.1 Installation Best Practices for Enterprise Manager High Availability

The following sections document best practices for installation and configuration of each Cloud Control component.

32.1.1 Configuring the Management Agent to Automatically Start on Boot and Restart on Failure

The Management Agent is started manually. It is important that the Management Agent be automatically started when the host is booted to insure monitoring of critical resources on the administered host. To that end, use any and all operating system mechanisms to automatically start the Management Agent. For example, on UNIX systems this is done by placing an entry in the UNIX `/etc/init.d` that calls the Management Agent on boot or by setting the Windows service to start automatically.

32.1.2 Configuring Restart for the Management Agent

Once the Management Agent is started, the watchdog process monitors the Management Agent and attempts to restart it in the event of a failure. The behavior of the watchdog is controlled by environment variables set before the Management

Agent process starts. The environment variables that control this behavior follow. All testing discussed here was done with the default settings.

- EM_MAX_RETRIES – This is the maximum number of times the watchdog will attempt to restart the Management Agent within the EM_RETRY_WINDOW. The default is to attempt restart of the Management Agent 3 times.
- EM_RETRY_WINDOW - This is the time interval in seconds that is used together with the EM_MAX_RETRIES environmental variable to determine whether the Management Agent is to be restarted. The default is 600 seconds.

The watchdog will not restart the Management Agent if the watchdog detects that the Management Agent has required restart more than EM_MAX_RETRIES within the EM_RETRY_WINDOW time period.

32.1.3 Installing the Management Agent Software on Redundant Storage

The Management Agent persists its intermediate state and collected information using local files in the `$AGENT_HOME/$HOSTNAME/sysman/emd` sub tree under the Management Agent home directory.

In the event that these files are lost or corrupted before being uploaded to the Management Repository, a loss of monitoring data and any pending alerts not yet uploaded to the Management Repository occurs.

At a minimum, configure these sub-directories on striped redundant or mirrored storage. Availability would be further enhanced by placing the entire $AGENT_HOME on redundant storage. The Management Agent home directory is shown by entering the command '`emctl getemhome`' on the command line, or from the Management Services and Repository tab and Agents tab in the Cloud Control console.

32.1.4 Install the Management Service Shared File Areas on Redundant Storage

The Management Service contains results of the intermediate collected data before it is loaded into the Management Repository. The loader receive directory contains these files and is typically empty when the Management Service is able to load data as quickly as it is received. Once the files are received by the Management Service, the Management Agent considers them committed and therefore removes its local copy. In the event that these files are lost before being uploaded to the Management Repository, data loss will occur. At a minimum, configure these sub-directories on striped redundant or mirrored storage. When Management Services are configured for the Shared Filesystem Loader, all services share the same loader receive directory. It is recommended that the shared loader receive directory be on a clustered file system like NetApps Filer.

32.2 Managing Multiple Hosts and Deploying a Remote Management Repository

Installing all the Cloud Control components on a single host is an effective way to initially explore the capabilities and features available to you when you centrally manage your Oracle environment.

A logical progression from the single-host environment is to a more distributed approach, where the Management Repository database is on a separate host and does not compete for resources with the Management Service. The benefit in such a configuration is scalability; the workload for the Management Repository and

Management Service can now be split. This configuration also provides the flexibility to adjust the resources allocated to each tier, depending on the system load.

In this more distributed configuration, data about your managed targets travels along the following paths so it can be gathered, stored, and made available to administrators by way of the Grid Control console:

1. Administrators use the Grid Control console to monitor and administer the targets just as they do in the single-host scenario described in Section 32.3, "Deploying Cloud Control Components on a Single Host".
2. Management Agents are installed on each host on the network, including the Management Repository host and the Management Service host. The Management Agents upload their data to the Management Service by way of the Management Service upload URL, which is defined in the `emd.properties` file in each Management Agent home directory. The upload URL uploads the data directly through the Oracle HTTP Server.
3. The Management Repository is installed on a separate machine that is dedicated to hosting the Management Repository database. The Management Service uses JDBC connections to load data into the Management Repository database and to retrieve information from the Management Repository so it can be displayed in the Grid Control console. This remote connection is defined in the Administration Server and can be accessed and changed through emctl commands.
4. The Management Service communicates directly with each remote Management Agent over HTTP by way of the Management Agent URL. The Management Agent URL is defined by the `EMD_URL` property in the `emd.properties` file of each Management Agent home directory. As described in Section 32.3, "Deploying Cloud Control Components on a Single Host", the Management Agent includes a built-in HTTP listener so no Oracle HTTP Server is required on the Management Agent host.

32.3 Deploying Cloud Control Components on a Single Host

In Enterprise Manager release 12*c*, the installation does not create a new database. You must install a database which should be on the same host as Enterprise Manager.

When you install all the Cloud Control components on a single host, the management data travels along the following paths:

1. Administrators use the Grid Control console to monitor and administer the managed targets that are discovered by the Management Agents on each host. The Grid Control console uses the following URL to connect to the Oracle HTTP Server:

   ```
   https://host1.acme.com:7799/em
   ```

   The Management Service retrieves data from the Management Repository as it is requested by the administrator using the Grid Control console.
2. The Management Agent loads its data (which includes management data about all the managed targets on the host, including the Management Service and the Management Repository database) by way of the Oracle HTTP Server upload URL. The Management Agent uploads data directly to Oracle HTTP Server. The default port for the upload URL is 1159 (if it is available during the installation procedure). The upload URL is defined by the `REPOSITORY_URL` property in the following configuration file in the Management Agent home directory:

   ```
   AGENT_HOME/sysman/config/emd.properties (UNIX)
   AGENT_HOME\sysman\config\emd.properties (Windows)
   ```

See Also: For more information about the Oracle Enterprise Manager directory structure (AGENT_HOME directory in particular), see the *Oracle® Enterprise Manager Cloud Control Advanced Installation and Configuration Guide*.

3. The Management Service uses JDBC connections to load data into the Management Repository database and to retrieve information from the Management Repository so it can be displayed in the Grid Control console. The Management Repository connection details can be listed and changed by using the following emctl commands:

```
emctl config oms -list_repos_details
emctl config oms -store_repos_details
```

See Also: "Reconfiguring the Oracle Management Service" on page 34-9 for more information on modifying the Management Repository connection information.

4. The Management Service sends data to the Management Agent by way of HTTP. The Management Agent software includes a built-in HTTP listener that listens on the Management Agent URL for messages from the Management Service. As a result, the Management Service can bypass the Oracle HTTP Server and communicate directly with the Management Agent. If the Management Agent is on a remote system, no Oracle HTTP Server is required on the Management Agent host.

   The Management Service uses the Management Agent URL to monitor the availability of the Management Agent, submit Enterprise Manager jobs, and other management functions.

   The Management Agent URL can be identified by the `EMD_URL` property in the following configuration file in the Management Agent home directory:

```
AGENT_HOME/sysman/config/emd.properties (UNIX)
AGENT_HOME\sysman\config\emd.properties (Windows)
```

   For example:

```
EMD_URL=https://host1.acme.com:3872/emd/main
```

   In addition, the name of the Management Agent as it appears in the Grid Control console consists of the Management Agent host name and the port used by the Management Agent URL.

32.4 Installing Multiple OMSs in Active/Active configuration

In larger production environments, you may find it necessary to add additional Management Service installations to help reduce the load on the Management Service and improve the efficiency of the data flow.

Note: When you add additional Management Service installations to your Cloud Control configuration, be sure to adjust the parameters of your Management Repository database. For example, you will likely need to increase the number of processes allowed to connect to the database at one time. Although the number of required processes will vary depending on the overall environment and the specific load on the database, as a general guideline, you should increase the number of processes by 40 for each additional Management Service.

For more information, see the description of the PROCESSES initialization parameter in the *Oracle Database Reference*.

Understanding the Flow of Management Data When Using Multiple Management Services

In a multiple Management Service configuration, the management data moves along the following paths:

1. Administrators can use one of two URLs to access the Grid Control console. Each URL refers to a different Management Service installation, but displays the same set of targets, all of which are loaded in the common Management Repository. Depending upon the host name and port in the URL, the Grid Control console obtains data from the Management Service (by way of the Oracle HTTP Server) on one of the Management Service hosts.
2. Each Management Agent uploads its data to a specific Management Service, based on the upload URL in its `emd.properties` file. That data is uploaded directly to the Management Service by way of Oracle HTTP Server.

   Whenever more than one Management Service is installed, it is a best practice to have the Management Service upload to a shared directory. This allows all Management Service processes to manage files that have been uploaded from any Management Agent. This protects from the loss of any one Management Server causing a disruption in upload data from Management Agents.

   Configure this functionality from the command line of each Management Service process as follows:

   ```
   emctl config oms loader -shared <yes|no> -dir <load
   directory>
   ```

Important: By adding a load balancer, you can avoid the following problems:

- Should the Management Service fail, any Management Agent connected to it cannot upload data.
- Because user accounts only know about one Management Service, users lose connectivity should the Management Service go down even if the other Management Service is up.

See Section 33.3, "High Availability Configurations" for information regarding load balancers.

Note: If the software library is being used in this environment, it should be configured to use shared storage in the same way as the shared Management Service loader. To modify the location for the software library:

1. Click the **Deployments** tab on the Enterprise Manager Home page.
2. Click the **Provisioning** subtab.
3. On the Provisioning page, click the **Administration** subtab.
4. In the Software Library Configuration section, click **Add** to set the Software Library Directory Location to a shared storage that can be accessed by any host running the Management Service.

3. Each Management Service communicates by way of JDBC with a common Management Repository, which is installed in a database on a dedicated Management Repository host. Each Management Service uses the same database connection information, defined in the `emgc.properties` file, to load data from its Management Agents into the Management Repository. The Management Service uses the same connection information to pull data from the Management Repository as it is requested by the Grid Control console.
4. Any Management Service in the system can communicate directly with any of the remote Management Agents defined in the common Management Repository. The Management Services communicate with the Management Agents over HTTP by way of the unique Management Agent URL assigned to each Management Agent.

   The Management Agent URL is defined by the `EMD_URL` property in the `emd.properties` file of each Management Agent home directory. Each Management Agent includes a built-in HTTP listener so no Oracle HTTP Server is required on the Management Agent host.

32.4.1 Configuring the First Management Service for High Availability

Once you configure the Management Repository, the next step is to install and configure the Enterprise Manager Cloud Control mid-tier, the Management Services, for greater availability. Before discussing steps that add mid-tier redundancy and scalability, note that the Management Service itself has a built-in restart mechanism based on the Oracle Weblogic Node Manager and the Oracle Process Management and Notification Service (OPMN). These services will attempt to restart a Management Service that is down. It is advisable to run OPMN and Node Manager as operating system services, so that they restart automatically if their host machine is restarted.

32.4.1.1 Management Service Install Location

If you are managing a large environment with multiple Management Services and Management Repository nodes, then consider installing the Management Services on hardware nodes that are different from Management Repository nodes (Figure 33–3). This allows you to scale out the Management Services in the future.

Figure 32–1 Management Service and Management Repository on Separate Hardware

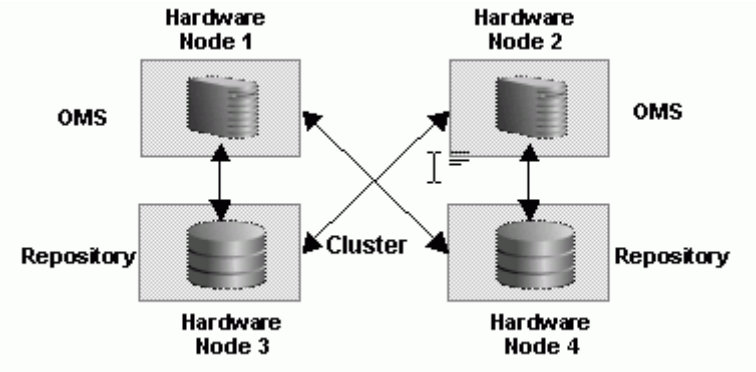

Also consider the network latency between the Management Service and the Management Repository while determining the Management Service install location. The distance between the Management Service and the Management Repository may be one of the factors that affect network latency and hence determine Enterprise Manager performance.

If the network latency between the Management Service and Management Repository tiers is high or the hardware available for running Enterprise Manager is limited, then the Management Service can be installed on the same hardware as the Management Repository (Figure 33–4). This allows for Enterprise Manager high availability, as well as keep the costs down.

Figure 32–2 Management Service and Management Repository on Same Hardware

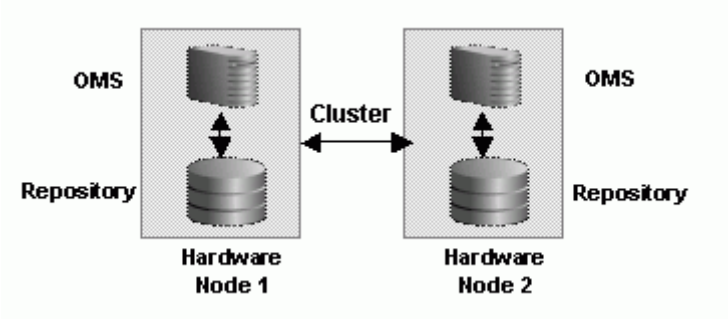

If you plan ahead, you can configure your Enterprise Manager deployment for high availability by choosing the correct options during the first Management Service install. You can also retrofit the MAA best practices into an existing Enterprise Manager deployment configured initially using the default install options.

32.4.2 Configuring Additional Management Services

Once your first Management Service is setup for high availability, there are two paths to setting up your additional Management Services for high availability:

- Installing a fresh additional Management Service as per MAA best practices
- Retrofitting MAA best practices on existing Additional Management Service

In either of the two cases, the following considerations should be noted:

- The additional Management Service should be hosted in close network proximity to the Management Repository database for network latency reasons.
- Configure the same path on all Management Services for the directory used for the shared filesystem loader.
- Additional Management Services should be installed using the same OS user and group as the first Management Service. Proper user equivalence should be setup so that files created by each Management Service on the shared loader directory can be accessed and modified by the other Management Service processes.
- Adjust the parameters of your Management Repository database. For example, you will likely need to increase the number of processes allowed to connect to the database at one time. Although the number of required processes will vary depending on the overall environment and the specific load on the database, as a general guideline, you should increase the number of processes by 40 for each additional Management Service.

32.4.2.1 Installing a Fresh Additional Management Service According MAA Best Practices

Install the additional Management Service using the OUI installer. The additional Management Service will inherit most of the HA configuration from the first Management Service. Post installation, do the following step to complete the HA configuration:

- Update the SLB configuration by adding the additional Management Service to the different pools on the SLB. Setup monitors for the new Management Service.

32.4.2.2 Retrofitting MAA Best Practices on Existing Additional Management Service

Once you have the additional Management Service installed, use the following steps to copy over the configuration from the first Management Service.

1. Export the configuration from the first Management Service using `emctl exportconfig oms -dir <location for the export file>`
2. Copy over the exported file to the additional Management Service
3. Shutdown the additional Management Service
4. Import the exported configuration on the additional Management Service using `emctl importconfig oms -file <full path of the export file>`
5. Restart the additional Management Service
6. Setup EMCLI using `emcli setup -url=https://slb.example.com/em -username sysman -password <sysman password> -nocertvalidate`
7. Resecure the Management Agent that is installed with the additional Management Service to upload to SLB using `emctl secure agent -emdWalletSrcUrl https://slb.example.com:<upload port>/em`
8. Update the SLB configuration by adding the additional Management Service to the different pools on the SLB. Setup monitors for the new Management Service. Modify the ssl.conf file to set the Port directive to the SLB virtual port used for UI access.

32.4.3 Configuring Shared File System Loader

The Management Service for Cloud Control has a high availability feature called the Shared Filesystem Loader. In the Shared Filesystem Loader, management data files received from Management Agents are stored temporarily on a common shared location called the shared receive directory. All Management Services are configured to use the same storage location for the shared receive directory. The Management Services coordinate internally and distribute among themselves the workload of uploading files into the Management Repository. Should a Management Service go down, its workload is taken up by surviving Management Services. You must choose a shared receive directory that is accessible by all the Management Services using redundant file storage.

During the first Management Service installation, the shared receive directory can be configured out of the box by passing SHARED_RECEIVE_DIRECTORY_ LOCATION=<shared recv directory> option to runInstaller (setup.exe on Windows). Oracle recommends that this location be outside the Instance Home and Middleware Home locations.

If not configured during installation, the Shared Filesystem Loader can also be configured post-installation by running the following emctl command on every Management Service:

```
emctl config oms loader -shared yes -dir <shared recv directory>
```

> **Note:** If shared filesystem Loader is configured on the first Management Service, any additional management service that is installed later will inherit the shared filesystem loader configuration. Therefore, ensure that the shared recv directory is available on the additional Management Service prior to running install.

Consider the following while configuring Shared Filesystem Loader on Windows.

- On Windows platforms, the Enterprise Manager install may configure the Management Service to run as a service using 'LocalSystem' account. This is a local account and will typically not have access to the network drive for the shared filesystem that requires domain authentication. To resolve this issue, configure the Management Service to run as a domain user as follows:
  1. Go to the Control Panel and then open the Services panel.
  2. Double click the appropriate service (Oracleoms11gProcessManager).
  3. Change the 'Log on as' user from the 'Local System Account' to **This Account**.
  4. Specify the domain user that has access to shared network drive.
  5. Click **OK**.
- Do not use local drive letters for mapped network shares while configuring the shared filesystem on Windows. Use UNC locations instead.

  ```
  emctl config oms loader -shared yes -dir \\\\<host>\\<share-name>\\<recv-dir>
  for example
  emctl config oms loader -shared yes -dir \\\\hostA\\vol1\\recv
  ```

  Note the use of double backslashes while specifying the directory location.

Note: User equivalence should be set up properly across OMS so that files created by one OMS on the loader directory are modifiable by other OMS.

32.4.4 Configuring Software Library

Since software library location has to be accessed by all Management Services, considerations similar to shared fileystem loader directory apply here too. The configuration of software library is not covered during install. It needs to be configured post-install using the Enterprise Manager Console:

1. On the Enterprise Manager home page, click the **Deployments** tab.
2. Click the **Provisioning** subtab.
3. On the Provisioning page, click the **Administration** subtab.
4. In the Software Library Configuration section, click **Add** to set the Software Library Directory Location to a shared storage that can be accessed by any host running the Management Service.

32.4.5 Configuring a Load Balancer

This section describes the guidelines for setting up a Server Load Balancer (SLB) to balance the agent and browser traffic to multiple Management Services. This is a two step process:

1. Configure the SLB
2. Make needed changes on the Management Services

32.4.5.1 SLB Setup

Use the following table as reference for setting up the SLB with Cloud Control Management Services.

Table 32–1 Management Service Ports

Cloud Control Service	TCP Port	Monitor Name	Persistence	Pool Name	Load Balancing	Virtual Server Name	Virtual Server Port
Secure Upload	1159	mon_gcsu1159	None	pool_ gcsu1159	Round Robin	vs_ gcsu1159	1159
Agent Registration	4889	mon_gcar4889	Active Cookie Insert	pool_ gcar4889	Round Robin	vs_ gcar4889	4889
Secure Console	7799	mon_gcsc7799	Source IP	pool_ gcsc7799	Round Robin	vs_ gcsc443	443
Unsecure Console (optional)	7788	mon_gcuc7788	Source IP	pool_ gcuc7788	Round Robin	vs_gcuc80	80

Use the administration tools that are packaged with your SLB. A sample configuration follows. This example assumes that you have two Management Services running on host A and host B using the default ports as listed in Table 33–1.

1. Create Pools

A *pool* is a set of servers grouped together to receive traffic on a specific TCP port using a load balancing method. Each pool can have its own unique characteristic for a persistence definition and the load-balancing algorithm used.

Table 32–2 Pools

Pool Name	Usage	Members	Persistence	Load Balancing
pool_gcsu1159	Secure upload	HostA:1159 HostB:1159	None	Round Robin
pool_gcar4889	Agent registration	HostA:4889 HostB:4889	Active cookie insert; expiration 60 minutes	Round Robin
pool_gcsc7799	Secured console access	HostA:7799 HostB:7799	Source IP; expiration 60 minutes	Round Robin
pool_gcuc7788 (optional)	Unsecured console access	HostA:7788 HostB:7788	Source IP; expiration 60 minutes	Round Robin

2. Create Virtual Servers

A *virtual server*, with its virtual IP Address and port number, is the client addressable hostname or IP address through which members of a load balancing pool are made available to a client. After a virtual server receives a request, it directs the request to a member of the pool based on a chosen load balancing method.

Table 32–3 Virtual Servers

Virtual Server Name	Usage	Virtual Server Port	Pool
vs_gcsu1159	Secure upload	1159	pool_gcsu1159
vs_gcar4889	Agent registration	4889	pool_gcar4889
vs_gcsc443	Secure console access	443	pool_gcsc7799
vs_gcuc80 (optional)	Unsecure console access	80	pool_gcuc7788

3. Create Monitors

Monitors are used to verify the operational state of pool members. Monitors verify connections and services on nodes that are members of load-balancing pools. A monitor is designed to check the status of a service on an ongoing basis, at a set interval. If the service being checked does not respond within a specified timeout period, the load balancer automatically takes it out of the pool and will choose the other members of the pool. When the node or service becomes available again, the monitor detects this and the member is automatically accessible to the pool and able to handle traffic.

Table 32–4 Monitors

Monitor Name	Configuration	Associate With
mon_gcsu1159	Type: https Interval: 60 Timeout: 181 Send String: GET /em/upload Receive String: Http Receiver Servlet active!	HostA:1159 HostB:1159
mon_gcar4889	Type: http Interval: 60 Timeout: 181 Send String: GET /em/genwallet Receive String: GenWallet Servlet activated	HostA:4889 HostB:4889
mon_gcsc7799	Type: https Interval: 5 Timeout: 16 Send String: GET /em/console/home HTTP/1.0\n Receive String: /em/console/logon/logon;jsessionid=	HostA:7799 HostB:7788
mon_gcuc7788 (optional)	Type: https Interval: 5 Timeout: 16 Send String: GET /em/console/home HTTP/1.0\n Receive String: /em/console/logon/logon;jsessionid=	HostA:7788 HostB:7788

Note: If you have SSO configured, use the following alternate definitions for the mon_gcsc7799 and mon_gcuc7788 monitors.

Table 32–5 Monitors for SSO Configuration

Monitor Name	Configuration	Associate With
mon_gcsc7799	Type: https Interval: 5 Timeout: 16 Send String: GET /em/genwallet Receive String: GenWallet Servlet activated	HostA:7799 HostB:7788
mon_gcuc7788 (optional)	Type: https Interval: 5 Timeout: 16 Send String: GET /em/genwallet Receive String: GenWallet Servlet activated	HostA:7788 HostB:7788

Note: F5 SLB monitors expect the "Receive String" within the first 5120 characters of a response. For SSO environments, the "Receive String" may be returned at some point beyond the 5120 limit. The monitor will not function in this situation.

32.4.5.2 Enterprise Manager Side Setup

Perform the following steps:

1. Resecure the Management Service

   By default, the service name on the Management Service-side certificate uses the name of the Management Service host. Management Agents do not accept this certificate when they communicate with the Management Service through a load balancer. You must run the following command to regenerate the certificate on the first Management Service:

```
emctl secure
  -oms -sysman_pwd <sysman_pwd>
  -reg_pwd <agent_reg_password>
  -host slb.example.com
  -secure_port 1159
  -slb_port 1159
  -slb_console_port 443
  [-lock]  [-lock_console]
```

2. Resecure all Management Agents

   Management Agents that were installed prior to SLB setup, including the Management Agent that comes with the Management Service install, would be uploading directly to the Management Service. These Management Agents will not be able to upload after SLB is setup. Resecure these Management Agents to upload to the SLB by running the following command on each Management Agent:

   ```
   emctl secure agent -emdWalletSrcUrl https://slb.example.com:<upload port>/em
   ```

32.4.6 Reconfiguring the Oracle Management Agent

The following sections describe reconfiguration and tuning changes you can make to the Management Agent after you have installed Enterprise Manager. Refer to the following sections for more information:

- Configuring the Management Agent to Use a New Management Service
- Changing the Management Agent Port
- Controlling the Amount of Disk Space Used by the Management Agent
- About the Management Agent Watchdog Process
- Setting the Management Agent Time Zone
- Adding Trust Points to the Management Agent Configuration

32.4.6.1 Configuring the Management Agent to Use a New Management Service

When you install the Management Agent on a managed host, you associate the Management Agent with a particular Management Service. The Management Agent uses the Management Service URL address and port to identify and communicate with the Management Service.

After you install the Management Agent, you can later reconfigure the Management Agent so it is associated with a different Management Service. Reconfiguring the Management Agent requires no changes to the Management Service. The reconfigured Management Agent will begin communicating with the new Management Service after the Management Agent is restarted.

If you are associating the Management Agent with a new Management Service that is locked in secure mode, then first associate the Management Agent with the new Management Service and then secure the Management Agent.

To associate the Management Agent with a new Management Service after you have installed the Management Agent:

1. Stop the Management Agent.
2. Locate the `emd.properties` file in the Management Agent home directory:

   ```
   AGENT_HOME/sysman/config/emd.properties
   ```

3. Use a text editor to open the file and locate the REPOSITORY_URL property.
4. Modify the value for the `REPOSITORY_URL` property so it references the new Management Service. For example:

   ```
   REPOSITORY_URL=http://mgmthost2.acme.com:4889/em/upload
   ```

5. Modify the value for the `emdWalletSrcUrl` property so it references the new Management Service. For example, if the new Management Service is on a host called `mgmthost2.acme.com`, modify the property as follows:

   ```
   emdWalletSrcUrl=http://mgmthost2.acme.com:4889/em/wallets/emd
   ```

6. Save your changes and close the `emd.properties` file.
7. To ensure that the Management Agent is no longer holding any specific data or settings from the previous Management Service, delete all the files in the following directories:

   ```
   AGENT_HOME/sysman/emd/upload/
   AGENT_HOME/sysman/emd/state/
   AGENT_HOME/sysman/emd/collection/*
   AGENT_HOME/sysman/emd/lastupld.xml
   AGENT_HOME/sysman/emd/agntstmp.txt
   AGENT_HOME/sysman/emd/blackouts.xml
   AGENT_HOME/sysman/emd/protocol.ini
   ```

   Note that this action removes all user-defined metrics (UDM)s and custom changes to metric and policy collections.

   Note: You can use the `emctl clearstate agent` command to delete the files in the state directory.

8. Restart the Management Agent.

32.4.6.2 Securing the Management Agent

To secure the Management Agent of the new Management Service, use the following command:

```
emctl secure agent [registration password] [-emdWalletSrcUrl <url>]
```

32.4.6.3 Changing the Management Agent Port

The Management Agent uses a predefined port number to receive requests from the Management Service. This port number is defined by default when you install the Management Agent on a managed host. If you later need to modify this port, you can use the following procedure. You might need to modify this port number if you have existing software that uses the default Management Agent port.

To change the Management Agent port:

1. Stop the Management Agent.
2. Locate the `emd.properties` file in the Management Agent home directory:

   ```
   AGENT_HOME/sysman/config/emd.properties
   ```

3. Use a text editor to open the file and locate the EMD_URL property.

   For example:

```
EMD_URL=http://managed_host1.acme.com:1813/emd/main
```

4. Modify the port number in the EMD_URL property so the Management Agent uses a new unused port on the managed host.

   For example:

   ```
   EMD_URL=http://managed_host1.acme.com:1913/emd/main
   ```

   You can also use the `netstat` command to check for the unused port:

   On Windows:

   ```
   netstat -an | findstr <new port number>
   ```

   On UNIX:

   ```
   netstat -an | grep <new port number>
   ```

5. Start the Management Agent.

> **Note:** After the changed URL is processed, the old Management Agent should not have any targets. Oracle recommends that you remove the old Management Agent from the Management Service to ensure that there are no unwanted targets in the Cloud Control console .

32.4.6.4 Controlling the Amount of Disk Space Used by the Management Agent

Oracle designed the Management Agent to work within a set of disk space limits. These limits prevent the Management Agent from using too much disk space and causing performance or resource issues on your enterprise systems. However, if disk space becomes an issue, you can adjust the default settings that are used to control the amount of disk space used by the Management Agent.

As the Management Agent on a particular host gathers management data about the targets on the host, it saves the collected data on the local disk until the data is uploaded to the Management Repository. The Management Agent saves this collected data and metadata in the following directory:

```
AGENT_HOME/sysman/emd/upload
```

By default, the Management Agent will save up to 50MB of collected data in the upload directory. If the amount of collected data exceeds 50MB, data collection is stopped temporarily until the data is uploaded to the repository and more disk space becomes available.

To verify how much space is available:

- Use the `emctl status agent` command. For example:

  ```
  Available disk space on upload filesystem    : 1.18%
  Collection Status                            : Disabled by Upload Manager
  Last successful heartbeat to OMS             : 2007-07-31 11:22:07
  ```

- Investigate the <AGENT_HOME>/sysman/log/emagent.trc file. The file will have errors such as :

  ```
  24.995519 MB Data. 34.06% of disk used. Disabling collections.
  2006-10-19 10:41:23 Thread-19 WARN collector: Disable collector
  ```

In addition, the Management Agent checks to be sure that the percentage of disk space currently in use on the local disk does not exceed 98 percent. If this value is exceeded, the Management Agent stops collecting data and stops saving information to the Management Agent log and trace files.

You can modify these default settings as follows:

1. Stop the Management Agent.
2. Locate the `emd.properties` file in the Management Agent home directory:

   `AGENT_HOME/sysman/config/emd.properties`

3. Use a text editor to open the file and modify the entries shown in Table 34–1.
4. Save your changes and exit the file.
5. Restart the Management Agent.

Table 32–6 Properties for Controlling the Disk Space Used by the Management Agent

Property	Explanation
UploadMaxBytesXML	Use this property in the emd.properties file to specify the maximum number of megabytes (MB) used by the collected data in the Management Agent upload directory. When this limit is exceeded, the Management Agent will stop collecting additional management data until the next upload to the Management Repository reduces the amount of collected data in the upload directory.
UploadMaxDiskUsedPct	Use this property in the `emd.properties` file to specify the maximum percentage of disk space that can be in use on the local disk before the Management Agent temporarily stops collecting additional data and stops saving information to the Management Agent log and trace files. The Management Agent will begin collecting data again when the percentage of disk space in use falls to less than the percentage specified in the `UploadMaxDiskUsedPctFloor` property in the `emd.properties` file.
UploadMaxDiskUsedPctFloor	Use this property in the `emd.properties` file to specify the amount (%) of disk space that can be used on the EMD filesystem before the following items are re-enabled after being disabled: ▪ Collection of data (upload manager) ▪ Logging and tracing ▪ Diagnosability tracing does minimal dumps above this limit
UploadMaxNumberXML	Use this property in the `emd.properties` files to specify the maximum number of files the upload manager will support in the upload directory. When this limit is exceeded, the Management Agent will temporarily disable collections, logging, and tracing.

32.4.6.5 About the Management Agent Watchdog Process

The Management Agent is the Enterprise Manager component that gathers the data you need to manage your enterprise efficiently. As a result, Enterprise Manager includes software that keeps track of the Management Agent processes and makes sure the Management Agent stays running.

For example, if the Management Agent quits unexpectedly, this self-monitoring process—referred to as the watchdog process—will restart the Management Agent automatically.

In most situations, the watchdog process works in the background and requires no configuration or maintenance. The watchdog process is controlled by the `emwd.pl` script located in the following directory of the Management Agent home directory:

`AGENT_HOME/bin`

You can identify the watchdog process by using the following commands:

```
$PROMPT> ps -ef | grep emwd
```

32.4.6.6 Setting the Management Agent Time Zone

In today's global economy, it is not uncommon for the systems you manage to reside in multiple locations throughout the world. For example, if your company headquarters are in New Hampshire, USA, you may need to manage systems that reside in California, Canada, and in Europe.

As Enterprise Manager collects monitoring data from Management Agents running on these remote systems, it is important that the data is correlated accurately. A software failure on a machine in Ontario, Canada might be the cause of a performance problem on a machine in Hoboken, New Jersey.

To correlate this data, it is important that Enterprise Manager obtains the correct time zone for each Management Agent that you install. The following sections describe how the Management Agent obtains the time zone and how to correct the problem if the time zone for a Management Agent is incorrect:

- Understanding How the Management Agent Obtains Time Zone Information
- Resetting the Time Zone of the Management Agent Due to Inconsistency of Time Zones
- Troubleshooting Management Agent Time Zone Problems

32.4.6.6.1 Understanding How the Management Agent Obtains Time Zone Information When you install the Management Agent, the software attempts to obtain the current time zone of the host computer. If successful, the installation procedure updates the `agentTZRegion` property setting in the following configuration file:

```
AGENT_HOME/sysman/config/emd.properties
```

The `agentTZRegion` property can be set to any of the values listed in the following file, which is installed in the Management Agent home directory:

```
AGENT_HOME/sysman/admin/supportedtzs.lst
```

To reconfigure a different time zone, perform the following steps. These steps assume that the original time zone used was EST and the target time zone is CST.

1. Set the environment correctly.
   - On Windows XP

     - From the Start Menu, access the Control Panel. Click **Date and Time**, then click the **Time Zone** tab.

     - Select GMT-06:00 Central Time (US & Canada) from the list.

     - Click **OK**.

     - Open a command line screen (cmd.exe).

     - Set the following environment variables:

     ```
     SET TZ=CST
     SET ORACLE_HOME=< your oracle home directory >
     SET PATH=%ORACLE_HOME%\bin;%PATH%
     ```

   - On UNIX

- Login to your UNIX server as the Oracle user.

- Set the following environment variables:

```
$ export TZ=CST
$ export ORACLE_HOME=< your oracle home directory >
$ export PATH=%ORACLE_HOME%\bin;%PATH%
```

2. Execute the following commands:

   - On Windows

     ```
     %ORACLE_HOME%\bin\emctl config agent getTZ
     %ORACLE_HOME%\bin\emctl stop iasconsole
     %ORACLE_HOME%\bin\emctl resetTZ agent
     ```

   - On UNIX

     ```
     $ORACLE_HOME/bin/emctl config agent getTZ
     $ORACLE_HOME/bin/emctl stop iasconsole
     $ORACLE_HOME/bin/emctl resetTZ agent
     ```

3. Delete all the files under the following directory:

   - On Windows

     ```
     %ORACLE_HOME%\sysman\logs
     ```

   - On UNIX

     ```
     $ORACLE_HOME$/sysman/logs
     ```

4. Start the console again.

   - On Windows

     ```
     %ORACLE_HOME%\bin\emctl start iasconsole
     %ORACLE_HOME%\bin\emctl status iasconsole
     %ORACLE_HOME%\bin\emctl status agent
     %ORACLE_HOME%\bin\emctl config agent getTZ
     ```

   - On UNIX

     ```
     $ORACLE_HOME\bin\emctl start iasconsole
     $ORACLE_HOME\bin\emctl status iasconsole
     $ORACLE_HOME\bin\emctl status agent
     $ORACLE_HOME\bin\emctl config agent getTZ
     ```

5. Check the timestamp in the log file.

6. Check the em.properties file. The agentTZRegion parameter should now look like this:

   - On Windows

     ```
     %ORACLE_HOME%\sysman\config\em.properties
     agentTZRegion=America/Chicago
     ```

   - On UNIX

     ```
     $ORACLE_HOME/sysman/config/em.properties
     agentTZRegion=America/Chicago
     ```

32.4.6.6.2 Resetting the Time Zone of the Management Agent Due to Inconsistency of Time Zones You need to reset the time zone of the Management Agent when *both* of the following situations are true:

- The Management Agent has been running with a particular time zone
- Subsequently a change occurs to the time zone of the host where the Management Agent is running

To propagate the time zone change to the `emd.properties` file, perform the following:

1. Execute the following script:

   ```
   ORACLE_HOME/bin/emctl resetTZ agent
   ```

   This script updates `ORACLE_HOME/sysman/config/emd.properties` so that the value of `agentTZRegion` matches that of the current time zone setting of the machine.

   Note: The location of the `emd.properties` file depends on the Control Console being used:

   - For the Database Control Console, the location is usually: ORACLE_HOME/<host>_<sid>/sysman/config
   - For the Application Server Control Console, the location is: ORACLE_HOME/sysman/config
   - For the Cloud Control Management Agent, the location is ORACLE_HOME/sysman/config
   - For the Real Application Cluster central Management Agent, the location is usually: ORACLE_HOME/<host>/sysman/config

2. In addition, this command prompts you to run a script against the Enterprise Manager Repository. You must log in to the database as the Enterprise Manager repository user and run the script `mgmt_target.set_agent_tzrgn`. An example follows:

   ```
   SQL> exec mgmt_target.set_agent_tzrgn('em.oracle.com:1830','PST8PDT');
   SQL> commit;
   SQL> exit
   ```

   `em.oracle.com:1830` represents the name of the emd target.

32.4.6.6.3 Troubleshooting Management Agent Time Zone Problems Sometimes, during the Management Agent installation, the time zone detected by the Management Agent configuration tool is not recognized by the Management Agent. In other words, the time zone obtained by the configuration tool is not listed in the Management Agent list of supported time zones.

This problem prevents the Management Agent from starting and results in an error similar to the following:

```
Could not determine agent time zone. Please refer to the file:
ORACLE_HOME/sysman/admin/supportedtzs.lst and pick a timezone region with a
standard offset of +5:0 from GMT and update the property 'agentTZRegion' in the
file: ORACLE_HOME/sysman/config/emd.properties
```

This error appears in one of the log files shown in Table 34–2, depending upon which Enterprise Manager product you are using.

Table 32–7 Location of Time Zone Error in the Enterprise Manager Log Files

If you are using...	Look for the Time Zone Error in This File...
Grid Control console	`emagent.nohup`
Application Server Control Console	`em.nohup`
Database Control Console	`emdb.nohup`

To configure the Management Agent to use a valid time zone:

1. Enter the following command in the Management Agent home directory to identify the time zone currently being used by the host computer:

   ```
   AGENT_HOME/bin/emctl config agent getTZ
   ```

2. Note the time zone that is returned by the `emctl config agent getTZ` command.

   This is the time zone of the host computer.

3. Use a text editor to open the following file in the Management Agent home directory:

   ```
   AGENT_HOME/sysman/admin/supportedtzs.lst
   ```

   This file contains a list of all the time zones supported by the Management Agent.

4. Browse the contents of the `supportedtzs.lst` file and note the supported time zone closest to the time zone of the host computer.

5. Use a text editor to open the following Management Agent configuration file:

   ```
   AGENT_HOME/sysman/config/emd.properties
   ```

6. Locate the following property near the end of the `emd.properties` file:

   ```
   agentTZRegion=
   ```

7. Set the value of this property to the time zone you identified as closest to the host time zone in the `supportedtzs.lst` file.

   For example:

   ```
   agentTZRegion=Europe/Warsaw
   ```

8. Save your changes and close the `emd.properties` file.

   You should now be able to start the Management Agent without generating the error in the log file.

32.4.6.7 Adding Trust Points to the Management Agent Configuration

Perform these steps to add the relevant security certificate:

1. Obtain the certificate, which is in Base64encoded X.509 (.CER) format, in the `b64SiteCertificate.txt` file. (The file name may be different in your configuration.) An example of the contents of the file is as follows:

   ```
   ------BEGIN CERTIFICATE--------------
   MIIDBzCCAnCgAw...
   ...... base 64 certificate content .....
   ```

```
------END CERTIFICATE----------------
```

2. In the Oracle Home of the Management Agent monitoring the wallet, run the following command to add the certificate to the Management Agent:

```
${ORACLE_HOME}/bin/mkwallet -i welcome
${ORACLE_HOME}/sysman/config/monwallet
${ORACLE_HOME}/sysman/config/b64SiteCertificate.txt NZDST_CLEAR_PTP
```

32.5 Configuring Standby Management Service

Consider the following before installing the standby Management Services.

- Oracle recommends that this activity be done during a lean period or during a planned maintenance window. When new Management Services are installed on the standby site, they are initially configured to connect to the Management Repository database on the primary site. Some workload will be taken up by the new Management Service. This could result in temporary loss in performance if the standby site Management Services are located far away from the primary site Management Repository database. However there would be no data loss and the performance would recover once the standby Management Services are shutdown post configuration.
- The shared storage used for the shared filesystem loader and software library must be made available on the standby site using the same paths as the primary site.

32.5.1 Installing the First Standby Management Service

Install the first standby Management Service using the following steps:

1. Copy the emkey to the Management Repository by running the following command on the first Management Service on the primary site:

   ```
   emctl config emkey -copy_to_repos
   ```

2. Perform a software-only install by running the installer with the following arguments:

   ```
   runInstaller -noconfig -validationaswarnings
   ```

3. Apply one-off patches

   ```
   <OMS Oracle Home>/install/oneoffs/apply_NewOneoffs.pl <OMS
   Oracle Home> OC9321514,9207217
   ```

4. Configure the Management Service by running `omsca` in standby mode. Choose a different domain name for the standby. For example, if the primary WebLogic domain is GCDomain, choose GCDomainStby.

   ```
   omsca standby -EM_DOMAIN_NAME GCDomainStby -nostart
   ```

   When prompted for Management Repository details, provide the Primary database details.

5. Configure the virtualization add on by running the following command:

   ```
   addonca -oui -omsonly -name vt -install gc
   ```

6. Configure the Management Agent that comes with the Management Service install by running:

```
<Agent Oracle Home>/bin/agentca -f
```

7. Export the configuration from the primary Management Service using:

   ```
   emctl exportconfig oms -dir <location for the export file>
   ```

8. Copy over the exported file to the standby Management Service
9. Import the exported configuration on the standby Management Service using:

   ```
   emctl importconfig oms -file <full path of the export file>
   ```

32.5.2 Installing Additional Standby Management Services

Install addional standby Management Services as follows:

Specify the primary database details and the standby administration server details on the installer screens. Post installation, do the following steps to complete the HA configuration.

1. Do a software only install by running the installer with following arguments:

   ```
   runInstaller -noconfig -validationaswarnings
   ```

2. Apply one-off patches

   ```
   <OMS Oracle Home>/install/oneoffs/apply_NewOneoffs.pl <OMS
   Oracle Home> OC9321514,9207217
   ```

3. Configure the Management Service by running `omsca`. When prompted for Management Repository details, provide the Primary database details. When prompted for Administration Server details, provide the standby administration server details.

   ```
   omsca add –nostart
   ```

4. Configure the virtualization add on by running the following command

   ```
   addonca -oui -omsonly -install gc
   ```

5. Configure the Management Agent that comes with the Management Service install by running:

   ```
   <Agent Oracle Home>/bin/agentca -f
   ```

6. Export the configuration from the primary Management Service using:

   ```
   emctl exportconfig oms -dir <location for the export file>
   ```

7. Copy over the exported file to the standby Management Service
8. Import the exported configuration on the standby Management Service using:

   ```
   emctl importconfig oms -file <full path of the export file>
   ```

32.5.3 Validating Your Installation and Complete the Setup

Validate your installation and complete the setup as follows:

1. Update the standby SLB configuration by adding the standby Management Services to the different pools on the SLB. Setup monitors for the new Management Service.
2. Make the standby Management Services point to the standby Management Repository database by running the following command on the first standby Management Service:

```
emctl config oms -store_repos_details -repos_conndesc
<connect descriptor of standby database> -repos_user sysman
-no_check_db
```

3. Shut down all Management Services by running the following command on each Management Service:

```
emctl stop oms -all
```

32.6 How to Configure Cloud Control OMS in Active/Passive Environment for High Availability Failover Using Virtual Host Names

This section provides a general reference for Cloud Control administrators who want to configure Enterprise Manager Cloud Control in Cold Failover Cluster (CFC) environments.

32.6.1 Overview and Requirements

The following conditions must be met for Cloud Control to fail over to a different host:

- The installation must be done using a Virtual Host Name and an associated unique IP address.
- Install on a shared disk/volume which holds the binaries, the configuration and the runtime data (including the recv directory).
- Configuration data and metadata must also failover to the surviving node.
- Inventory location must failover to the surviving node.
- Software owner and time zone parameters must be the same on all cluster member nodes that will host this Oracle Management Service (OMS).

32.6.2 Installation and Configuration

To override the physical host name of the cluster member with a virtual host name, software must be installed using the parameter ORACLE_HOSTNAME. For inventory pointer, the software must be installed using the command line parameter *-invPtrLoc* to point to the shared inventory location file, which includes the path to the shared inventory location.

If you are using an NFS mounted volume for the installation, please ensure that you specify rsize and wsize in your mount command to prevent running into I/O issues.

For example:

```
oms.acme.com:/u01/app/share1 /u01/app/share1 nfs
rw,bg,rsize=32768,wsize=32768,hard,nointr,tcp,noac,vers=3,timeo=
600 0 0
```

Note: Any reference to shared failover volumes could also be true for non-shared failover volumes which can be mounted on active hosts after failover.

32.6.3 Setting Up the Virtual Host Name/Virtual IP Address

You can set up the virtual host name and virtual IP address by either allowing the clusterware to set it up, or manually setting it up yourself before installation and startup of Oracle services. The virtual host name must be static and resolvable

consistently on the network. All nodes participating in the setup must resolve the virtual IP address to the same host name. Standard TCP tools such as nslookup and traceroute can be used to verify the host name. Validate using the following commands:

`nslookup` <*virtual hostname*>

This command returns the virtual IP address and full qualified host name.

`nslookup` <*virtual IP*>

This command returns the virtual IP address and fully qualified host name.

Be sure to try these commands on every node of the cluster and verify that the correct information is returned.

32.6.4 Setting Up Shared Storage

Storage can be managed by the clusterware that is in use or you can use any shared file system (FS) volume as long as it is not an unsupported type, such as OCFS V1. The most common shared file system is NFS.

Note: If the OHS directory is on a shared storage, the LockFile directive in the httpd.conf file should be modified to point to a local disk, otherwise there is a potential for locking issues.

32.6.5 Setting Up the Environment

Some operating system versions require specific operating system patches be applied prior to installing 11*g*R1. The user installing and using the 11*g*R1 software must also have sufficient kernel resources available. Refer to the operating system's installation guide for more details.

Before you launch the installer, certain environment variables need to be verified. Each of these variables must be identically set for the account installing the software on ALL machines participating in the cluster:

- OS variable TZ

  Time zone setting. You should unset this variable prior to installation

- PERL variables

  Variables such as PERL5LIB should also be unset to avoid association to the incorrect set of PERL libraries

32.6.6 Synchronizing Operating System IDs

The user and group of the software owner should be defined identically on all nodes of the cluster. This can be verified using the 'id' command:

```
$ id -a
uid=550(oracle) gid=50(oinstall) groups=501(dba)
```

32.6.7 Setting Up Shared Inventory

Use the following steps to set up shared inventory:

1. Create your new ORACLE_HOME directory.
2. Create the Oracle Inventory directory under the new oracle home:

```
$ cd <shared oracle home>

$ mkdir oraInventory
```

3. Create the oraInst.loc file. This file contains the Inventory directory path information needed by the Universal Installer.

   a. `vi oraInst.loc`

   b. Enter the path information to the Oracle Inventory directory and specify the group of the software owner as the oinstall user. For example:

   ```
   inventory_loc=/app/oracle/product/11.1/oraInventory

   inst_group=oinstall
   ```

32.6.8 Installing the Software

Refer to the following steps when installing the software:

1. Create the shared disk location on both the nodes for the software binaries
2. Install WebLogic Server. For information on installing WebLogic Server, refer to *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.
3. Point to the inventory location file oraInst.loc (under the ORACLE_BASE in this case), as well as specifying the host name of the virtual group. For example:

```
$ export ORACLE_HOSTNAME=lxdb.acme.com
$ runInstaller -invPtrloc /app/oracle/share1/oraInst.loc
ORACLE_HOSTNAME=lxdb.acme.com -debug
```

1. Install Oracle Management Services on cluster member *Host1* using the option, "EM install using the existing DB"
2. Continue the remainder of the installation normally.
3. Once completed, copy the files *oraInst.loc* and *oratab* to */etc*. Also copy */opt/oracle* to all cluster member hosts (*Host2*, *Host3*, and so on).

32.6.8.1 Windows Specific Configuration Steps

On Windows environments, an additional step is required to copy over service and keys required by the Oracle software. Note that these steps are required if your clustering software does not provide a shared windows registry.

1. Using regedit on the first host, export each Oracle service from under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services.
2. Using regedit on the first host, export HKEY_LOCAL_MACHINE\SOFTWARE\ORACLE.
3. Use regedit to import the files created in step 1 and 2 to the failover host.

32.6.9 Starting Up Services

Ensure that you start your services in the proper order. Use the order listed below:

1. Establish IP address on the active node
2. Start the TNS listener (if it is part of the same failover group)
3. Start the database (if it is part of the same failover group)
4. Start Cloud Control using `emctl start oms`

5. Test functionality

In case of failover, refer to the following steps:

1. Establish IP on failover box
2. Start TNS listener using the command `lsnrctl start` if it is part of the same failover group
3. Start the database using the command `dbstart` if it is part of the same failover group
4. Start Cloud Control using the command `emctl start oms`
5. Test the functionality

32.6.10 Summary

The OMS mid-tier component of Cloud Control can now be deployed in a CFC environments that utilize a floating host name..

32.7 Configuring the Management Repository

Before installing Enterprise Manager, you should prepare the database, which will be used for setting up Management Repository. Install the database using Database Configuration Assistant (DBCA) to make sure that you inherit all Oracle install best practices.

- Configure Database
  - For both high availability and scalability, you should configure the Management Repository in the latest certified database version, with the RAC option enabled. Check for the latest version of database certified for Enterprise Manager from the Certify tab on the My Oracle Support website.
  - Choose Automatic Storage Management (ASM) as the underlying storage technology.
  - When the database installation is complete:

    Go to $ORACLE_HOME/rbdms/admin directory of the database home and execute the 'dbmspool.sql'

    This installs the DBMS_SHARED_POOL package, which will help in improving throughput of the Management Repository.
- Install Enterprise Manager

  While installing Enterprise Manager using Oracle Universal Installer (OUI), you will be presented with the option for configuring the Management Repository using an existing database.

32.7.1 Post Management Service - Install Management Repository Configuration

There are some parameters that should be configured during the Management Repository database install (as previously mentioned) and some parameters that should be set after the Management Service has been installed.

Start by installing Management Agents on each Management Repository RAC node. Once the Management Agents are installed and the Management Repository database is discovered as a target, the Enterprise Manager console can be used to configure these best practices in the Management Repository.

These best practices fall in the area of:

- Configuring Storage
- Configuring Oracle Database 11*g* with RAC for High Availability and Fast Recoverability
  - Enable ARCHIVELOG Mode
  - Enable Block Checksums
  - Configure the Size of Redo Log Files and Groups Appropriately
  - Use a Flash Recovery Area
  - Enable Flashback Database
  - Use Fast-Start Fault Recovery to Control Instance Recovery Time
  - Enable Database Block Checking
  - Set DISK_ASYNCH_IO

  The details of these settings are available in *Oracle Database High Availability Best Practices*.

  Use the MAA Advisor for additional high availability recommendations that should be applied to the Management Repository. To access the MAA Advisor:

  1. On the Database Target Home page, locate the High Availability section.
  2. Click **Details** next to the Console item.
  3. In the Availability Summary section of the High Availability Console page, click **Details** located next to the MAA Advisor item.

32.8 Converting the Enterprise Manager Repository from Single Instance to RAC

The scenario used in the following conversion process provisions a two-node cluster, creates an additional physical standby database that uses Oracle ASM or some type of shared storage on the new server, and converts the standby database to an Oracle RAC database. Finally, a switchover operation is performed to move the new Oracle RAC standby database to the primary database role.

> **Note:** The MAA recommended best practice is to add a standby database during this process even if there is already a standby database in place. This is to maintain the same level of protection from disasters during all phases of this process, with no compromise to the existing availability solution. Thus, if the starting configuration contained a standby database, the ending configuration would contain an Oracle RAC primary database and two single-instance standby databases. After the switchover is performed, either of the single-instance physical standby databases can be removed with no negative impact.

Conversion Tasks

Task 1: Provision Oracle Clusterware, ASM

Task 2: Create a Single Instance Physical Standby Database on the New Cluster , and Oracle RAC

Task 3: Prepare the Environment Prior to Conversion to Oracle RAC

Task 4: Convert a Physical Standby Database to Oracle RAC

Task 5: Perform a switchover and Enable Additional Threads

32.8.1 Task 1: Provision Oracle Clusterware, Oracle ASM, and Oracle RAC

This section describes using Deployment Procedures to provision a two-node cluster running Oracle Clusterware, Oracle ASM, and Oracle RAC. If you have already provisioned Oracle RAC software onto the new server, skip to Task 2 to create the physical standby database.

The process described in this section uses gold images of Oracle Clusterware, Oracle ASM, and Oracle RAC that were pre-created and stored in the Software Library.

Also, see "Provisioning Oracle RAC Using Gold Image" in the Oracle Enterprise Manager Administrator's Guide for Software and Server Provisioning and Patching for additional prerequisite information.

Perform the following steps:

1. In Cloud Control, click the Deployments tab.
2. On the Deployments page, in the Deployment Procedure Manager section, click RAC Provisioning Procedures.
3. On the Deployment Procedure Manager page, in the Procedure subtab, select to run the Deployment Procedure for Oracle Clusterware and Oracle RAC that you created previously in the "Install Software and Upload Components" section.

   Note: Ensure the customized copy of the Deployment Procedure uses the appropriate commands (for example, sudo) for step execution.

   Click Schedule Deployment. Enterprise Manager Cloud Control displays the Select Source page of the Deployment Procedure.

4. On the Select Source page, do the following:

a. In the Select Source section, select Select from Software Library.

b. In the Source for Clusterware section, click the torch icon and select the Component that has the gold image of Oracle Clusterware.

c. In the Source for RAC section, click the torch icon and select the Component that has the gold image of Oracle RAC.

d. In the Source for ASM section, choose whether or not you want to deploy Oracle ASM. The MAA team chose to deploy Oracle ASM using the same component for the ASM Oracle home as was used for the Oracle RAC Oracle home.

Note: Ensure that you select only Components that are in "Active" status. Once you select the component name, the application automatically displays the component location.

5. On the Select Hosts page, perform the following:

a. In the Hosts to Include in Cluster section, click Add and select the target hosts that should form the cluster.

By default, the Show Suitable Hosts option is selected and the table lists only those hosts that are best suited for provisioning. If you do not find the host you want to add, then select Show All Hosts to view a complete list of hosts.

By default, the procedure automatically pre-fills the Private Host Name and Virtual Host Name fields with values. Ensure the correct Virtual Host Name for the VIP is used. If necessary, edit them to specify values that match your environment. Optionally, you can also specify IP addresses. For example:

b. In the Hosts to Include in Cluster section, configure the private and public network interfaces by clicking Select Interfaces. By default, the interfaces that have the same name and subnet for the selected target hosts are displayed. You can also choose Show Interface List to view all the interfaces for the selected target hosts. Select one of the existing interfaces or specify a completely new one. Click OK.

c. In the Network Interface Configuration section, review the details of the private and public interfaces.

Click Next.

6. On the Credentials/Schedule page:

a. In the Hosts section, specify the Host Credentials (username and password that will be used to access the server during installation) to be used for the target.

You can opt to retain the default selection (Preferred) so that the preferred credentials stored in the Management Repository are used, or override the preferred credentials to explicitly specify the host credentials.

b. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

c. Click Next.

7. On the Configure Cluster page:

a. In the Cluster Name and Location section, review the default name and location details provided for Oracle Clusterware, Oracle RAC Database, and Oracle ASM and edit them if necessary. For example:

In this example, the default cluster name is based on the host cluster name you provided in the Agent Deploy application in Enterprise Manager Cloud Control, while deploying Management Agents on a cluster. The scratch location is a temporary location on the target host where temporary files are placed before provisioning and configuring Oracle RAC.

For security purposes, the clusterware configuration sets the ownership of Oracle Clusterware home and all its parent directories to be owned by root. Hence, Oracle recommends that you install Oracle Clusterware outside the Oracle base of the Oracle RAC home.

b. In the Database Details section, by default, a starter database is created. However, because this environment will be used to create a standby database, no starter database is needed. Deselect the Create Starter Database checkbox.

c. In the ASM Instance Details section (that appears only if you had selected to deploy Oracle ASM), provide the password for the SYS user and the Oracle ASM disk string. For example:

Click Next.

8. On the Storage page, do the following:

a. On the Shared Storage Configuration section, specify locations in the partition name and the mount location fields for the Oracle Cluster Registry (OCR), voting disks, and data files. Also, select the mount format and a storage device for storing data. While partition name is the path to the location where the device is installed, mount location is the mount point that represents the partition location.

If you are using raw devices, you can select the Clear raw devices checkbox under the table to clear the devices as a part of the installation. Doing so ensures the installation does not fail due to information remaining on the devices from previous installations.

b. In the Advanced Options section, select a checkbox for ASM redundancy: None, Normal, or High.

Click Next.

9. On the Advance Configuration page, do the following:

a. In the Configure the Bonding Interface (Private Interconnect) section, if necessary. select Configure Bonding Interface to configure the bonding interface. For more information about the individual settings, click Help in the Enterprise Manager console.

b. In the Sysctl File Configuration section, select Configure Sysctl file if you want to configure the sysctl.conf file. Specify the mode of editing the system configuration file and the location of the reference system configuration file used for modifying the kernel parameters. For more information about the individual settings, click Help in the Enterprise Manager console.

10. On the "Configure Oracle Home" page, you can optionally choose to install and initiate the configuration manager to receive security updates:

• If the host where the database is being provisioned has a direct connection to the Internet, then specify an e-mail address and My Oracle Support password to install and initiate the configuration manager. An e-mail address is required so that security updates and install updates can be sent from My Oracle Support.

• If the host where the database is being provisioned has an indirect connection to the Internet through a proxy server, then specify an e-mail address and My Oracle Support password, and then in the Connection Details section, specify the proxy server details.

Click Next.

11. On the Review page, review the details you have provided for provisioning Oracle RAC and click Finish to submit the job to Enterprise Manager.

After the Deployment Procedure ends successfully you can verify the cluster configuration on the Target tab. In the following MAA example you can verify the clusterware home and the hosts included in the cluster:

32.8.2 Task 2: Create a Single-Instance Physical Standby Database on the New Cluster

The scenario used in the following steps creates a new physical standby database using the Oracle RAC home on the new cluster. This new standby database will then be converted to Oracle RAC during Task 4. During standby creation, create the database files on shared storage to facilitate conversion to Oracle RAC later. The MAA recommendation is to use Oracle ASM managed storage.

Follow the instructions in <link to creating a standy database> to create a single-instance physical standby database on some form of shared storage using the

newly installed Oracle RAC home that you created in Task 1. In this example, the new physical standby database is called *racsby*. The standby database is created on Oracle ASM managed storage.

32.8.3 Task 3: Prepare the Environment Prior to Conversion

Prepare the environment prior to conversion to Oracle RAC by performing the following tasks:

1. Verify that the STANDBY_FILE_MANAGEMENT initialization parameter is set to AUTO on both the primary database and on the standby database that is to be converted.

2. Run the $ORACLE_HOME/rdbms/admin/catclust.sql script on the primary database to install the cluster database views into the environment.

Note: This step is required only if the configuration does not already include a RAC database.

3. Manually create a second undo tablespace on the primary database to support the new database instance being created. The following SQL statement is an example:

```
create undo tablespace undotbs2 datafile
'/u01/app/oracle/oradata/gcprim/undotbs02.dbf' size 500M autoextend on
retention guarantee;
```

Note: This step is required only if there are not enough undo tablespaces already created to support the new Oracle RAC database. You must create an undo tablespace for each new database instance to be added. The Convert to Cluster wizard (executed in Task 4) will display an error if there are not enough undo tablespaces created.

4. Remove the standby database to be converted from the Data Guard broker configuration.

Note: It is necessary to temporarily remove the new standby database from broker management in order to update the broker configuration file initialization parameters. The next step describes how to change the parameters to re-create the broker configuration files on Oracle ASM shared storage (instead of the current location on file system storage).

Perform the following steps to prepare the standby database for conversion to an Oracle RAC standby database. In the following examples, the standby database to be converted is called racsby.

1. In Cloud Control, click the Targets tab, and then click the Databases subtab.

2. On the Databases page, select the primary database (gcprim) from the table.

3. On the gcprim database home page, click Primary in the High Availability section.

4. On the Data Guard page for gcprim, select the standby database that is to be converted (this is the racsby database in the MAA examples), and click Remove. For example:

You should click the check box to “Preserve the destination corresponding to this standby” to continue shipping redo data to the standby.

The Remove Standby Database wizard asks you to confirm the removal and then proceeds to remove the racsby standby database from the broker configuration.

Note: Removing the standby database from the broker configuration only removes the entries from the broker configuration files. The removal does not delete the database from the Data Guard configuration. In fact, the current state of the physical standby database does not change.

5. Modify the following broker initialization parameters for the racsby standby database by using Initialization Parameters option on the Server tab:

• The broker configuration files must be on shared storage when used in conjunction with a RAC database. Edit the DG_BROKER_CONFIG_FILE1 and DG_BROKER_CONFIG_FILE2 parameter values to set them to a location on Oracle ASM shared storage. The recommended Oracle ASM location is in the same diskgroup as the datafiles for the database. Note this directory must exist prior to adding the standby database back into the broker configuration.

• Set the DG_BROKER_START initialization parameter to FALSE to disable the broker. This is necessary for the parameter changes for the broker configuration file locations to take effect.

Click Save to File.

6. Add the racsby standby database back into the broker configuration. This step completes the action of relocating the broker configuration files into the new Oracle ASM shared storage locations.

a. In Cloud Control click Targets, and then click Databases.

b. On the Databases page, select the primary database (gcprim in this example).

c. On the Database Instance page for the primary database, click Primary in the list under the High Availability section of the page.

d. On the Data Guard page for the primary database, click Add Standby Database to invoke the Add Standby Database wizard.

Note: The following steps automatically re-enable the broker management of the racsby standby database.

The following sequence of steps is similar to the process described in Module 1: Create a Physical Standby Database except you will be enabling broker management only for an existing standby database rather than creating a new standby. Respond to the wizard dialog as follows:

• On the Add Standby Database page, select “Manage an existing standby database” and click Continue.

• On the Select Existing Standby Database page, select the standby database from the table. In this example, this is the racsby database. Click Next.

• On the Configuration page, click Next.

• On the Review page, review the configuration settings and click Finish to add the standby database back into the broker configuration. The following screenshot shows the broker configuration running with the gcprim primary database and the two physical standby databases: gcsby and racsby.

32.8.4 Task 4: Convert the Physical Standby Database to an Oracle RAC Database

This task converts the newly created physical standby database into an Oracle RAC database, with no downtime occurring to the primary database.

Figure 4 shows the state of the configuration after you perform the steps to convert the new racsby physical standby database to an Oracle RAC standby database. The configuration contains a single-instance primary database, a single-instance physical standby database, and an Oracle RAC physical standby database.

Figure 4: Intermediate State 2 for Module 2 After Conversion to Oracle RAC

Perform the following steps to convert the racsby standby database to Oracle RAC:

1. In Cloud Control click Targets, and then click Databases.

2. On the Databases page, select the standby database (racsby in our example) from the table and then on the standby database home page, click the Server subtab.

3. On the Server page in the Change Database section, click Convert to Cluster. The Convert to Cluster Database wizard starts.

4. Respond to the Convert to Cluster Database wizard, as follows:

a. On the Cluster Credentials page, specify the credentials for the Oracle RAC and Oracle ASM homes.

Note: In the Information section of the page, you can see the wizard recognizes that the racsby database is a standby database for the gcprim primary database.

On this page, only login credentials were needed to be supplied because the database was already configured to use Oracle RAC and Oracle ASM homes.

Click Next.

b. On the Hosts page, select the hosts from the table on which you want to run the converted Oracle RAC database. Note: The current host for the standby database is always selected.

Click Next.

c. On the Options page, select to use either an existing listener or create a new listener, and specify a prefix to be used to name the cluster database instance (ORACLE_SID).

d. On the Shared Storage page, specify the data file and FLASH_RECOVERY_AREA storage locations. If the database is to be converted to Oracle RAC in-place (that is, the files are already located on shared storage), then you can use the existing locations. Otherwise specify the target disk groups for the data files and the FLASH_ RECOVERY_AREA.

e. Review the settings and click Submit to run the Convert Cluster Database job in Enterprise Manager.

5. After the Convert Cluster Database job has completed successfully, remove the original (non Oracle RAC) standby database definition (racsby) from Cloud Control.

Warning: When you remove the old database from Cloud Control, all the monitoring history is deleted. Remove a database only when this data is no longer needed.

a. In Cloud Control, click Targets.

b. On the Databases page, notice that there are two entries for same standby database: one is for the original single-instance standby database and the other is for the new clustered standby database. Select the single-instance standby database definition from the table and click Remove.

32.8.5 Task 5: Perform a Switchover and Enable Additional Threads

This section describes how to perform a switchover to complete the database conversion to Oracle RAC and to enable the new thread on the Oracle RAC database.

In the MAA example used in this section, the primary database is gcprim and has a single-instance standby database with the database unique name gcsby. The primary database is running on a file system. The new Oracle RAC database unique name is racsby and resides on.

Oracle ASM. After the switchover, racsby is the Oracle RAC primary database, and gcprim is a single-instance physical standby database.

Perform a Switchover

Perform the following steps to switchover the newly converted Oracle RAC standby database (racsby) to run in the primary database role.

1. In Cloud Control, click Targets and then click the All Targets subtab.

2. Select the gcprim database hyperlink.

3. Select Details in the High Availability section.

4. On the Data Guard page, select the Oracle RAC physical standby database that you want to switch to the primary database role, and click Switchover.

After the switchover completes, the Data Guard page shows the roles have been switched between racsby (now an Oracle RAC primary database), and gcprim (now a single-instance physical standby database).

Enable Additional Threads

Manually enable additional threads on the Oracle RAC primary database to make this a two-node Oracle RAC database. For example:

SQL> ALTER DATABASE ENABLE PUBLIC THREAD 2;

Database altered.

Note: Repeat this step for each additional thread added.

Figure 5 shows the configuration that results after you have completed the steps to perform a switchover and enable the additional database threads. The configuration contains an Oracle RAC primary database and two single-instance physical standby databases.

Figure 5: Intermediate State 3 for Module 2 after Switchover

At this point, you can optionally remove the additional physical standby database. Figure 6 shows the ending configuration containing an Oracle RAC primary database and one single-instance physical standby database.

32.9 Configuring Management Service to Management Repository Communication

Management Service processes need to be configured to communicate with each node of the RAC Management Repository in a redundant fashion.

Note that Real Application Cluster (RAC) nodes are referred to by their virtual IP (vip) names. The service_name parameter is used instead of the system identifier (SID) in

connect_data mode and failover is turned on. Refer to *Oracle Database Net Services Administrator's Guide* for details.

Configure the repository connect descriptor by running the emctl command from any Management Service:

```
emctl config oms -store_repos_details -repos_conndesc '"(DESCRIPTION=
(ADDRESS_LIST=(FAILOVER=ON)
(ADDRESS=(PROTOCOL=TCP)(HOST=node1-vip.example.com)(PORT=1521))
(ADDRESS=(PROTOCOL=TCP)(HOST=node2-vip.example.com)(PORT=1521)))
(CONNECT_DATA=(SERVICE_NAME=EMREP)))"' -repos_user sysman
```

After making the previous change, run the following command to make the same change to the monitoring configuration used for the Management Services and Repository target: `emctl config emrep -conn_desc`

32.10 Configuring Standby Database for the Enterprise Manager Repository

The starting point of this step is to have the primary site configured as per Cloud Control MAA guidelines. The following steps lay down the procedure for setting up the standby Management Repository database.

1. Prepare Standby Management Repository hosts for Data Guard

   Install a Management Agent on each of the standby Management Repository hosts. Configure the Management Agents to upload by the SLB on the primary site. Install CRS and Database software on the standby Management Repository hosts. The version used must be the same as that on the primary site.

2. Prepare Primary Management Repository database for Data Guard

   If the primary Management Repository database is not already configured, enable archive log mode, setup flash recovery area and enable flashback database on the primary Management Repository database.

3. Create Physical Standby Database

   In Enterprise Manager, the standby Management Repository database must be physical standbys. Use the Enterprise Manager Console to setup a physical standby database in the standby environment. Note that Enterprise Manager Console does not support creating a standby RAC database. If the standby database has to be RAC, configure the standby database using a single instance and then use the Convert to RAC option from Enterprise Manager Console to convert the single instance standby database to RAC. Also, note that during single instance standby creation, the database files should be created on shared storage to facilitate conversion to RAC later.

   Note that the Convert to RAC option is available for Oracle Database releases 10.2.0.5, 11.1.0.7, and above. Oracle Database release 11.1.0.7 requires patch 8824966 for the Convert to RAC option to work.

4. Add Static Service to Listener

   To enable Data Guard to restart instances during the course of broker operations, a service with a specific name must be statically registered with the local listener of each instance. The value for the GLOBAL_DBNAME attribute must be set to a concatenation of <db_unique_name>_DGMGRL.<db_domain>. For example, in the LISTENER.ORA file:

   ```
   SID_LIST_LISTENER=(SID_LIST=(SID_DESC=(SID_NAME=sid_name)
   ```

```
(GLOBAL_DBNAME=db_unique_name_DGMGRL.db_domain)
(ORACLE_HOME=oracle_home)))
```

5. Enable Flashback Database on the Standby Database
6. Verify Physical Standby

   Verify the Physical Standby database through the Enterprise Manager Console. Click the **Log Switch** button on the Data Guard page to switch log and verify that it is received and applied to the standby database.

32.11 Disaster Recovery

While high availability typically protects against local outages such as application failures or system-level problems, disaster tolerance protects against larger outages such as catastrophic data-center failure due to natural disasters, fire, electrical failure, evacuation, or pervasive sabotage. For Maximum Availability, the loss of a site cannot be the cause for outage of the management tool that handles your enterprise.

Maximum Availability Architecture for Enterprise Manager mandates deploying a remote failover architecture that allows a secondary datacenter to take over the management infrastructure in the event that disaster strikes the primary management infrastructure.

Setting up disaster recovery for Enterprise Manager essentially consists of installing a standby Management Repository Database, a standby Management Service and a standby Server Load Balancer and configuring them to automatically startup when the primary components fail.

Prerequisites

The primary site must be configured as per Cloud Control MAA guidelines described in previous sections. This includes Management Services fronted by an SLB and all Management Agents configured to upload to Management Services by the SLB.

· The standby site must be similar to the primary site in terms of hardware and network resources to ensure there is no loss of performance when failover happens.

· There must be sufficient network bandwidth between the primary and standby sites to handle peak redo data generation.

· Configure storage used by the software library to be replicated at the primary and standby site. In the event of a site outage, the contents of this storage must be made available on the standby site using hardware vendor disk level replication technologies.

· For complete redundancy in a disaster recovery environment, a second load balancer must be installed at the standby site. The secondary SLB must be configured in the same fashion as the primary. Some SLB vendors (such as F5 Networks) offer additional services that can be used to pass control of the Virtual IP presented by the SLB on the primary site to the SLB on the standby site in the event of a site outage. This can be used to facilitate automatic switching of Management Agent traffic from the primary site to the standby site.

32.11.1 Setup Standby Management Service

Consider the following before installing the standby Management Services.

· Oracle recommends that this activity be done during a lean period or during a planned maintenance window. When new Management Services are installed on the

standby site, they are initially configured to connect to the Management Repository database on the primary site. Some workload will be taken up by the new Management Service. This could result in temporary loss in performance if the standby site Management Services are located far away from the primary site Management Repository database. However there would be no data loss and the performance would recover once the standby Management Services are shutdown post configuration.

· The shared storage used for the software library must be made available on the standby site using the same paths as the primary site.

32.11.1.1 Installing the First Standby Management Service

Install the first standby Management Service using the following steps:

1. Copy the emkey to the Management Repository by running the following command on the first Management Service on the primary site:

   ```
   emctl config emkey -copy_to_repos
   ```

2. Export the configuration from the first Management Service on the primary site using:

   ```
   emctl exportconfig oms -dir <location for the export file>
   ```

   After the configuration is exported, do not make any configuration changes to the primary site still the standby management service is configured.

3. Install a Management Agent on the standby host if one does not already exist.

4. Perform a software-only install of the Enterprise Manager software using a modified version of the "Add Management Service" Deployment Procedure.

   Navigate to **Enterprise** -> **Provisioning and Patching** -> **Procedure Library**.

   Select **Add Management Service procedure** and click on the "Create Like" button.

   Go to the **Procedure Steps** tab and select and disable the steps - "Configure Management Service", "Targets Discovery" and "Post Configuration Tasks".

   Save the modified deployment procedure and use it to install the Enterprise Manager software on the standby OMS host.

   After the Deployment Procedure completes, delete the file emInstanceMapping.properties from <OMS Oracle Home>/sysman/config on the standby OMS host.

5. Configure the Management Service by running omsca in standby mode. Choose a different domain name for the standby. For example, if the primary WebLogic domain is GCDomain, choose GCDomainStby.

   ```
   omsca standby –EM_DOMAIN_NAME GCDomainStby –NM_USER
   nodemanager –AS_USERNAME weblogic –nostart
   ```

   When prompted for the Administration Server host and EM Instance host, enter the standby OMS hostname (or accept the default).

   When prompted for the passwords, provide the same passwords as the primary site.

   When prompted for Management Repository details, provide the Primary database details.

6. Configure the required plugins by running the following command:

```
pluginca -action deploy -isFirstOMS true -plugins
<plugin-list> -oracleHome <oms oracle home> -middlewareHome
<wls middleware home>
```

where plugin-list is the list of plugins returned by the SQL query

```
SELECT epv.plugin_id, epv.version FROM   em_plugin_version
epv, em_current_deployed_plugin ecp WHERE epv.plugin_type NOT
IN ( 'BUILT_IN_TARGET_TYPE' , 'INSTALL_HOME') AND ecp.dest_
type='2' AND epv.plugin_version_id = ecp.plugin_version_id;
```

and is a comma separate list in the following format:
`<plugin-id>=<plugin-version>,<plugin-id>=<plugin-version>,…`

Example:

```
"oracle.sysman.empa=12.1.0.1.0,oracle.sysman.mos=12.1.0.1.0,o
racle.sysman.emas=12.1.0.1.0,oracle.sysman.emfa=12.1.0.1.0,or
acle.sysman.db=12.1.0.1.0,oracle.sysman.emct=12.1.0.1.0,oracl
e.sysman.vt=12.1.0.1.0,oracle.sysman.ssa=12.1.0.1.0"
```

7. Copy over the configuration exported from the Primary Management Service in step 2 above to the standby Management Service host. Import the exported configuration on the standby Management Service using:

   ```
   emctl importconfig oms -file <full path of the export file>
   ```

   Note this command emits a warning about a failed export and prompts for confirmation to proceed. The warning can be ignored by entering "y" to proceed.

   Note this command will start the Management Service.

8. Stop the Management Service but leave the Administration Server running using:

   ```
   emctl stop oms
   ```

9. Add the standby Weblogic Domain and associated targets:

   The standby Weblogic Domain and associated targets can be added using the Guided Discovery process via the Setup -> Add Target -> Add Targets Manually page, selecting 'Oracle Fusion Middleware' from the 'Target Types' drop-down. Use the secure port (typically 7101) and, under 'Advanced', set the JMX Protocol to “t3s”.

   Note that the Weblogic targets except the Administration Server will be shown as down as the standby OMS is down at this stage.

10. If you have Single Sign On configured on the primary site, follow the same steps to configure SSO on the standby OMS.

11. If you have Real User Experience Insight, AD4J Manager or BI Publisher configured on the primary site, follow the same steps to configure them on the standby OMS.

32.11.1.2 Installing Additional Standby Management Services

It is recommended that your standby site be similar in configuration as your primary site. This means configuring multiple OMS on your standby site, similar to your primary site. Install additional standby Management Services as per the procedure listed below under “Additional Standby Management Services”.

If, however, you choose to start with a single OMS on the standby site initially, you may skip this step and continue with the next section “Validating your installation and Complete the Setup”. If you decide to add an additional standby OMS later after having run the “Validating your installation and Complete the Setup” steps, the steps

listed under "Additional Standby Management Services" can be followed after executing the following additional steps:

Start up the standby Administration Server by running the following command on the first standby Management Service:

```
emctl start oms –admin_only
```

Additional Standby Management Services

1. Export the configuration from the first Management Service on the primary site using:

   ```
   emctl exportconfig oms -dir <location for the export file>
   ```

   After the configuration is exported, do not make any configuration changes to the primary site still the standby management service is configured.

2. Install a Management Agent on the standby host.

3. Perform a software-only install of the Enterprise Manager software using a modified version of "Add Management Service" Deployment Procedure.

   Navigate to Enterprise -> Provisioning and Patching -> Procedure Library.

   Select Add Management Service procedure and click on "Create Like" button.

   Go to the Procedure Steps tab and select and disable the steps - "Configure Management Service", "Targets Discovery" and "Post Configuration Tasks".

   Save the modified deployment procedure and use it to install the Enterprise Manager software on the standby OMS host.

   After the Deployment Procedure completes, delete the file emInstanceMapping.properties from <OMS Oracle Home>/sysman/config on the standby OMS host.

4. Configure the Management Service by running omsca.

   ```
   omsca add –nostart
   ```

   When prompted for Management Repository details, provide the Primary database details.

   When prompted for Administration Server details, provide the standby administration server details.

5. Configure the required plugins by running the following command:

   ```
   pluginca -action deploy –isFirstOMS false -plugins
   <plugin-list> -oracleHome <oms oracle home> -middlewareHome
   <wls middleware home>
   ```

   where plugin-list is the list of plugins returned by the SQL query above and is a comma separate list in the following format:
   `<plugin-id>=<plugin-version>,<plugin-id>=<plugin-version>,…`

   Example

   ```
   "oracle.sysman.empa=12.1.0.1.0,oracle.sysman.mos=12.1.0.1.0,o
   racle.sysman.emas=12.1.0.1.0,oracle.sysman.emfa=12.1.0.1.0,or
   acle.sysman.db=12.1.0.1.0,oracle.sysman.emct=12.1.0.1.0,oracl
   e.sysman.vt=12.1.0.1.0,oracle.sysman.ssa=12.1.0.1.0"
   ```

6. Copy over the configuration exported from the Primary Management Service in step 1 above to the standby Management Service host. Import the exported configuration on the standby Management Service using:

```
emctl importconfig oms -file <full path of the export file>
```

Note this command emits a warning about a failed export and prompts for confirmation to proceed. The warning can be ignored by entering "y" to proceed.

Note this command will start the Management Service.

7. Stop the Management Service using:

```
emctl stop oms
```

8. Refresh the standby domain target from the console. This will present a guided workflow to discover and add the new managed server and associated targets.
9. If you have Single Sign On configured on the primary site, follow the same steps to configure SSO on the standby OMS.
10. If you have Real User Experience Insight, AD4J Manager or BI Publisher configured on the primary site, follow the same steps to configure them on the standby OMS.

Validating Your Installation and Complete the Setup

Update the standby SLB configuration by adding the standby Management Services to the different pools on the SLB. Setup monitors for the new Management Service.

32.11.2 Setup Standby Database

The starting point of this step is to have the primary site configured as per Cloud Control MAA guidelines. Please refer to this document <need a link here to detailed steps of configuring a Physical Standby> for details steps. Note the following points.

- The Standby Management Repository database must be a Physical Standby. Logical standby are not supported.
- Use the Enterprise Manager itself to setup a physical standby database in the standby environment. Note that Enterprise Manager Console does not support creating a standby RAC database. If the standby database has to be RAC, configure the standby database using a single instance and then use the Convert to RAC option from Enterprise Manager Console to convert the single instance standby database to RAC. Also, note that during single instance standby creation, the database files should be created on shared storage to facilitate conversion to RAC later.

  Note that the Convert to RAC option is available for Oracle Database releases 10.2.0.5, 11.1.0.7, and above. Oracle Database release 11.1.0.7 requires patch 8824966 for the Convert to RAC option to work.

- Add Static Service to Listener

  To enable Data Guard to restart instances during the course of broker operations, a service with a specific name must be statically registered with the local listener of each instance. The value for the GLOBAL_DBNAME attribute must be set to a concatenation of <db_unique_name>_DGMGRL.<db_domain>. For example, in the LISTENER.ORA file:

```
SID_LIST_LISTENER=(SID_LIST=(SID_DESC=(SID_NAME=sid_name)
  (GLOBAL_DBNAME=db_unique_name_DGMGRL.db_domain)
  (ORACLE_HOME=oracle_home)))
```

- To allow re-instate of an old primary database as a standby database after a failover, flashback database must be enabled. Hence do so for both the primary and the standby databases.
- To allow Enterprise Manager to monitor a Physical Standby database (which is typically in a mounted state), specify sysdba monitoring privileges. This can be specified either during the Standby creation wizard itself or post creation by modifying the Monitoring Configuration for the standby database target.

33

High Availability Solutions

Highly Available systems are critical to the success of virtually every business today. It is equally important that the management infrastructure monitoring these mission-critical systems are highly available. The Enterprise Manager Cloud Control architecture is engineered to be scalable and available from the ground up. It is designed to ensure that you concentrate on managing the assets that support your business, while it takes care of meeting your business Service Level Agreements.

When you configure Cloud Control for high availability, your aim is to protect each component of the system, as well as the flow of management data in case of performance or availability problems, such as a failure of a host or a Management Service.

Maximum Availability Architecture (MAA) provides a highly available Enterprise Manager implementation by guarding against failure at each component of Enterprise Manager.

The impacts of failure of the different Enterprise Manager components are:

- Management Agent failure or failure in the communication between Management Agents and Management Services

  Results in targets no longer monitored by Enterprise Manager, though the Enterprise Manager console is still available and one can view historical data from the Management Repository.

- Management Service failure

  Results in the unavailability of Enterprise Manager console, as well as unavailability of almost all Enterprise Manager services.

- Management Repository failure

  Results in failure on the part of Enterprise Manager to save the uploaded data by the Management Agents as well as unavailability of almost all Enterprise Manager services.

Overall, failure in any component of Enterprise Manager can result in substantial service disruption. Therefore it is essential that each component be hardened using a highly available architecture.

This chapter covers the following topics:

- Implementing High Availability Levels

33.1 Latest High Availability Information

Because of rapidly changing technology, and the fact that high availability implementations extend beyond the realm of Oracle Enterprise Manager, the following resources should be checked regularly for the latest information on third-party. integration with Oracle's high availability solutions (F5 or third-party cluster ware, for example).

- Oracle Maximum Availability Architecture Website

  http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm

- Support Note 330072.1: "How To Configure Grid Control Components for High Availability "

33.2 Defining High Availability

Oracle Enterprise Manager's flexible, distributed architecture permits a wide range of deployment configurations, allowing it to meet the monitoring and management needs of your business, as well as allowing for expansion as business needs dictate.

For this reason, high availability for Enterprise Manager cannot be narrowly defined as a singular implementation, but rather a range of protection levels based on your available resources, Oracle technology and best practices that safeguard the investment in your IT infrastructure. Depending on your Enterprise Manager deployment and business needs, you can implement the level of high availability necessary to sustain your business. High availably for Enterprise Manager can be categorized into four levels, each level building on the previous and increasing in implementation cost and complexity, but also incrementally increasing the level of availability.

33.2.1 Levels of High Availability

Each high availability solution level is driven by your business requirements and available IT resources. However, it is important to note that the levels represent a subset of possible deployments that are useful in presenting the various options available. Your IT organization will likely deploy its own configuration which need not exactly match one of the levels.

The following table summarizes the four example high availability levels for Oracle Enterprise Manager installations as well as general resource requirements.

Table 33–1 Enterprise Manager High Availability Levels

Level	Description	Minimum Number of Nodes	Recommended Number of Nodes	Load Balancer Requirements
Level 1	OMS and repository database each on their own host with no failover.	1	2	None
Level 2	OMS installed on shared storage with a VIP based failover database using Local Data Guard.	2	4	None

Table 33–1 (Cont.) Enterprise Manager High Availability Levels

Level	Description	Minimum Number of Nodes	Recommended Number of Nodes	Load Balancer Requirements
Level 3	OMS in Active/Active configuration database using RAC + Local Data Guard	2	5	Local Load Balancer
Level 4	OMS in Active/Active configuration on the primary site standby RAC database (DataGuard) at the disaster recovery site. Multiple standby OMS's at remote site. Data Guard RAC database at the primary site **Note**: Level 4 is a MAA Best Practice, achieving highest availability in the most cost effective, simple architecture.	4	8	**Required**: Local Load Balancer for each site. **Optional**: Global Load Balancer

33.3 Determining Your High Availability Needs

As previously mentioned, the availability level you choose depends on factors such as the hardware resources available and the business need of your organization. However, developing your high availability plan in a way that objectively encompasses all aspects of your high availability needs (hardware, business processes, effort, cost) can be problematic. The solution is to define high availability needs in terms of Recovery Time Objective (RTO) and Recovery Point Objective (RPO).

- **Recovery Time Objective** - The period of time within which your business process or technological resources must be restored after failure. Key Question: How fast do your business processes/resources need to be running again before the bottom line is impacted?
- **Recovery Point Objective** - The period of time between the time of failure and the last backup. Key Question: How much data are you willing to loose?

Defining your high availability needs in terms of RTO and RPO allows you to effectively meet the demands of users. Both values should be determined using the worst-case scenarios.

33.4 Comparing Availability Levels

Given the broad range of factors that must be taken into consideration when implementing a highly available Enterprise Manager environment, your ultimate decision will be based on the interrelationship between RTO, RPO and the cost involved with implementing one of the availability levels. The following table shows the interrelationship between these factors.

Table 33–2 Comparison of High Availability Levels

Level	RTO	RPO	Build Time	Cost
Level 1	98.0%	Hours	Hours to Days	$
Level 2	98.8%	Minutes	Hours to Days	$$

Table 33–2 (Cont.) Comparison of High Availability Levels

Level	RTO	RPO	Build Time	Cost
Level 3	99.9%	Minutes to Seconds	Days	$$$
Level 4	99.9%	Minutes to Seconds	Days	$$$$

The table is not a prescriptive recommendation for choosing a high availability level, but instead should be used to aid your decision making process based on your business needs. For example, you have an uptime requirement of 95% and a desired mean time to recovery of seconds, the you should select level four.

What is not reflected in the table are such factors as survivability and scalability. Hence, although the differences between level three and level four seem outwardly insignificant, there are differences. If you need survivability in the event of a primary site loss you need to go with a Level 4 architecture. If you need equalized performance in the event of site loss it's essential. A level three architecture with DG that's asymmetrically scaled will mean degradation in performance when activated.

If you need to maintain performance levels you will need for level 4 with a symmetrically sized architecture on both sites. This is particularly true if you want to run through planned failover routines where you actively run on the primary or secondary site for extended periods of time. For example, some finance institutions mandate this as part of operating procedures.

The following tables compare the protection levels and recovery times for the various high-availability levels.

Table 33–3 High Availability Levels of Protection

Level	OMS Host Failure	OMS Storage Failure	Database Host Failure	Database Storage Failure	Site Failure
Level 1	No	No	No	No	No
Level 2	Yes	No	Yes	Yes	No
Level 3	Yes	Yes	Yes	Yes	No
Level 4	Yes	Yes	Yes	Yes	Yes

Table 33–4 High Availability Level Recovery Times

Level	Node Failure	Local Storage Failure	Site Failure	Cost
Level 1	Hours-Days	Hours-Days	Hours-Days	$
Level 2	Minutes	Hours-Days	Hours-Days	$$
Level 3	No Outage	Minutes	Hours-Days	$$$
Level 4	No Outage	Minutes	Minutes	$$$$

33.5 Implementing High Availability Levels

Once you have determined the high availability requirements for your enterprise, you are ready to begin implementing one of the high availability levels that is suitable for

your environment. Use the following information roadmap to find implementation instructions for each level.

Level	Where to find information
Level 1	*Oracle Enterprise Manager Basic Installation Guide* and the *Oracle Enterprise Manager Advanced Installation and Configuration Guide*
Level 2	*Oracle Enterprise Manager Basic Installation Guide* and the *Oracle Enterprise Manager Advanced Installation and Configuration Guide* PLUS ■ Installing Multiple OMSs in Active/Active configuration ■ Configuring Standby Database for the Enterprise Manager Repository
Level 3	*Oracle Enterprise Manager Basic Installation Guide* and the *Oracle Enterprise Manager Advanced Installation and Configuration Guide* PLUS ■ Installing Multiple OMSs in Active/Active configuration ■ Configuring the First Management Service for High Availability ■ Configuring Additional Management Services ■ Configuring Software Library ■ Configuring a Load Balancer ■ Reconfiguring the Oracle Management Agent ■ Converting the Enterprise Manager Repository from Single Instance to RAC ■ Configuring Standby Database for the Enterprise Manager Repository
Level 4	*Oracle Enterprise Manager Basic Installation Guide* and the *Oracle Enterprise Manager Advanced Installation and Configuration Guide* PLUS ■ Installing Multiple OMSs in Active/Active configuration ■ Configuring the First Management Service for High Availability ■ Configuring Additional Management Services ■ Configuring Shared File System Loader ■ Configuring Software Library ■ Configuring a Load Balancer ■ Reconfiguring the Oracle Management Agent ■ Configuring Standby Management Service ■ Converting the Enterprise Manager Repository from Single Instance to RAC ■ Configuring Standby Database for the Enterprise Manager Repository

34

Configuring Monitoring for Enterprise Manager High Availability

34.1 Configuration With Cloud Control

Cloud Control comes preconfigured with a series of default rules to monitor many common targets. These rules can be extended to monitor the Cloud Control infrastructure as well as the other targets on your network to meet specific monitoring needs.

34.1.1 Console Warnings, Alerts, and Notifications

The following list is a set of recommendations that extend the default monitoring performed by Enterprise Manager. Use the Incident Rules link to adjust the default rules provided on the Configuration/Rules page:

- Ensure the Agent Unreachable rule is set to alert on all Management Agents unreachable and Management Agents clear errors.
- Ensure the Repository Operations Availability rule is set to notify on any unreachable problems with the Management Service or Management Repository nodes. Also modify this rule to alert on the Targets Not Providing Data condition and any database alerts that are detected against the database serving as the Management Repository.

Modify the Agent Upload Problems Rule to alert when the Management Service status has hit a warning or clear threshold.

34.1.2 Configure Additional Error Reporting Mechanisms

Enterprise Manager provides error reporting mechanisms through e-mail notifications, PL/SQL packages, and SNMP alerts. Configure these mechanisms based on the infrastructure of the production site. If using e-mail for notifications, configure the incident rule through the Cloud Control console to notify administrators using multiple SMTP servers if they are available. This can be done by modifying the default e-mail server setting on the Notification Methods option under Setup.

34.1.3 Component Backup

Backup procedures for the database are well established standards. Configure backup for the Management Repository using the RMAN interface provided in the Cloud

Control console. Refer to the RMAN documentation or the Maximum Availability architecture document for detailed implementation instructions.

In addition to the Management Repository, the Management Service and Management Agent should also have regular backups. Backups should be performed after any configuration change.

34.1.4 Troubleshooting

In the event of a problem with Cloud Control, the starting point for any diagnostic effort is the console itself. The Management System tab provides access to an overview of all Management Service operations and current alerts. Other pages summarize the health of Management Service processes and logged errors These pages are useful for determining the causes of any performance problems as the summary page shows at a historical view of the amount of files waiting to be loaded to the Management Repository and the amount of work waiting to be completed by Management Agents.

34.1.4.1 Upload Delay for Monitoring Data

When assessing the health and availability of targets through the Cloud Control console, information is slow to appear in the UI, especially after a Management Service outage. The state of a target in the Cloud Control console may be delayed after a state change on the monitored host. Use the Management System page to gauge backlog for pending files to be processed.

34.1.4.2 Notification Delay of Target State Change

The model used by the Management Agent to assess the state of health for any particular monitored target is poll based. Management Agents immediately post a notification to the Management Service as soon as a change in state is detected. This infers that there is some potential delay for the Management Agent to actually detect a change in state.

35

Backing Up Enterprise Manager

35.1 Backing Up Your Deployment

Although Enterprise Manager functions as a single entity, technically, it is built on a distributed, multi-tier software architecture composed of the following software components:

- Oracle Management Services (OMS)
- Oracle Management Agent (Agent)
- Oracle Management Repository (Repository)

Each component, being uniquely different in composition and function, requires different approaches to backup and recovery. For this reason, the backup strategies are discussed on a per-tier basis in this chapter. For an overview of Enterprise Manager architecture, refer to the Oracle® Enterprise Manager Cloud Control Basic Installation Guide.

Oracle Configuration Manager

Oracle Configuration Manager (OCM) is used to collect client configuration information and upload it to the Oracle repository. When the client configuration data is uploaded on a regular basis, customer support representatives can analyze this data and provide better service to the customers.

35.1.1 Repository Backup

The Repository is the storage location where all the information collected by the Agent gets stored. It consists of objects such as database jobs, packages, procedures, views, and tablespaces. Because it is configured in an Oracle Database, the backup and recovery strategies for the repository are essentially the same as those for the Oracle Database. Backup procedures for the database are well established standards and can be implemented using the RMAN backup utility, which can be accessed via the Enterprise Manager console.

35.1.1.1 Repository Backup

Oracle recommends using High Availability Best Practices for protecting the Repository database against unplanned outages. As such, use the following standard database backup strategies.

- Database should be in *archivelog* mode. Not running the repository database in *archivelog* mode leaves the database vulnerable to being in an unrecoverable condition after a media failure.
- Perform regular hot backups with RMAN using the *Recommended Backup Strategy* option via the Enterprise Manager console. Other utilities such as DataGuard and RAC can also be used as part of a comprehensive strategy to prevent data loss.

Adhering to these strategies will create a full backup and then create incremental backups on each subsequent run. The incremental changes will then be rolled up into the baseline, creating a new full backup baseline.

Using the *Recommended Backup Strategy* also takes advantage of the capabilities of Enterprise Manager to execute the backups: Jobs will be automatically scheduled through the Job sub-system of Enterprise Manager. The history of the backups will then be available for review and the status of the backup will be displayed on the repository database target home page. This backup job along with archiving and flashback technologies will provide a restore point in the event of the loss of any part of the repository. This type of backup, along with archive and online logs, allows the repository to be recovered to the last completed transaction.

You can view when the last repository backup occurred on the Management Services and Repository Overview page under the Repository details section.

A thorough summary of how to configure backups using Enterprise Manager is available in the *Oracle Database 2 Day DBA* guide. For additional information on Database high availability best practices, review the *Oracle Database High Availability Best Practices* documentation.

35.1.2 Oracle Management Service Backup

The Oracle Management Service (OMS) orchestrates with Management Agents to discover targets, monitor and manage them, and store the collected information in a repository for future reference and analysis. The OMS also renders the Web interface for the Enterprise Manager console. For Enterprise Manager version 11.1, the OMS architecture has changed.

35.1.2.1 Backing Up the OMS

The OMS is generally stateless. Some configuration data is stored on the OMS file system.

A snapshot of OMS configuration can be taken using the `emctl exportconfig oms` command.

```
$ <OMS_HOME>/bin/emctl exportconfig oms [-sysman_pwd <sysman password>]
[-dir <backup dir>] Specify directory to store backup file
[-keep_host] Specify this parameter if the OMS was installed using a virtual
hostname (using
ORACLE_HOSTNAME=<virtual_hostname>)
```

Running exportconfig captures a snapshot of the OMS at a given point in time, thus allowing you to back up the most recent OMS configuration on a regular basis. If required, the most recent snapshot can then be restored on a fresh OMS installation on the same or different host.

Backup strategies for the OMS components are as follows:

- **Software Homes**

omposed of Fusion Middeware Home, the OMS Oracle Home and the WebTier (OHS) Oracle Home and multiple Management Plug-in Oracle Homes.

Software Homes change when patches or patchsets are applied or updates are applied through the new Self Update feature. For this reason, filesystem-level backups should be taken after each patch/patchset application or application of updates through Self Update. You should back up the Oracle inventory files along with the Software Homes and save the output of opatch lsinventory –detail to make it easy to determine which patches are applied to the backed up Oracle Homes.

Note: If you do not have filesystem-level backups, you can also reinstall the software homes using the "Installing Software Only" install method.

Important: The location of the OMS Oracle Home must be the same for all OMS instances in your Cloud Control deployment.

- **Instance Home**

  The gc_inst directory, composed of WebLogic, OMS and WebTier configuration files.

  The Instance Home can be backed up using the `emctl exportconfig oms` command.

- **AdminServer**

  Beginning with Enterprise Manager version 11.1, the OMS's WebLogic architecture introduces the concept of an AdminServer. The AdminServer operates as the central control entity for the configuration of the entire OMS(s) domain. The AdminServer is an integral part of the first OMS installed in your Cloud Control deployment and shares the Software Homes and Instance Home.

  The AdminServer is backed up at the same time as the Instance Home, the emctl exportconfig oms command.

Note: The `exportconfig oms` command is only available with Enterprise Manager version 10.2.0.5 or newer.

Running `exportconfig` captures a snapshot of the OMS at a given point in time, thus allowing you to back up the most recent OMS configuration on a regular basis. If required, the most recent snapshot can then be restored on a fresh OMS installation on the same or different host.

Backup strategies for the OMS components are as follows:

- **Software Homes**

  Composed of three WebLogic components – Middeware Home, the OMS Oracle Home and the WebTier (OHS) Oracle Home. Software Homes only change when patches or patchsets are applied. For this reason, filesystem-level backups should be taken after each patch/patchset application. You should back up the Oracle inventory files along with the Software Homes. .

Important: Beginning with Enterprise Manager version 11.1, the location of the OMS Oracle Home must be the same for all OMS's in your monitored environment.

- **Instance Home**

  Composed of WebLogic, OMS and WebTier configuration files. The Instance Home can be backed up using the `emctl exportconfig oms` command.

- **Software Library**

  Composed of components used by Enterprise Manager patching and provisioning functions. Oracle Database Filesystem (DBFS) is recommended for software library backup. DBFS technology allows an Oracle database tablespace to be exposed to applications as a mounted filesystem. Internally, all the files are stored as secure files in the Oracle database. Storing the software library in the Enterprise Manager repository database using DBFS lets you leverage the comprehensive capabilities of the Oracle database to take consistent backups of the software library along with the Enterprise Manager repository. For more information about DBFS, see the *Oracle® Database SecureFiles and Large Objects Developer's Guide*.

- **Shared Loader RECV Directory**

  The shared loader receive (RECV) directory temporarily stores metric data uploaded from Agents before the data is loaded into the repository. Use a high availability storage technology to protect the receive directory.

- **AdminServer**

  Beginning with Enterprise Manager version 11.1, the OMS's WebLogic architecture introduces the concept of an AdminServer. The AdminServer operates as the central control entity for the configuration of the entire OMS(s) domain. The AdminServer is an integral part of the first OMS installed in your Cloud Control deployment and shares the Software Homes and Instance Home.

35.1.3 Agent Backup

The Agent is an integral software component that is deployed on each monitored host. It is responsible for monitoring all the targets running on those hosts, communicating that information to the middle-tier OMS and managing and maintaining the hosts and its targets.

35.1.3.1 Backing Up Agents

There are no special considerations for backing up Agents. As a best practice, reference Agent installs should be maintained for different platforms and kept up-to-date in terms of customizations in the emd.properties file and patches applied. Use Deployment options from the Cloud Control console to install and maintain reference Agent installs.

36

Enterprise Manager Outages

Outages can be planned as might be the case when performing upgrades or periodic maintenance, or unplanned as can happen in the event of hardware/software failure, or perhaps some environmental catastrophe. Regardless of the type of outage, you want to ensure that your IT infrastructure can be restored an running as soon as possible.

This chapter covers the following:

- Enterprise Manager Recovery
- Recovering from a Simultaneous OMS-Repository Failure
- Switchover
- Failover
- Automatic Failover
- How to Configure Cloud Control OMS in Active/Passive Environment for High Availability Failover Using Virtual Host Names
- Configuring Targets for Failover in Active/Passive Environments

36.1 Enterprise Manager Recovery

Recovering Enterprise Manager means restoring any of the three fundamental components of the Enterprise Manager architecture. Specifically, restoration involves restoring the

- Management Repository
- Management Service
- Management Agent

36.1.1 Repository Recovery

Recovery of the Repository database must be performed using RMAN since Cloud Control will not be available when the repository database is down. There are two recovery cases to consider:

- **Full Recovery**: No special consideration is required for Enterprise Manager.
- **Point-in-Time/Incomplete Recovery**: Recovered repository may be out of sync with Agents because of lost transactions. In this situation, some metrics may show up incorrectly in the Cloud Control console unless the repository is synchronized with the latest state available on the Agents.

A repository resync feature (Enterprise Manager version 10.2.0.5 and later) allows you to automate the process of synchronizing the Enterprise Manager repository with the latest state available on the Agents.

> **Note:** resync requires Agents version 10.2.0.5 or later. Older Agents must be synchronized manually. See "Manually Resynchronizing Agents" on page 1-2.

To resynchronize the repository with the Agents, you use Enterprise Manager Command-line utility (emctl) `resync repos` command:

```
emctl resync repos -full -name "<descriptive name for the operation>"
```

You must run this command from the OMS Oracle Home after restoring the repository but BEFORE starting the OMS. After submitting the command, start up all OMS's and monitor the progress of repository resychronization from the Enterprise Manager console's Repository Resynchronization page, as shown in Figure 1–1.

Figure 36–1 Repository Synchronization Page

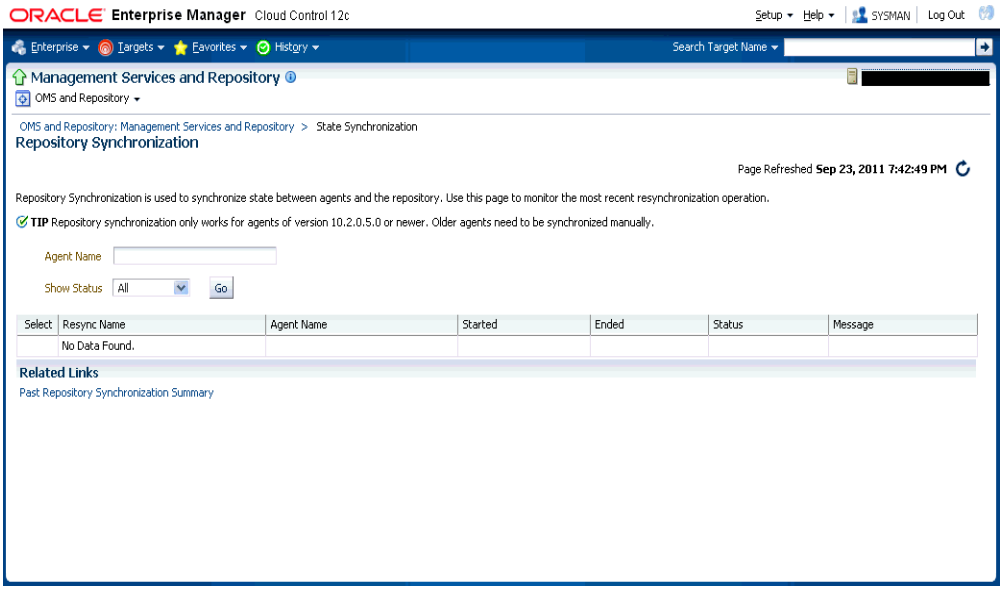

Repository recovery is complete when the resynchronization jobs complete on all Agents.

Oracle strongly recommends that the repository database be run in *archivelog* mode so that in case of failure, the database can be recovered to the latest transaction. If the database cannot be recovered to the last transaction, *Repository Synchronization* can be used to restore monitoring capabilities for targets that existed when the last backup was taken. Actions taken after the backup will not be recovered automatically. Some examples of actions that will not be recovered automatically by *Repository Synchronization* are:

- Incident Rules
- Preferred Credentials
- Groups, Services, Systems

- Jobs/Deployment Procedures
- Custom Reports
- New Agents

Manually Resynchronizing Agents

The Enterprise Manager Repository Synchronization feature can only be used for Agents 10.2.0.5 or later. Older Agents must be resynchronized manually by shutting down the Agent using the following procedure:

1. Shut down the Agent.
2. Delete the agentstmp.txt, lastupld.xml, state/* and upload/* files from the `<AGENT_HOME>/sysman/emd` directory.
3. Restart the Agent.

36.1.2 Recovery Scenarios

A prerequisite for repository (or any database) recovery is to have a valid, consistent backup of the repository. Using Enterprise Manager to automate the backup process ensures regular, up-to-date backups are always available if repository recovery is ever required. Recovery Manager (RMAN) is a utility that backs up, restores, and recovers Oracle Databases. The RMAN recovery job syntax should be saved to a safe location. This allows you to perform a complete recovery of the Enterprise Manager repository database. In its simplest form, the syntax appears as follows:

```
run {
restore database;
recover database;
}
```

Actual syntax will vary in length and complexity depending on your environment. For more information on extracting syntax from an RMAN backup and recovery job, or using RMAN in general, see the *Oracle Database Backup and Recovery Advanced User's Guide*.

The following scenarios illustrate various repository recovery situations along with the recovery steps.

36.1.2.1 Full Recovery on the Same Host

Repository database is running in *archivelog* mode. Recent backup, archive log files and redo logs are available. The repository database disk crashes. All datafiles and control files are lost.

Resolution:

1. Stop all OMS instances using `emctl stop oms`.
2. Recover the database using RMAN
3. Bring the site up using the command `emctl start oms` on all OMS instances.
4. Verify that the site is fully operational.

36.1.2.2 Incomplete Recovery on the Same Host

Repository database is running in *noarchivelog* mode. Full offline backup is available. The repository database disk crashes. All datafiles and control files are lost.

Resolution:

1. Stop the OMS(s) using `emctl stop oms`.
2. Recover the database using RMAN.
3. Initiate Repository Resync using `emctl resync repos -full -name` "<resync name>" from one of the OMS Oracle Home.
4. Start the OMS(s) using `emctl start oms`.
5. Log into Cloud Control. Navigate to **Management Services and Repository Overview** page. Click on **Repository Synchronization** under **Related Links**. Monitor the status of resync jobs. Resubmit failed jobs, if any, after fixing the error.
6. Verify that the site is fully operational.

36.1.2.3 Full Recovery on a Different Host

The repository database is running on host "A" in *archivelog* mode. Recent backup, archive log files and redo logs are available. The repository database crashes. All datafiles and control files are lost.

Resolution:

1. Stop the OMS instances using the command `emctl stop oms`.
2. Recover the database using RMAN on a different host (host "B").
3. Correct the connect descriptor for the repository in credential store by running

```
$emctl config oms -store_repos_details -repos_conndesc <connect descriptor>
-repos_user sysman
```

4. Start the OMS(s) using the command `emctl start oms`.
5. Relocate the repository database target to the Agent running on host "B" by running the following command from the OMS:

```
$emctl config repos -host <hostB> -oh <OH of repository on hostB>  -conn_desc
"<TNS connect descriptor>"
```

> **Note:** This command can only be used to relocate the repository database under the following conditions:
>
> - An Agent is already running on this machine.
> - No database on host "B" has been discovered.

6. Change the monitoring configuration for the OMS and Repository target: by running the following command from the OMS:

```
$emctl config emrep -conn_desc "<TNS connect descriptor>"
```

7. Verify that the site is fully operational.

36.1.2.4 Incomplete Recovery on a Different Host

The repository database is running on host "A" in *noarchivelog* mode. Full offline backup is available. Host "A" is lost due to hardware failure. All datafiles and control files are lost.

Resolution:

1. Stop the OMS(s) using `emctl stop oms`.

2. Recover the database using RMAN on a different host (host "B").
3. Correct the connect descriptor for the repository in credential store.

   ```
   $emctl config oms -store_repos_details -repos_conndesc <connect descriptor>
   -repos_user sysman
   ```

4. Initiate Repository Resync:

   ```
   $emctl resync repos -full -name "<resync name>"
   ```

   from one of the OMS Oracle Homes.
5. Start the OMS using the command `emctl start oms.`
6. Run the command to relocate the repository database target to the Agent running on host "B":

   ```
   $emctl config repos -agent <agent on host B> -host <hostB>
   -oh <OH of repository on hostB> -conn_desc "<TNS connect
   descriptor>"
   ```

7. Run the command to change monitoring configuration for the OMS and Repository target:

   ```
   emctl config emrep -conn_desc "<TNS connect descriptor>"
   ```

8. Manually fix all pre-10.2.0.5 Agents by shutting down the Agents, deleting the agentstmp.txt, lastupld.xml, state/* and upload/* files under the <AGENT_HOME>/sysman/emd directory and then restarting the Agents.
9. Log in to Cloud Control. Navigate to **Management Services and Repository Overview** page. Choose on **Repository Synchronization** under **Related Links**. Monitor the status of resync jobs. Resubmit failed jobs, if any, after fixing the error mentioned.
10. Verify that the site is fully operational.

36.1.3 Recovering the OMS

If an OMS is lost, recovering an OMS essentially consists of two steps, recovering the Software Homes and then configuring the Instance Home. When restoring on the same host, the software homes can be restored from filesystem backup. In case a backup does not exist, or if installing to a different host, the software homes can be reconstructed using the “Install Software Only" option from the Cloud Control software distribution. Care should be taken to select and install all Management Plugins that existed in your environment prior to crash. The following SQL command can be run against the repository database as the “sysman” user to get the list of plugins already deployed:

```
select epv.display_name, gcp.plugin_id||':'||gcp.version "plugin-version" from GC_
CURRENT_DEPLOYED_PLUGIN gcp, MGMT_OMS_PARAMETERS omsp, EM_PLUGIN_VERSION epv where
gcp.DESTINATION_TYPE='OMS' and gcp.DESTINATION_NAME = omsp.host_url and omsp.NAME
= 'HOST_NAME' and gcp.plugin_id = epv.PLUGIN_ID;
```

Note that some plugins might have not shipped with Cloud Control and might not be present in the install media. Such plugins should be downloaded from OTN and their location passed to the Oracle Installer. Choose all plugins returned by the SQL query above in the Plugins page of the Installer. Recovery will fail if all the required plugins are not selected.

After running the installer in software only mode, all patches that were installed prior to the crash must be re-applied. Assuming the repository is intact, the post scripts that

run SQLs against the repository can be skipped as the repository already has those patches applied.

As stated earlier, the location of the OMS Oracle Home is fixed and cannot be changed. Hence, ensure that the OMS Oracle Home is restored in the same location that was used previously.

Once the Software Homes are recovered, the instance home can be reconstructed using the omsca command in recovery mode:

```
omsca recover –as –ms -nostart –backup_file <exportconfig file>
```

Use the export file generated by the `emctl exportconfig` command shown in the previous section.

36.1.4 OMS Recovery Scenarios

The following scenarios illustrate various OMS recovery situations along with the recovery steps.

Important: A prerequisite for OMS recovery is to have recent, valid OMS configuration backups available. Oracle recommends that you back up the OMS using the `emctl exportconfig oms` command whenever an OMS configuration change is made. This command must be run on the primary OMS running the WebLogic AdminServer.

Alternatively, you can run this command on a regular basis using the Enterprise Manager Job system.

Each of the following scenarios cover the recovery of the Software homes using either a filesystem backup (when available and only when recovering to the same host) or using the Software only option from the installer. In either case, the best practice is to recover the instance home (gc_inst) using the omsca recover command, rather than from a filesystem backup. This guarantees that the instance home is valid and up to date.

36.1.4.1 Single OMS, No Server Load Balancer (SLB), OMS Restored on the same Host

Site hosts a single OMS. No SLB is present. The OMS configuration was backed up using the `emctl exportconfig oms` command on the primary OMS running the AdminServer. The OMS Oracle Home is lost.

Resolution:

1. Perform cleanup on failed OMS host.

   Make sure there are no processes still running from the Middleware home using a command similar to the following:

   ```
   ps -ef | grep -i -P "(Middleware|gc_inst)" | grep -v grep | awk '{print $2}' |
   xargs kill -9
   ```

   If they exist, remove the ‘Middleware’ and ‘gc_inst’ directories.

2. Ensure that software library locations are still accessible.

3. Restore the software homes.

If restoring from a filesystem backup, delete the file <OMS_HOME>/sysman/config/emInstanceMapping.properties and any gc_inst directory that may have been restored, if they exist.

Alternatively, if a backup does not exist, use the software only install method to reconstruct the software homes:

1. Select the 'Install Software Only' option from the 'Install Types' step page within the Cloud Control software installer.
2. Ensure all previously deployed plug-ins are selected on the 'Select Plug-ins' step page.

   It is possible to determine which plugins were deployed previously by running the following SQL against the repository database:

   ```
   select epv.display_name, gcp.plugin_id||':'||gcp.version "plugin-version"
   from GC_CURRENT_DEPLOYED_PLUGIN gcp, MGMT_OMS_PARAMETERS omsp, EM_PLUGIN_
   VERSION epv where gcp.DESTINATION_TYPE='OMS' and gcp.DESTINATION_NAME =
   omsp.host_url and omsp.NAME = 'HOST_NAME' and gcp.plugin_id = epv.PLUGIN_
   ID;
   ```

   Note: At the end of the Software only installation, do NOT run *ConfigureGC.pl* when told to do so by the installer. This step should only be performed as part of a fresh install, not as part of a recovery operation.

3. Apply any patches that were previously applied to the OMS software homes.

4. Run omsca in recovery mode specifying the export file taken earlier to configure the OMS:

   ```
   <OMS_HOME>/bin/omsca recover -as -ms -nostart -backup_file <exportconfig file>
   ```

   Note: The -backup_file to be passed must be the latest file generated from emctl exportconfig oms command.

5. Configure and deploy required plugins

   ```
   <OMS_HOME>/bin/pluginca -action midtierconfig -plugins <plugin-list>
   -oraclehome <oms oracle home> -middlewarehome <wls middleware home>
   ```

   where plugin-list is a comma separated list of "plugin-version" value for each row returned by the query:

   ```
   select epv.display_name, gcp.plugin_id||':'||gcp.version "plugin-version" from
   GC_CURRENT_DEPLOYED_PLUGIN gcp, MGMT_OMS_PARAMETERS omsp, EM_PLUGIN_VERSION epv
   where gcp.DESTINATION_TYPE='OMS' and gcp.DESTINATION_NAME = omsp.host_url and
   omsp.NAME = 'HOST_NAME' and gcp.plugin_id = epv.PLUGIN_ID;
   ```

   For exmple:

   ```
   /u01/app/oracle/Middleware/oms/bin/pluginca -action midtierconfig -plugins
   "oracle.sysman.mos=12.1.0.0.0,oracle.sysman.emas=12.1.0.0.0,oracle.sysman.db=12
   .1.0.0.0" -oraclehome /u01/app/oracle/Middleware/oms -middlewarehome
   /u01/app/oracle/Middleware
   ```

6. Start the OMS.

   ```
   <OMS_HOME>/bin/emctl start oms
   ```

7. Recover the Agent (if necessary).

   If the Agent software home was recovered along with the OMS software homes (as is likely in a single OMS install recovery where the agent and agent_inst directories are commonly under the Middleware home), the Agent instance directory should be recreated to ensure consistency between the Agent and OMS.

   Remove the agent_inst directory if it was restored from backup

   Use agentDeploy.sh to configure the agent:

   ```
   <AGENT_HOME>/core/12.1.0.0.0/sysman/install/agentDeploy.sh AGENT_BASE_
   DIR=<AGENT_BASE_DIR> AGENT_INSTANCE_HOME=<AGENT_INSTANCE_HOME> ORACLE_
   HOSTNAME=<AGENT_HOSTNAME> AGENT_PORT=<AGENT_PORT> -configOnly OMS_HOST=<oms
   host> EM_UPLOAD_PORT=<OMS_UPLOAD_PORT> AGENT_REGISTRATION_PASSWORD=<REG_
   PASSWORD>
   ```

   The OMS automatically blocks the Agent. Resync the Agent from the Agent homepage.

   If the Agent software home was not recovered along with the OMS but the Agent still needs to be recovered, follow the instructions in section *Agent Reinstall Using the Same Port*.

   Note: This is only likely to be needed in the case where a filesystem recovery has been performed that did not include a backup of the Agent software homes. If the OMS software homes were recovered using the Software only install method, this step will not be required because a Software only install installs an Agent software home under the Middleware home.

8. Verify that the site is fully operational.

36.1.4.2 Single OMS, No SLB, OMS Restored on a Different Host

Site hosts a single OMS. The OMS is running on host "A." No SLB is present. The OMS configuration was backed up using the `emctl exportconfig oms` command. Host "A" is lost.

Resolution:

1. Ensure that software library locations are accessible from “Host B”.
2. Restore the software homes on “Host B”.

   Oracle does not support restoring OMS Oracle Homes from filesystem backup across different hosts. Use the software-only install method to reconstruct the software homes:

   1. Select the ‘Install Software Only’ option from the ‘Install Types’ step page within the Cloud Control software installer.
   2. Ensure all previously deployed plug-ins are selected on the ‘Select Plug-ins’ step page.

      It is possible to determine which plugins were deployed previously by running the following SQL against the repository database:

      ```
      select epv.display_name, gcp.plugin_id||':'||gcp.version "plugin-version"
      from GC_CURRENT_DEPLOYED_PLUGIN gcp, MGMT_OMS_PARAMETERS omsp, EM_PLUGIN_
      VERSION epv where gcp.DESTINATION_TYPE='OMS' and gcp.DESTINATION_NAME =
      omsp.host_url and omsp.NAME = 'HOST_NAME' and gcp.plugin_id = epv.PLUGIN_
      ```

```
ID;
```

Note: At the end of the Software only installation, do NOT run *ConfigureGC.pl* when told to do so by the installer. This step should only be performed as part of a fresh install, not as part of a recovery operation.

3. Apply any patches that were previously applied to the OMS software homes.

3. Run omsca in recovery mode specifying the export file taken earlier to configure the OMS:

```
<OMS_HOME>/bin/omsca recover –as –ms –nostart –backup_file <exportconfig file>
```

Note: The -backup_file to be passed must be the latest file generated from emctl exportconfig oms command.

4. Configure and deploy required plugins.

```
<OMS_HOME>/bin/pluginca -action midtierconfig -plugins <plugin-list>
-oraclehome <oms oracle home> -middlewarehome <wls middleware home>
```

where plugin-list is a comma separated list of “plugin-version” value for each row returned by the query:

```
select epv.display_name, gcp.plugin_id||':'||gcp.version "plugin-version" from
GC_CURRENT_DEPLOYED_PLUGIN gcp, MGMT_OMS_PARAMETERS omsp, EM_PLUGIN_VERSION epv
where gcp.DESTINATION_TYPE='OMS' and gcp.DESTINATION_NAME = omsp.host_url and
omsp.NAME = 'HOST_NAME' and gcp.plugin_id = epv.PLUGIN_ID;
```

For example:

```
/u01/app/oracle/Middleware/oms/bin/pluginca -action midtierconfig -plugins
"oracle.sysman.mos=12.1.0.0.0,oracle.sysman.emas=12.1.0.0.0,oracle.sysman.db=12
.1.0.0.0" -oraclehome /u01/app/oracle/Middleware/oms -middlewarehome
/u01/app/oracle/Middleware
```

5. Start the OMS.

```
<OMS_HOME>/bin/emctl start oms
```

An agent is installed as part of the Software only install and needs to be configured using the agentDeploy.sh command:

6. Configure the Agent.

```
<AGENT_HOME>/core/12.1.0.0.0/sysman/install/agentDeploy.sh AGENT_BASE_
DIR=<AGENT_BASE_DIR> AGENT_INSTANCE_HOME=<AGENT_INSTANCE_HOME> ORACLE_
HOSTNAME=<AGENT_HOSTNAME> AGENT_PORT=<AGENT_PORT> -configOnly OMS_HOST=<oms
host> EM_UPLOAD_PORT=<OMS_UPLOAD_PORT> AGENT_REGISTRATION_PASSWORD=<REG_
PASSWORD>
```

The OMS automatically blocks the Agent. Resync the Agent from the Agent homepage

7. Relocate the oracle_emrep target to the Agent of the new OMS host using the following commands:

```
<OMS_HOME>/bin/emcli login –username=sysman
<OMS_HOME>/bin/emcli sync
```

```
<OMS_HOME>/bin/emctl config emrep -agent <agent on host "B", e.g
myNewOMSHost.example.com:3872>
```

8. In the Cloud Control console, locate the 'WebLogic Domain' target for the Cloud Control Domain. Go to 'Monitoring Credentials' and update the adminserver host to host B. Then do a Refresh Weblogic Domain to reconfigure the domain with new hosts.
9. Locate duplicate targets from the Management Services and Repository Overview page of the Enterprise Manager console. Click the Duplicate Targets link to access the Duplicate Targets page. To resolve duplicate target errors, the duplicate target must be renamed on the conflicting Agent. Relocate duplicate targets from Agent "A" to Agent "B".
10. Change the OMS to which all Agents point and then resecure all Agents.

    Because the new machine is using a different hostname from the one originally hosting the OMS, all Agents in your monitored environment must be told where to find the new OMS. On each Agent, run the following command:

    ```
    <AGENT_INST_DIR>/bin/emctl secure agent -emdWalletSrcUrl "http://hostB:<http_
    port>/em"
    ```

11. Assuming the original OMS host is no longer in use, remove the Host target (including all remaining monitored targets) from Cloud Control by selecting the host on the Targets > Hosts page and clicking 'Remove'. You will be presented with an error that informs you to remove all monitored targets first. Remove those targets then repeat the step to remove the Host target successfully.
12. Verify that the site is fully operational.

36.1.4.3 Single OMS, No SLB, OMS Restored on a Different Host using the Original Hostname

Site hosts a single OMS. The OMS is running on host "A." No SLB is present. The OMS configuration was backed up using the `emctl exportconfig oms` command. Host "A" is lost. Recovery is to be performed on "Host B" but retaining the use of "Hostname A".

Resolution:

1. Ensure that loader receive directory and software library locations are accessible from Host "B".

   Oracle does not support restoring OMS Oracle Homes from filesystem backup across different hosts. Use the software-only install method to reconstruct the software homes:

   1. Select the 'Install Software Only' option from the 'Install Types' step page within the Cloud Control software installer.
   2. Ensure all previously deployed plug-ins are selected on the 'Select Plug-ins' step page.

      It is possible to determine which plugins were deployed previously by running the following SQL against the repository database:

      ```
      select epv.display_name, gcp.plugin_id||':'||gcp.version "plugin-version"
      from GC_CURRENT_DEPLOYED_PLUGIN gcp, MGMT_OMS_PARAMETERS omsp, EM_PLUGIN_
      VERSION epv where gcp.DESTINATION_TYPE='OMS' and gcp.DESTINATION_NAME =
      omsp.host_url and omsp.NAME
      ```

> **Note:** At the end of the Software only installation, do NOT run *ConfigureGC.pl* when told to do so by the installer. This step should only be performed as part of a fresh install, not as part of a recovery operation.

3. Apply any patches that were previously applied to the OMS software homes.

2. Modify the network configuration such that "Host B" also responds to hostname of "Host A". Specific instructions on how to configure this are beyond the scope of this document. However, some general configuration suggestions are:

   Modify your DNS server such that both "Hostname B" and "Hostname A" network addresses resolve to the physical IP of "Host B".

   Multi-home "Host B". Configure an additional IP on "Host B" for the IP address that "Hostname A" resolves to. For example, on "Host B" run the following commands:

   ```
   ifconfig eth0:1 <IP assigned to "Hostname A"> netmask <netmask>
   /sbin/arping -q -U -c 3 -I eth0 <IP of HostA>
   ```

3. Run omsca in recovery mode specifying the export file taken earlier to configure the OMS:

   ```
   <OMS_HOME>/bin/omsca recover -as -ms -nostart -backup_file <exportconfig file>
   -AS_HOST <hostA> -EM_INSTANCE_HOST <hostA>
   ```

   > **Note:** The -backup_file to be passed must be the latest file generated from emctl exportconfig oms command.

4. Configure and deploy required plug-ins.

   ```
   <OMS_HOME>/bin/pluginca -action midtierconfig -plugins <plugin-list>
   -oraclehome <oms oracle home> -middlewarehome <wls middleware home>
   ```

   where plugin-list is a comma separated list of "plugin-version" value for each row returned by the query:

   ```
   select epv.display_name, gcp.plugin_id||':'||gcp.version "plugin-version" from
   GC_CURRENT_DEPLOYED_PLUGIN gcp, MGMT_OMS_PARAMETERS omsp, EM_PLUGIN_VERSION epv
   where gcp.DESTINATION_TYPE='OMS' and gcp.DESTINATION_NAME = omsp.host_url and
   omsp.NAME = 'HOST_NAME' and gcp.plugin_id = epv.PLUGIN_ID;
   ```

   For example:

   ```
   /u01/app/oracle/Middleware/oms/bin/pluginca -action midtierconfig -plugins
   "oracle.sysman.mos=12.1.0.0.0,oracle.sysman.emas=12.1.0.0.0,oracle.sysman.db=12
   .1.0.0.0" -oraclehome /u01/app/oracle/Middleware/oms -middlewarehome
   /u01/app/oracle/Middleware
   ```

5. Start the OMS

   ```
   <OMS_HOME>/bin/emctl start oms
   ```

6. Configure the agent.

   An agent is installed as part of the Software only install and needs to be configured using the agentDeploy.sh command:

```
<AGENT_HOME>/core/12.1.0.0.0/sysman/install/agentDeploy.sh AGENT_BASE_
DIR=<AGENT_BASE_DIR> AGENT_INSTANCE_HOME=<AGENT_INSTANCE_HOME> ORACLE_
HOSTNAME=<AGENT_HOSTNAME> AGENT_PORT=<AGENT_PORT> -configOnly OMS_HOST=<oms
host> EM_UPLOAD_PORT=<OMS_UPLOAD_PORT> AGENT_REGISTRATION_PASSWORD=<REG_
PASSWORD>
```

7. The OMS automatically blocks the Agent. Resync the Agent from the Agent homepage.

   Run the command to relocate Management Services and Repository target to Agent "B":

   ```
   emctl config emrep -agent <agent on host B>
   ```

8. Locate duplicate targets from the Management Services and Repository Overview page of the Enterprise Manager console. Click the Duplicate Targets link to access the Duplicate Targets page. To resolve duplicate target errors, the duplicate target must be renamed on the conflicting Agent. Relocate duplicate targets from Agent "A" to Agent "B".
9. Verify that the site is fully operational.

36.1.4.4 Multiple OMS, Server Load Balancer, Primary OMS Recovered on the Same Host

Site hosts multiple OMSs. All OMSs are fronted by a Server Load Balancer. OMS configuration backed up using the emctl exportconfig oms command on the primary OMS running the WebLogic AdminServer. The primary OMS is lost.

Resolution:

1. Perform cleanup on failed OMS host.

   Make sure there are no processes still running from the Middleware home using a command similar to the following:

   ```
   ps -ef | grep -i -P "(Middleware|gc_inst)" | grep -v grep | awk '{print $2}' |
   xargs kill -9
   ```

   If recovering the software homes using the software only install method, first de-install the existing Oracle Homes using the Cloud Control software distribution installer. This is required even if the software homes are no longer available as it is necessary to remove any record of the lost Oracle Homes from the Oracle inventory.

   If they exist, remove the 'Middleware' and 'gc_inst' directories.
2. Ensure that software library locations are still accessible.
3. Restore the software homes.

   If restoring from a filesystem backup, delete the file <OMS_HOME>/sysman/config/emInstanceMapping.properties and any gc_inst directory that may have been restored, if they exist.

   Alternatively, if a backup does not exist, use the software only install method to reconstruct the software homes:

   1. Select the 'Install Software Only' option from the 'Install Types' step page within the Cloud Control software installer.
   2. Ensure all previously deployed plug-ins are selected on the 'Select Plug-ins' step page.

It is possible to determine which plugins were deployed previously by running the following SQL against the repository database:

```
select epv.display_name, gcp.plugin_id||':'||gcp.version "plugin-version"
from GC_CURRENT_DEPLOYED_PLUGIN gcp, MGMT_OMS_PARAMETERS omsp, EM_PLUGIN_
VERSION epv where gcp.DESTINATION_TYPE='OMS' and gcp.DESTINATION_NAME =
omsp.host_url and omsp.NAME = 'HOST_NAME' and gcp.plugin_id = epv.PLUGIN_
ID;
```

Note: At the end of the Software only installation, do NOT run *ConfigureGC.pl* when told to do so by the installer. This step should only be performed as part of a fresh install, not as part of a recovery operation.

3. Apply any patches that were previously applied to the OMS software homes.

4. Run omsca in recovery mode specifying the export file taken earlier to configure the OMS:

```
<OMS_HOME>/bin/omsca recover –as –ms –nostart –backup_file <exportconfig file>
```

Note: The -backup_file to be passed must be the latest file generated from emctl exportconfig oms command.

5. Configure and deploy required plug-ins.

```
<OMS_HOME>/bin/pluginca -action midtierconfig -plugins <plugin-list>
-oraclehome <oms oracle home> -middlewarehome <wls middleware home>
```

where plugin-list is a comma separated list of “plugin-version” value for each row returned by the query:

```
select epv.display_name, gcp.plugin_id||':'||gcp.version "plugin-version" from
GC_CURRENT_DEPLOYED_PLUGIN gcp, MGMT_OMS_PARAMETERS omsp, EM_PLUGIN_VERSION epv
where gcp.DESTINATION_TYPE='OMS' and gcp.DESTINATION_NAME = omsp.host_url and
omsp.NAME = 'HOST_NAME' and gcp.plugin_id = epv.PLUGIN_ID;
```

For example:

```
/u01/app/oracle/Middleware/oms/bin/pluginca -action midtierconfig -plugins
"oracle.sysman.mos=12.1.0.0.0,oracle.sysman.emas=12.1.0.0.0,oracle.sysman.db=12
.1.0.0.0" -oraclehome /u01/app/oracle/Middleware/oms -middlewarehome
/u01/app/oracle/Middleware
```

6. Start the OMS.

```
<OMS_HOME>/bin/emctl start oms
```

7. Recover the Agent.

If the Agent software home was recovered along with the OMS software homes (as is likely in a Primary OMS install recovery where the agent and agent_inst directories are commonly under the Middleware home), the Agent instance directory should be recreated to ensure consistency between the Agent and OMS.

Remove the agent_inst directory if it was restored from backup

Use agentDeploy.sh to configure the Agent:

```
<AGENT_HOME>/core/12.1.0.0.0/sysman/install/agentDeploy.sh AGENT_BASE_
```

```
DIR=<AGENT_BASE_DIR> AGENT_INSTANCE_HOME=<AGENT_INSTANCE_HOME> ORACLE_
HOSTNAME=<AGENT_HOSTNAME> AGENT_PORT=<AGENT_PORT> -configOnly OMS_HOST=<oms
host> EM_UPLOAD_PORT=<OMS_UPLOAD_PORT> AGENT_REGISTRATION_PASSWORD=<REG_
PASSWORD>
```

The OMS automatically blocks the Agent. Resync the Agent from the Agent homepage.

If the Agent software home was not recovered along with the OMS but the Agent still needs to be recovered, follow the instructions in section *Agent Reinstall Using the Same Port*.

Note: This is only likely to be needed in the case where a filesystem recovery has been performed that did not include a backup of the Agent software homes. If the OMS software homes were recovered using the Software only install method, this step will not be required because a Software only install installs an Agent software home under the Middleware home.

---

8. Re-enroll the additional OMS, if any, with the recovered Administration Server by running emctl enroll oms on each additional OMS.
9. Verify that the site is fully operational.

36.1.4.5 Multiple OMS, Server Load Balancer configured, Primary OMS Recovered on a Different Host

Site hosts multiple OMSs. OMSs fronted by a Server Load Balancer. OMS Configuration backed up using emctl exportconfig oms command. Primary OMS on host "A" is lost and needs to be recovered on Host "B".

1. If necessary, perform cleanup on failed OMS host.

   Make sure there are no processes still running from the Middleware home using a command similar to the following:

   ```
   ps -ef | grep -i -P "(Middleware|gc_inst)" | grep -v grep | awk '{print $2}' |
   xargs kill -9
   ```

2. Ensure that software library locations are accessible from “Host B”.
3. Restore the software homes on “Host B”.

   Oracle does not support restoring OMS Oracle Homes from filesystem backup across different hosts. Use the software-only install method to reconstruct the software homes:

   1. Select the ‘Install Software Only’ option from the ‘Install Types’ step page within the Cloud Control software installer.
   2. Ensure all previously deployed plug-ins are selected on the ‘Select Plug-ins’ step page.

      It is possible to determine which plugins were deployed previously by running the following SQL against the repository database:

      ```
      select epv.display_name, gcp.plugin_id||':'||gcp.version "plugin-version"
      from GC_CURRENT_DEPLOYED_PLUGIN gcp, MGMT_OMS_PARAMETERS omsp, EM_PLUGIN_
      VERSION epv where gcp.DESTINATION_TYPE='OMS' and gcp.DESTINATION_NAME =
      omsp.host_url and omsp.NAME = 'HOST_NAME' and gcp.plugin_id = epv.PLUGIN_
      ID;
      ```

> **Note:** At the end of the Software only installation, do NOT run *ConfigureGC.pl* when told to do so by the installer. This step should only be performed as part of a fresh install, not as part of a recovery operation.

3. Apply any patches that were previously applied to the OMS software homes.

4. Run omsca in recovery mode specifying the export file taken earlier to configure the OMS:

```
<OMS_HOME>/bin/omsca recover -as -ms -nostart -backup_file <exportconfig file>
```

> **Note:** The -backup_file to be passed must be the latest file generated from emctl exportconfig oms command.

5. Configure and deploy required plug-ins.

```
<OMS_HOME>/bin/pluginca -action midtierconfig -plugins <plugin-list>
-oraclehome <oms oracle home> -middlewarehome <wls middleware home>
```

where plugin-list is a comma separated list of "plugin-version" value for each row returned by the query:

```
select epv.display_name, gcp.plugin_id||':'||gcp.version "plugin-version" from
GC_CURRENT_DEPLOYED_PLUGIN gcp, MGMT_OMS_PARAMETERS omsp, EM_PLUGIN_VERSION epv
where gcp.DESTINATION_TYPE='OMS' and gcp.DESTINATION_NAME = omsp.host_url and
omsp.NAME = 'HOST_NAME' and gcp.plugin_id = epv.PLUGIN_ID;
```

For example:

```
/u01/app/oracle/Middleware/oms/bin/pluginca -action midtierconfig -plugins
"oracle.sysman.mos=12.1.0.0.0,oracle.sysman.emas=12.1.0.0.0,oracle.sysman.db=12
.1.0.0.0" -oraclehome /u01/app/oracle/Middleware/oms -middlewarehome
/u01/app/oracle/Middleware
```

6. Start the OMS.

```
<OMS_HOME>/bin/emctl start oms
```

7. Configure the Agent.

An agent is installed as part of the Software only install and needs to be configured using the agentDeploy.sh command:

```
<AGENT_HOME>/core/12.1.0.0.0/sysman/install/agentDeploy.sh AGENT_BASE_
DIR=<AGENT_BASE_DIR> AGENT_INSTANCE_HOME=<AGENT_INSTANCE_HOME> ORACLE_
HOSTNAME=<AGENT_HOSTNAME> AGENT_PORT=<AGENT_PORT> -configOnly OMS_HOST=<oms
host> EM_UPLOAD_PORT=<OMS_UPLOAD_PORT> AGENT_REGISTRATION_PASSWORD=<REG_
PASSWORD>
```

The OMS automatically blocks the Agent. Resync the Agent from the Agent homepage

8. Add the new OMS to the SLB virtual server pools and remove the old OMS.

9. Relocate the oracle_emrep target to the Agent of the new OMS host using the following commands:

```
<OMS_HOME>/bin/emcli sync
<OMS_HOME>/bin/emctl config emrep -agent <agent on host "B", e.g
```

```
myNewOMSHost.example.com:3872>
```

10. In the Cloud Control console, locate the 'WebLogic Domain' target for the Cloud Control Domain. Go to 'Monitoring Credentials' and update the adminserver host to host B. Then do a Refresh Weblogic Domain to reconfigure the domain with new hosts.
11. Locate duplicate targets from the Management Services and Repository Overview page of the Enterprise Manager console. Click the Duplicate Targets link to access the Duplicate Targets page. To resolve duplicate target errors, the duplicate target must be renamed on the conflicting Agent. Relocate duplicate targets from Agent "A" to Agent "B".
12. Assuming the original OMS host is no longer in use, remove the Host target (including all remaining monitored targets) from Cloud Control by selecting the host on the Targets > Hosts page and clicking 'Remove'. You will be presented with an error that informs you to remove all monitored targets first. Remove those targets then repeat the step to remove the Host target successfully.
13. Verify that the site is fully operational.

36.1.4.6 Multiple OMS, SLB configured, additional OMS recovered on same or different host

Multi OMS site. OMSs fronted by SLB. OMS configuration backed up using emctl exportconfig oms command on the first OMS. Additional OMS is lost and needs to be recovered on the same or a different host.

1. If recovering to the same host, ensure cleanup of the failed OMS has been peformed:

   Make sure there are no processes still running from the Middleware home using a command similar to the following:

   ```
   ps -ef | grep -i -P "(Middleware|gc_inst)" | grep -v grep | awk '{print $2}' |
   xargs kill -9
   ```

   If recovering the software homes using the software only install method, first de-install the existing Oracle Homes using the Cloud Control software distribution installer. This is required even if the software homes are no longer available as it is necessary to remove any record of the lost Oracle Homes from the Oracle inventory.

   If they exist, remove the 'Middleware' and 'gc_inst' directories.
2. Ensure that shared software library locations are accessible.
3. Install an Agent on the required host (same or different as the case may be).
4. Use the Additional OMS deployment procedure to configure a new additional OMS.
5. Verify that the site is fully operational.

36.1.5 Recovering Agents

If an Agent is lost, it should be reinstalled by cloning from a reference install. Cloning from a reference install is often the fastest way to recover an Agent install as it is not necessary to track and reapply customizations and patches. Care should be taken to reinstall the Agent using the same port. Using the Enterprise Manager's Agent Resynchronization feature, a reinstalled Agent can be reconfigured using target information present in the repository. When the Agent is reinstalled using the same

port, the OMS detects that it has been re-installed and blocks it temporarily to prevent the auto-discovered targets in the re-installed Agent from overwriting previous customizations.

> **Blocked Agents:** A Blocked Agent is a condition where the OMS rejects all heartbeat or upload requests from the blocked Agent. Hence, a blocked Agent will not be able to upload any alerts or metric data to the OMS. However, blocked Agents continue to collect monitoring data.

The Agent can be resynchronized and unblocked from the Agent homepage by clicking on the **Resynchronize Agent** button. Resynchronization pushes all targets from the repository to the Agent and then unblocks the Agent.

36.1.6 Agent Recovery Scenarios

The following scenarios illustrate various Agent recovery situations along with the recovery steps. Agent recovery is supported for Agent versions 10.2.0.5 and later. The Agent resynchronization feature requires that a reinstalled Agent use the same port as the previous Agent that crashed.

36.1.6.1 Agent Reinstall Using the Same Port

An Agent is monitoring multiple targets. The Agent installation is lost.

1. Deinstall Agent OracleHome using the Oracle Universal Installer.

   > **Note:** This step is necessary in order to clean up the inventory.

2. Install a new Agent or use the Agent clone option to reinstall the Agent though Enterprise Manager. Specify the same port that was used by the crashed Agent. The location of the install need not be same as the previous install.

   The OMS detects that the Agent has been re-installed and blocks the Agent.

3. Initiate Agent Resynchronization from the Agent homepage.

   All targets in the repository are pushed to the new Agent. The Agent is instructed to clear backlogged files and then do a clearstate. The Agent is then unblocked.

4. Reconfigure User-defined Metrics if the location of User-defined Metric scripts have changed.

5. Verify that the Agent is operational and all target configurations have been restored restored using the following emctl commands:

   ```
   emctl status agent
   emctl upload agent
   ```

   There should be no errors and no XML files in the backlog.

36.1.6.2 Agent Restore from Filesystem Backup

An Agent is monitoring multiple targets. File system backup for the Agent Oracle Home exists. The Agent install is lost.

1. Deinstall Agent OracleHome using OUI.

Note: This step is necessary in order to clean up the inventory.

2. Restore the Agent from the filesystem backup then start the Agent.

   The OMS detects that the Agent has been restored from backup and blocks the Agent.

3. Initiate Agent Resynchronization from the Agent homepage.

   All targets in the repository are pushed to the new Agent. The Agent is instructed to clear backlogged files and performs a clearstate. The Agent is unblocked.

4. Verify that the Agent is functional and all target configurations have been restored using the following emctl commands:

   ```
   emctl status agent
   emctl upload agent
   ```

   There should be no errors and no XML files in the backlog.

36.2 Recovering from a Simultaneous OMS-Repository Failure

When both OMS and repository fail simultaneously, the recovery situation becomes more complex depending upon factors such as whether the OMS and repository recovery has to be performed on the same or different host, or whether there are multiple OMSs fronted by an SLB. In general, the order of recovery for this type of compound failure should be repository first, followed by OMS(s) following the steps outlined in the appropriate recovery scenarios discussed earlier. The following scenarios illustrate two OMS-Repository failures and the requisite recovery steps.

36.2.1 Collapsed Configuration: Incomplete Repository Recovery, Primary OMS on the Same Host

Repository and the primary OMS are installed on same host (host "A"). The repository database is running in noarchivelog mode. Full cold backup is available. A recent OMS backup file exists (emctl exportconfig oms). The repository, OMS and the Agent crash.

1. Follow the repository recovery procedure shown in Incomplete Recovery on the Same Host with the following exception:

   Since the OMS OracleHome is not available and repository resynchronization has to be initiated before starting an OMS against the restored repository, submit "resync" via the following PL/SQL block. Log into the repository as SYSMAN using SQLplus and run:

   ```
   begin emd_maintenance.full_repository_resync('<resync name>'); end;
   ```

2. Follow the OMS recovery procedure shown in Section 1.2.2.1, "Single OMS, No Server Load Balancer (SLB), OMS Restored on the same Host"

3. Verify that the site is fully operational.

36.2.2 Distributed Configuration: Incomplete Repository Recovery, Primary OMS and additional OMS on Different Hosts, SLB Configured

The Repository, primary OMS, and additional OMS all reside on the different hosts. Repository database was running in noarchivelog mode. OMS backup file from a

recent backup exists (emctl exportconfig oms). Full cold backup of the database exists. All three hosts are lost.

1. Follow the repository recovery procedure shown in Incomplete Recovery on the Same Host" with the following exception:

   Since OMS OracleHome is not yet available and Repository resync has to be initiated before starting an OMS against the restored repository, submit resync via the following PL/SQL block. Log into the repository as SYSMAN using SQLplus and run the following:

   ```
   begin emd_maintenance.full_repository_resync('resync name'); end;
   ```

2. Follow the OMS recovery procedure shown in "Multiple OMS, Server Load Balancer configured, Primary OMS Recovered on a Different Host" with the following exception:

   Override the repository connect description present in the backup file by passing the additional omsca parameter:

   ```
   -REPOS_CONN_STR <restored repos descriptor>
   ```

   This needs to be added along with other parameters listed in Multiple OMS, Server Load Balancer configured, Primary OMS Recovered on a Different Host.

3. Follow the OMS recovery procedure shown in *Multiple OMS, SLB configured, additional OMS recovered on same or different host*.

4. Verify that the site is fully operational.

36.3 Switchover

Switchover is a planned activity where operations are transferred from the Primary site to a Standby site. This is usually done for testing and validation of Disaster Recovery (DR) scenarios and for planned maintenance activities on the primary infrastructure. This section describes the steps to switchover to the standby site. The same procedure is applied to switchover in either direction.

Enterprise Manager Console cannot be used to perform switchover of the Management Repository database. Use the Data Guard Broker command line tool DGMGRL instead.

1. Prepare the Standby Database

   Verify that recovery is up-to-date. Using the Enterprise Manager Console, you can view the value of the ApplyLag column for the standby database in the Standby Databases section of the Data Guard Overview Page.

2. Shut down the Primary Enterprise Manager Application Tier.

   Shutdown all the Management Service instances in the primary site by running the following command on each Management Service:

   ```
   emctl stop oms -all
   ```

3. Verify Shared Loader Directory / Software Library Availability

   Ensure all files from the primary site are available on the standby site.

4. Switch over to the Standby Database

   Use DGMGRL to perform a switchover to the standby database. The command can be run on the primary site or the standby site. The switchover command verifies the states of the primary database and the standby database, affects

switchover of roles, restarts the old primary database, and sets it up as the new standby database.

```
SWITCHOVER TO <standby database name>;
```

Verify the post switchover states. To monitor a standby database completely, the user monitoring the database must have SYSDBA privileges. This privilege is required because the standby database is in a mounted-only state. A best practice is to ensure that the users monitoring the primary and standby databases have SYSDBA privileges for both databases.

```
SHOW CONFIGURATION;
SHOW DATABASE <primary database name>;
SHOW DATABASE <standby database name>;
```

5. Make the standby Management Services point to the Standby Database which is now the new Primary by running the following on each standby Management Service.

   ```
   emctl config oms -store_repos_details -repos_conndesc
   <connect descriptor of new primary database> -repos_user
   sysman
   ```

6. Startup the Enterprise Manager Application Tier

   Startup all the Management Services on the standby site:

   ```
   emctl start oms
   ```

7. Relocate Management Services and Management Repository target

   The Management Services and Management Repository target is monitored by a Management Agent on one of the Management Services on the primary site. To ensure that the target is monitored after switchover/failover, relocate the target to a Management Agent on the standby site by running the following command on one of the Management Service standby sites.

   ```
   emctl config emrep -agent <agent name> -conn_desc
   ```

8. Switchover to Standby SLB.

   Make appropriate network changes to failover your primary SLB to standby SLB that is, all requests should now be served by the standby SLB without requiring any changes on the clients (browser and Management Agents).

9. Establish the old primary Management Services as the new standby Management Services to complete the switchover process.

   Start the Administration Server on old primary site

   ```
   emctl start oms -admin_only
   ```

   Point the old primary Management Services to the new Primary Repository database by running the following command on each Management Service on the old primary site.

   ```
   emctl config oms -store_repos_details -repos_conndesc
   <connect descriptor of new primary database> -repos_user
   sysman
   ```

This completes the switchover operation. Access and test the application to ensure that the site is fully operational and functionally equivalent to the primary site.

Repeat the same procedure to switchover in the other direction.

Keeping the Standby Site in Sync

After the initial setup of the standby site, the standby site has to be kept in sync with the changes done on primary site. Transactions on the Primary Management Repository get propagated to the Standby Management Repository automatically through Data Guard but the OMS side changes have to be redone manually on the standby site. The following sections describe this procedure for typical activities.

Applying patches

When patches are applied on the primary site Management Services, they have to be applied on the standby site Management Services too. Note that patches typically update the Oracle Homes (via the opatch apply command) and optionally might require scripts to be run against the Management Repository. On the standby site, it is sufficient to update the Oracle Homes (via the opatch apply command) and skip the running of scripts on the Management Repository because database changes are automatically propagated to the standby site using Data Guard.

Managing Plugins

When new Plugins are deployed on the Primary Site or existing Plugins upgraded or un-deployed on the Primary Site, the following procedures needs to be run on the standby site too to keep the Standby Management Services in sync. Note if the Standby Management Services are not kept in sync, they would fail to start when a switchover or failover operation is attempted.

The procedure below assume that the standby site was setup as per the documented process and the standby management services are currently down and point to the standby repository. The plugin(s) deployment on the Primary site has been completed successfully.

Deploying a new Plugin or upgrading a Plugin on Standby Site

1. Extract the Plugin archives from the Primary site

   Go to the Self Update Home, click on Plugins, select the required plugin and select export from the Action table menu. Note the emcli command from the popup that gets displayed.

   ```
   emcli export_update -id=<update id> -deep -host=<standby OMS host>
   -dir=<directory to export archives> <host credential options>
   ```

   Note that an additional option "-deep" is required. This command would create 4 files on the destination directory specified. The filename <version>_OMS_ <platform>_<revision>.zip is the one to be used in next step.

2. Startup the Standby Administration Server, if it is down.

   ```
   emctl start oms -admin_only
   ```

3. Install the OMS archive to First Standby OMS Oracle Home

   ```
   pluginia –archives <path to plugin archive>
   ```

4. Configure the Plug-in on First Standby OMS Oracle Home

   ```
   pluginca -action deploy -isFirstOMS true -plugins <plugin-list> -oracleHome
   <oms oracle home> -middlewareHome <wls middleware home>
   ```

   where <plugin-list> is the plugin name in the format <plugin-id>=<plugin-version>

5. Repeat steps 3 and 4 for each Standby additional OMS

```
pluginia -archives <path to plugin archive>
pluginca -action deploy -isFirstOMS false -plugins <plugin-list> -oracleHome
<oms oracle home> -middlewareHome <wls middleware home>
```

This completes the plugin deployment on Standby site.

36.4 Failover

A standby database can be converted to a primary database when the original primary database fails and there is no possibility of recovering the primary database in a timely manner. This is known as a manual failover. There may or may not be data loss depending upon whether your primary and target standby databases were synchronized at the time of the primary database failure.

This section describes the steps to failover to a standby database, recover the Enterprise Manager application state by resynchronizing the Management Repository database with all Management Agents, and enabling the original primary database as a standby using flashback database.

The word *manual* is used here to contrast this type of failover with a fast-start failover described later in Section 1.7, "Automatic Failover".

1. Verify Shared Loader Directory and Software Library Availability

   Ensure all files from the primary site are available on the standby site.

2. Failover to Standby Database.

   Shutdown the database on the primary site. Use DGMGRL to connect to the standby database and execute the FAILOVER command:

   ```
   FAILOVER TO <standby database name>;
   ```

   Verify the post failover states:

   ```
   SHOW CONFIGURATION;
   SHOW DATABASE <primary database name>;
   SHOW DATABASE <standby database name>;
   ```

   Note that after the failover completes, the original primary database cannot be used as a standby database of the new primary database unless it is re-enabled.

3. Make the standby Management Services point to the Standby Database which is now the new Primary by running the following on each standby Management Service.

   ```
   emctl config oms -store_repos_details -repos_conndesc <connect descriptor of
   new primary database> -repos_user sysman
   ```

4. Resync the New Primary Database with Management Agents.

   Skip this step if you are running in Data Guard Maximum Protection or Maximum Availability level as there is no data loss on failover. However, if there is data loss, synchronize the new primary database with all Management Agents.

   On any one Management Service on the standby site, run the following command:

   ```
   emctl resync repos -full -name "<name for recovery action>"
   ```

   This command submits a resync job that would be executed on each Management Agent when the Management Services on the standby site are brought up.

   Repository resynchronization is a resource intensive operation. A well tuned Management Repository will help significantly to complete the operation as

quickly as possible. Specifically if you are not routinely coalescing the IOTs/indexes associated with Advanced Queueing tables as described in My Oracle Support note 271855.1, running the procedure before resync will significantly help the resync operation to complete faster.

5. Startup the Enterprise Manager Application Tier

   Startup all the Management Services on the standby site by running the following command on each Management Service.

   ```
   emctl start oms
   ```

6. Relocate Management Services and Management Repository target.

   The Management Services and Management Repository target is monitored by a Management Agent on one of the Management Services on the primary site. To ensure that target is monitored after switchover/failover, relocate the target to a Management Agent on the standby site by running the following command on one of the standby site Management Service.

   ```
   emctl config emrep -agent <agent name> -conn_desc
   ```

7. Switchover to the Standby SLB.

   Make appropriate network changes to failover your primary SLB to the standby SLB, that is, all requests should now be served by the standby SLB without requiring any changes on the clients (browser and Management Agents).

8. Establish Original Primary Database as Standby Database Using Flashback

   Once access to the failed site is restored and if you had flashback database enabled, you can reinstate the original primary database as a physical standby of the new primary database.

   1. Shutdown all the Management Services in original primary site.

      ```
      emctl stop oms -all
      ```

   2. Restart the original primary database in mount state:

      ```
      shutdown immediate;

      startup mount;
      ```

   3. Reinstate the Original Primary Database

      Use DGMGRL to connect to the old primary database and execute the REINSTATE command

      ```
      REINSTATE DATABASE <old primary database name>;
      ```

   4. The newly reinstated standby database will begin serving as standby database to the new primary database.

   5. Verify the post reinstate states.

      ```
      SHOW CONFIGURATION;

      SHOW DATABASE <primary database name>;

      SHOW DATABASE <standby database name>;
      ```

9. Establish Original Primary Management Service as the standby Management Service.

   Start the Administration Server on old primary site

   ```
   emctl start oms -admin_only
   ```

Point the old primary Management Service to the new Primary Repository database by running the following command on each Management Service on the old primary site.

```
emctl config oms -store_repos_details -repos_conndesc <connect descriptor of
new primary database> -repos_user sysman
```

10. Monitor and complete Repository Resynchronization

    Navigate to the Management Services and Repository Overview page of Cloud Control Console. Under Related Links, click Repository Synchronization. This page shows the progress of the resynchronization operation on a per Management Agent basis. Monitor the progress.

    Operations that fail should be resubmitted manually from this page after fixing the error mentioned. Typically, communication related errors are caused by Management Agents being down and can be fixed by resubmitting the operation from this page after restarting the Management Agent.

    For Management Agents that cannot be started due to some reason, for example, old decommissioned Management Agents, the operation should be stopped manually from this page. Resynchronization is deemed complete when all the jobs have a completed or stopped status.

This completes the failover operation. Access and test the application to ensure that the site is fully operational and functionally equivalent to the primary site.

Perform a switchover procedure if the site operations have to be moved back to the original primary site.

36.5 Automatic Failover

This section details the steps to achieve complete automation of failure detection and failover procedure by utilizing Fast-Start Failover and Observer process. At a high level the process works as follows:

- Fast-Start Failover (FSFO) determines that a failover is necessary and initiates a failover to the standby database automatically
- When the database failover has completed the DB_ROLE_CHANGE database event is fired
- The event causes a trigger to be fired which calls a script that configures and starts Enterprise Manager Application Tier

Perform the following steps:

1. Develop Enterprise Manager Application Tier Configuration and Startup Script

   Develop a script that will automate the Enterprise Manager Application configuration and startup process. See the sample shipped with Cloud Control in the OH/sysman/ha directory. A sample script for the standby site is included here and should be customized as needed. Make sure ssh equivalence is setup so that remote shell scripts can be executed without password prompts. Place the script in a location accessible from the standby database host. Place a similar script on the primary site.

   ```
   #!/bin/sh
   # Script: /scratch/EMSBY_start.sh
   # Primary Site Hosts
   # Repos: earth, OMS: jupiter1, jupiter2
   # Standby Site Hosts
   ```

```
# Repos: mars, # OMS: saturn1, saturn2
LOGFILE="/net/mars/em/failover/em_failover.log"
OMS_ORACLE_HOME="/scratch/OracleHomes/em/oms11"
CENTRAL_AGENT="saturn1.example.com:3872"

#log message
echo "##############################" >> $LOGFILE
date >> $LOGFILE
echo $OMS_ORACLE_HOME >> $LOGFILE
id >>  $LOGFILE 2>&1

#switch all OMS to point to new primary and startup all OMS
ssh orausr@saturn1 "$OMS_ORACLE_HOME/bin/emctl oms -store_repos_details -repos_
conndesc <connect descriptor of new primary database> -repos_user sysman
-repos_pwd <password>" >> $LOGFILE 2>&1
ssh orausr@saturn1 "$OMS_ORACLE_HOME/bin/emctl start oms" >>  $LOGFILE 2>&1

#Repeat the above two lines for each OMS in a multiple OMS setup. E.g.
ssh orausr@saturn2 "$OMS_ORACLE_HOME/bin/emctl oms -store_repos_details -repos_
conndesc <connect descriptor of new primary database> -repos_user sysman
-repos_pwd <password>" >> $LOGFILE 2>&1
ssh orausr@saturn2 "$OMS_ORACLE_HOME/bin/emctl start oms" >>  $LOGFILE 2>&1

#relocate Management Services and Repository target
#to be done only once in a multiple OMS setup
#allow time for OMS to be fully initialized
ssh orausr@saturn1 "$OMS_ORACLE_HOME/bin/emctl config emrep -agent $CENTRAL_
AGENT -conn_desc -sysman_pwd <password>" >> $LOGFILE 2>&1

#always return 0 so that dbms scheduler job completes successfully
exit 0
```

2. Automate Execution of Script by Trigger

Create a database event "DB_ROLE_CHANGE" trigger, which fires after the database role changes from standby to primary. See the sample shipped with Cloud Control in OH/sysman/ha directory.

```
--
--
-- Sample database role change trigger
--
--
CREATE OR REPLACE TRIGGER FAILOVER_EM
AFTER DB_ROLE_CHANGE ON DATABASE
DECLARE
    v_db_unique_name varchar2(30);
    v_db_role varchar2(30);
BEGIN
    select upper(VALUE) into v_db_unique_name
    from v$parameter where NAME='db_unique_name';
    select database_role into v_db_role
    from v$database;

    if v_db_role = 'PRIMARY' then

      -- Submit job to Resync agents with repository
      -- Needed if running in maximum performance mode
      -- and there are chances of data-loss on failover
      -- Uncomment block below if required
      -- begin
```

```
      --  SYSMAN.setemusercontext('SYSMAN', SYSMAN.MGMT_USER.OP_SET_
IDENTIFIER);
      --  SYSMAN.emd_maintenance.full_repository_resync('AUTO-FAILOVER to '||v_
db_unique_name||' - '||systimestamp, true);
      --  SYSMAN.setemusercontext('SYSMAN', SYSMAN.MGMT_USER.OP_CLEAR_
IDENTIFIER);
      -- end;

      -- Start the EM mid-tier
      dbms_scheduler.create_job(
          job_name=>'START_EM',
          job_type=>'executable',
          job_action=> '<location>' || v_db_unique_name|| '_start_oms.sh',
          enabled=>TRUE
      );
    end if;
EXCEPTION
WHEN OTHERS
THEN
    SYSMAN.mgmt_log.log_error('LOGGING', SYSMAN.MGMT_GLOBAL.UNEXPECTED_ERR,
SYSMAN.MGMT_GLOBAL.UNEXPECTED_ERR_M || 'EM_FAILOVER: ' ||SQLERRM);
END;
/
```

Note: Based on your deployment, you might require additional steps to synchronize and automate the failover of SLB and shared storage used for software library. These steps are vendor specific and beyond the scope of this document. One possibility is to invoke these steps from the Enterprise Manager Application Tier startup and configuration script.

3. Configure Fast-Start Failover and Observer.

   Use the Fast-Start Failover configuration wizard in Enterprise Manager Console to enable FSFO and configure the Observer.

   This completes the setup of automatic failover.

36.6 How to Configure Cloud Control OMS in Active/Passive Environment for High Availability Failover Using Virtual Host Names

This section provides a general reference for Cloud Control administrators who want to configure Enterprise Manager 11*g*R1 Cloud Control in Cold Failover Cluster (CFC) environments.

36.6.1 Overview and Requirements

The following conditions must be met for Cloud Control to fail over to a different host:

- The installation must be done using a Virtual Host Name and an associated unique IP address.
- Install on a shared disk/volume which holds the binaries, the configuration and the runtime data (including the recv directory).
- Configuration data and metadata must also failover to the surviving node.

- Inventory location must failover to the surviving node.
- Software owner and time zone parameters must be the same on all cluster member nodes that will host this Oracle Management Service (OMS).

36.6.2 Installation and Configuration

To override the physical host name of the cluster member with a virtual host name, software must be installed using the parameter ORACLE_HOSTNAME. For inventory pointer, the software must be installed using the command line parameter *-invPtrLoc* to point to the shared inventory location file, which includes the path to the shared inventory location.

If you are using an NFS mounted volume for the installation, please ensure that you specify rsize and wsize in your mount command to prevent running into I/O issues.

For example:

```
oms.acme.com:/u01/app/share1 /u01/app/share1 nfs
rw,bg,rsize=32768,wsize=32768,hard,nointr,tcp,noac,vers=3,timeo=
600 0 0
```

Note: Any reference to shared failover volumes could also be true for non-shared failover volumes which can be mounted on active hosts after failover.

36.6.3 Setting Up the Virtual Host Name/Virtual IP Address

You can set up the virtual host name and virtual IP address by either allowing the clusterware to set it up, or manually setting it up yourself before installation and startup of Oracle services. The virtual host name must be static and resolvable consistently on the network. All nodes participating in the setup must resolve the virtual IP address to the same host name. Standard TCP tools such as nslookup and traceroute can be used to verify the host name. Validate using the following commands:

`nslookup` <*virtual hostname*>

This command returns the virtual IP address and full qualified host name.

`nslookup` <*virtual IP*>

This command returns the virtual IP address and fully qualified host name.

Be sure to try these commands on every node of the cluster and verify that the correct information is returned.

36.6.4 Setting Up Shared Storage

Storage can be managed by the clusterware that is in use or you can use any shared file system (FS) volume as long as it is not an unsupported type, such as OCFS V1. The most common shared file system is NFS.

Note: If the OHS directory is on a shared storage, the LockFile directive in the httpd.conf file should be modified to point to a local disk, otherwise there is a potential for locking issues.

36.6.5 Setting Up the Environment

Some operating system versions require specific operating system patches be applied prior to installing 11*g*R1. The user installing and using the 11*g*R1 software must also have sufficient kernel resources available. Refer to the operating system's installation guide for more details.

Before you launch the installer, certain environment variables need to be verified. Each of these variables must be identically set for the account installing the software on ALL machines participating in the cluster:

- OS variable TZ

  Time zone setting. You should unset this variable prior to installation

- PERL variables

  Variables such as PERL5LIB should also be unset to avoid association to the incorrect set of PERL libraries

36.6.6 Synchronizing Operating System IDs

The user and group of the software owner should be defined identically on all nodes of the cluster. This can be verified using the 'id' command:

```
$ id -a
uid=550(oracle) gid=50(oinstall) groups=501(dba)
```

36.6.7 Setting Up Shared Inventory

Use the following steps to set up shared inventory:

1. Create your new ORACLE_HOME directory.
2. Create the Oracle Inventory directory under the new oracle home:

   ```
   $ cd <shared oracle home>
   $ mkdir oraInventory
   ```

3. Create the oraInst.loc file. This file contains the Inventory directory path information needed by the Universal Installer.

   a. `vi oraInst.loc`

   b. Enter the path information to the Oracle Inventory directory and specify the group of the software owner as the oinstall user. For example:

   ```
   inventory_loc=/app/oracle/product/11.1/oraInventory
   inst_group=oinstall
   ```

36.6.8 Installing the Software

Refer to the following steps when installing the software:

1. Create the shared disk location on both the nodes for the software binaries
2. Install WebLogic Server. For information on installing WebLogic Server, refer to *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.
3. Point to the inventory location file oraInst.loc (under the ORACLE_BASE in this case), as well as specifying the host name of the virtual group. For example:

   ```
   $ export ORACLE_HOSTNAME=lxdb.acme.com
   ```

```
$ runInstaller -invPtrloc /app/oracle/share1/oraInst.loc
ORACLE_HOSTNAME=lxdb.acme.com -debug
```

1. Install Oracle Management Services on cluster member *Host1* using the option, "EM install using the existing DB"
2. Continue the remainder of the installation normally.
3. Once completed, copy the files *oraInst.loc* and *oratab* to */etc*. Also copy */opt/oracle* to all cluster member hosts (*Host2*, *Host3*, and so on).

36.6.8.1 Windows Specific Configuration Steps

On Windows environments, an additional step is required to copy over service and keys required by the Oracle software. Note that these steps are required if your clustering software does not provide a shared windows registry.

1. Using regedit on the first host, export each Oracle service from under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services.
2. Using regedit on the first host, export HKEY_LOCAL_MACHINE\SOFTWARE\ORACLE.
3. Use regedit to import the files created in step 1 and 2 to the failover host.

36.6.9 Starting Up Services

Ensure that you start your services in the proper order. Use the order listed below:

1. Establish IP address on the active node
2. Start the TNS listener (if it is part of the same failover group)
3. Start the database (if it is part of the same failover group)
4. Start Cloud Control using `emctl start oms`
5. Test functionality

In case of failover, refer to the following steps:

1. Establish IP on failover box
2. Start TNS listener using the command `lsnrctl start` if it is part of the same failover group
3. Start the database using the command `dbstart` if it is part of the same failover group
4. Start Cloud Control using the command `emctl start oms`
5. Test the functionality

36.6.10 Summary

The OMS mid-tier component of Cloud Control can now be deployed in a CFC environments that utilize a floating host name.

36.7 Configuring Targets for Failover in Active/Passive Environments

This section provides a general reference for Cloud Control administrators who want to relocate Cold Failover Cluster (CFC) targets from one existing Management Agent to another. Although the targets are capable of running on multiple nodes, these targets run only on the active node in a CFC environment.

CFC environments generally use a combination of cluster software to provide a virtual host name and IP address along with interconnected host and storage systems to share information and provide high availability for applications. Automating failover of the virtual host name and IP, in combination with relocating the Enterprise Manager targets and restarting the applications on the passive node, requires the use of Oracle Enterprise Manager command-line interface (EM CLI) and Oracle Clusterware (running Oracle Database release 10*g* or 11*g*) or third-party cluster software. Several Oracle partner vendors provide clusterware solutions in this area.

The Enterprise Manager Command Line Interface (EM CLI) allows you to access Enterprise Manager Cloud Control functionality from text-based consoles (terminal sessions) for a variety of operating systems. Using EM CLI, you can perform Enterprise Manager Cloud Control console-based operations, like monitoring and managing targets, jobs, groups, blackouts, notifications, and alerts. See the *Oracle Enterprise Manager Command Line Interface* manual for more information.

36.7.1 Target Relocation in Active/Passive Environments

Beginning with Oracle Enterprise Manager 10g release 10.2.0.5, a single Oracle Management Agent running on each node in the cluster can monitor targets configured for active / passive high availability. Only one Management Agent is required on each of the physical nodes of the CFC cluster because, in case of a failover to the passive node, Enterprise Manager can move the HA monitored targets from the Management Agent on the failed node to another Management Agent on the newly activated node using a series of EMCLI commands. See the *Oracle® Enterprise Manager Command Line Interface* manual for more information.

If your application is running in an active/passive environment, the clusterware brings up the applications on the passive node in the event that the active node fails. For Enterprise Manager to continue monitoring the targets in this type of configuration, the existing Management Agent needs additional configuration.

The following sections describe how to prepare the environment to automate and restart targets on the new active node. Failover and fallback procedures are also provided.

36.7.2 Installation and Configuration

The following sections describe how to configure Enterprise Manager to support a CFC configuration using the existing Management Agents communicating with the Oracle Management Service processes:

- Prerequisites
- Configuration Steps

36.7.2.1 Prerequisites

Prepare the Active/Passive environments as follows:

- Ensure the operating system clock is synchronized across all nodes of the cluster. (Consider using Network Time Protocol (NTP) or another network synchronization method.)
- Use the EM CLI RELOCATE_TARGETS command only with Enterprise Manager Release 10.2.0.5 (and higher) Management Agents.

36.7.2.2 Configuration Steps

The following steps show how to configure Enterprise Manager to support a CFC configuration using the existing Management Agents that are communicating with the OMS processes. The example that follows is based on a configuration with a two-node cluster that has one failover group. For additional information about targets running in CFC active/passive environments, see My Oracle Support note 406014.1.

1. Configure EM CLI

   To set up and configure target relocation, use the Oracle Enterprise Manager command-line interface (EM CLI). See the *Oracle Enterprise Manager Command Line Interface* manual and the *Oracle Enterprise Manager Extensibility* manual for information about EM CLI and Management Plug-Ins.

2. Install Management Agents

   Install the Management Agent on a local disk volume on each node in the cluster. Once installed, the Management Agents are visible in the Cloud Control console.

3. Discover Targets

   After the Active / Passive targets have been configured, use the Management Agent discovery screen in the Cloud Control console to add the targets (such as database, listener, application server, and so on). Perform the discovery on the active node, which is the node that is currently hosting the new target.

36.7.3 Failover Procedure

To speed relocation of targets after a node failover, configure the following steps using a script that contains the commands necessary to automatically initiate a failover of a target. Typically, the clusterware software has a mechanism with which you can automatically execute the script to relocate the targets in Enterprise Manager. Also, see Section 1.9.6, "Script Examples" for sample scripts.

1. Shut down the target services on the failed active node.

   On the active node where the targets are running, shut down the target services running on the virtual IP.

2. If required, disconnect the storage for this target on the active node.

   Shut down all the applications running on the virtual IP and shared storage.

3. Enable the target's IP address on the new active node.

4. If required, connect storage for the target on the currently active node.

5. Relocate the targets in Cloud Control using EM CLI.

   To relocate the targets to the Management Agent on the new active node, issue the EM CLI RELOCATE TARGET command for each target type (listener, application servers, and so on) that you must relocate after the failover operation. For example:

```
emcli relocate_targets
-src_agent=<node 1>:3872
-dest_agent=<node 2>:3872
-target_name=<database_name>
-target_type=oracle_database
-copy_from_src
-force=yes
```

In the example, port 3872 is the default port for the Management Agent. To find the appropriate port number for your configuration, use the value for the EMD_URL parameter in the emd.properties file for this Management Agent.

Note: In case of a failover event, the source agent will not be running. However, there is no need to have the source Management Agent running to accomplish the RELOCATE operation. EM CLI is an OMS client that performs its RELOCATE operations directly against the Management Repository.

36.7.4 Fallback Procedure

To return the HA targets to the original active node or to any other cluster member node:

1. Repeat the steps in Section 1.9.3, "Failover Procedure" to return the HA targets to the active node.
2. Verify the target status in the Cloud Control console.

36.7.5 EM CLI Parameter Reference

Issue the same command for each target type that will be failed over to (or be switched over) during relocation operations. For example, issue the same EM CLI command to relocate the listener, the application servers, and so on. Table 1–1 shows the EM CLI parameters you use to relocate targets:

Table 36–1 EM CLI Parameters

EM CLI Parameter	Description
-src_agent	Management Agent on which the target was running before the failover occurred.
-dest_agent	Management Agent that will be monitoring the target after the failover.
-target_name	Name of the target to be failed over.
-target_type	Type of target to be failed over (internal Enterprise Manager target type). For example, the Oracle database (for a standalone database or an Oracle RAC instance), the Oracle listener for a database listener, and so on.
-copy_from_src	Use the same type of properties from the source Management Agent to identify the target. This is a **MANDATORY** parameter! If you do not supply this parameter, you can corrupt your target definition!
-force	Force dependencies (if needed) to failover as well.

36.7.6 Script Examples

The following sections provide script examples:

- Relocation Script
- Start Listener Script
- Stop Listener Script

36.7.6.1 Relocation Script

```
#! /bin/ksh
```

```
#get the status of the targets

emcli get_targets -
targets="db1:oracle_database;listener_db1:oracle_listener" -noheader

  if [[ $? != 0 ]]; then exit 1; fi

# blackout the targets to stop false errors.  This blackout is set to expire in 30
minutes

emcli create_blackout -name="relocating active passive test targets" -
add_targets="db1:oracle_database;listener_db1:oracle_listener" -
reason="testing failover" -
schedule="frequency:once;duration:0:30"
  if [[ $? != 0 ]]; then exit 1; fi

# stop the listener target.  Have to go out to a OS script to use the 'lsnrctl set
current_listener' function

emcli execute_hostcmd -cmd="/bin/ksh" -osscript="FILE" -
input_file="FILE:/scratch/oraha/cfc_test/listener_stop.ksh" -
credential_set_name="HostCredsNormal" -
targets="host1.us.oracle.com:host"
  if [[ $? != 0 ]]; then exit 1; fi

# now, stop the database

emcli execute_sql -sql="shutdown abort" -
targets="db1:oracle_database" -
credential_set_name="DBCredsSYSDBA"
  if [[ $? != 0 ]]; then exit 1; fi

# relocate the targets to the new host

emcli relocate_targets -
src_agent=host1.us.oracle.com:3872 -
dest_agent=host2.us.oracle.com:3872 -
target_name=db1 -target_type=oracle_database -
copy_from_src -force=yes  -
changed_param=MachineName:host1vip.us.oracle.com
  if [[ $? != 0 ]]; then exit 1; fi

emcli relocate_targets -
src_agent=host1.us.oracle.com:3872 -
dest_agent=host2.us.oracle.com:3872 -
target_name=listener_db1 -target_type=oracle_listener -
copy_from_src -force=yes  -
changed_param=MachineName:host1vip.us.oracle.com
  if [[ $? != 0 ]]; then exit 1; fi

# Now, restart database and listener on the new host

emcli execute_hostcmd -cmd="/bin/ksh" -osscript="FILE" -
input_file="FILE:/scratch/oraha/cfc_test/listener_start.ksh" -
credential_set_name="HostCredsNormal" -
targets="host2.us.oracle.com:host"
  if [[ $? != 0 ]]; then exit 1; fi

emcli execute_sql -sql="startup" -
targets="db1:oracle_database" -
```

```
credential_set_name="DBCredsSYSDBA"
  if [[ $? != 0 ]]; then exit 1; fi

# Time to end the blackout and let the targets become visible

emcli stop_blackout -name="relocating active passive test targets"
  if [[ $? != 0 ]]; then exit 1; fi

# and finally, recheck the status of the targets

emcli get_targets -
targets="db1:oracle_database;listener_db1:oracle_listener" -noheader
  if [[ $? != 0 ]]; then exit 1; fi
```

36.7.6.2 Start Listener Script

```
#!/bin/ksh

export
ORACLE_HOME=/oradbshare/app/oracle/product/11.1.0/db
export PATH=$ORACLE_HOME/bin:$PATH

lsnrctl << EOF
set current_listener listener_db1
start
exit
EOF
```

36.7.6.3 Stop Listener Script

```
#!/bin/ksh
export
ORACLE_HOME=/oradbshare/app/oracle/product/11.1.0/db
export PATH=$ORACLE_HOME/bin:$PATH

lsnrctl << EOF
set current_listener listener_db1
stop
exit
EOF
```

Index

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

R

S

T

U

V

W

X