Oracle® Enterprise Manager

Monitoring Guide

24ai Release 1 (24.1)

F97211-04

May 2025

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

F97211-04

Copyright © 2022, 2025, Oracle and/or its affiliates.

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

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

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

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

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

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

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

Contents

Preface

Part I Enterprise Monitoring

Part II Discovery

1 Discovering and Adding Host and Non-Host Targets

2 Discovering and Adding Database Targets

3 Discovering and Adding Middleware Targets

Part III Monitoring and Managing Targets

5 Using Incident Management

10 Using Administration Groups

15 Utilizing the Job System and Corrective Actions

21 Managing Storage

26 Managing Oracle SuperCluster

27 Monitoring Oracle Operating Systems

29 Monitoring Oracle VM Server for SPARC

30 Provisioning Zones with Oracle Database on Database Domains

Part VI Appendix

Index

Preface

This guide describes how to use Oracle Enterprise Manager Cloud Control 13*c* core functionality.

The preface covers the following:

- Audience
- Documentation Accessibility
- Related Documents
- Conventions

Audience

This document is intended for Enterprise Manager administrators and developers who want to manage their Enterprise Manager infrastructure.

Documentation Accessibility

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

Access to Oracle Support

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

Related Documents

For the latest releases of Enterprise Manager Cloud Control and other Oracle documentation, see:

`http://docs.oracle.com/en/enterprise-manager/`

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.

Convention	Meaning
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

Enterprise Monitoring

This chapter covers the following topics:

- Monitoring Overview
- Monitoring: Basics
- Monitoring: Advanced Setup
- Notifications
- Managing Events, Incidents, and Problems
- Accessing Monitoring Information

Monitoring Overview

Enterprise Manager monitoring functionality permits unattended monitoring of your IT environment. Enterprise Manager comes with a comprehensive set of performance and health metrics that allows monitoring of key components in your environment, such as applications, application servers, databases, as well as the back-end components on which they rely (such as hosts, operating systems, storage).

The Management Agent on each monitored host monitors the status, health, and performance of all managed components (targets) on that host. If a target goes down, or if a performance metric crosses a warning or critical threshold, an event is triggered and sent to Enterprise Manager. Administrators or any interested party can be notified of the triggered event through the Enterprise Manager notification system.

Adding targets to monitor is simple. Enterprise Manager provides you with the option of either adding targets manually or automatically discovering all targets on a host. Enterprise Manager can also automatically and intelligently apply monitoring settings for newly added targets. For more information, see Administration Groups and Template Collections). While Enterprise Manager provides a comprehensive set of metrics used for monitoring, you can also use metric extensions (see Metric Extensions: Customizing Monitoring) to monitor conditions that are specific to your environment. As your data center grows, it will become more challenging to manage individual targets separately, thus you can use Enterprise Manager's group management functionality to organize large sets of targets into groups, allowing you to monitor and manage many targets as one.

Comprehensive Out-of-Box Monitoring

Monitoring begins as soon as you install Enterprise Manager. Enterprise Manager's Management Agents automatically start monitoring their host's systems (including hardware and software configuration data on these hosts) as soon as they are deployed and started. Enterprise Manager provides auto-discovery scripts that enable these Agents to automatically discover all Oracle components and start monitoring them using a comprehensive set of metrics at Oracle-recommended thresholds.

This monitoring functionality includes other components of the Oracle ecosystem such as NetApp Filer, BIG-IP load balancers, Checkpoint Firewall, and IBM WebSphere. Metrics from all monitored components are stored and aggregated in the Management Repository, providing administrators with a rich source of diagnostic information and trend analysis data. When critical alerts are detected, notifications are sent to administrators for rapid resolution.

Out-of-box, Enterprise Manager monitoring functionality provides:

- In-depth monitoring with Oracle-recommended metrics and thresholds.
- Monitoring of all components of your IT infrastructure (Oracle and non-Oracle) as well as the applications and services that are running on them.
- Access to real-time performance charts.
- Collection, storage, and aggregation of metric data in the Management Repository. This allows you to perform strategic tasks such as trend analysis and reporting.
- E-mail and pager notifications for detected critical events.

Enterprise Manager can monitor a wide variety of components (such as databases, hosts, and routers) within your IT infrastructure.

Some examples of monitored metrics are:

- Archive Area Used (Database)
- Component Memory Usage (Application Server)
- Segments Approaching Maximum Extents Count (Database)
- Network Interface Total I/O Rate (Host)

Monitoring Without Management Agents

When it is not practical to have a Management Agent present to monitor specific components of your IT infrastructure, as might be the case with an IP traffic controller or remote Web application, Enterprise Manager provides Extended Network and Critical URL Monitoring functionality. This feature allows the Beacon functionality of the Agent to monitor remote network devices and URLs for availability and responsiveness without requiring an Agent to be physically present on that device. You simply select a specific Beacon, and add key network components and URLs to the *Network and URL Watch Lists*. Enterprise Manager monitoring concepts and the underlying subsystems that support this functionality are discussed in the following sections.

Monitoring: Basics

Enterprise Manager comes with a comprehensive set of predefined performance and health metrics that enables automated monitoring of key components in your environment, such as applications, application servers, databases, as well as the back-end components on which they rely, such as hosts, operating systems, storage. While Enterprise Manager can monitor for many types of conditions (events), the most common use of its monitoring capability centers around the basics of monitoring for violation of acceptable performance boundaries defined by metric values. The following sections discuss the basic concepts and Enterprise Manger functionality that supports monitoring of targets.

Metric Thresholds: Determining When a Monitored Condition is an Issue

Some metrics have associated predefined limiting parameters called thresholds that cause metric alerts (specific type of event) to be triggered when collected metric values exceed these limits. Enterprise Manager allows you to set metric threshold values for two levels of alert severity:

- **Warning** - Attention is required in a particular area, but the area is still functional.
- **Critical** - Immediate action is required in a particular area. The area is either not functional or indicative of imminent problems.

Hence, thresholds are boundary values against which monitored metric values are compared. For example, for each disk device associated with the Disk Utilization (%) metric, you might define a warning threshold at 80% disk space used and critical threshold at 95%.

> **Note:**
>
> Not all metrics need a threshold: If the values do not make sense, or are not needed in a particular environment, they can be removed or simply not set.

While the out-of-box predefined metric threshold values will work for most monitoring conditions, your environment may require that you customize threshold values to more accurately reflect the operational norms of your environment. Setting accurate threshold values, however, may be more challenging for certain categories of metrics such as performance metrics.

For example, what are appropriate warning and critical thresholds for the *Response Time Per Transaction* database metric? For such metrics, it might make more sense to be alerted when the monitored values for the performance metric deviates from normal behavior. Enterprise Manager provides features to enable you to capture normal performance behavior for a target and determine thresholds that are deviations from that performance norm.

> **Note:**
>
> Enterprise Manager administrators must be granted *Manage Target Metrics* or greater privilege on a target in order to perform any metric threshold changes.

Preventing False Alerts:Setting the Number of Occurrences after which an Alert has been Triggered

To prevent false alerts due to spikes in metric values, the *Number of Occurrences* determines the period of time a collected metric value must remain above or below the threshold value before an alert is triggered or cleared. For example, if a metric value is collected every 5 minutes, and the *Number of Occurrences* is set to 6, the metric values (collected successively) must stay above the threshold value for 30 minutes before an alert is triggered. Also, after the alert is triggered, the same metric value needs to stay below its threshold for the same number of occurrences before the alert is cleared. For server-generated alerts, the evaluation frequency is determined by the Oracle Database internals. Refer to the Oracle Database documentation on Oracle Database Server-Generated Alerts for additional details.

Events: Defining What Conditions are of Interest

When a metric threshold value is reached, a metric alert is raised. A metric alert is a type of event. An event is a significant occurrence that indicates a potential problem; for example, either a warning or critical threshold for a monitored metric has been crossed. Other examples of events include: database instance is down, a configuration file has been changed, job executions ended in failure, or a host exceeded a specified percentage CPU utilization. Two of the most important event types used in enterprise monitoring are:

- Metric Alert
- Target Availability

For more information on events and available event types for which you can monitor, see Using Incident Management .

Corrective Actions: Resolving Issues Automatically

Corrective actions allow you to specify automated responses to metric alerts, saving administrator time and ensuring issues are dealt with before they noticeably impact users. For example, if Enterprise Manager detects that a component, such as the SQL*Net listener is down, a corrective action can be specified to automatically start it back up. A corrective action is, therefore, any task you specify that will be executed when a metric triggers a warning or critical alert severity. In addition to performing a corrective task, a corrective action can be used to gather more diagnostic information, if needed. By default, the corrective action runs on the target on which the event has been raised.

A corrective action can also consist of multiple tasks, with each task running on a different target. Administrators can also receive notifications for the success or failure of corrective actions. A corrective action can also consist of multiple tasks, with each task running on a different target.

Corrective actions for a target can be defined by all Enterprise Manager administrators who have been granted *Manage Target Metrics* or greater privilege on the target. For any metric, you can define different corrective actions when the metric triggers at warning severity or at critical severity.

Corrective actions must run using the credentials of a specific Enterprise Manager administrator. For this reason, whenever a corrective action is created or modified, the credentials that the modified action will run with must be specified. You specify these credentials when you associate the corrective action with elements such as incident or event rules.

Metric Extensions: Customizing Monitoring

Metric Extensions let you extend Enterprise Manager's monitoring capabilities to cover conditions specific to your IT environment, thus providing you with a complete and comprehensive view of your monitored environment.

Metric extensions allow you to define new metrics on any target type that utilize the same full set of data collection mechanisms used by Oracle provided metrics. For example, some target types you can create metrics on are:

- Hosts
- Databases
- IBM Websphere
- Oracle Exadata Databases and Storage Servers
- Oracle Business Intelligence Components

Once these new metrics are defined, they are used like any other Enterprise Manager metric. For more information about metric extensions, see Using Metric Extensions .

User-Defined Metrics (Pre-12c)

If you upgraded your Enterprise Manager 12*c* site from an older version of Enterprise Manager, then all user-defined metrics defined in the older version will also be migrated to Enterprise Manager 12*c*. These user-defined metrics will continue to work, however they will no longer be supported a future release. 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. For information about the migration process, see *Converting User-defined Metrics to Metric Extensions* in Using Metric Extensions

Blackouts and Notification Blackouts

Blackouts allow you to support planned outage periods to perform scheduled or emergency maintenance. When a target is put under blackout, monitoring is suspended, thus preventing unnecessary alerts from being sent when you bring down a target for scheduled maintenance operations such as database backup or hardware upgrade. Blackout periods are automatically excluded when calculating a target's overall availability.

A blackout period can be defined for individual targets, a group of targets or for all targets on a host. The blackout can be scheduled to run immediately or in the future, and to run indefinitely or stop after a specific duration. Blackouts can be created on an as-needed basis, or scheduled to run at regular intervals. If, during the maintenance period, you discover that you need more (or less) time to complete maintenance tasks, you can easily extend (or stop) the blackout that is currently in effect. Blackout functionality is available from both the Enterprise Manager console as well as via the Enterprise Manager command-line interface (EM CLI). EM CLI is often useful for administrators who would like to incorporate the blacking out of a target within their maintenance scripts. When a blackout ends, the Management Agent automatically re-evaluates all metrics for the target to provide current status of the target post-blackout.

If an administrator inadvertently performs scheduled maintenance on a target without first putting the target under blackout, these periods would be reflected as target downtime instead of planned blackout periods. This has an adverse impact on the target's availability records. In such cases, Enterprise Manager allows Super Administrators to go back and define the blackout period that should have happened at that time. The ability to create these retroactive blackouts provides Super Administrators with the flexibility to define a more accurate picture of target availability.

Notification Blackouts

Beginning with Enterprise Manager 13*c*, you can stop notifications only. These are called Notification Blackouts and are intended solely for suppressing event notifications on targets. Because the Agent continues to monitor the target during the Notification Blackout duration, the OMS will continue to show the actual target status along with an indication that the target is currently under Notification Blackout.

Monitoring: Advanced Setup

Enterprise Manager greatly simplifies managing your monitored environment and also allows you to customize and extend Enterprise Manager monitoring capabilities. However, the primary advantage Enterprise Manager monitoring provides is the ability to monitor and manage large-scale, heterogeneous environments. Whether you are monitoring an environment with 10 targets or 10,000 targets, the following Enterprise Manager advanced features allow you to implement and maintain your monitored environment with the equal levels of convenience and simplicity.

Monitoring Templates

Monitoring Templates simplify the task of standardizing monitoring settings across your enterprise by allowing you to specify your standards for monitoring in a template once and apply them to monitored targets across your organization. This makes it easy for you to apply specific monitoring settings to specific classes of targets throughout your enterprise. For example, you can define one monitoring template for test databases and another monitoring template for production databases.

A monitoring template defines all Enterprise Manager parameters you would normally set to monitor a target, such as:

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

When a change is made to a template, you can reapply the template across affected targets in order to propagate the new changes. The apply operation can be automated using Administration Groups and Template Collections. For any target, you can preserve custom monitoring settings by specifying metric settings that can never be overwritten by a template.

Enterprise Manager comes with an array of Oracle-certified templates that provide recommended metric settings for various Oracle target types.

For more information about monitoring templates, see Using Monitoring Templates .

Administration Groups and Template Collections

Monitored environments are rarely static—new targets are constantly being added from across your ecosystem. Enterprise Manager allows you to maintain control of this dynamic environment through administration groups. Administration groups automate the process of setting up targets for management in Enterprise Manager by automatically applying management settings such as monitoring settings or compliance standards. Typically, these settings are manually applied to individual targets, or perhaps semi-automatically using monitoring templates (see Monitoring Templates) or custom scripts. Administration groups combine the convenience of applying monitoring settings using monitoring templates with the power of automation.

Template collections contain the monitoring settings and other management settings that are meant to be applied to targets as they join the administration group. Monitoring settings for targets are defined in monitoring templates. Monitoring templates are defined on a per target type basis, so you will need to create monitoring templates for each of the different target types in your administration group. You will most likely create multiple monitoring templates to define the appropriate monitoring settings for an administration group.

Every target added to Enterprise Manager possesses innate attributes called *target properties*. Enterprise Manager uses these target properties to add targets to the correct administration group. Administration group membership is based on target properties as membership criteria so target membership is dynamic. Once added to the administration group, Enterprise Manager automatically applies the requisite monitoring settings using monitoring templates that are part of the associated template collection .

Administration groups use the following target properties to define membership criteria:

- Contact
- Cost Center
- Customer Support Identifier
- Department
- Lifecycle Status
- Line of Business
- Location
- Target Version
- Target Type

Customizing Alert Messages

Whenever a metric threshold is reached, an alert is raised along with a metric-specific message. These messages are written to address generic metric alert conditions. Beginning with Enterprise Manager Release 12.1.0.4, you can customize these messages to suit the specific requirements of your monitored environment.

Customizing an alert message allows you to tailor the message to suit your monitoring needs. You can tailor the message to include their operational context specific to your environment such as IT error codes used in your data center, or add additional information collected by Enterprise Manager such as:

- Metric name for which the alert has been triggered
- Severity level of the alert or violation
- Threshold value for which warning or critical violation has been triggered
- Number of Occurrences after which alert has been triggered

  To prevent false alerts due to spikes in metric values, the *Number of Occurrences* determines the period of time a collected metric value must remain above or below the threshold value before an alert is triggered or cleared. For example, if a metric value is collected every 5 minutes, and the *Number of Occurrences* is set to 6, the metric values (collected successively) must stay above the threshold value for 30 minutes before an alert is triggered. Also, after the alert is triggered, the same metric value needs to stay below its threshold for the same number of occurrences before the alert is cleared. For server-generated alerts, the evaluation frequency is determined by the Oracle Database internals. Refer to the Oracle Database documentation on Oracle Database Server-Generated Alerts for additional details.

Alert message customization allows for more efficient alert management by increasing message usability.

To customize a metric alert message:

1. Navigate to a target homepage.

2. From the *target* menu (host target type is shown in the graphic), select **Monitoring** and then **Metric and Collection Settings**.

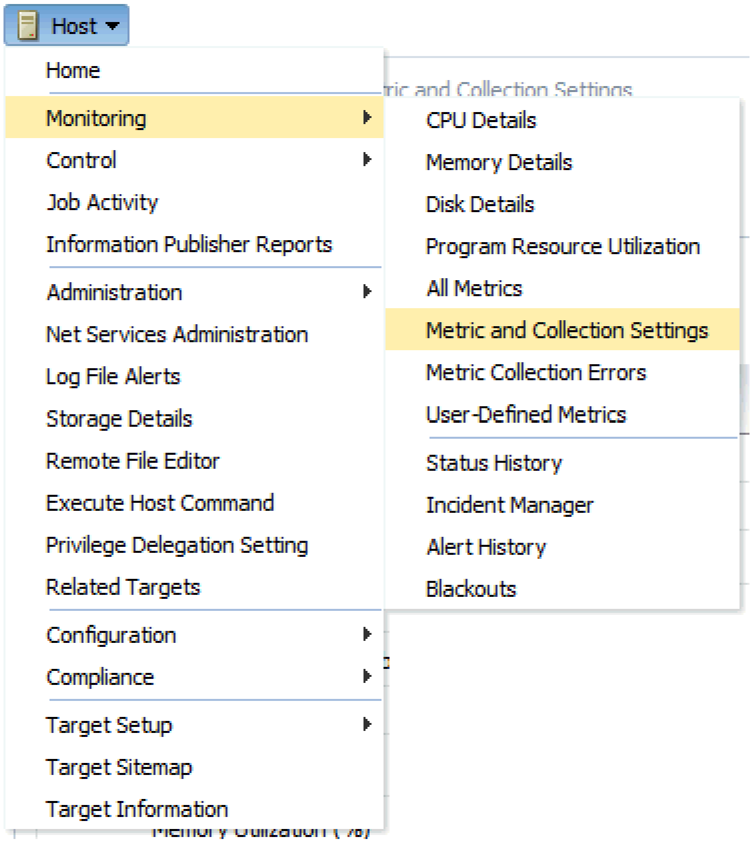

The *Metric and Collection Settings* page displays.

3. In the metric table, find the specific metric whose message you want to change and click the edit icon (pencil).

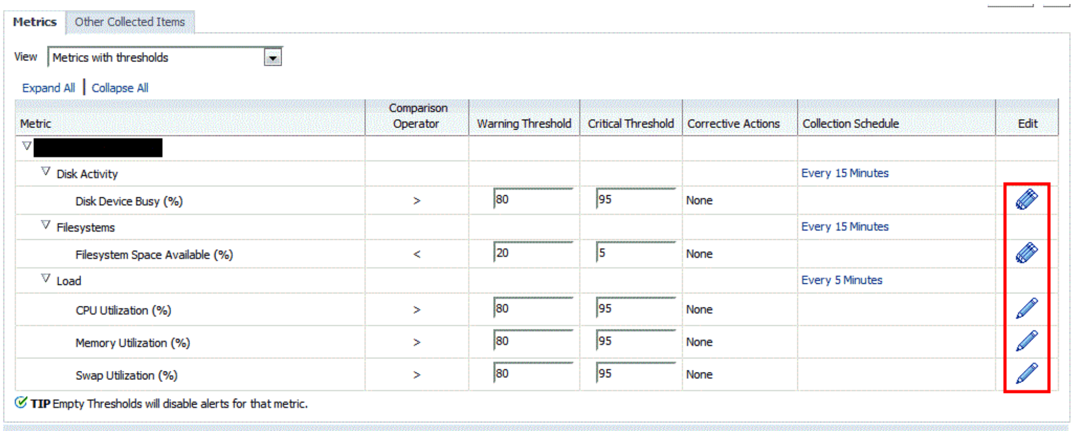

Metrics | Other Collected Items

View: Metrics with thresholds

Expand All | Collapse All

Metric	Comparison Operator	Warning Threshold	Critical Threshold	Corrective Actions	Collection Schedule	Edit
▽ [redacted]						
▽ Disk Activity					Every 15 Minutes	
Disk Device Busy (%)	>	80	95	None		✎
▽ Filesystems					Every 15 Minutes	
Filesystem Space Available (%)	<	20	5	None		✎
▽ Load					Every 5 Minutes	
CPU Utilization (%)	>	80	95	None		✎
Memory Utilization (%)	>	80	95	None		✎
Swap Utilization (%)	>	80	95	None		✎

✓ **TIP** Empty Thresholds will disable alerts for that metric.

The *Edit Advanced Settings* page displays.

4. In the *Monitored Objects* region, click **Edit Alert Message**.

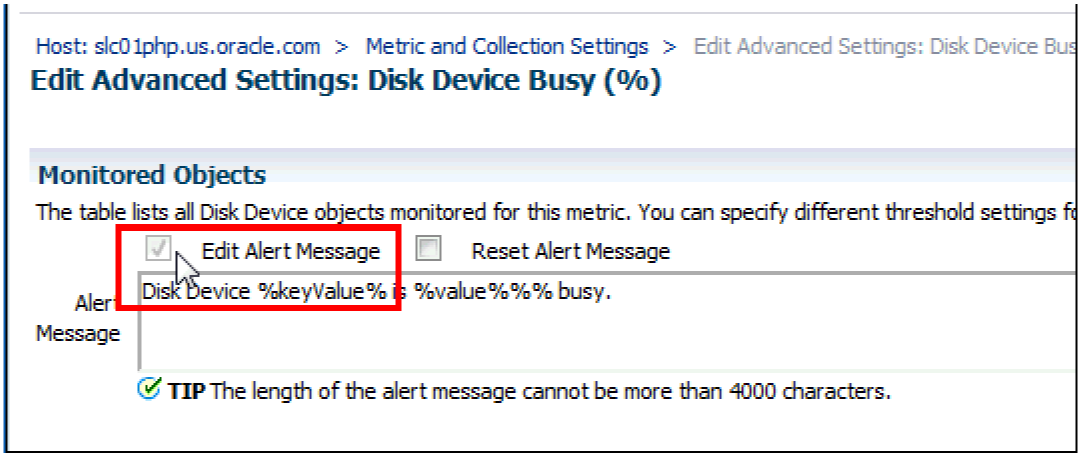

Host: slc01php.us.oracle.com > Metric and Collection Settings > Edit Advanced Settings: Disk Device Bus

Edit Advanced Settings: Disk Device Busy (%)

Monitored Objects

The table lists all Disk Device objects monitored for this metric. You can specify different threshold settings fo

☑ Edit Alert Message □ Reset Alert Message

Alert Message: Disk Device %keyValue% is %value%%% busy.

TIP The length of the alert message cannot be more than 4000 characters.

5. Modify the alert message as appropriate.

> **Note:**
>
> To change your revised message back to the original Oracle-defined message at any time, click **Reset Alert Message**.

6. Click **Continue** to return to the *Metric and Collection Settings* page.
7. To modify additional metric alert messages, repeat steps three through six.
8. Once you are finished, click **OK** to save all changes to the Enterprise Manager Repository. Enterprise Manager will display a message indicating the updates have succeeded.
9. Click **OK** to dismiss the message and return to the target homepage.

Notifications

For a typical monitoring scenario, when a target becomes unavailable or if thresholds for performance are crossed, events are raised and notifications are sent to the appropriate administrators. Enterprise Manager supports notifications via email, pager, SNMP traps, Webhooks, Slack or by running custom scripts and allows administrators to control these notification mechanisms through:

- Notification Methods
- Rules and Rule Sets
- Notification Blackouts

Notification Methods

A notification method represents a specific way to send notifications. Besides e-mail, there are the following types of notification methods: OS Command, PL/SQL, SNMP Traps, Webhooks and Slack. When configuring a notification method, you need to specify the particulars associated with a specific notification mechanism such as which SMTP gateway(s) to use for e-mail or which custom OS script to run. Super Administrators perform a one-time setup of the various types of notification methods available for use.

Rules

A rule instructs Enterprise Manager to take specific action when events or incidents (entity containing one important event or related events) occur, such as notifying an administrator or opening a helpdesk ticket (see Managing Events, Incidents, and Problems). For example, you can define a rule that specifies e-mail should be sent to you when CPU Utilization on any host target is at critical severity, or another rule that notifies an administrator's supervisor if an incident is not acknowledged within 24 hours.

Notification Blackouts

Notification Blackouts allow you to stop notifications while at the same time allowing the Agents to continue monitoring your targets. This allows Enterprise Manager to more accurately collect target availability information. For more information, see "Blackouts and Notification Blackouts."

Customizing Notifications

Notifications that are sent to Administrators can be customized based on message type and on-call schedule. Message customization is useful for administrators who rely on both e-mail and paging systems as a means for receiving notifications. The message formats for these systems typically vary—messages sent to e-mail can be lengthy and can contain URLs, and messages sent to a pager are brief and limited to a finite number of characters. To support these types of mechanisms, Enterprise Manager allows administrators to associate a long or short message format with each e-mail address. E-mail addresses that are used to send regular e-mails can be associated with the *long* format; pages can be associated with the *short* format. The *long* format contains full details about the event/incident; the *short* format contains the most critical pieces of information.

Notifications can also be customized based on an administrator's on-call schedule. An administrator who is on-call might want to be contacted by both his pager and work email address during business hours and only by his pager address during off hours. Enterprise Manager offers a flexible notification schedule to support the wide variety of on-call schedules. Using this schedule, an administrator defines his on-call schedule by specifying the email addresses by which they should be contacted when they are on-call. For periods where they are not on-call, or do not wish to receive notifications for incidents, they simply leave that part of the schedule blank. All alerts that are sent to an administrator automatically adhere to his specified schedule.

Managing Events, Incidents, and Problems

Enterprise Manager's monitoring functionality is built upon the precept of monitoring by exception. This means it monitors and raises events when exception conditions exist in your IT environment and allowing administrators to address them in a timely manner. As discussed earlier, the two most commonly used event types to monitor for are metric alert and target availability. Although these are the most common event types for which Enterprise Manager monitors, there are many others. Available event types include:

- Target Availability
- Metric Alert
- Metric Evaluation Errors
- Job Status Changes
- Compliance Standard Rule Violations
- Compliance Standard Score Violations

- High Availability
- Service Level Agreement Alerts
- User-reported
- JVM Diagnostics Threshold Violation

By definition, an incident is a unit containing a single, or closely correlated set of events that identify an issue that needs administrator attention within your managed environment. So an incident might be as simple as a single event indicating available space in a tablespace has fallen below a specified limit, or more complex such as an incident consisting of multiple events relating to potential performance issue when a server is running out of resources. Such an incident would contain events relating to the usage of CPU, I/O , and memory resources. Managing by incident gives you the ability to address issues that may consist of any number of causal factors. For an in-depth discussion on incidents and events, see Using Incident Management .

Although incidents can correspond to a single events, incidents more commonly correspond to groups of related events. A large number of discrete events can quickly become unmanageable, but handled as an assemblage of related events, incidents allow you to manage large numbers of event occurrences more effectively.

Once an incident is created, Enterprise Manager makes available a rich set of incident management workflow features that let you to manage and track the incident through its complete lifecycle. Incident management features include:

- Assign incident ownership.
- Track the incident resolution status.
- Set incident priority.
- Set incident escalation level.
- Ability to provide a manual summary.
- Ability to add user comments.
- Ability to suppress/unsuppress
- Ability to manually clear the incident.
- Ability to create a ticket manually.

Problems pertain to the diagnostic incidents and problems stored in Automatic Diagnostic Repository (ADR), which are automatically raised by Oracle software when it encounters critical errors in the software. When problems are raised for Oracle software, Oracle has determined that the recommended recourse is to open a Service Request (SR), send support the diagnostic logs, and eventually provide a solution from Oracle. A problem represents the underlying root cause of a set of incidents. Enterprise Manager provides features to track and manage the lifecycle of a problem.

Incident Manager

Enterprise Manager simplifies managing incidents through an intuitive UI called Incident Manager. Incident Manager provides and easy-to-use interface that allows you to search, view, manage, and resolve incidents and problems impacting your environment. To access Incident Manager, from the **Enterprise** menu, select **Monitoring**, and then **Incident Manager**.

Figure 1 Incident Manager

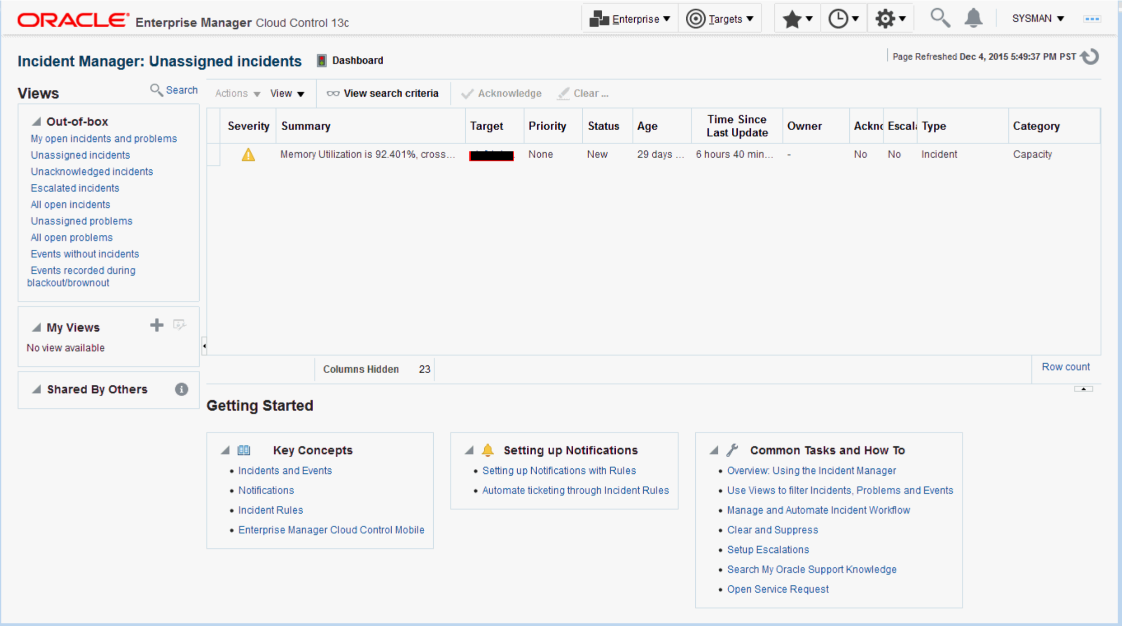

ORACLE Enterprise Manager Cloud Control 13c

Incident Manager: Unassigned incidents — Dashboard — Page Refreshed Dec 4, 2015 5:49:37 PM PST

Views

- Out-of-box
  - My open incidents and problems
  - Unassigned incidents
  - Unacknowledged incidents
  - Escalated incidents
  - All open incidents
  - Unassigned problems
  - All open problems
  - Events without incidents
  - Events recorded during blackout/brownout
- My Views
  - No view available
- Shared By Others

Severity	Summary	Target	Priority	Status	Age	Time Since Last Update	Owner	Ackno	Escala	Type	Category
⚠	Memory Utilization is 92.401%, cross...	[redacted]	None	New	29 days ...	6 hours 40 min...	-	No	No	Incident	Capacity

Columns Hidden 23

Getting Started

Key Concepts
- Incidents and Events
- Notifications
- Incident Rules
- Enterprise Manager Cloud Control Mobile

Setting up Notifications
- Setting up Notifications with Rules
- Automate ticketing through Incident Rules

Common Tasks and How To
- Overview: Using the Incident Manager
- Use Views to filter Incidents, Problems and Events
- Manage and Automate Incident Workflow
- Clear and Suppress
- Setup Escalations
- Search My Oracle Support Knowledge
- Open Service Request

From the Incident Manager UI, you can:

- Filter incidents, problems, and events by using custom views.
- Respond and work on an incident.
- Manage incident lifecycle including assigning, acknowledging, tracking its status, prioritization, and escalation
- Access (in context) My Oracle Support knowledge base articles and other Oracle documentation to help resolve the incident.
- Access direct in-context diagnostic/action links to relevant Enterprise Manager functionality allowing you to quickly diagnose or resolve the incident.

Incident Rules and Rule Sets

An *incident rule* specifies criteria and actions that determine when a notification should be sent and how it should be sent whenever an event or incident is raised. The criteria defined within a rule can apply to attributes such as the target type, events and severity states (clear, warning or critical) and the notification method that should be used when an incident is raised that matches the rule criteria. Rule actions can be conditional in nature. For example, a rule action can be defined to page a user when an incident severity is critical or just send e-mail if it is warning.

A *rule set* is a collection of rules that apply to a common set of targets such as hosts, databases, groups, jobs, metric extensions, or self updates and take appropriate actions to automate the business processes underlying incident. Incident rule sets can be made public for sharing across administrators. For example, administrators can subscribe to the same rule set if they are interested in receiving notifications for the same criteria defined in the rule. Alternatively, an Enterprise Manager Super Administrator can assign incident rule sets to other administrators so that they receive notifications for incidents as defined in the rule.

In addition to being used by the notification system (see *Rules* in Notifications), rule sets can also instruct Enterprise Manager to perform other actions, such as creating incidents, updating incidents, or call into a trouble ticketing system as discussed in Connectors.

Event Compression

Depending on what you are monitoring for, the number of incident alerts triggered by underlying events could potentially be high, with many of these notifications being related to the same underlying issue. Enterprise Manager lets you manage a potential flood of events using *Event Compression* so that you can focus on a smaller set of meaningful incidents. Event Compression allows you to group two or more related events into a single incident so that you can receive one notification about the incident instead of multiple notifications about each event that is part of the incident. There are two types of Event Compression you can use:

- **Event Compression Policies**: These out-of-the box event compression policies will be enabled by default for all newly created event rules. The policies work with your incident rule sets to decide whether sets of related events can compress into a single incident. For more information, see Event Compression Policies.

  > **Note:**
  >
  > Event Compression Policies are available with Enterprise Manager 13c Release 5 Update 8 or later.

  You can compare how Event Compression Policies would have affected the number of incidents in the past by running an Event Compression Analysis. For more information, see Assessing the Benefits of Using Event Compression Policies.

- **Rule-based Event Compression**: This compression functionality is enabled at rule set level and implemented in two steps:

  1. Create an event rule that compresses related events into a single incident.
  2. Create an incident rule to send a notification (email, ticket creation, etc.) when an incident is created.

  For more information, see Rule-based Event Compression.

Connectors

An Oracle Management Connector integrates third-party management systems with Enterprise Manager. There are two types of connectors: event connectors and ticketing connectors.

Using the event connector, you can configure Enterprise Manager to share events with non-Oracle event management systems. The connector monitors all events sent from Oracle Enterprise Manager and automatically updates alert information in the third-party management system. Event connectors support the following functions:

- Sharing of event information from Oracle Enterprise Manager to the third-party management system.
- Customization of event to alert mappings between Oracle Enterprise Manager and the third-party management system.
- Synchronization of event changes in Oracle Enterprise Manager with the alerts in the third-party management system.

Using the ticketing connector, you can configure Enterprise Manager to create, update, or close a ticket for any event created in Enterprise Manager. The ticket generated by the connector contains the relevant information about the Enterprise Manager incident, including a link to the Enterprise Manager console to enable helpdesk analysts leverage Enterprise Manager's diagnostic and resolution features to resolve the incident. In Enterprise Manger, the ticket ID, ticket status, and link to the third-party ticketing system is the shown in the context of the incident. This provides Enterprise Manager administrators with ticket status information and an easy way to quickly access the ticket.

Available connectors include:

- ServiceNow Connector
- Jira Service Management Connector
- PagerDuty Connector
- Microsoft Systems Center Operations Manager (SCOM) Connector
- IBM Tivoli Netcool/OMNIbus Connector

For more information about Oracle-built connectors, see the Enterprise Manager Plug-ins Exchange.

http://www.oracle.com/goto/emextensibility

Accessing Monitoring Information

Enterprise Manager provides multiple ways to access monitoring information. The primary focal point for incident management is the Incident Manager console, however Enterprise Manager also provides other ways to access monitoring information. The following figures show the various locations within Enterprise Manager that display target monitoring information. The following figure shows the Enterprise Manager Overview page that conveniently displays target status rollup and rollup of incidents.

Figure 2   Enterprise Manager Console

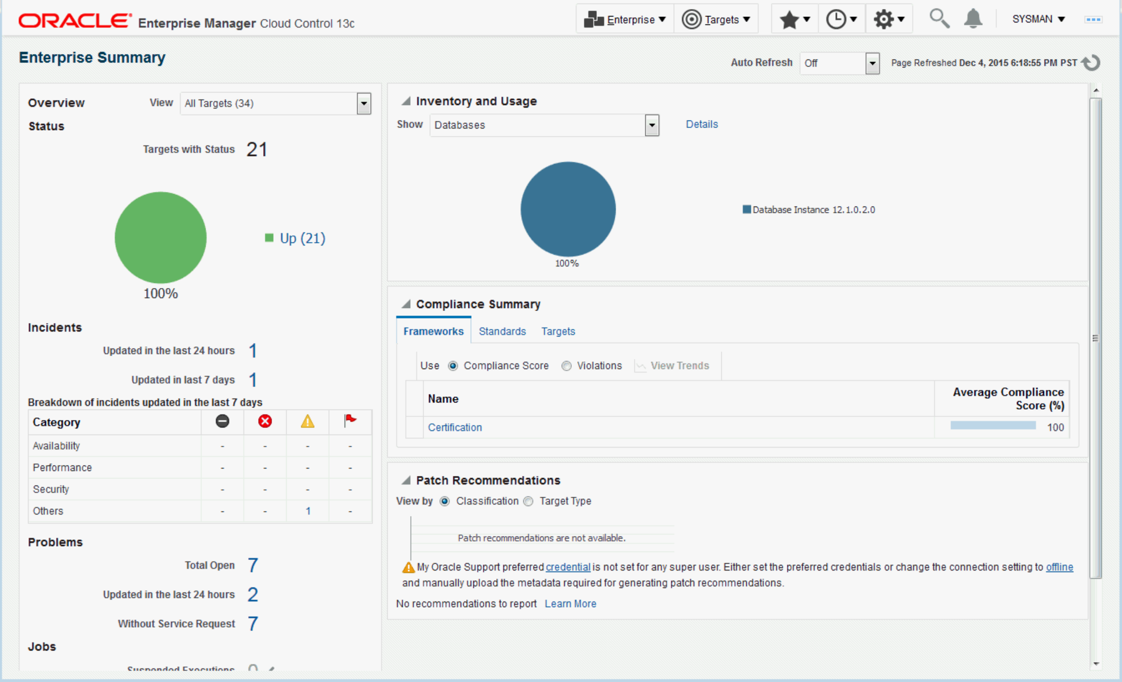

ORACLE Enterprise Manager Cloud Control 13c
Enterprise ▾
Targets ▾
SYSMAN ▾
100%
Up (21)
100%
Database Instance 12.1.0.2.0

Enterprise Summary

Auto Refresh: Off

Page Refreshed Dec 4, 2015 6:18:55 PM PST

Overview

View: All Targets (34)

Status

Targets with Status 21

Incidents

Updated in the last 24 hours 1

Updated in last 7 days 1

Breakdown of incidents updated in the last 7 days

Category	⊖	⊗	⚠	⚑
Availability	-	-	-	-
Performance	-	-	-	-
Security	-	-	-	-
Others	-	-	1	-

Problems

Total Open 7

Updated in the last 24 hours 2

Without Service Request 7

Jobs

Suspended Executions 0

Inventory and Usage

Show: Databases   Details

Compliance Summary

Frameworks   Standards   Targets

Use ● Compliance Score ○ Violations   View Trends

Name	Average Compliance Score (%)
Certification	100

Patch Recommendations

View by ● Classification ○ Target Type

Patch recommendations are not available.

⚠ My Oracle Support preferred credential is not set for any super user. Either set the preferred credentials or change the connection setting to offline and manually upload the metadata required for generating patch recommendations.

No recommendations to report   Learn More

The next figure shows the Incident Manager home page which displays incidents for a system or target.

Figure 3 Incident Manager (in context of a system or target)

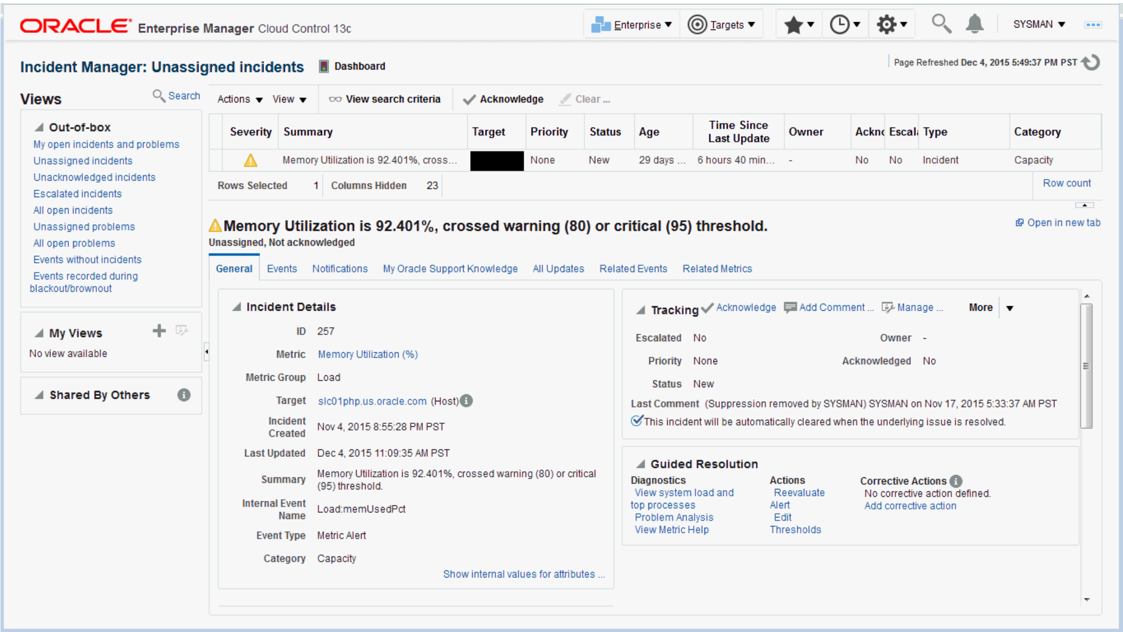

Incident Manager: Unassigned incidents   Dashboard

Page Refreshed Dec 4, 2015 5:49:37 PM PST

Views

Out-of-box
- My open incidents and problems
- Unassigned incidents
- Unacknowledged incidents
- Escalated incidents
- All open incidents
- Unassigned problems
- All open problems
- Events without incidents
- Events recorded during blackout/brownout

My Views

No view available

Shared By Others

Severity	Summary	Target	Priority	Status	Age	Time Since Last Update	Owner	Ackno	Escal	Type	Category
⚠	Memory Utilization is 92.401%, cross...	[redacted]	None	New	29 days ...	6 hours 40 min...	-	No	No	Incident	Capacity

Rows Selected 1   Columns Hidden 23

⚠Memory Utilization is 92.401%, crossed warning (80) or critical (95) threshold.

Unassigned, Not acknowledged

General   Events   Notifications   My Oracle Support Knowledge   All Updates   Related Events   Related Metrics

Incident Details

ID	257
Metric	Memory Utilization (%)
Metric Group	Load
Target	slc01php.us.oracle.com (Host)
Incident Created	Nov 4, 2015 8:55:28 PM PST
Last Updated	Dec 4, 2015 11:09:35 AM PST
Summary	Memory Utilization is 92.401%, crossed warning (80) or critical (95) threshold.
Internal Event Name	Load:memUsedPct
Event Type	Metric Alert
Category	Capacity

Show internal values for attributes ...

Tracking

Escalated	No	Owner	-
Priority	None	Acknowledged	No
Status	New		

Last Comment (Suppression removed by SYSMAN) SYSMAN on Nov 17, 2015 5:33:37 AM PST

✓This incident will be automatically cleared when the underlying issue is resolved.

Guided Resolution

Diagnostics	Actions	Corrective Actions
View system load and top processes Problem Analysis View Metric Help	Reevaluate Alert Edit Thresholds	No corrective action defined. Add corrective action

Monitoring information is also displayed on target home pages. In the following figure, you can see target status as well as a rollup of incidents.

Figure 4 Target Home Pages

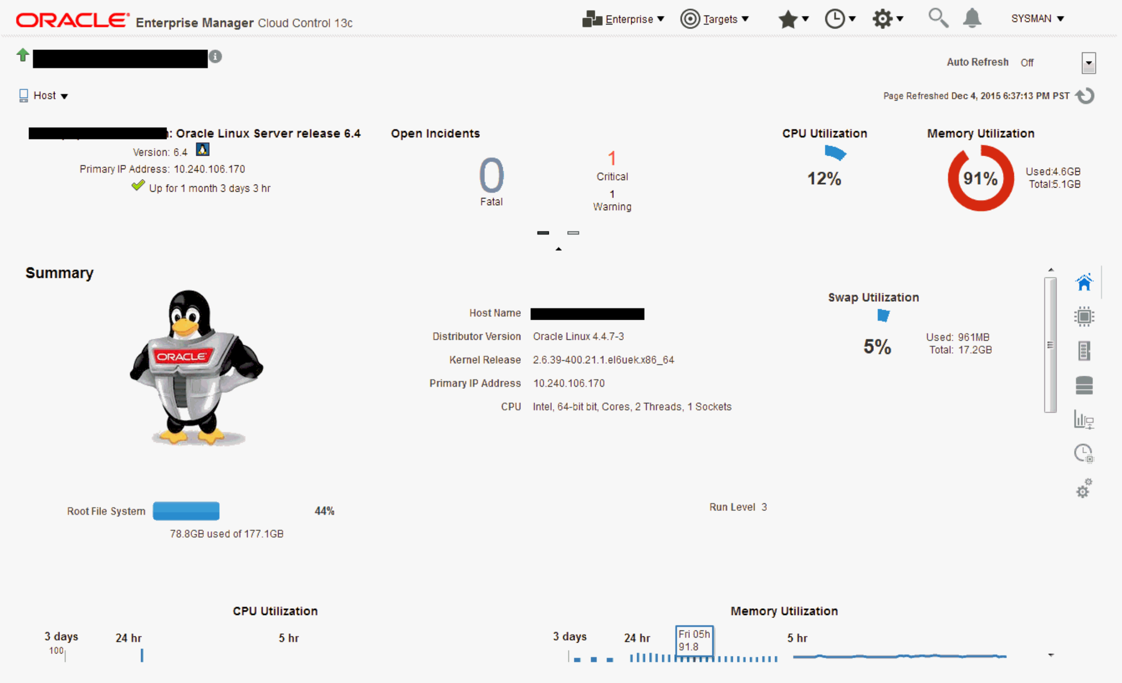

[redacted]

Host

Auto Refresh Off

Page Refreshed Dec 4, 2015 6:37:13 PM PST

[redacted]: Oracle Linux Server release 6.4

Version: 6.4

Primary IP Address: 10.240.106.170

✓ Up for 1 month 3 days 3 hr

Open Incidents

0 Fatal   1 Critical   1 Warning

CPU Utilization 12%

Memory Utilization 91% Used:4.6GB Total:5.1GB

Summary

Host Name	[redacted]
Distributor Version	Oracle Linux 4.4.7-3
Kernel Release	2.6.39-400.21.1.el6uek.x86_64
Primary IP Address	10.240.106.170
CPU	Intel, 64-bit bit, Cores, 2 Threads, 1 Sockets

Swap Utilization 5% Used: 961MB Total: 17.2GB

Root File System 44%

78.8GB used of 177.1GB

Run Level 3

Part II
Discovery

This section contains the following chapters:

1

Discovering and Adding Host and Non-Host Targets

Enterprise Manager enables you to discover, promote, add, and then monitor software deployments across your network, using a single GUI-rich console. This chapter introduces you to the concepts of discovery and promotion, and describes how you can perform these tasks using Enterprise Manager.

In particular, this chapter covers the following:

- Overview of Discovering and Adding Targets
- Discovering and Adding Host Targets
- Discovering and Adding Non-Host Targets
- Discovering and Promoting Oracle Homes
- Retrieving Deleted Targets

Overview of Discovering and Adding Targets

This section introduces you to basic concepts of discovery, promotion, and monitoring. It familiarizes you with the different methods of discovering and monitoring targets using Enterprise Manager. In particular, this section covers the following:

- Understanding Discovery Terminology
- Options for Discovering Targets
- Discovery and Monitoring in Enterprise Manager Lifecycle
- Discovery and Monitoring Process

Understanding Discovery Terminology

This section describes the following:

- What are Targets and Managed Targets?
- What is Discovery?
- What is Promotion?

What are Targets and Managed Targets?

Targets are entities such as host machines, databases, Fusion Middleware components, server targets (hardware), that can be managed and monitored in Enterprise Manager.

Managed targets are entities that are actively being monitored and managed by Enterprise Manager.

What is Discovery?

Discovery refers to the process of identifying unmanaged hosts and targets in your environment. You can discover hosts and targets automatically or manually.

Figure 1-1 illustrates the discovery process.

Figure 1-1 Discovery

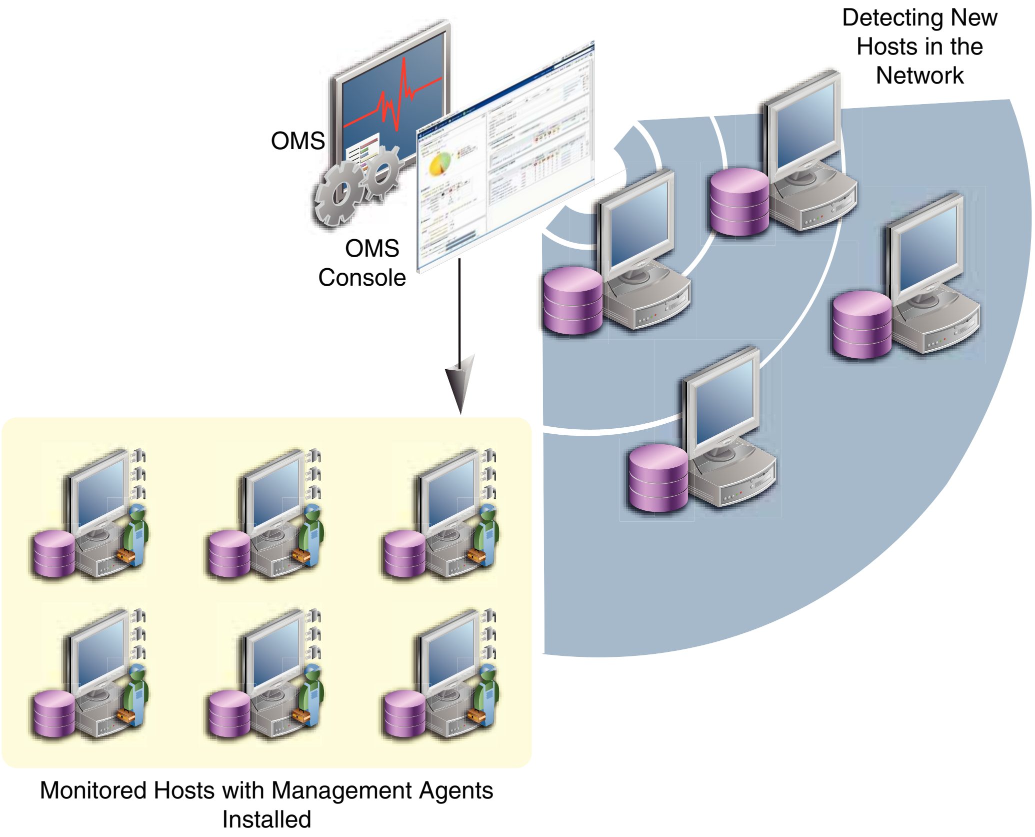
Detecting New Hosts in the Network
OMS
OMS Console
Monitored Hosts with Management Agents Installed

What is Promotion?

Promotion refers to the process of converting unmanaged hosts and targets, which have been discovered in your network, to managed hosts and targets in Enterprise Manager so that they can be monitored and managed efficiently. While conversion of unmanaged hosts to managed hosts involves deployment of a Management Agent on those hosts, conversion of unmanaged targets running on those hosts to managed targets involves only adding the targets as manageable entities in Enterprise Manager without deploying any additional component on the hosts.

Figure 1-2 illustrates the promotion process.

Figure 1-2 Promotion

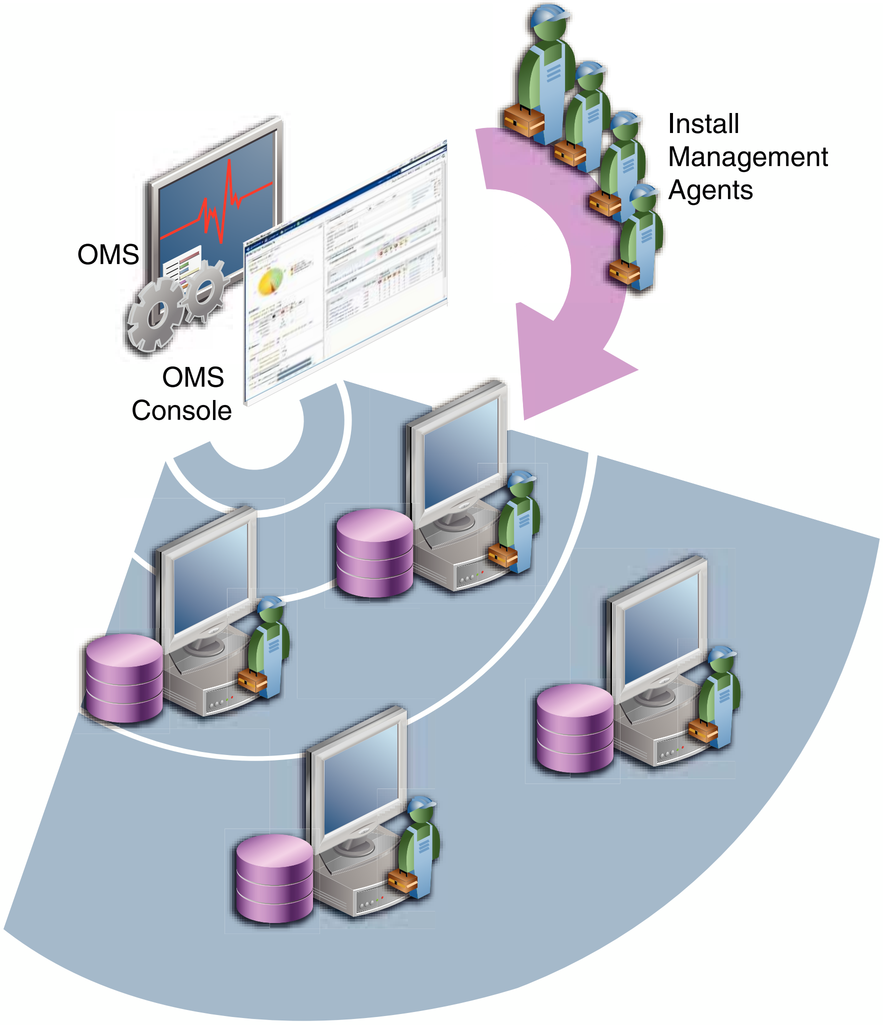

Options for Discovering Targets

You can discover targets using either of the following methods:

Autodiscovery Process

For discovery of a host, the autodiscovery process enables a Management Agent running on the host to run an Enterprise Manager job that scans for unmanaged hosts. You then convert these unmanaged hosts to managed hosts by deploying Management Agents on these hosts. Next, you search for targets such as databases or other deployed components or applications on these managed hosts, and finally you promote these targets to managed status.

For discovery of targets, the autodiscovery process enables you to search for targets on the host and then add these targets using Enterprise Manager.

The benefit of using this process is that as new components are added to your infrastructure, they can be found and brought under management on a regularly-scheduled basis.

Guided Discovery Process

The guided discovery process enables you to explicitly add a target to bring under management. The discovery wizard guides you through the process and most of the specifications required are filled by default.

The benefits of using this process are as follows:

- You can find targets with less effort.
- You can find a new database that has been added recently even if autodiscovery has not been run.
- You can find a non-promoted database that already exists in autodiscovery results, but has a change in details. For example, the port.
- You eliminate unnecessary consumption of resources on the Management Agent when discovery is not needed.

Declarative Process

Declaring target monitoring properties enables you to manually specify all the details required to discover the database target, such as the host name and location, target name and location, and other specific information. This process is generally used when the autodiscovery process and guided discovery process fails to discover the target that you want to add.

Discovery and Monitoring in Enterprise Manager Lifecycle

Figure 1-3 illustrates the lifecycle process of discovering and monitoring targets in Enterprise Manager.

Figure 1-3 Discovery and Monitoring in Enterprise Manager Lifecycle

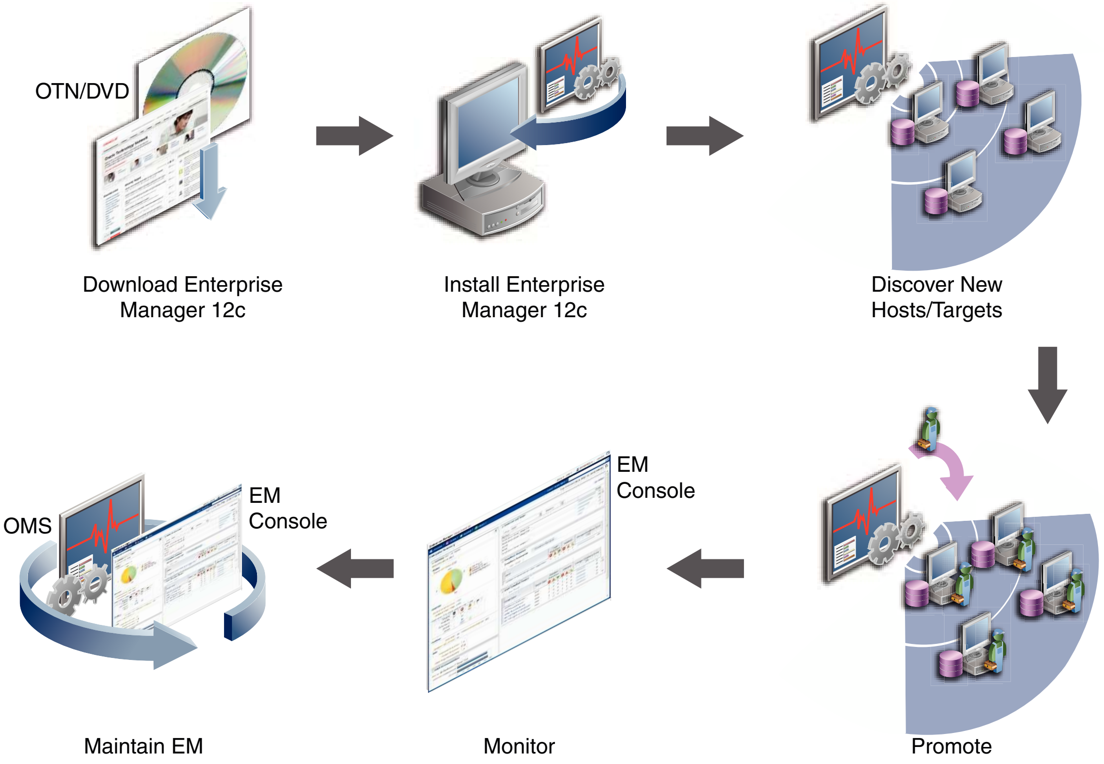
OTN/DVD
Download Enterprise Manager 12c
Install Enterprise Manager 12c
Discover New Hosts/Targets
Promote
EM Console
Monitor
OMS
EM Console
Maintain EM

Discovery and Monitoring Process

Figure 1-4 illustrates the high level process of discovering and monitoring targets:

Figure 1-4   Discovery and Monitoring Process

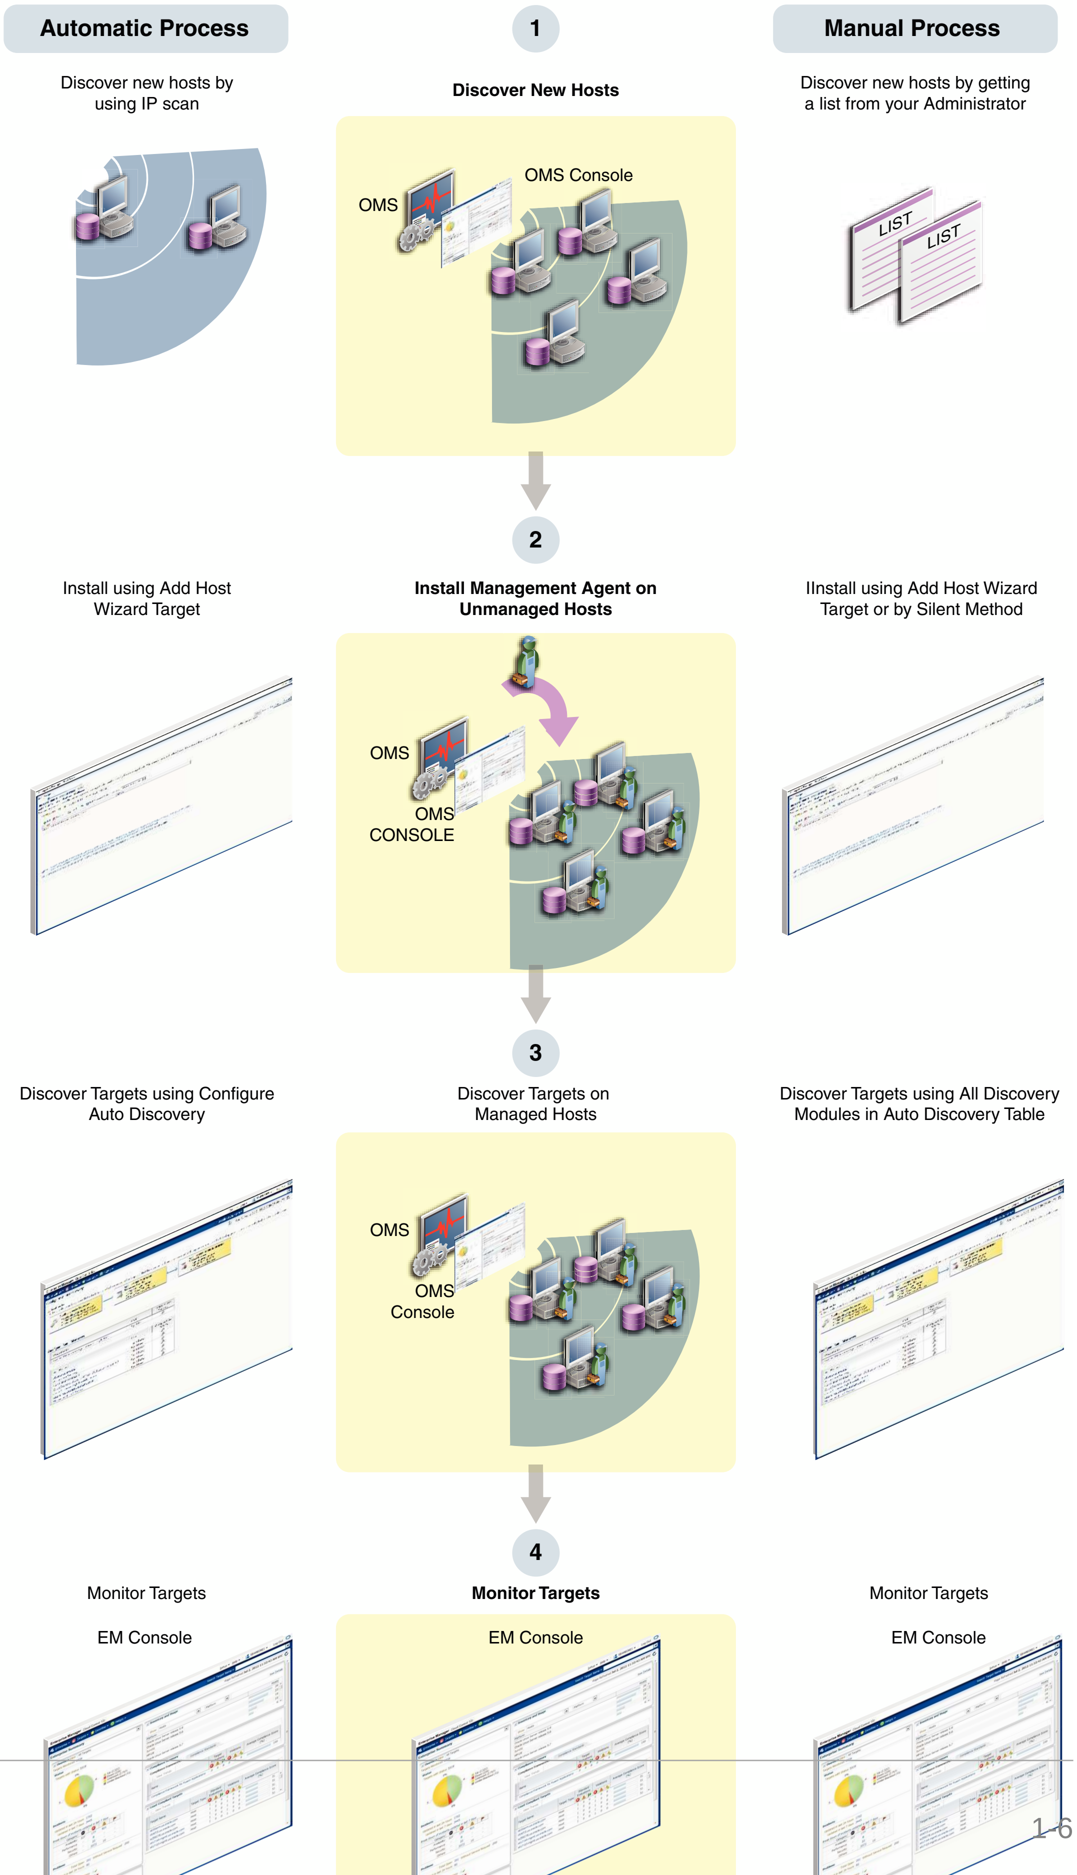
Automatic Process
1
Manual Process
Discover new hosts by using IP scan
Discover New Hosts
Discover new hosts by getting a list from your Administrator
OMS
OMS Console
LIST
LIST
2
Install using Add Host Wizard Target
Install Management Agent on Unmanaged Hosts
IInstall using Add Host Wizard Target or by Silent Method
OMS
OMS CONSOLE
3
Discover Targets using Configure Auto Discovery
Discover Targets on Managed Hosts
Discover Targets using All Discovery Modules in Auto Discovery Table
OMS
OMS Console
4
Monitor Targets
Monitor Targets
Monitor Targets
EM Console
EM Console
EM Console

Discovering and Adding Host Targets

This section covers the following:

- Adding Remote Hosts to the Remote Agent
- Adding Host Targets Using the Manual Guided Discovery Process

Adding Remote Hosts to the Remote Agent

Prerequisites

- SSH connection is enabled from remote agent host to all remote hosts.
- Password-less sudo configuration is required for the remote host target.

> **Note:**
>
> - Remote host addition to remote agent is only supported on Linux x86-64, in this release.
> - PowerBroker option is not supported for adding remote host target to remote agent, in this release.

Adding Host Target to Remote Agent Using the Console

From **Setup**, go to **Add Targets Manually**, click **Add Remote Host Targets** and fill the form with the following inputs and click **Submit**.

Table 1-1 UI Input

Input field	Description
Session Name	Pre-populated by default with the task name, user name and date stamp, to be useful when reusing from previous session, but can be changed as needed.
Remote agent	Select the remote agent to be used from the list.
Host Monitoring Credential	Used for monitoring and will not need escalation to root. All metric collections will use this credential. Supports SSH keys and username/password.
Host Admin Credential	Root credentials used for pushing sbin content as part of Cache Maintenance operations Require ability of escalating to root without a password using PDP settings (sudo)
SSH Port	Pre-populated with the default value, `22`, can be changed as needed.
Privileged Delegation Setting	Pre-populated with the default path, `/usr/bin/sudo -u %RUNAS% %COMMAND%`, can be changed as needed.

Table 1-1 (Cont.) UI Input

Input field	Description
Remote cache	Cache of scripts, archives, monitoring utilities and application metadata maintained on remote host. Minimum 1 GB space required on remote host.
Remote cache Sbin	Cache contains utilities like nmosudo and nmgsshe root owned contents on remote host. Remote agent may need certain root owned executables for monitoring.

Adding Host Target to Remote Agent Using emcli

Use the `submit_add_remote_host` emcli command line option to add remote host to the remote agent:

```
emcli submit_add_remote_host
    -host_names=<host name>
    -remote_cache_root=<path to remote cache>
    -remote_cache_sbin_root=<path to remote cache utilities>
    -agent=<agent name>
    -host_cred_mon=<credentials type, e.g. SSH>
    -host_cred_admin=<root credentials, e.g. SSH_ROOT>
    -privilege_delegation_setting="/usr/bin/sudo -u %RUNAS% %COMMAND%"
    -wait_for_completion
```

Example:

```
emcli submit_add_remote_host
    -host_names=host1.example.com
    -remote_cache_root=/u01/app/rcache
    -remote_cache_sbin_root=/u01/app/rcacheroo
    -agent=host1.example.com:1838
    -host_cred_mon=SSH
    -host_cred_admin=SSH_ROOT
    -privilege_delegation_setting="/usr/bin/sudo -u %RUNAS% %COMMAND%"
    -wait_for_completion
```

Adding Host Targets Using the Manual Guided Discovery Process

To add host targets manually, refer to the instructions outlined in the Installing Oracle Management Agents.

Discovering and Adding Non-Host Targets

This section covers the following:

- Configuring Autodiscovery of Non-Host Targets
- Adding Non-Host Targets Using the Guided Discovery Process
- Adding Non-Host Targets By Using the Declarative Process

Configuring Autodiscovery of Non-Host Targets

To discover targets on managed hosts, follow these steps:

1. From the **Setup** menu, select **Add Target,** and then select **Configure Auto Discovery.**
2. On the Setup Discovery page, in the **Targets on Hosts** tab, 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.

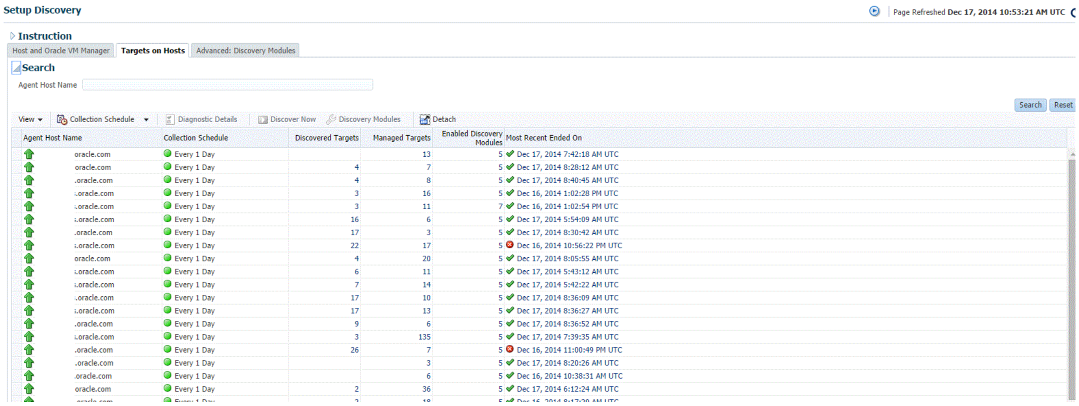

Setup Discovery

Page Refreshed **Dec 17, 2014 10:53:21 AM UTC**

▷ Instruction

Host and Oracle VM Manager | **Targets on Hosts** | Advanced: Discovery Modules

Search

Agent Host Name

Agent Host Name	Collection Schedule	Discovered Targets	Managed Targets	Enabled Discovery Modules	Most Recent Ended On
oracle.com	● Every 1 Day		13	5	✓ Dec 17, 2014 7:42:18 AM UTC
oracle.com	● Every 1 Day	4	7	5	✓ Dec 17, 2014 8:28:12 AM UTC
.oracle.com	● Every 1 Day	4	8	5	✓ Dec 17, 2014 8:40:45 AM UTC
.oracle.com	● Every 1 Day	3	16	5	✓ Dec 16, 2014 1:02:28 PM UTC
.oracle.com	● Every 1 Day	3	11	7	✓ Dec 16, 2014 1:02:54 PM UTC
.oracle.com	● Every 1 Day	16	6	5	✓ Dec 17, 2014 5:54:09 AM UTC
.oracle.com	● Every 1 Day	17	3	5	✓ Dec 17, 2014 8:30:42 AM UTC
.oracle.com	● Every 1 Day	22	17	5	✗ Dec 16, 2014 10:56:22 PM UTC
oracle.com	● Every 1 Day	4	20	5	✓ Dec 17, 2014 8:05:55 AM UTC
.oracle.com	● Every 1 Day	6	11	5	✓ Dec 17, 2014 5:43:12 AM UTC
.oracle.com	● Every 1 Day	7	14	5	✓ Dec 17, 2014 5:42:22 AM UTC
.oracle.com	● Every 1 Day	17	10	5	✓ Dec 17, 2014 8:36:09 AM UTC
.oracle.com	● Every 1 Day	17	13	5	✓ Dec 17, 2014 8:36:27 AM UTC
.oracle.com	● Every 1 Day	9	6	5	✓ Dec 17, 2014 8:36:52 AM UTC
.oracle.com	● Every 1 Day	3	135	5	✓ Dec 17, 2014 7:39:35 AM UTC
.oracle.com	● Every 1 Day	26	7	5	✗ Dec 16, 2014 11:00:49 PM UTC
.oracle.com	● Every 1 Day		3	5	✓ Dec 17, 2014 8:20:26 AM UTC
.oracle.com	● Every 1 Day		6	5	✓ Dec 16, 2014 10:38:31 AM UTC
oracle.com	● Every 1 Day	2	36	5	✓ Dec 17, 2014 6:12:24 AM UTC

3. On the Target Discovery (Agent-based) page, 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.
4. To search for a specific Management Agent, click **Search.** The table lists all the Management Agents and filters the list based on what you search for.
5. Select the host in the table and click **Discovery Modules.**
6. On the Discovery Modules page, select the target types that you want to discover on the host. 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 Domain: 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.

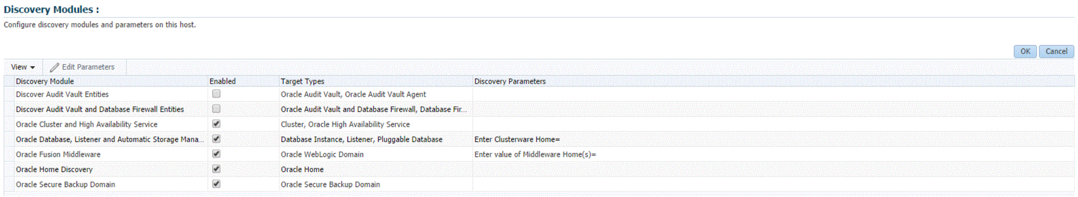

Discovery Modules :

Configure discovery modules and parameters on this host.

Discovery Module	Enabled	Target Types	Discovery Parameters
Discover Audit Vault Entities	□	Oracle Audit Vault, Oracle Audit Vault Agent	
Discover Audit Vault and Database Firewall Entities	□	Oracle Audit Vault and Database Firewall, Database Fir...	
Oracle Cluster and High Availability Service	☑	Cluster, Oracle High Availability Service	
Oracle Database, Listener and Automatic Storage Mana...	☑	Database Instance, Listener, Pluggable Database	Enter Clusterware Home=
Oracle Fusion Middleware	☑	Oracle WebLogic Domain	Enter value of Middleware Home(s)=
Oracle Home Discovery	☑	Oracle Home	
Oracle Secure Backup Domain	☑	Oracle Secure Backup Domain	

Click **OK** when finished. Target discovery has been configured on this host.

7. On the Setup Discovery page, in the Targets on Host tab, select the hosts you want to set the schedule at which discovery will be run. Click **Collection Schedule,** and then select **For all hosts,** or **For selected hosts.** In the Collection Schedule dialog box, enable or disable collection for the hosts that you have selected. If you have enabled collection, then select the frequency of collection. This schedule will be applied to all selected hosts. By default the discovery will run every 24 hours. Click **OK.**
8. Repeat these steps for each additional host on which you want to configure discovery.
9. Click **Discover Now** to discover targets immediately. The discovery will also run at the scheduled interval.
10. To check for discovered targets from the **Setup** menu, select **Add Target,** then select **Auto Discovery Results.**
11. Select a target to promote, then click **Promote**. A wizard specific to the target type you are promoting opens. Supply the required values.

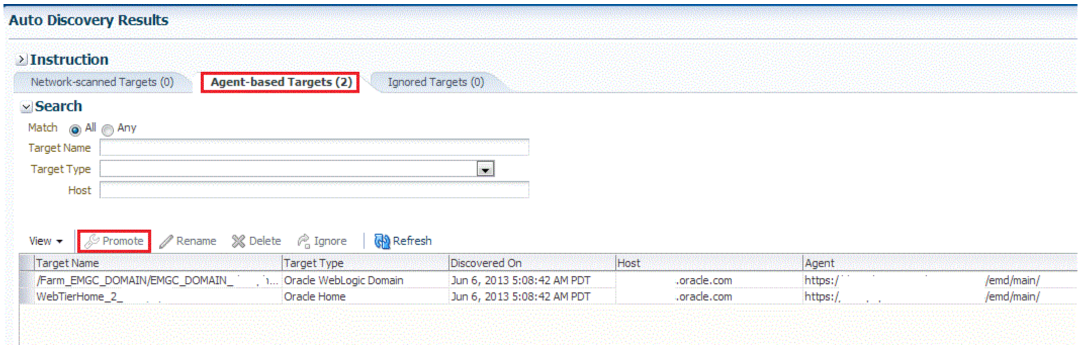

Auto Discovery Results

Instruction

Network-scanned Targets (0) | **Agent-based Targets (2)** | Ignored Targets (0)

Search

Match ● All ○ Any

Target Name

Target Type

Host

Target Name	Target Type	Discovered On	Host	Agent
/Farm_EMGC_DOMAIN/EMGC_DOMAIN_ ...	Oracle WebLogic Domain	Jun 6, 2013 5:08:42 AM PDT	.oracle.com	https:/ /emd/main/
WebTierHome_2_	Oracle Home	Jun 6, 2013 5:08:42 AM PDT	.oracle.com	https:/ /emd/main/

12. Click the **Agent-based Targets** tab.You can choose one or several targets to promote.
13. Note that you can optionally click **Ignore** for a discovered target. Ignoring a target puts it into a list of targets that you do not want to manage.

    Ignored targets will be displayed in the Ignored Targets tab, and will remain in Enterprise Manager as un-managed targets until you decide to either promote or remove them. If you delete a target, it would be rediscovered the next time discovery runs.
14. 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:**
>
> - When you promote a discovered target to managed status, the plug-in required for the target is automatically deployed to the Management Agent, which monitors the host where the target has been discovered. For the plug-in to be deployed, the Management Agent must be secure. Therefore, before promoting the discovered targets to managed status, ensure that the Management Agent is secure. You can always unsecure it after the discovered target is promoted to managed status, that is, after the required plug-in is deployed.
>
>   To verify the secure status of a Management Agent, and to secure it if required, use any one of the following methods:
>
>   - From the **Setup** menu, select **Manage the Manager,** and then click **Agents**. Click the required Management Agent. Verify whether the Management Agent is secure. If it is not secure, from the **Agent** menu, click **Secure** to secure it.
>   - Run the following command to verify if the Management Agent is secure:
>
>     ```
>     <EMSTATE>/bin/emctl status agent
>     ```
>
>   If the Management Agent is secure, the Management Agent URL displayed in the output of the previous command is an HTTPS URL. However, if the Management Agent URL displayed is an HTTP URL, secure the Management Agent by running the following command:
>
>     ```
>     <EMSTATE>/bin/emctl secure agent
>     ```
>
> - Enterprise Manager supports simultaneous promotion of multiple targets only for some target types. 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.

Adding Non-Host Targets Using the Guided Discovery Process

To add non-host targets using the guided process, follow these steps:

> **Note:**
>
> When you add a target using the guided process, some scripts and automated processes are run that are particular for the target type that you select. You may have to input credentials in order to run the guided process.

1. From the **Setup** menu, select **Add Target**, then select **Add Targets Manually**. Enterprise Manager displays the Add Targets Manually page.
2. On the Add Targets Manually page, select **Add Using Guided Process.**
3. In the Add Using Guided Process dialog box, select the target type such as **Exalogic Elastic Cloud, Oracle Cluster and High Availability Service, System Infrastructure server ILOM,** or **Oracle WebLogic Domain.** Click **Add...**

4. After you select the target type, a wizard specific to the target type guides you through the process of manually adding 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.

> **Note:**
>
> When you manually add a non-host target to Enterprise Manager, the plug-in required for the target is automatically deployed to the Management Agent, which monitors the host where the non-host target exists. For the plug-in to be deployed, the Management Agent must be secure. Therefore, before manually adding a non-host target to Enterprise Manager, ensure that the Management Agent is secure. You can always unsecure it after the target is added to Enterprise Manager, that is, after the required plug-in is deployed.
>
> To verify the secure status of a Management Agent, and to secure it if required, use any one of the following methods:
>
> - From the **Setup** menu, select **Manage the Manager** and then, click **Agents**. Click the required Management Agent. Verify whether the Management Agent is secure. If it is not secure, from the **Agent** menu, click **Secure** to secure it.
> - Run the following command to verify if the Management Agent is secure:
>
>   ```
>   <EMSTATE>/bin/emctl status agent
>   ```
>
>   If the Management Agent is secure, the Management Agent URL displayed in the output of the previous command is an HTTPS URL. However, if the Management Agent URL displayed is an HTTP URL, secure the Management Agent by running the following command:
>
>   ```
>   <EMSTATE>/bin/emctl secure agent
>   ```

Adding Non-Host Targets By Using the Declarative Process

To add a target on a managed host by specifying the target monitoring properties, follow these steps:

1. From the **Setup** menu, select **Add Target**, then select **Add Targets Manually**. Enterprise Manager displays the Add Targets Manually page.
2. On the Add Targets Manually page, select **Add Target Declaratively.**
3. In the Add Target Declaratively dialog box, choose one of the target types to add from the Target Types list, such as **ADF Business Components for Java, Cluster Database,** or **Oracle HTTP Server.**
4. Specify the Management Agent that will be used to monitor the target, or click on the Search icon to search for and select the Management Agent. Click **Add...**
5. After you select the target type, a wizard specific to the target type guides you through the process of manually adding 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.

Note:

When you manually add a non-host target to Enterprise Manager, the plug-in required for the target is automatically deployed to the Management Agent, which monitors the host where the non-host target exists. For the plug-in to be deployed, the Management Agent must be secure. Therefore, before manually adding a non-host target to Enterprise Manager, ensure that the Management Agent is secure. You can always unsecure it after the target is added to Enterprise Manager, that is, after the required plug-in is deployed.

To verify the secure status of a Management Agent, and to secure it if required, use any one of the following methods:

- From the **Setup** menu, select **Manage the Manager**, and then, click **Agents**. Click the required Management Agent. Verify whether the Management Agent is secure. If it is not secure, from the **Agent** menu, click **Secure** to secure it.
- Run the following command to verify if the Management Agent is secure:

  ```
  <EMSTATE>/bin/emctl status agent
  ```

  If the Management Agent is secure, the Management Agent URL displayed in the output of the previous command is an HTTPS URL. However, if the Management Agent URL displayed is an HTTP URL, secure the Management Agent by running the following command:

  ```
  <EMSTATE>/bin/emctl secure agent
  ```

Discovering and Promoting Oracle Homes

When you deploy an Oracle software component outside of the deployment procedures provided by Enterprise Manager, the Oracle home is not automatically discovered and promoted as targets. You will have to manually discover and promote the Oracle home target.

To discover and promote an Oracle home target, follow these steps:

1. From the **Enterprise** menu, select **Job,** and then select **Activity.**
2. On the Job Activity page, from the drop-down list in the table, select **Discover Promote Oracle Home Target.**

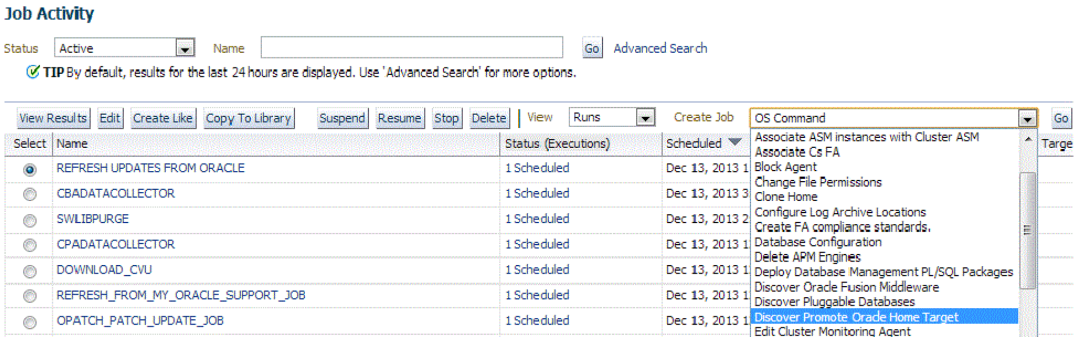

Job Activity

Status: Active
Name:
TIP By default, results for the last 24 hours are displayed. Use 'Advanced Search' for more options.

View: Runs
Create Job: OS Command

Select	Name	Status (Executions)	Scheduled
●	REFRESH UPDATES FROM ORACLE	1 Scheduled	Dec 13, 2013 1
○	CBADATACOLLECTOR	1 Scheduled	Dec 13, 2013 3
○	SWLIBPURGE	1 Scheduled	Dec 13, 2013 2
○	CPADATACOLLECTOR	1 Scheduled	Dec 13, 2013 1
○	DOWNLOAD_CVU	1 Scheduled	Dec 13, 2013 1
○	REFRESH_FROM_MY_ORACLE_SUPPORT_JOB	1 Scheduled	Dec 13, 2013 1
○	OPATCH_PATCH_UPDATE_JOB	1 Scheduled	Dec 13, 2013 1

Go
Advanced Search
View Results
Edit
Create Like
Copy To Library
Suspend
Resume
Stop
Delete
Go
Targe
Associate ASM instances with Cluster ASM
Associate Cs FA
Block Agent
Change File Permissions
Clone Home
Configure Log Archive Locations
Create FA compliance standards.
Database Configuration
Delete APM Engines
Deploy Database Management PL/SQL Packages
Discover Oracle Fusion Middleware
Discover Pluggable Databases
Discover Promote Oracle Home Target
Edit Cluster Monitoring Agent

Click **Go.**

3. On the 'Create Discover Promote Oracle Home Target' Job page, in the General tab, specify the name of the discovery.

   For example: `OHDiscovery`

   You can optionally add a description for the discovery.

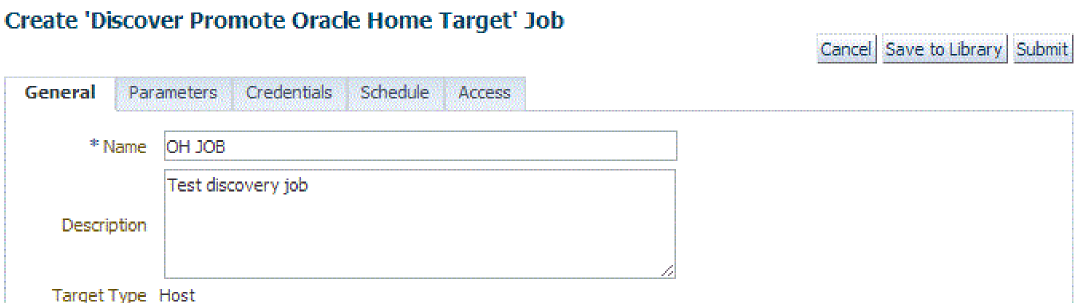

Create 'Discover Promote Oracle Home Target' Job

* Name: OH JOB

Description: Test discovery job

Target Type: Host

   Click **Add.**

4. In the Search and Select: Target dialog box, select Target Type as **Host**, and then select all the host targets listed by clicking **Select All.**

   Click **Select.**

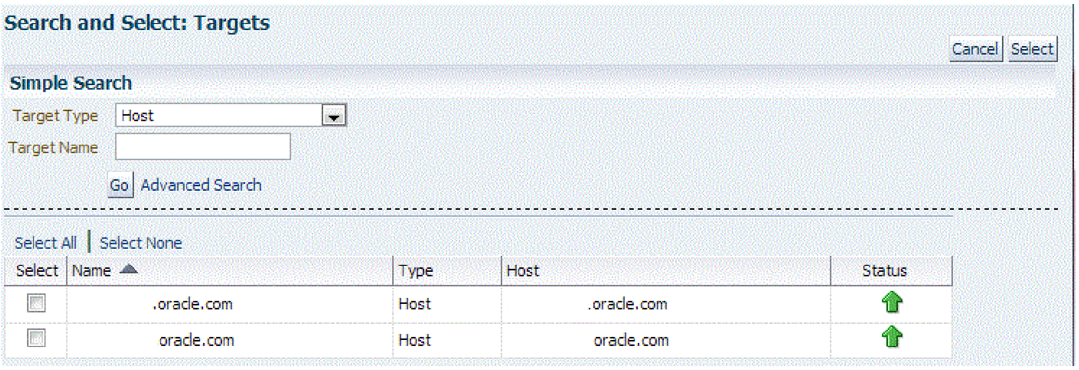

Search and Select: Targets

Simple Search

Target Type: Host

Target Name:

Select	Name	Type	Host	Status
☐	.oracle.com	Host	.oracle.com	⇧
☐	oracle.com	Host	oracle.com	⇧

5. On the'Create Discover Promote Oracle Home Target' Job page, the host targets that you selected are displayed in the table.

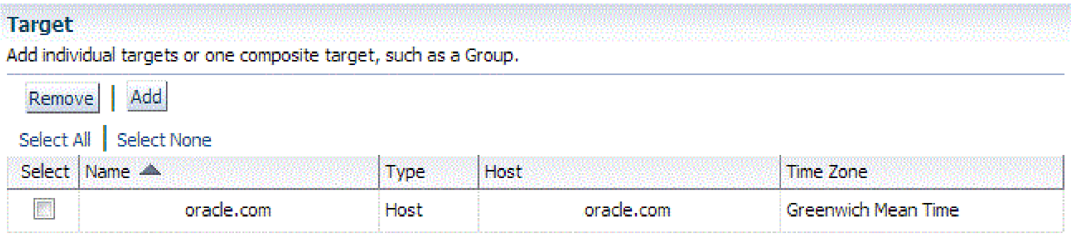

Target

Add individual targets or one composite target, such as a Group.

Select	Name	Type	Host	Time Zone
☐	oracle.com	Host	oracle.com	Greenwich Mean Time

6. Select the Parameters tab, and then do one of the following:

   - To discover a single Oracle Home, specify the path to the home, and then select **Oracle Home** as the manage entity.

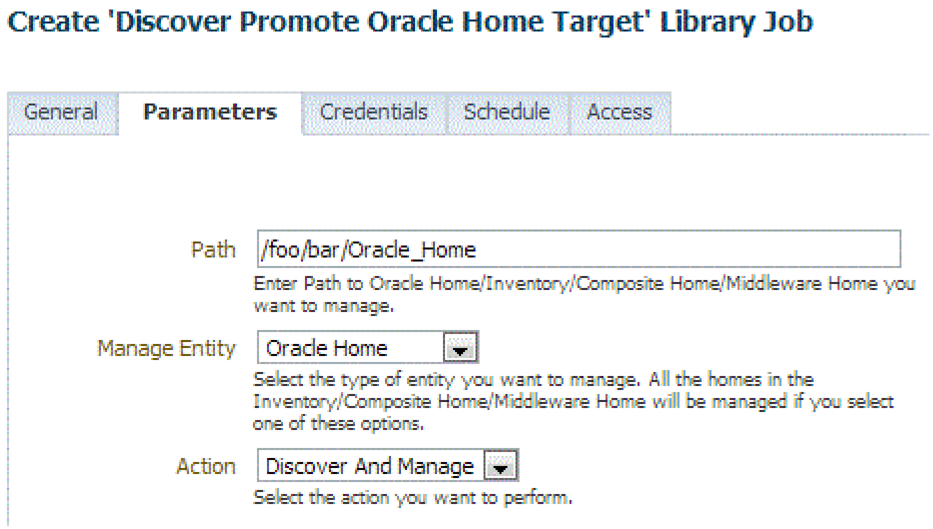

Create 'Discover Promote Oracle Home Target' Library Job

Path: /foo/bar/Oracle_Home
Enter Path to Oracle Home/Inventory/Composite Home/Middleware Home you want to manage.

Manage Entity: Oracle Home
Select the type of entity you want to manage. All the homes in the Inventory/Composite Home/Middleware Home will be managed if you select one of these options.

Action: Discover And Manage
Select the action you want to perform.

- To discover all Homes in an inventory, specify the path to the inventory, and then select **Inventory** as the manage entity.

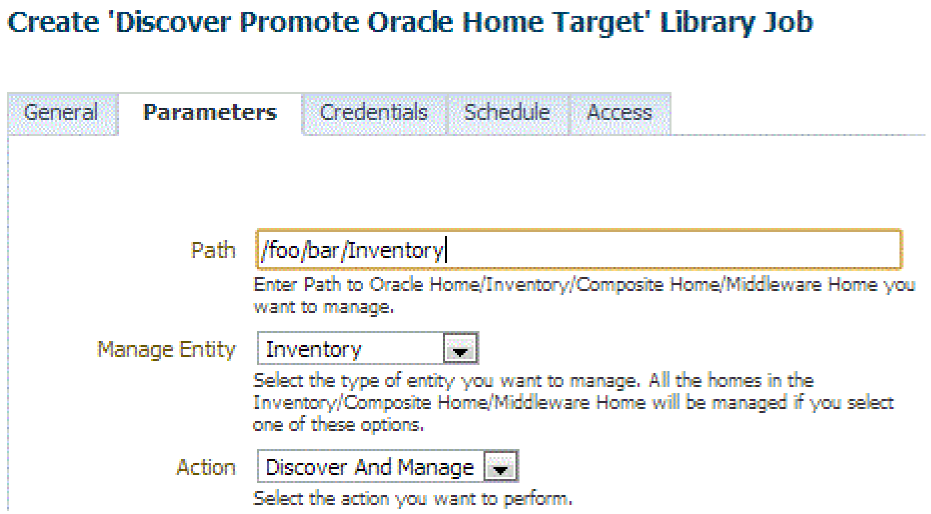

Create 'Discover Promote Oracle Home Target' Library Job

Path: /foo/bar/Inventory
Enter Path to Oracle Home/Inventory/Composite Home/Middleware Home you want to manage.

Manage Entity: Inventory
Select the type of entity you want to manage. All the homes in the Inventory/Composite Home/Middleware Home will be managed if you select one of these options.

Action: Discover And Manage
Select the action you want to perform.

- To discover all Homes in a Middleware Home, specify the path to the Middleware Home and select **Middleware Home** as the manage entity.

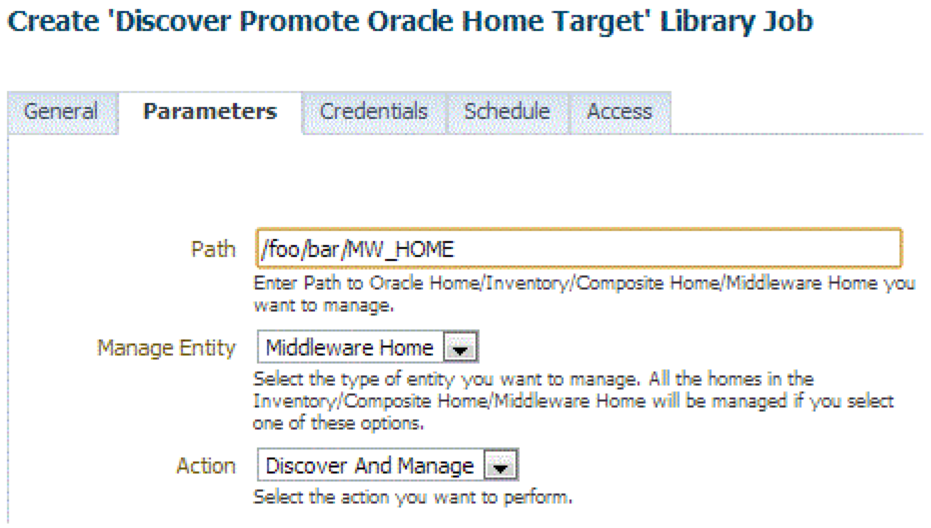

Create 'Discover Promote Oracle Home Target' Library Job

Path: /foo/bar/MW_HOME
Enter Path to Oracle Home/Inventory/Composite Home/Middleware Home you want to manage.

Manage Entity: Middleware Home
Select the type of entity you want to manage. All the homes in the Inventory/Composite Home/Middleware Home will be managed if you select one of these options.

Action: Discover And Manage
Select the action you want to perform.

7. To save the job for later, click **Save to Library.** To submit it, click **Submit.** When the discovery is successful, a confirmation is displayed on the Job Activity page.

> **Note:**
>
> If you submit the discovery job without specifying a path, a discovery of the whole host will be performed. In order for a Home to be discoverable by the Management Agent, it needs to be registered in an inventory that the Management Agent recognizes. The default inventory is the central inventory, which in Unix systems is found in `/etc/oraInst.loc`. Any Home registered here will automatically be discovered.
>
> If there are other inventories in the host, they need to be added to the inventory list of the Management Agent. A line must be added to `$EMSTATE/sysman/config/OUIinventories.add`.
>
> If the inventory is not found here, the Management Agent will not know of its existence, and hence any Home registered there will not be discovered.

Retrieving Deleted Targets

This sections covers the following:

- Retrieving Deleted Target Types
- Retrieving Deleted Host and Corresponding Management Agent Targets

Retrieving Deleted Target Types

If you have deleted one or more targets (such as a database target or a weblogic domain, or any other target), you can retrieve them and add them back to the Enterprise Manager Console. If autodiscovery is configured on the host where the targets were present, the targets are automatically discovered during the next scheduled autodiscovery operation. Once they are autodiscovered, you can promote them and add them to the console. If autodiscovery is not configured on the host where the targets were present, you have to discover the targets using one of the following methods:

- By enabling autodiscovery as described in Configuring Autodiscovery of Non-Host Targets to automatically discover and promote the targets in the next scheduled autodiscovery operation.
- By using the guided discovery process as described in Adding Non-Host Targets Using the Guided Discovery Process to manually discover and add the discovered targets to the console.
- By specifying the target monitoring properties for each target as described in Adding Non-Host Targets By Using the Declarative Process to manually discover and add the discovered targets to the console.
- By using the following EM CLI verb:

```
$ emcli add_target
      -name="name"
      -type="type"
      -host="hostname"
      [-properties="pname1:pval1;pname2:pval2;..."]
      [-separator=properties="sep_string"]
      [-subseparator=properties="subsep_string"]
      [-credentials="userpropname:username;pwdpropname:password;..."]
      [-input_file="parameter_tag:file_path"]
      [-display_name="display_name"]
      [-groups="groupname1:grouptype1;groupname2:grouptype2;..."]
      [-timezone_region="gmt_offset"]
      [-monitor_mode="monitor_mode"]
      [-instances="rac_database_instance_target_name1:target_type1;..."]
      [-force]
      [-timeout="time_in_seconds"]

[ ] indicates that the parameter is optional
```

For more information, see Verb Reference in the *Oracle Enterprise Manager Command Line Interface Guide*.

Retrieving Deleted Host and Corresponding Management Agent Targets

If you have deleted a host target and the corresponding Management Agent target, you can retrieve both of them. To do so, follow these steps:

Discover and add the host and the Management Agent by running the following command from the agent instance home of the corresponding host:

```
$ emctl config agent addInternalTargets
```

Once the host and the Management Agent are discovered and added to the console, add each target on that host as targets to be monitored in the console, by running the following EM CLI verb:

```
$ emcli add_target
      -name="name"
      -type="type"
      -host="hostname"
      [-properties="pname1:pval1;pname2:pval2;..."]
      [-separator=properties="sep_string"]
      [-subseparator=properties="subsep_string"]
      [-credentials="userpropname:username;pwdpropname:password;..."]
      [-input_file="parameter_tag:file_path"]
      [-display_name="display_name"]
      [-groups="groupname1:grouptype1;groupname2:grouptype2;..."]
      [-timezone_region="gmt_offset"]
      [-monitor_mode="monitor_mode"]
      [-instances="rac_database_instance_target_name1:target_type1;..."]
      [-force]
      [-timeout="time_in_seconds"]

[ ] indicates that the parameter is optional
```

For more information, see Verb Reference in the *Oracle Enterprise Manager Command Line Interface Guide*.

2

Discovering and Adding Database Targets

This chapter describes how you can discover and add database targets to be managed by Enterprise Manager. In particular, this chapters covers the following:

Discovering and Monitoring using Remote Agents

Starting with Enterprise Manager 24ai, databases can be monitored and managed using Remote Agents. This involves deploying the remote agent, then adding all of the hosts of your database system targets to the remote agent. Next discover your database targets in the same way that you would do if you were using a local agent on each host.

> **Note:**
>
> When choosing the remote agent for database monitoring, the general best practice is to choose the same remote agent for all members of your database system (e.g. database, listener, cluster, ASM, etc.). However, if the database target is not in close network proximity to the remote agent, use a different remote agent instead.
>
> For example, if you have primary and standby database targets in different datacenters, choose a remote agent in each respective datacenter that is in close network proximity to the primary or standby database.

Monitoring Using Remote Agents

Follow these steps to remotely discover and monitor a database system using the Remote Agent.

1. Deploy the agent.
   Deploy the remote agent to a host dedicated to remote agent use. For more information, see Installing Oracle Remote Management Agents.

2. Add all the hosts of your database system.
   Add all of the hosts of your database system as remotely monitored hosts to the same remote agent (best practice).

   Remember to choose a remote agent that is in close network proximity to your database system.

   For example, for a database system with a two node RAC, add both hosts of your two node RAC to the remote agent.

   These hosts will be marked as monitored by a Remote Agent. For more information, see Adding Remote Hosts to the Remote Agent.
3. Discover the database targets
   Discover your database targets using the same discovery workflow that you would use if you were using a local agent.

   All database discovery options, such as supported Auto-discovery, guided discovery, manual discovery, emcli add_target, emcli discover_db, are supported

> **Note:**
>
> Engineered Systems targets are not supported with Remote Agents.
>
> The remote agents use remote protocols to monitor targets and SSH for running commands on hosts.

Enabling Autodiscovery of Database Targets

Autodiscovery of database targets is enabled by default. If autodiscovery has been disabled, you can enable it, by following these steps:

1. From the **Setup** menu, select **Add Target,** then select **Configure Auto Discovery.**
2. On the Configure Auto Discovery page, in the Agent-based Auto Discovery table, select **Oracle Database, Listener, and Automatic Storage Management.**

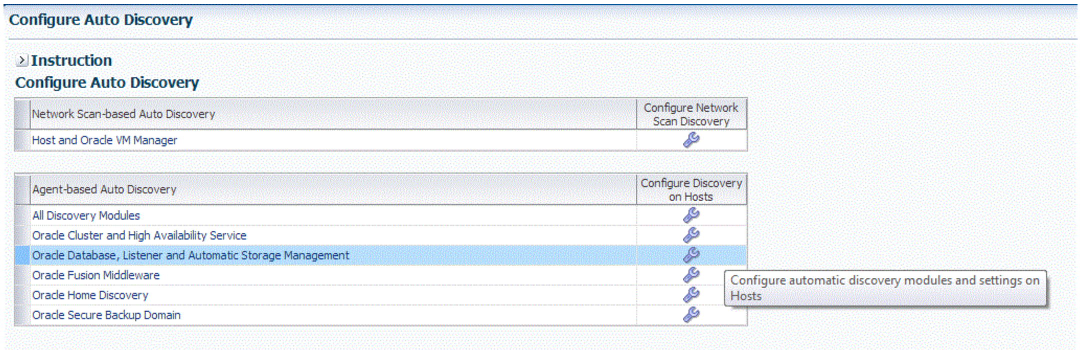

Configure Auto Discovery

Instruction

Configure Auto Discovery

Network Scan-based Auto Discovery	Configure Network Scan Discovery
Host and Oracle VM Manager	🔧

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

3. On the Configure Target Discovery page, click **Add Host.**

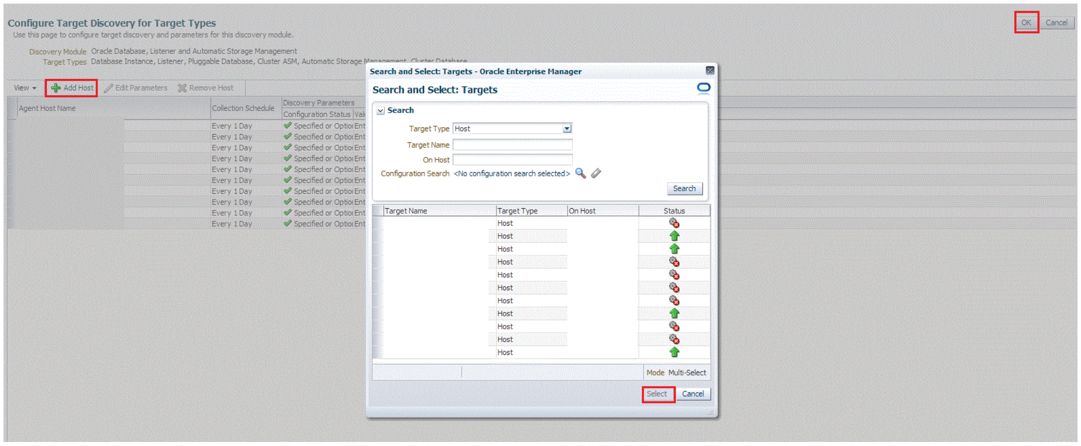

Configure Target Discovery for Target Types

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

Discovery Module Oracle Database, Listener and Automatic Storage Management
Target Types Database Instance, Listener, Pluggable Database, Cluster ASM, Automatic Storage Management, Cluster Database

Agent Host Name	Collection Schedule	Discovery Parameters	
		Configuration Status	Val
	Every 1 Day	✓ Specified or Optio	Ent
	Every 1 Day	✓ Specified or Optio	Ent
	Every 1 Day	✓ Specified or Optio	Ent
	Every 1 Day	✓ Specified or Optio	Ent
	Every 1 Day	✓ Specified or Optio	Ent
	Every 1 Day	✓ Specified or Optio	Ent
	Every 1 Day	✓ Specified or Optio	Ent
	Every 1 Day	✓ Specified or Optio	Ent
	Every 1 Day	✓ Specified or Optio	Ent
	Every 1 Day	✓ Specified or Optio	Ent

Search and Select: Targets - Oracle Enterprise Manager

Search and Select: Targets

Search

Target Type Host
Target Name
On Host
Configuration Search <No configuration search selected>

Target Name	Target Type	On Host	Status
	Host		⊗
	Host		↑
	Host		↑
	Host		⊗
	Host		⊗
	Host		⊗
	Host		⊗
	Host		↑
	Host		⊗
	Host		⊗
	Host		↑

Mode Multi-Select

4. From the Search and Select Targets dialog box, select a target that you want to be configured, and click **Select.**

5. You can click **Edit Parameters** to edit the parameters of the target.

   Click **OK.**

Discovering and Adding Container Database and Pluggable Database Targets

This section describes the different methods in which you can discover, promote, and add container database (CDB) and pluggable database (PDB) targets in Enterprise Manager. In particular, this section covers the following:

- Discovering CDB and PDB Targets Using Autodiscovery
- Adding CDB and PDB Targets Using the Guided Discovery Process
- Adding CDB and PDB Targets By Using the Declarative Process

Discovering CDB and PDB Targets Using Autodiscovery

Autodiscovery of databases is enabled by default. If autodiscovery has been disabled, follow the steps described in Enabling Autodiscovery of Database Targets.

> **Note:**
>
> A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
> A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

To promote a CDB target and its associated PDB targets using automatic discovery, follow these steps:

> **Note:**
>
> By default, promoting a CDB target also promotes all its associated discovered PDB targets. Also, by default, Enterprise Manager runs a background job (every 24 hours) to automatically discover and promote newly created PDB targets present on managed hosts. Hence, to discover and promote PDB targets, you only need to discover and promote the associated CDB target, as described in this section.

1. From the **Setup** menu, select **Add Target,** then select **Auto discovery Results.** Click **Agent-based Targets.**
2. Search for and select the Database Instance target that you want to promote, then click **Promote.**

   > **Note:**
   >
   > You can also discover Application Root PDBs. However, they are not explicitly shown as Application Root Pluggable Database, but just as Pluggable Database.

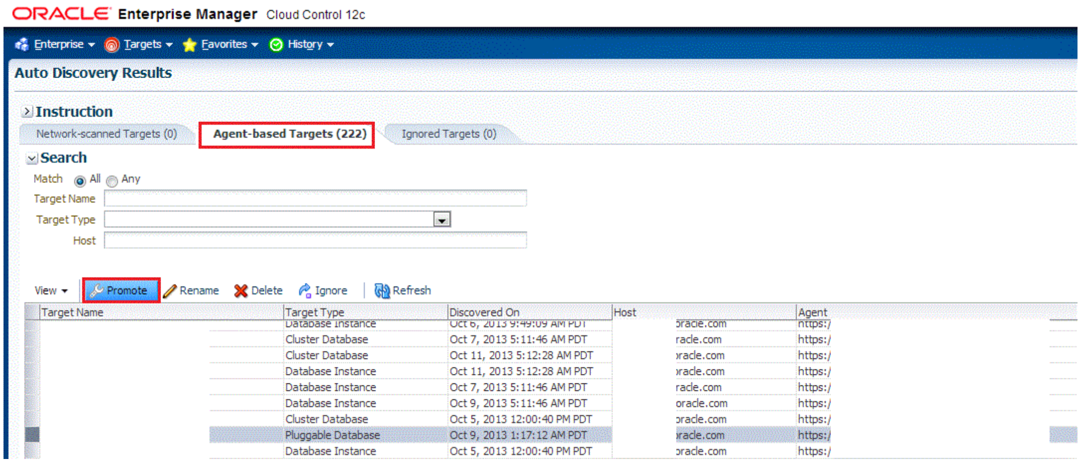
ORACLE Enterprise Manager Cloud Control 12c
Enterprise
Targets
Favorites
History
Network-scanned Targets (0)
Agent-based Targets (222)
Ignored Targets (0)
View
Promote
Rename
Delete
Ignore
Refresh

Auto Discovery Results

Instruction

Search

Match ● All ○ Any

Target Name

Target Type

Host

Target Name	Target Type	Discovered On	Host	Agent
	Database Instance	Oct 6, 2013 9:49:09 AM PDT	oracle.com	https:/
	Cluster Database	Oct 7, 2013 5:11:46 AM PDT	racle.com	https:/
	Cluster Database	Oct 11, 2013 5:12:28 AM PDT	racle.com	https:/
	Database Instance	Oct 11, 2013 5:12:28 AM PDT	racle.com	https:/
	Database Instance	Oct 7, 2013 5:11:46 AM PDT	racle.com	https:/
	Database Instance	Oct 9, 2013 5:11:46 AM PDT	oracle.com	https:/
	Cluster Database	Oct 5, 2013 12:00:40 PM PDT	oracle.com	https:/
	Pluggable Database	Oct 9, 2013 1:17:12 AM PDT	racle.com	https:/
	Database Instance	Oct 5, 2013 12:00:40 PM PDT	oracle.com	https:/

3. On the Promote Target: Results page, under the Databases section, select the CDB target.

   By default, selecting a CDB target for promotion also selects all its associated and discovered PDB (and Application Root PDB) targets for promotion. If you want to add or remove a target from the ones selected for promotion, select the CDB target, then click **Configure.**

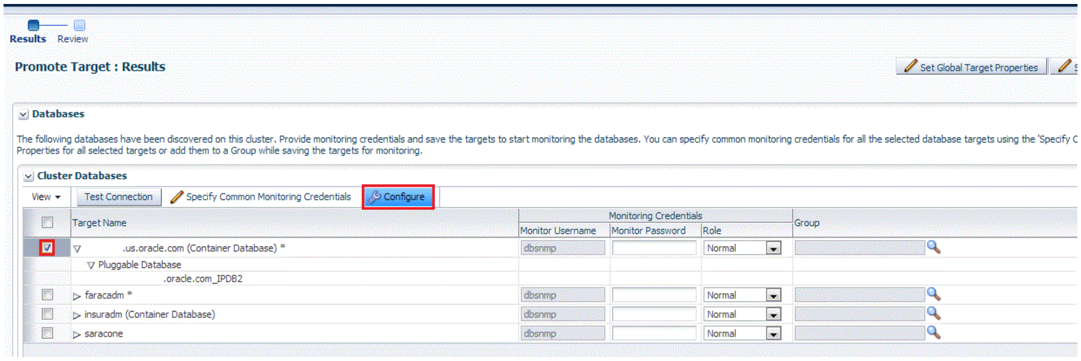
Results
Review
Set Global Target Properties
View
Test Connection
Specify Common Monitoring Credentials
Configure

Promote Target : Results

Databases

The following databases have been discovered on this cluster. Provide monitoring credentials and save the targets to start monitoring the databases. You can specify common monitoring credentials for all the selected database targets using the 'Specify C Properties for all selected targets or add them to a Group while saving the targets for monitoring.

Cluster Databases

□	Target Name	Monitoring Credentials			Group
		Monitor Username	Monitor Password	Role	
☑	▽ .us.oracle.com (Container Database) *	dbsnmp		Normal	
	▽ Pluggable Database				
	.oracle.com_IPDB2				
□	▷ faracadm *	dbsnmp		Normal	
□	▷ insuradm (Container Database)	dbsnmp		Normal	
□	▷ saracone	dbsnmp		Normal	

Select the **Pluggable Databases** tab, then click **Add** or **Remove.** Click **Save.**

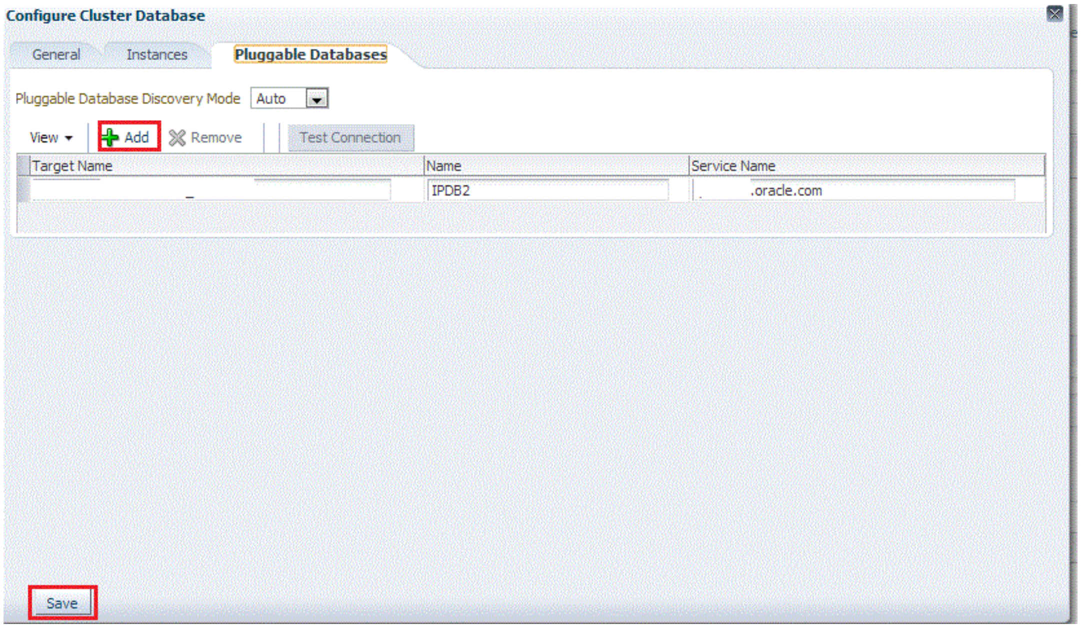

Pluggable Database Discovery Mode: Auto

Target Name	Name	Service Name
	IPDB2	.oracle.com

Enterprise Manager runs a background job (every 24 hours) to automatically discover and promote newly created PDB targets present on managed hosts. If you do not want Enterprise Manager to automatically promote the PDB targets associated with a particular CDB target, and instead want to promote them manually, select the CDB target on the Promote Target: Results page, then click **Configure.** Select the **Pluggable Databases** tab, then select Manual for **Pluggable Database Discovery Mode.** Click **Save.**

4. Specify the monitoring credentials for the selected CDB target, that is, the user name, password, and role. Also, if you want the selected target to be added to a group, specify a value for **Group.**

   If you specify Normal for **Role**, then the user name can be dbsnmp or a DB monitoring user. If you specify SYSDBA for Role, then you can provide any SYSDBA user.

5. Click **Test Connection** to test the connection made to the CDB target using the specified monitoring credentials.

6. To specify global target properties for all the targets you have selected on the Promote Target: Results page, click **Set Global Target Properties,** specify the required properties, then click **OK.**

   To specify a common group for all the targets you have selected on the Promote Target: Results page, click **Specify Group for Targets,** select a group, then click **Select.**

7. Click **Next.**

8. Review the displayed information, then click **Submit.**

Adding CDB and PDB Targets Using the Guided Discovery Process

To add a CDB target and its associated PDB (and Application Root PDB) targets using a guided discovery process, follow these steps:

> **Note:**
>
> - You can also add Application Root PDBs. However, they are not explicitly shown as Application Root Pluggable Database, but just as Pluggable Database.
> - A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
>   A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

1. From the **Setup** menu, select **Add Target,** then select **Add Targets Manually.**
2. On the Add Targets Manually page, select **Add Using Guided Process.**
3. In the Add Using Guided Process dialog box, select **Oracle Database, Listener, and Automatic Storage Management.** Click **Add...**
4. Select **Add Targets Using Guided Process (Also Adds Related Targets).** For **Target Types,** select **Oracle Database, Listener, and Automatic Storage Management.** Click **Add Using Guided Process.**
5. On the Database Discovery: Search Criteria page, for **Specify Host or Cluster,** specify the CDB host.

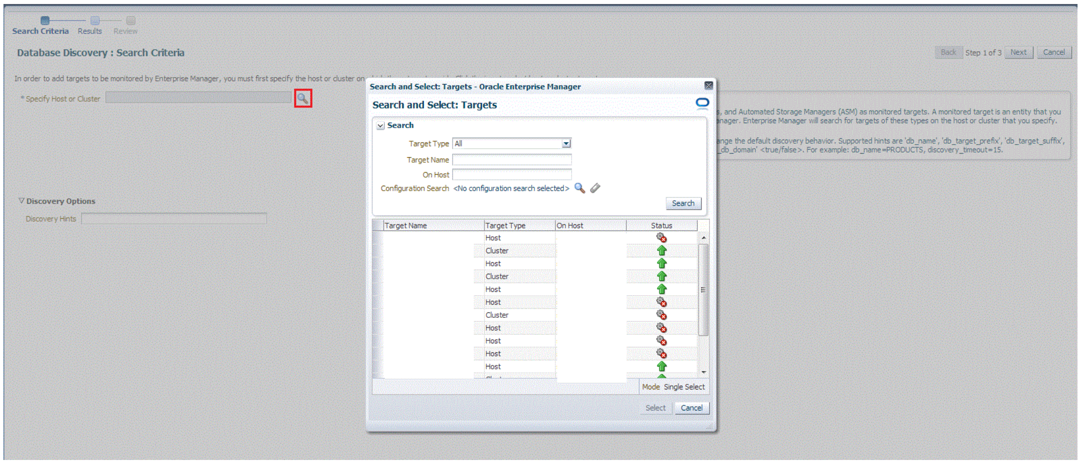

Database Discovery : Search Criteria

In order to add targets to be monitored by Enterprise Manager, you must first specify the host or cluster on

s, and Automated Storage Managers (ASM) as monitored targets. A monitored target is an entity that you
anager. Enterprise Manager will search for targets of these types on the host or cluster that you specify.

ange the default discovery behavior. Supported hints are 'db_name', 'db_target_prefix', 'db_target_suffix',
_db_domain' <true/false>. For example: db_name=PRODUCTS, discovery_timeout=15.

* Specify Host or Cluster

Discovery Options

Discovery Hints

Search and Select: Targets - Oracle Enterprise Manager

Search and Select: Targets

Search

Target Type: All

Target Name

On Host

Configuration Search <No configuration search selected>

Target Name	Target Type	On Host	Status
	Host		⊗
	Cluster		⬆
	Host		⬆
	Cluster		⬆
	Host		⬆
	Host		⊗
	Cluster		⊗
	Host		⊗
	Host		⊗
	Host		⊗
	Host		⬆

Mode Single Select

If the host you select is a member of a cluster, then use this page to also indicate whether the PDB targets should be searched only on the selected host or on all hosts within the cluster.

After you click the search icon, a Select Targets dialog box appears. Specify the target name or select a target from the target table. You can filter the search to search for all down targets or up targets, by clicking the filter arrows in the Status column.

You can also configure your search by clicking the Configuration Search icon. After you select the CDB, click **Select.**

You get an option to choose if you want to search for pluggable database targets only on the current host that you selected, or on all the hosts in the cluster that the host target is a member of.

This option enables you to save time, if you want to search for PDB targets only on the current host selected.

The Discovery Options section enables you to specify discovery hints. You can specify additional discovery options to change the default discovery behavior. Supported hints are `db_name, db_target_prefix, db_target_suffix, discovery_timeout <in seconds per host>`, and `no_db_domain.`

For example,

```
db_name=PRODUCTS, discovery timeout=15.
```

6. Click **Next.**
7. On the Database Discovery: Results page, under the Databases section, select the CDB target.

   By default, selecting a CDB target for promotion also selects all its associated and discovered PDB (and Application Root PDB) targets for promotion. If you want to add or remove a PDB target from the ones selected for promotion, select the CDB target, then click **Configure.** Select the **Pluggable Databases** tab, then click **Add** or **Remove.** Click **Save.**

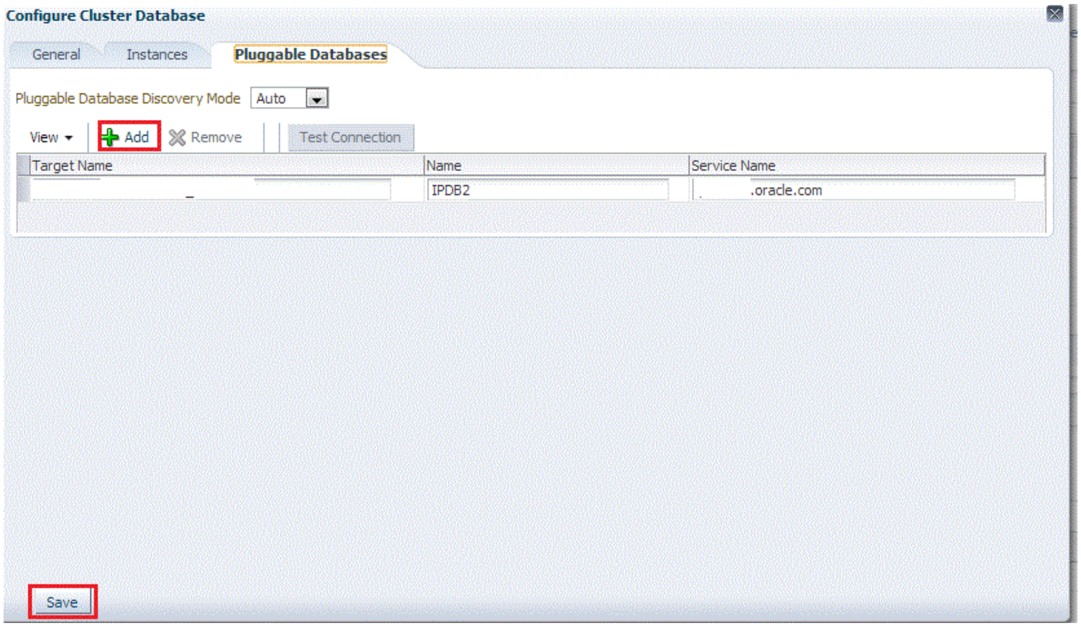

Configure Cluster Database

Pluggable Database Discovery Mode: Auto

Target Name	Name	Service Name
	IPDB2	.oracle.com

Enterprise Manager runs a background job (every 24 hours) to automatically discover and promote newly created PDB targets present on managed hosts. If you do not want Enterprise Manager to automatically promote the PDB targets associated with a particular CDB target, and instead want to promote them manually, select the CDB target on the Promote Target: Results page, then click **Configure.** Select the **Pluggable Databases** tab, then select Manual for **Pluggable Database Discovery Mode.** Click **Save.**

8. Specify the monitoring credentials for the selected CDB target, that is, the user name, password, and role. Also, if you want the selected target to be added to a group, specify a value for **Group.**

   If you specify Normal for **Role**, then the user name can be dbsnmp or a DB monitoring user. If you specify SYSDBA for Role, then you can provide any SYSDBA user.

9. Click **Test Connection** to test the connection made to the CDB target using the specified monitoring credentials.

10. To specify global target properties for all the targets you have selected on the Promote Target: Results page, click **Set Global Target Properties,** specify the required properties, then click **OK.**

    To specify a common group for all the targets you have selected on the Promote Target: Results page, click **Specify Group for Targets,** select a group, then click **Select.**

    If you have selected multiple databases and you want to set the same monitoring properties for all of them, select **Specify Common Monitoring Credentials.** Enter the monitoring credentials, monitoring password, and role. Click **Apply.**

11. Click **Next.**

12. Review the displayed information, then click **Submit.**

Adding CDB and PDB Targets By Using the Declarative Process

To add a CDB target and its associated PDB (and Application Root) targets by specifying target monitoring properties, follow these steps:

> **Note:**
>
> A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
> A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

1. From the **Setup** menu, select **Add Target,** then select **Add Targets Manually.**

2. On the Add Targets Manually page, select **Add Target Declaratively.**

3. In the Add Target Declaratively dialog box, select **Pluggable Database,** and click **Add...**

   For **Monitoring Agent,** specify the Management Agent present on the CDB host. Click **Add Manually.**

4. On the Add Pluggable Database: Specify Container Database page, specify the CDB or click the search icon to select the CDB to which the target will be added.

   Click **Continue.**

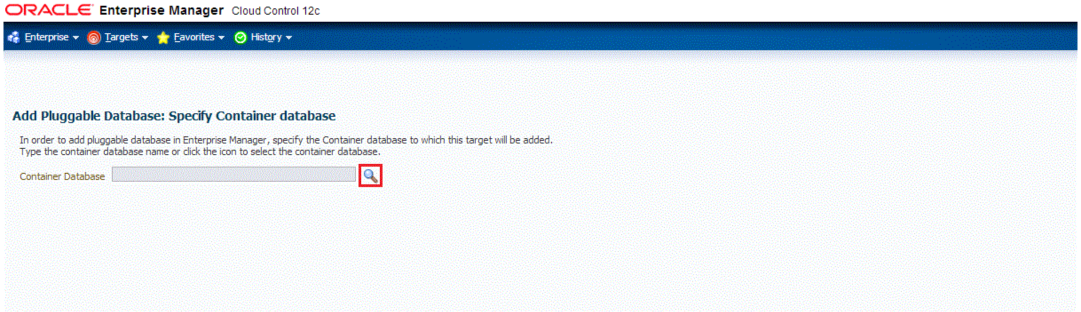

Add Pluggable Database: Specify Container database

In order to add pluggable database in Enterprise Manager, specify the Container database to which this target will be added.
Type the container database name or click the icon to select the container database.

Container Database

5. On the **Add Pluggable Database: Properties** page, specify a unique name for the target, the name, the default service name, the OMS preferred string connection, and the Agent preferred connect string.

> **Note:**
>
> In Enterprise Manager Release Update 12 and later, Service Name is now supported through Preferred Connection Strings. For more information, please view the following Preferred Connect Strings.

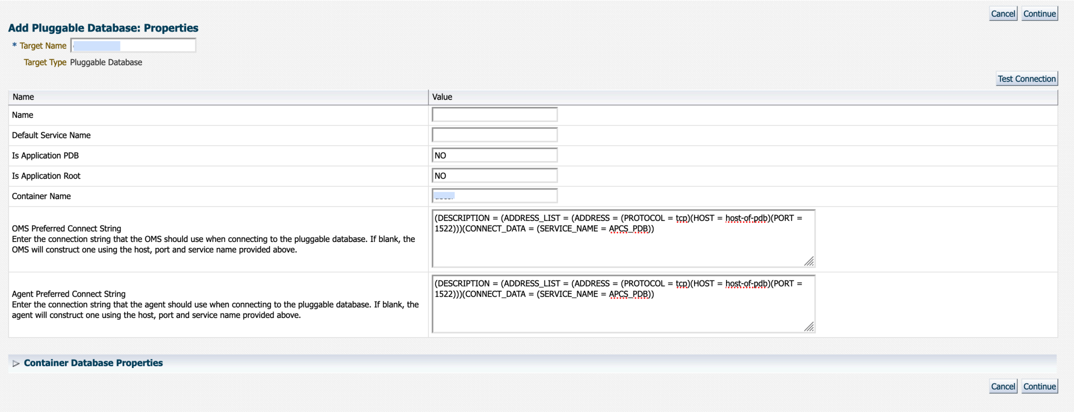

Add Pluggable Database: Properties

* Target Name

Target Type Pluggable Database

Name	Value
Name	
Default Service Name	
Is Application PDB	NO
Is Application Root	NO
Container Name	
OMS Preferred Connect String Enter the connection string that the OMS should use when connecting to the pluggable database. If blank, the OMS will construct one using the host, port and service name provided above.	(DESCRIPTION = (ADDRESS_LIST = (ADDRESS = (PROTOCOL = tcp)(HOST = host-of-pdb)(PORT = 1522)))(CONNECT_DATA = (SERVICE_NAME = APCS_PDB))
Agent Preferred Connect String Enter the connection string that the agent should use when connecting to the pluggable database. If blank, the agent will construct one using the host, port and service name provided above.	(DESCRIPTION = (ADDRESS_LIST = (ADDRESS = (PROTOCOL = tcp)(HOST = host-of-pdb)(PORT = 1522)))(CONNECT_DATA = (SERVICE_NAME = APCS_PDB))

▷ Container Database Properties

Expand the Container Database section to verify the properties of the CDB.

Click **Test Connection** to test the connection made to the CDB target using the specified monitoring credentials.

Click **Continue.**

6. Specify the required information on each page and then click **Next,** until you reach the Review page.
7. Review the displayed information, and then click **Submit.**

Discovering and Adding Cluster Database Targets

This section describes the different methods in which you can discover and add cluster database targets. In particular, this section cover the following:

- Discovering Cluster Database Targets Using Autodiscovery
- Adding Cluster Database Targets Using the Guided Discovery Process
- Adding Cluster Database Targets By Using the Declarative Process

Discovering Cluster Database Targets Using Autodiscovery

Autodiscovery of cluster databases is enabled by default. If autodiscovery has been disabled, follow the steps described in Enabling Autodiscovery of Database Targets.

Oracle Clusterware is a high availability framework that provides the infrastructure necessary to run the cluster systems called Oracle Real Application Clusters (RAC). It manages RAC resources, such as virtual IP (VIP) addresses, databases, listeners, and storage, and it enables them to communicate with each other, so that they appear to function as a single system (cluster).

Enterprise Manager discovers a cluster as a Cluster Database target. A database system and Oracle High Availability Services are automatically associated with a Cluster Database. A database system contains a primary database and related targets such as listeners, Automatic Storage Management (ASM) and Oracle Homes. The Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.

Starting with Oracle Clusterware 23ai, you can use the REST APIs to gather clusterware events information. Events are critical for monitoring and managing RAC, as they indicate changes in cluster state and resource health.

To promote cluster database targets using autodiscovery, follow these steps:

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

   From the results table, from the Agent-based targets tab, select the Cluster Database and Oracle High Availability Service that you want to add for monitoring, and click **Promote.**

2. The Promote Targets: Result page displays the databases discovered on the cluster. Select the cluster database

   On the Promote Target: Results page, under the Databases section, in the Cluster Databases section, select the cluster database target that you want to promote.

   By default, selecting the cluster database target for promotion also selects all its associated discovered database instance targets for promotion. If you want to add or remove a database instance target from the ones selected for promotion, select the cluster database target, then click **Configure.** Select the **Instances** tab, then click **Add** or **Remove.** Click **Save.**

3. In the Cluster Databases section, specify the monitoring credentials for the selected cluster database target, that is, the Monitor user name, Monitor password, and role. Also, if you want the selected target to be added to a group, specify a value for **Group.**

   If you specify Normal for **Role**, then the user name can be dbsnmp or a DB monitoring user. If you specify SYSDBA for Role, then you can provide any SYSDBA user.

> **Note:**
>
> Enterprise Manager lets you use the job system to automate database password change for the monitoring user. See Automate Monitoring and Non-monitoring User Password Management for more information.

4. Click **Test Connection** to test the connection made to the cluster database target using the specified monitoring credentials.
5. To specify global target properties for all the targets you have selected on the Promote Target: Results page, click **Set Global Target Properties,** specify the required properties, then click **OK.**

   To specify a common group for all the targets you have selected on the Promote Target: Results page, click **Specify Group for Targets,** select a group, then click **Select.**
6. If you have selected multiple databases and you want to set the same monitoring properties for all of them, select **Specify Common Monitoring Credentials.** Enter the monitoring credentials, monitoring password, and role. Click **Apply.**
7. Click **Next.**
8. Review the displayed information, then click **Submit.**

Adding Cluster Database Targets Using the Guided Discovery Process

Oracle Clusterware is a high availability framework that provides the infrastructure necessary to run the cluster systems called Oracle Real Application Clusters (RAC). It manages RAC resources, such as virtual IP (VIP) addresses, databases, listeners, and storage, and it enables them to communicate with each other, so that they appear to function as a single system (cluster).

Enterprise Manager discovers a cluster as a Cluster Database target. A database system and Oracle High Availability Services are automatically associated with a Cluster Database. A database system contains a primary database and related targets such as listeners, Automatic Storage Management (ASM) and Oracle Homes. The Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.

Starting with Oracle Clusterware 23ai, you can use the REST APIs to gather clusterware events information. Events are critical for monitoring and managing RAC, as they indicate changes in cluster state and resource health.

Prerequisites:

For Clusterware 23ai, to use REST based methods, using the command below in the root folder, set the password for the events user, which will be used to monitor the clusterware events using the clusterware REST API.

```
./srvctl modify cdp -passfile_events <location of password file>
```

To add cluster database targets using the guided process, follow these steps:

1. From the Setup menu, select **Add Target,** then select **Add Targets Manually.**
2. On the Add Targets Manually page, select **Add Using Guided Process.**
3. In the Add Using Guided Process dialog box, select the cluster to be discovered, and select **Add...**

4. On the **Specify Host** page, select the search icon and select the host on which the target resides from the **Select Targets** dialog box and click **Select**.

   On the **Select Targets** dialog box, you can filter the search for all down targets or up targets, by clicking the filter arrows in the Status column.

   You can also configure your search by clicking the Configuration Search icon. After you select the host target, click **Select.**

5. Select **Discover Target**.
6. After the discovery cluster process is completed, if the target cluster is using Oracle Database 23ai, a **Clusterware REST API Properties** tab will be available. Setting the REST API properties is optional, but recommended.
   - Specify the password for the Clusterware events user. All other fields should be prepopulated. Select **Save**.
7. By default, selecting a cluster database target for promotion also selects all its associated discovered database instance targets for promotion. If you want to add or remove a database instance target from the ones selected for promotion, select the cluster database target, then click **Add** or **Remove**.
8. To specify global target properties for all the targets you have selected click **Set Global Target Properties,** specify the required properties, then click **OK**. Select **Save**.
9. Select **Discover Target**.

Adding Cluster Database Targets By Using the Declarative Process

Oracle Clusterware is a high availability framework that provides the infrastructure necessary to run the cluster systems called Oracle Real Application Clusters (RAC). It manages RAC resources, such as virtual IP (VIP) addresses, databases, listeners, and storage, and it enables them to communicate with each other, so that they appear to function as a single system (cluster).

Enterprise Manager discovers a cluster as a Cluster Database target. A database system and Oracle High Availability Services are automatically associated with a Cluster Database. A database system contains a primary database and related targets such as listeners, Automatic Storage Management (ASM) and Oracle Homes. The Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.

Starting with Oracle Clusterware 23ai, you can use the REST APIs to gather clusterware events information. Events are critical for monitoring and managing RAC, as they indicate changes in cluster state and resource health.

Prerequisites:

For Clusterware 23ai, to use REST based methods, using the command below in the root folder, set the password for the events user, which will be used to monitor the clusterware events using the clusterware REST API.

```
./srvctl modify cdp -passfile_events <location of password file>
```

To add a cluster database target declaratively by specifying target monitoring properties, follow these steps:

1. From the Setup menu, select **Add Target,** then select **Add Targets Manually.**
2. On the Add Targets Manually page, select **Add Target Manually.**

3. In the **Add Target Manually** dialog box, select the search icon and select the host on which the target resides from the **Select Host** dialog box and click **Select**. Select **Cluster** from the **Target Type** list, and select **Add...**.
4. Specify the Cluster Target Properties: Target Name, Oracle Home, SCAN Name, SCAN Port, ONS Port.
5. If the target cluster is using Oracle Database 23ai, a **Clusterware REST API Properties** tab will be available. Setting the REST API properties is optional, but recommended.
   - Specify the cdp port and password for the Clusterware events user. All other fields should be prepopulated. Select **Save**.
6. By default, selecting a cluster database target for promotion also selects all its associated discovered database instance targets for promotion. If you want to add or remove a database instance target from the ones selected for promotion, select the cluster database target, then click **Add** or **Remove**.
7. To specify global target properties for all the targets you have selected click **Set Global Target Properties,** specify the required properties, then click **OK**. Select **Save**.
8. Select **Discover Target**.

Discovering and Adding Single Instance Database Targets

This section describes the different methods in which you can discover and add single instance database targets. In particular, this section covers the following:

- Discovering Single Instance Database Targets Using Autodiscovery
- Adding Single Instance Database Targets Using Guided Discovery Process
- Adding Single Instance Database Targets By Using the Declarative Process

Discovering Single Instance Database Targets Using Autodiscovery

Autodiscovery of databases is enabled by default. If autodiscovery has been disabled, follow the steps described in Enabling Autodiscovery of Database Targets.

> **Note:**
>
> A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
> A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

To promote single instance database targets, follow these steps:

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

From the results table, from the Agent-based targets tab, select the discovered database instance target that you want to add for monitoring, and click **Promote.**

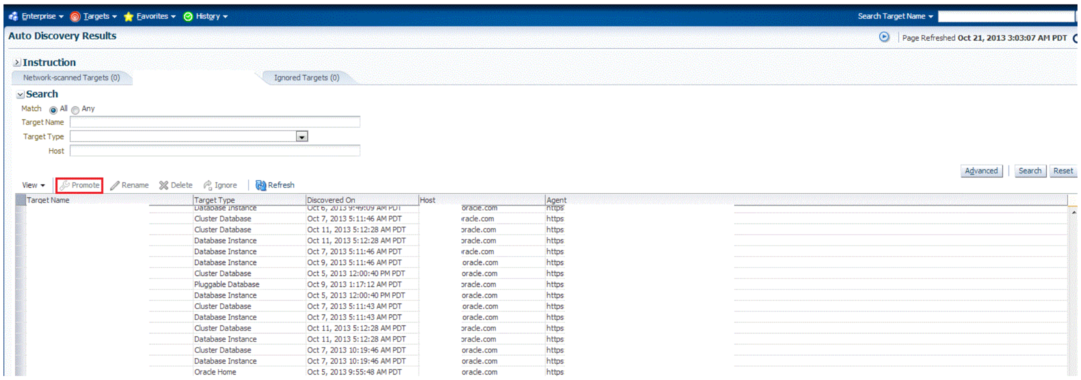

Auto Discovery Results

Search

Match ● All ○ Any

Target Name

Target Type

Host

Target Name	Target Type	Discovered On	Host	Agent
	Database Instance	Oct 6, 2013 9:49:09 AM PDT	oracle.com	https
	Cluster Database	Oct 7, 2013 5:11:46 AM PDT	racle.com	https
	Cluster Database	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Database Instance	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Database Instance	Oct 7, 2013 5:11:46 AM PDT	racle.com	https
	Database Instance	Oct 9, 2013 5:11:46 AM PDT	oracle.com	https
	Cluster Database	Oct 5, 2013 12:00:40 PM PDT	oracle.com	https
	Pluggable Database	Oct 9, 2013 1:17:12 AM PDT	oracle.com	https
	Database Instance	Oct 5, 2013 12:00:40 PM PDT	oracle.com	https
	Cluster Database	Oct 7, 2013 5:11:43 AM PDT	oracle.com	https
	Database Instance	Oct 7, 2013 5:11:43 AM PDT	oracle.com	https
	Cluster Database	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Database Instance	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Cluster Database	Oct 7, 2013 10:19:46 AM PDT	oracle.com	https
	Database Instance	Oct 7, 2013 10:19:46 AM PDT	oracle.com	https
	Oracle Home	Oct 5, 2013 9:55:48 AM PDT	oracle.com	https

2. The Promote Targets:Result page displays the databases Select the database instance.

   On the Promote Target: Results page, under the Databases section, in the Single Instance Databases section, select the database instance target that you want to promote.

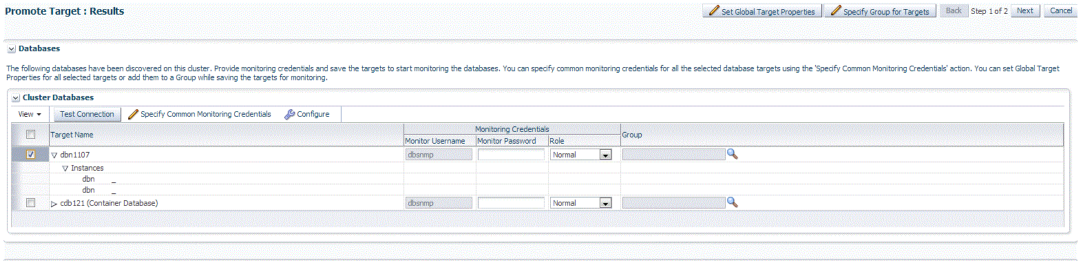
Set Global Target Properties
Specify Group for Targets
Back
Step 1 of 2
Next
Cancel
View
Test Connection
Specify Common Monitoring Credentials
Configure

Promote Target : Results

Databases

The following databases have been discovered on this cluster. Provide monitoring credentials and save the targets to start monitoring the databases. You can specify common monitoring credentials for all the selected database targets using the 'Specify Common Monitoring Credentials' action. You can set Global Target Properties for all selected targets or add them to a Group while saving the targets for monitoring.

Cluster Databases

□	Target Name	Monitoring Credentials			Group
		Monitor Username	Monitor Password	Role	
☑	▽ dbn1107	dbsnmp		Normal	
	▽ Instances				
	dbn _				
	dbn _				
□	▷ cdb121 (Container Database)	dbsnmp		Normal	

3. Specify the monitoring credentials for the selected database instance target, that is, the Monitor user name, Monitor password, and role. Also, if you want the selected target to be added to a group, specify a value for **Group.**

   If you specify Normal for **Role**, then the user name can be dbsnmp or a DB monitoring user. If you specify SYSDBA for Role, then you can provide any SYSDBA user.

   > **Note:**
   >
   > Enterprise Manager lets you use the job system to automate database password change for the monitoring user. See Automate Monitoring and Non-monitoring User Password Management for more information.

4. Click **Test Connection** to test the connection made to the database instance target using the specified monitoring credentials.

5. To specify global target properties for all the targets you have selected on the Promote Target: Results page, click **Set Global Target Properties,** specify the required properties, then click **OK.**

To specify a common group for all the targets you have selected on the Promote Target: Results page, click **Specify Group for Targets,** select a group, then click **Select.**

6. On the Configure Database Instance: Properties, specify a name and a database system name for the database instance target. Next, specify all the properties of the target.

> **Note:**
>
> In Enterprise Manager Release Update 12 and later, Service Name is now supported through Preferred Connection Strings. For more information, please view the following Preferred Connect Strings.

7. Click **Next.**
8. Review the displayed information, then click **Submit.**

Adding Single Instance Database Targets Using Guided Discovery Process

To add single instance database targets, follow these steps:

> **Note:**
>
> A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
> A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

1. From the Setup menu, select **Add Target,** then select **Add Targets Manually.**
2. On the Add Targets Manually page, select **Add Using Guided Process.**
3. In the Add Using Guided Process dialog box, select **Oracle Database, Listener, and Automatic Storage Management.** Click **Add...**
4. On the Database Discovery: Search Criteria page, specify the database instance host, or click on the **Specify Host or Cluster** search icon to select the host.

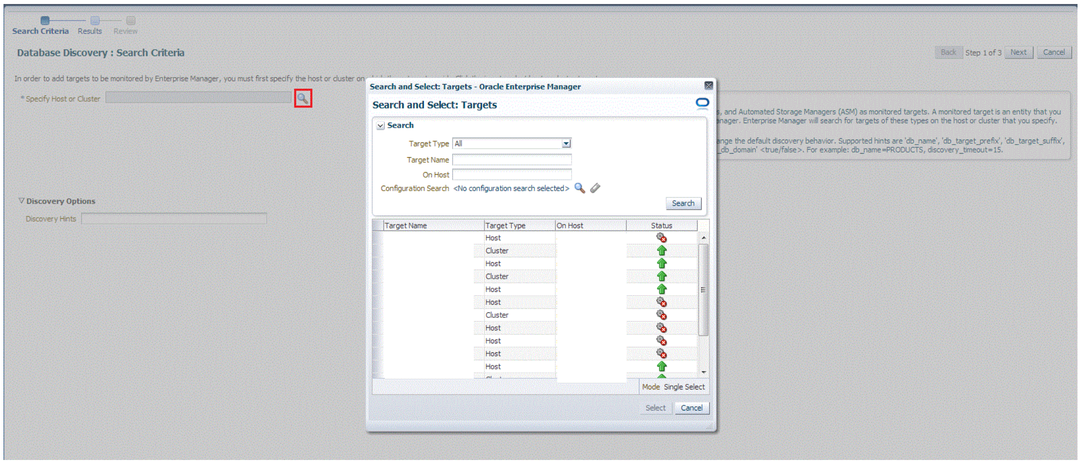

Database Discovery : Search Criteria

In order to add targets to be monitored by Enterprise Manager, you must first specify the host or cluster on

* Specify Host or Cluster

▽ Discovery Options

Discovery Hints

, and Automated Storage Managers (ASM) as monitored targets. A monitored target is an entity that you nager. Enterprise Manager will search for targets of these types on the host or cluster that you specify.

nge the default discovery behavior. Supported hints are 'db_name', 'db_target_prefix', 'db_target_suffix', _db_domain' <true/false>. For example: db_name=PRODUCTS, discovery_timeout=15.

Search and Select: Targets - Oracle Enterprise Manager

Search and Select: Targets

Search

Target Type: All

Target Name:

On Host:

Configuration Search <No configuration search selected>

Target Name	Target Type	On Host	Status
	Host		
	Cluster		↑
	Host		↑
	Cluster		↑
	Host		↑
	Host		
	Cluster		
	Host		
	Host		
	Host		
	Host		↑

Mode Single Select

If the host you select is a member of a cluster, then use this page to also indicate whether the databases should be searched only on the selected host or on all hosts within the cluster.

After you click the search icon, a Select Targets dialog box appears.Specify the target name or select a target from the target table. You can filter the search to search for all down targets or up targets, by clicking the filter arrows in the Status column.

You can also configure your search by clicking the Configuration Search icon. After you select the host, click **Select.**

You get an option to choose if you want to search for database targets only on the current host that you selected, or on all the hosts in the cluster that the host target is a member of.

This option enables you to save time, if you want to search for database targets only on the current host selected.

The Discovery Options section enables you to specify discovery hints. You can specify additional discovery options to change the default discovery behavior. Supported hints are `db_name, db_target_prefix, db_target_suffix, discovery_timeout <in seconds per host>`, and `no_db_domain`.

For example,

```
db_name=PRODUCTS, discovery timeout=15.
```

Click **Next** to proceed with the discovery process.

Click **Next.**

5. The Database Discovery:Result page displays the databases discovered on the cluster. Select the database

   On the Database Discovery: Results page, under the Databases section, in the Single Instance Databases section, select the database instance target that you want to promote.

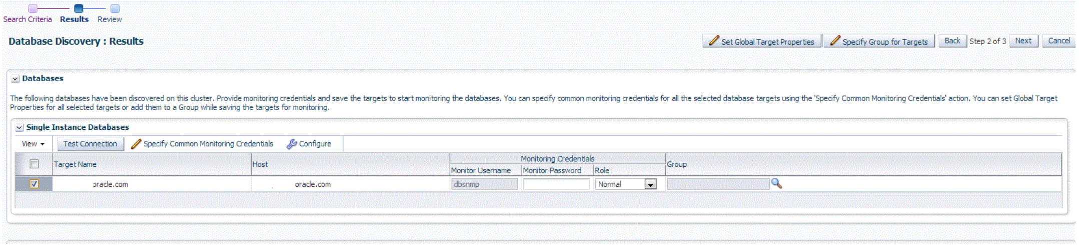
Search Criteria
Results
Review
Set Global Target Properties
Specify Group for Targets
Back
Step 2 of 3
Next
Cancel
View
Test Connection
Specify Common Monitoring Credentials
Configure

Database Discovery : Results

Databases

The following databases have been discovered on this cluster. Provide monitoring credentials and save the targets to start monitoring the databases. You can specify common monitoring credentials for all the selected database targets using the 'Specify Common Monitoring Credentials' action. You can set Global Target Properties for all selected targets or add them to a Group while saving the targets for monitoring.

Single Instance Databases

□	Target Name	Host	Monitoring Credentials			Group
			Monitor Username	Monitor Password	Role	
☑	oracle.com	oracle.com	dbsnmp		Normal	

6. On the Configure Database Instance: Properties, specify a name and a database system name for the database instance target. Next, specify all the properties of the target.

   > **Note:**
   >
   > In Enterprise Manager Release Update 12 and later, Service Name is now supported through Preferred Connection Strings. For more information, please view the following Preferred Connect Strings.

7. Click **Test Connection** to test the connection made to the database instance target using the specified monitoring credentials.
8. To specify global target properties for all the targets you have selected on the Promote Target: Results page, click **Set Global Target Properties,** specify the required properties, then click **OK.**

   To specify a common group for all the targets you have selected on the Promote Target: Results page, click **Specify Group for Targets,** select a group, then click **Select.**
9. If you have selected multiple databases and you want to set the same monitoring properties for all of them, select **Specify Common Monitoring Credentials.** Enter the monitoring credentials, monitoring password, and role. Click **Apply.**
10. Click **Next.**
11. Review the displayed information, then click **Submit.**

Adding Single Instance Database Targets By Using the Declarative Process

To add a single instance database target declaratively by specifying target monitoring properties, follow these steps:

> **Note:**
>
> A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
> A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

1. From the Setup menu, select **Add Target,** then select **Add Targets Manually.**
2. On the Add Targets Manually page, select **Add Target Declaratively.**
3. In the Add Target Declaratively dialog box, select **Database Instance.**
4. In the Monitoring Agent field, select the Management Agent monitoring the database.
5. Click **Add...**

6. On the Configure Database Instance: Properties, specify a name and a database system name for the database instance target. Next, specify all the properties of the target.

> **Note:**
>
> In Enterprise Manager Release Update 12 and later, Service Name is now supported through Preferred Connection Strings. For more information, please view the following Preferred Connect Strings.

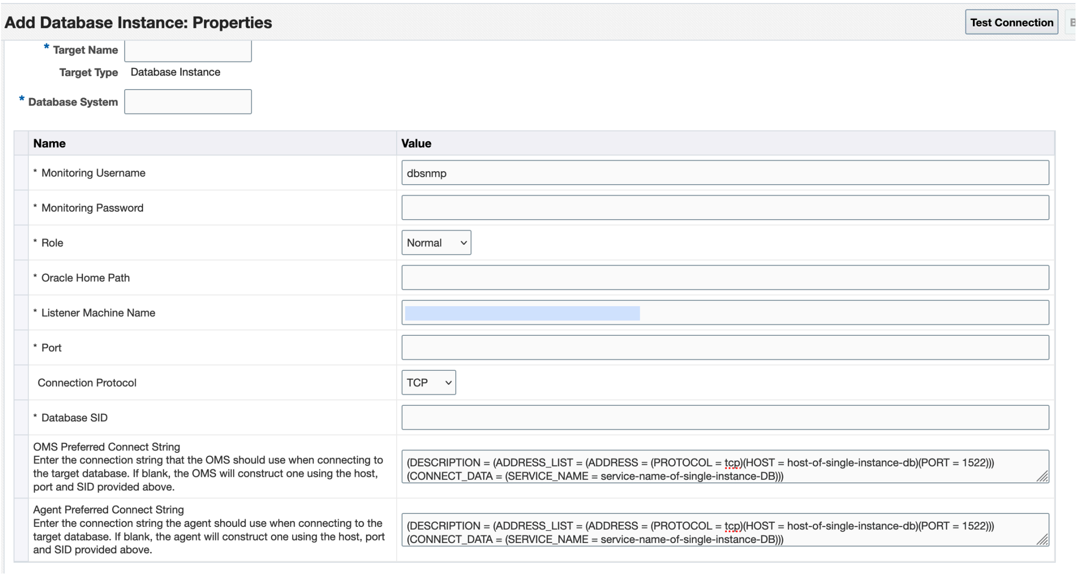

Add Database Instance: Properties

* Target Name	
Target Type	Database Instance
* Database System	

Name	Value
* Monitoring Username	dbsnmp
* Monitoring Password	
* Role	Normal
* Oracle Home Path	
* Listener Machine Name	
* Port	
Connection Protocol	TCP
* Database SID	
OMS Preferred Connect String Enter the connection string that the OMS should use when connecting to the target database. If blank, the OMS will construct one using the host, port and SID provided above.	(DESCRIPTION = (ADDRESS_LIST = (ADDRESS = (PROTOCOL = tcp)(HOST = host-of-single-instance-db)(PORT = 1522))) (CONNECT_DATA = (SERVICE_NAME = service-name-of-single-instance-DB)))
Agent Preferred Connect String Enter the connection string the agent should use when connecting to the target database. If blank, the agent will construct one using the host, port and SID provided above.	(DESCRIPTION = (ADDRESS_LIST = (ADDRESS = (PROTOCOL = tcp)(HOST = host-of-single-instance-db)(PORT = 1522))) (CONNECT_DATA = (SERVICE_NAME = service-name-of-single-instance-DB)))

Click **Next.**

7. Specify all the details required on the Install Packages, Credentials, and Parameters pages. Click **Next** after each page until you reach the Review page.

> **Note:**
>
> Enterprise Manager lets you use the job system to automate database password change for the monitoring user. See Automate Monitoring and Non-monitoring User Password Management for more information.

8. Review the displayed information, then click **Submit.**

Discovering and Adding Cluster Targets

This section describes the different methods in which you can discover and add cluster targets. In particular, this section covers the following:

- Discovering Cluster Targets Using Autodiscovery
- Adding Cluster Targets Using the Guided Discovery Process

- Adding Cluster Targets By Using the Declarative Process

Discovering Cluster Targets Using Autodiscovery

Autodiscovery of databases is enabled by default. If autodiscovery has been disabled, follow the steps described in Enabling Autodiscovery of Database Targets.

> **Note:**
>
> A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
> A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

To promote cluster targets, follow these steps:

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

   From the results table, from the Agent-based targets tab, select the discovered cluster target that you want to add for monitoring, and click **Promote.**

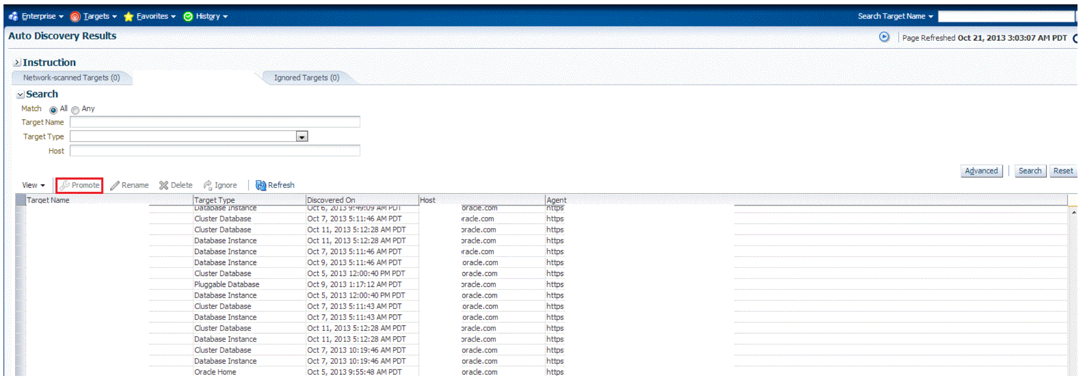

Auto Discovery Results

Match ● All ○ Any

Target Name

Target Type

Host

Target Name	Target Type	Discovered On	Host	Agent
	Database Instance	Oct 6, 2013 9:49:09 AM PDT	oracle.com	https
	Cluster Database	Oct 7, 2013 5:11:46 AM PDT	racle.com	https
	Cluster Database	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Database Instance	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Database Instance	Oct 7, 2013 5:11:46 AM PDT	racle.com	https
	Database Instance	Oct 9, 2013 5:11:46 AM PDT	oracle.com	https
	Cluster Database	Oct 5, 2013 12:00:40 PM PDT	oracle.com	https
	Pluggable Database	Oct 9, 2013 1:17:12 AM PDT	oracle.com	https
	Database Instance	Oct 5, 2013 12:00:40 PM PDT	oracle.com	https
	Cluster Database	Oct 7, 2013 5:11:43 AM PDT	oracle.com	https
	Database Instance	Oct 7, 2013 5:11:43 AM PDT	oracle.com	https
	Cluster Database	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Database Instance	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Cluster Database	Oct 7, 2013 10:19:46 AM PDT	oracle.com	https
	Database Instance	Oct 7, 2013 10:19:46 AM PDT	oracle.com	https
	Oracle Home	Oct 5, 2013 9:55:48 AM PDT	oracle.com	https

2. The Promote Targets:Result page displays the hosts discovered on the cluster. Select the host that you want to promote.

   > **Note:**
   >
   > A host can belong to only one cluster. If a particular host is not displayed, it can mean that the host belongs to another cluster.

On the Promote Target: Results page, in the Clusters section, select the cluster target that you want to promote.

By default, selecting the cluster target for promotion also selects all its associated discovered hosts for promotion. If you want to add or remove a host from the ones selected for promotion, select the cluster database target, then click **Configure.** Select the **Hosts** tab, then click **Add** or **Remove.** Click **Save.**

3. In the Clusters section, specify the monitoring credentials for the selected cluster target, that is, the Monitor user name, Monitor password, and role. Also, if you want the selected target to be added to a group, specify a value for **Group.**

   If you specify Normal for **Role**, then the user name can be *dbsnmp* or a DB monitoring user. If you specify SYSDBA for Role, then you can provide any SYSDBA user.

4. Click **Test Connection** to test the connection made to the cluster target using the specified monitoring credentials.
5. To specify global target properties for all the targets you have selected on the Promote Target: Results page, click **Set Global Target Properties,** specify the required properties, then click **OK.**

   To specify a common group for all the targets you have selected on the Promote Target: Results page, click **Specify Group for Targets,** select a group, then click **Select.**

6. If you have selected multiple databases and you want to set the same monitoring properties for all of them, select **Specify Common Monitoring Credentials.** Enter the monitoring credentials, monitoring password, and role. Click **Apply.**
7. Click **Next.**
8. Review the displayed information, then click **Submit.**

Adding Cluster Targets Using the Guided Discovery Process

To add cluster targets using the guided process, follow these steps:

> **Note:**
>
> A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
> A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

1. From the Setup menu, select **Add Target,** then select **Add Targets Manually.**
2. On the Add Targets Manually page, select **Add Using Guided Process.**
3. In the Add Using Guided Process dialog box, select **Oracle Database, Listener, and Automatic Storage Management.** Click **Add...**

4. On the Cluster Discovery: Specify Host page, specify the cluster host, or click on the **Specify Host or Cluster** search icon to select the cluster.

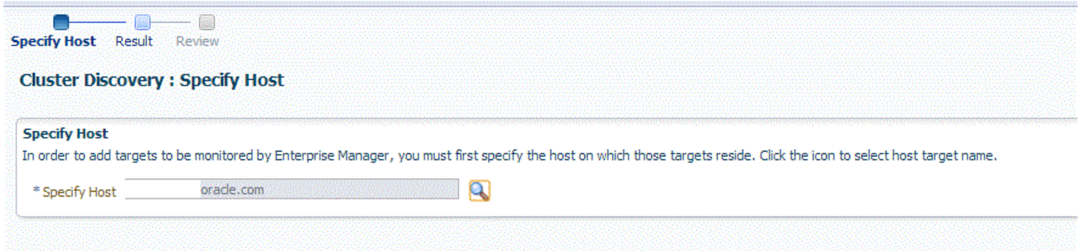

Cluster Discovery : Specify Host

Specify Host

In order to add targets to be monitored by Enterprise Manager, you must first specify the host on which those targets reside. Click the icon to select host target name.

* Specify Host oracle.com

If the host you select is a member of a cluster, then use this page to also indicate whether the clusters should be searched only on the selected host or on all hosts within the cluster.

After you click the search icon, a Select Targets dialog box appears.Specify the target name or select a target from the target table. You can filter the search to search for all down targets or up targets, by clicking the filter arrows in the Status column.

You can also configure your search by clicking the Configuration Search icon. After you select the CDB, click **Select.**

You get an option to choose if you want to search for cluster targets only on the current host that you selected, or on all the hosts in the cluster that the host target is a member of.

This option enables you to save time, if you want to search for cluster targets only on the current host selected.

The Discovery Options section enables you to specify discovery hints. You can specify additional discovery options to change the default discovery behavior. Supported hints are `db_name, db_target_prefix, db_target_suffix, discovery_timeout <in seconds per host>`, and `no_db_domain.`

For example,

`db_name=PRODUCTS, discovery timeout=15.`

Click **Next** to proceed with the discovery process.

Click **Next.**

5. The Cluster Discovery:Result page displays the hosts discovered on the cluster. Select the host that you want to add.

> **Note:**
>
> A host can belong to only one cluster. If a particular host is not displayed, it can mean that the host belongs to another cluster.

On the Cluster Discovery: Result page, in the Clusters Target Properties section, verify the properties of the cluster.

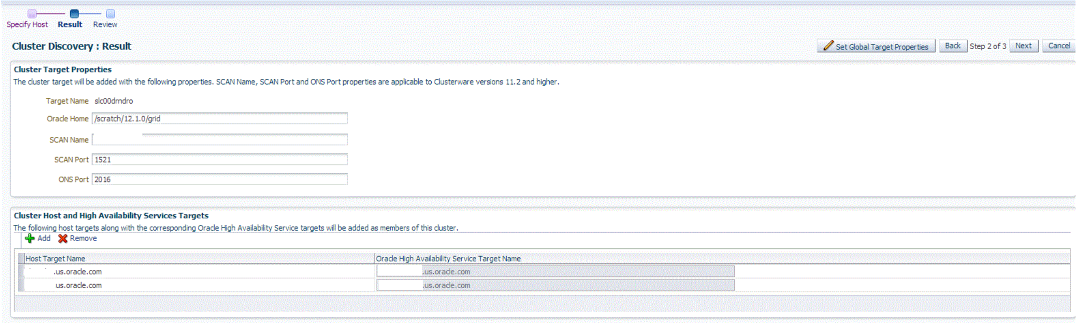

Cluster Discovery : Result

Cluster Target Properties

The cluster target will be added with the following properties. SCAN Name, SCAN Port and ONS Port properties are applicable to Clusterware versions 11.2 and higher.

Target Name slc00drndro

Oracle Home /scratch/12.1.0/grid

SCAN Name

SCAN Port 1521

ONS Port 2016

Cluster Host and High Availability Services Targets

The following host targets along with the corresponding Oracle High Availability Service targets will be added as members of this cluster.

Host Target Name	Oracle High Availability Service Target Name
.us.oracle.com	.us.oracle.com
us.oracle.com	.us.oracle.com

If you want to add or remove a host, select a host from the Cluster Host and High Availability Service Targets section. Click **Add** or **Remove.**

6. Click **Next.**
7. Review the displayed information, then click **Submit.**

Adding Cluster Targets By Using the Declarative Process

To add a cluster target declaratively by specifying target monitoring properties, follow these steps:

> **Note:**
>
> A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
> A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

1. From the Setup menu, select **Add Target,** then select **Add Targets Manually.**
2. On the Add Targets Manually page, select **Add Target Declaratively.**
3. In the Add Target Declaratively dialog box, select **Cluster.**
4. In the Monitoring Agent field, select the Management Agent monitoring the database.
5. Click **Add...**
6. On the Cluster Discovery:Result page, specify the target name, Oracle Home, SCAN name, SCAN port, and ONS port.

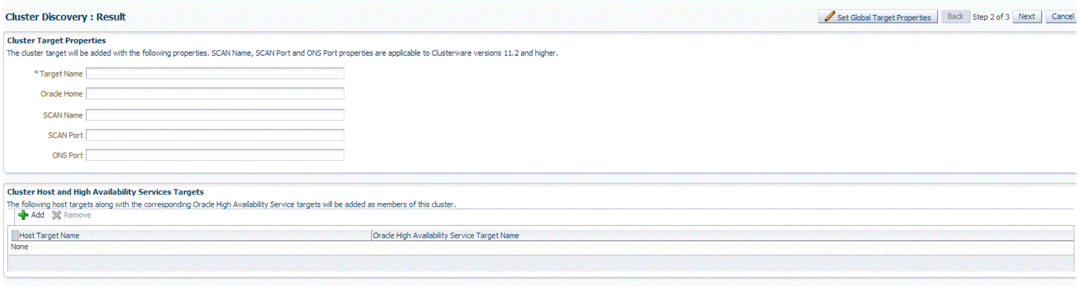

Cluster Discovery : Result

Cluster Target Properties

The cluster target will be added with the following properties. SCAN Name, SCAN Port and ONS Port properties are applicable to Clusterware versions 11.2 and higher.

* Target Name

Oracle Home

SCAN Name

SCAN Port

ONS Port

Cluster Host and High Availability Services Targets

The following host targets along with the corresponding Oracle High Availability Service targets will be added as members of this cluster.

Host Target Name	Oracle High Availability Service Target Name
None	

Note:

The SCAN name, SCAN port, and ONS port properties are applicable only for Clusterware versions 11.2 and higher.

7. You can add more hosts and high availability service targets to the cluster. If a particular host is not displayed when you click **Add,** it is possible that the host already belongs to another cluster.

   To specify global target properties for all the targets you have selected on the Promote Target: Results page, click **Set Global Target Properties,** specify the required properties, then click **OK.**

   Click **Next.**

8. Review the displayed information, then click **Submit.**

Discovering and Adding Single Instance High Availability Service Targets

This section describes the different methods in which you can discover and add single instance high availability service targets. In particular, this section covers the following:

- Discovering Single Instance High Availability Service Targets Using Autodiscovery
- Adding Single Instance High Availability Service Targets Using the Guided Discovery Process
- Adding Single Instance High Availability Service Targets By Using the Declarative Process

Discovering Single Instance High Availability Service Targets Using Autodiscovery

Autodiscovery of databases is enabled by default. If autodiscovery has been disabled, follow the steps described in Enabling Autodiscovery of Database Targets.

> **Note:**
>
> A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
> A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

To promote single instance high availability targets, follow these steps:

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

    > **Note:**
    >
    > If you do not see any results, then autodiscovery of targets has been disabled. To enable autodiscovery refer to Enabling Autodiscovery of Database Targets.

    From the results table, from the Agent-based targets tab, select the discovered High Availability instance target that you want to add for monitoring, and click **Promote.**

    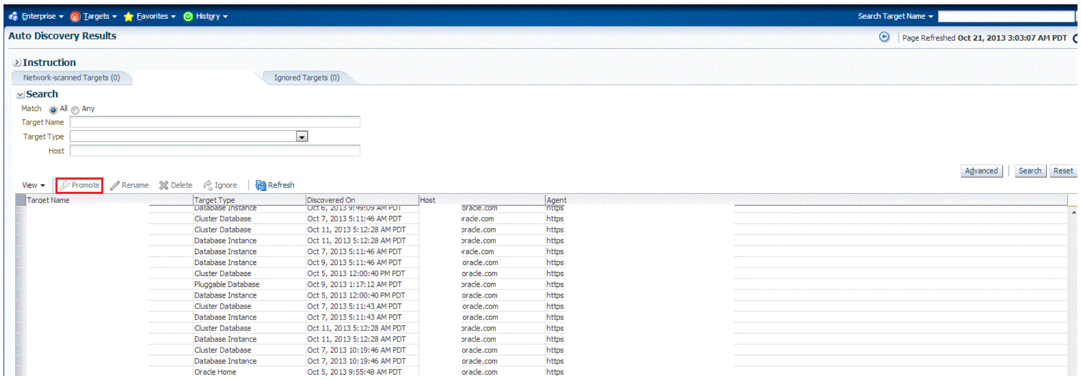

    Page Refreshed **Oct 21, 2013 3:03:07 AM PDT**

    Match ● All ○ Any

    Target Name

    Target Type

    Host

Target Name	Target Type	Discovered On	Host	Agent
	Database Instance	Oct 6, 2013 9:49:09 AM PDT	oracle.com	https
	Cluster Database	Oct 7, 2013 5:11:46 AM PDT	racle.com	https
	Cluster Database	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Database Instance	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Database Instance	Oct 7, 2013 5:11:46 AM PDT	racle.com	https
	Database Instance	Oct 9, 2013 5:11:46 AM PDT	oracle.com	https
	Cluster Database	Oct 5, 2013 12:00:40 PM PDT	oracle.com	https
	Pluggable Database	Oct 9, 2013 1:17:12 AM PDT	oracle.com	https
	Database Instance	Oct 5, 2013 12:00:40 PM PDT	oracle.com	https
	Cluster Database	Oct 7, 2013 5:11:43 AM PDT	oracle.com	https
	Database Instance	Oct 7, 2013 5:11:43 AM PDT	oracle.com	https
	Cluster Database	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Database Instance	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Cluster Database	Oct 7, 2013 10:19:46 AM PDT	oracle.com	https
	Database Instance	Oct 7, 2013 10:19:46 AM PDT	oracle.com	https
	Oracle Home	Oct 5, 2013 9:55:48 AM PDT	oracle.com	https

2. The Promote Targets:Result page displays the High Availability instances discovered. Select the database

    On the Promote Target: Results page, under the High Availability Services section, select the SIHA target that you want to promote.

3. Specify the monitoring credentials for the selected SIHA target, that is, the Monitor user name, Monitor password, and role. Also, if you want the selected target to be added to a group, specify a value for **Group.**

    If you specify Normal for **Role**, then the user name can be *dbsnmp* or a DB monitoring user. If you specify SYSDBA for **Role**, then you can provide any SYSDBA user.

4. Click **Test Connection** to test the connection made to the SIHA target using the specified monitoring credentials.
5. To specify global target properties for all the targets you have selected on the Promote Target: Results page, click **Set Global Target Properties,** specify the required properties, then click **OK.**

   To specify a common group for all the targets you have selected on the Promote Target: Results page, click **Specify Group for Targets,** select a group, then click **Select.**
6. If you have selected multiple databases and you want to set the same monitoring properties for all of them, select **Specify Common Monitoring Credentials.** Enter the monitoring credentials, monitoring password, and role. Click **Apply.**
7. Click **Next.**
8. Review the displayed information, then click **Submit.**

Adding Single Instance High Availability Service Targets Using the Guided Discovery Process

To add single instance high availability service targets using the guided process, follow these steps:

> **Note:**
>
> A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
> A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

1. From the Setup menu, select **Add Target,** then select **Add Targets Manually.**
2. On the Add Targets Manually page, select **Add Using Guided Process.**
3. In the Add Using Guided Process dialog box, select **Oracle Database, Listener, and Automatic Storage Management.** Click **Add...**
4. On the Cluster Discovery: Specify Host page, specify the single instance high availability service host, or click on the **Specify Host or Cluster** search icon to select the cluster.

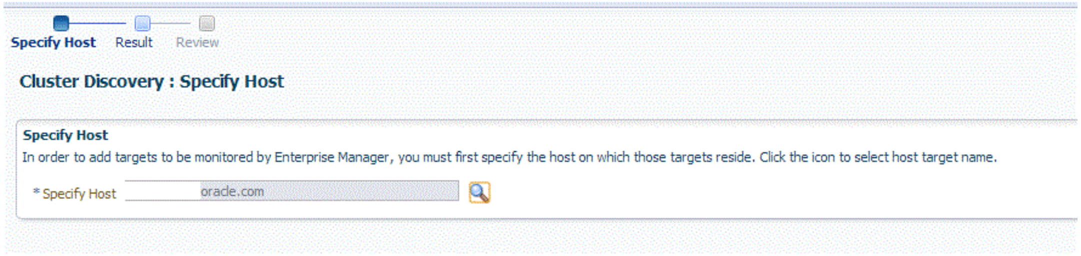

Cluster Discovery : Specify Host

Specify Host

In order to add targets to be monitored by Enterprise Manager, you must first specify the host on which those targets reside. Click the icon to select host target name.

* Specify Host: oracle.com

Click **Next.**

5. The Cluster Discovery:Result page displays the high availability service instances discovered.

   On the Cluster Discovery: Result page, under the High Availability Services section, select the high availability service instance target that you want to promote.

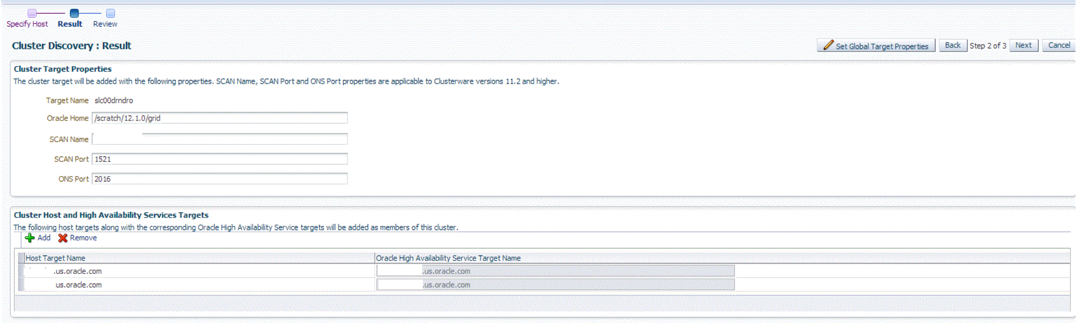

Cluster Discovery : Result

Cluster Target Properties

The cluster target will be added with the following properties. SCAN Name, SCAN Port and ONS Port properties are applicable to Clusterware versions 11.2 and higher.

Target Name slc00drndro

Oracle Home /scratch/12.1.0/grid

SCAN Name

SCAN Port 1521

ONS Port 2016

Cluster Host and High Availability Services Targets

The following host targets along with the corresponding Oracle High Availability Service targets will be added as members of this cluster.

Host Target Name	Oracle High Availability Service Target Name
.us.oracle.com	.us.oracle.com
us.oracle.com	.us.oracle.com

6. Click **Next.**
7. Review the displayed information, then click **Submit.**

Adding Single Instance High Availability Service Targets By Using the Declarative Process

To add a single instance High Availability Service target declaratively by specifying target monitoring properties, follow these steps:

> **Note:**
>
> A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
> A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

1. From the Setup menu, select **Add Target,** then select **Add Targets Manually.**
2. On the Add Targets Manually page, select **Add Target Declaratively.**
3. In the Add Target Declaratively dialog box, select **Oracle High Availability Service.**
4. In the Monitoring Agent field, select the Management Agent monitoring the database.
5. Click **Add...**

6. On the High Availability Service: Result page, specify a name for the target, the Oracle home, and the ONS port.

   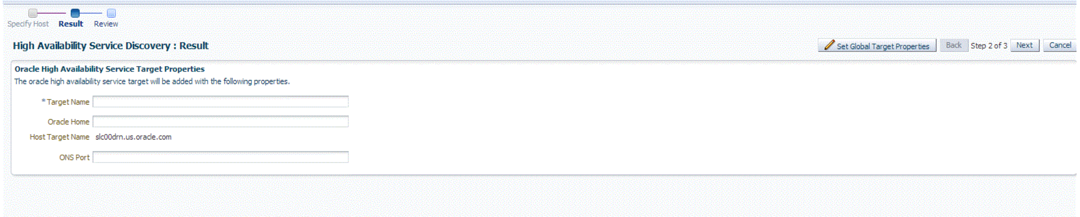

   High Availability Service Discovery : Result

   Oracle High Availability Service Target Properties

   The oracle high availability service target will be added with the following properties.

   * Target Name
   Oracle Home
   Host Target Name slc00drn.us.oracle.com
   ONS Port

   To specify global target properties for all the targets you have selected on the Promote Target: Results page, click **Set Global Target Properties,** specify the required properties, then click **OK.**

   Click **Next.**

7. Review the displayed information, and then click **Submit.**

Discovering and Adding Cluster Automatic Storage Management Targets

This section describes the different methods in which you can discover and add ASM cluster targets. In particular, this section covers the following:

- Discovering Cluster ASM Targets Using Autodiscovery
- Adding Cluster ASM Targets Using the Guided Discovery Process
- Adding Cluster ASM Targets By Using the Declarative Process

Discovering Cluster ASM Targets Using Autodiscovery

Autodiscovery of databases is enabled by default. If autodiscovery has been disabled, follow the steps described in Enabling Autodiscovery of Database Targets.

> **Note:**
>
> A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
> A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

To promote Cluster Automatic Storage Management targets using autodiscovery, follow these steps:

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

   > **Note:**
   >
   > If you do not see any results, then autodiscovery of targets has been disabled. To enable autodiscovery refer to Enabling Autodiscovery of Database Targets.

   From the results table, from the Agent-based targets tab, select the discovered cluster ASM target that you want to add for monitoring, and click **Promote.**

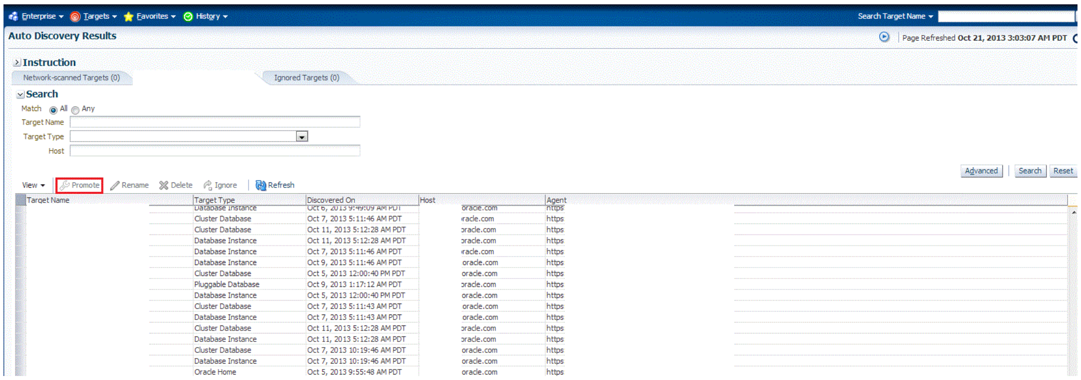

Target Name	Target Type	Discovered On	Host	Agent
	Database Instance	Oct 6, 2013 9:49:09 AM PDT	oracle.com	https
	Cluster Database	Oct 7, 2013 5:11:46 AM PDT	racle.com	https
	Cluster Database	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Database Instance	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Database Instance	Oct 7, 2013 5:11:46 AM PDT	racle.com	https
	Database Instance	Oct 9, 2013 5:11:46 AM PDT	oracle.com	https
	Cluster Database	Oct 5, 2013 12:00:40 PM PDT	oracle.com	https
	Pluggable Database	Oct 9, 2013 1:17:12 AM PDT	oracle.com	https
	Database Instance	Oct 5, 2013 12:00:40 PM PDT	oracle.com	https
	Cluster Database	Oct 7, 2013 5:11:43 AM PDT	oracle.com	https
	Database Instance	Oct 7, 2013 5:11:43 AM PDT	oracle.com	https
	Cluster Database	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Database Instance	Oct 11, 2013 5:12:28 AM PDT	oracle.com	https
	Cluster Database	Oct 7, 2013 10:19:46 AM PDT	oracle.com	https
	Database Instance	Oct 7, 2013 10:19:46 AM PDT	oracle.com	https
	Oracle Home	Oct 5, 2013 9:55:48 AM PDT	oracle.com	https

2. The Promote Targets:Result page displays the targets discovered on the cluster ASM.

   On the Promote Target: Results page, in the Cluster ASM section, select the target that you want to promote.

   By default, selecting the cluster ASM target for promotion also selects all its associated discovered targets for promotion. If you want to add or remove a target from the ones selected for promotion, select the cluster ASM target, then click **Configure.** Select the **Instances** tab, then click **Add** or **Remove.** Click **Save.**

3. In the cluster ASM section, specify the monitoring credentials for the selected cluster ASM target, that is, the Monitor user name, Monitor password, and role. Also, if you want the selected target to be added to a group, specify a value for **Group.**

   If you specify Normal for **Role**, then the user name can be *dbsnmp* or a DB monitoring user. If you specify SYSDBA for **Role**, then you can provide any SYSDBA user.

4. Click **Test Connection** to test the connection made to the cluster ASM target using the specified monitoring credentials.

5. To specify global target properties for all the targets you have selected on the Promote Target: Results page, click **Set Global Target Properties,** specify the required properties, then click **OK.**

   To specify a common group for all the targets you have selected on the Promote Target: Results page, click **Specify Group for Targets,** select a group, then click **Select.**

6. If you have selected multiple targets and you want to set the same monitoring properties for all of them, select **Specify Common Monitoring Credentials.** Enter the monitoring credentials, monitoring password, and role. Click **Apply.**

7. Click **Next.**

8. Review the displayed information, then click **Submit.**

Adding Cluster ASM Targets Using the Guided Discovery Process

To add cluster ASM targets using the guided process, follow these steps:

> **Note:**
>
> A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
> A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

1. From the Setup menu, select **Add Target,** then select **Add Targets Manually.**
2. On the Add Targets Manually page, select **Add Using Guided Process.**
3. In the Add Using Guided Process dialog box, select **Oracle Database, Listener, and Automatic Storage Management.** Click **Add...**
4. On the Database Discovery: Search Criteria page, specify the cluster ASM host, or click on the **Specify Host or Cluster** search icon to select the cluster.

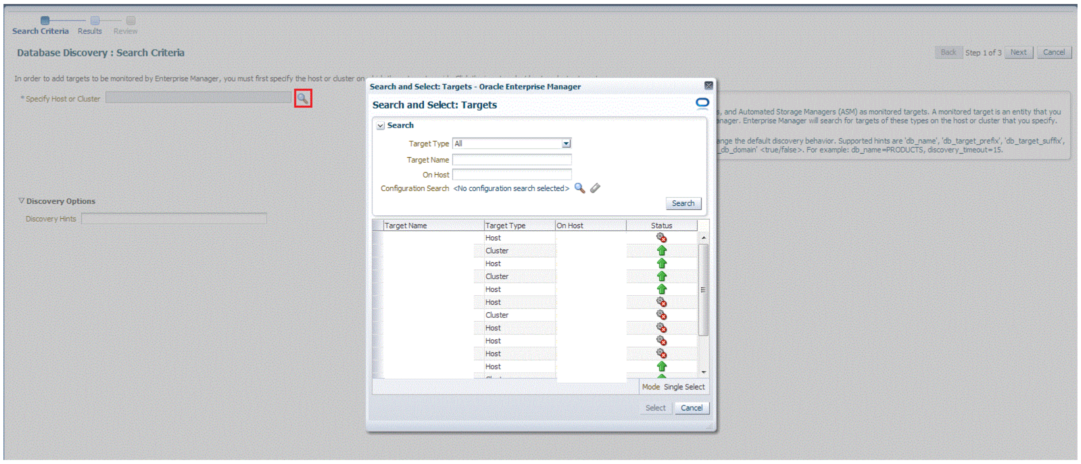

Database Discovery : Search Criteria

In order to add targets to be monitored by Enterprise Manager, you must first specify the host or cluster on

* Specify Host or Cluster

…, and Automated Storage Managers (ASM) as monitored targets. A monitored target is an entity that you …nager. Enterprise Manager will search for targets of these types on the host or cluster that you specify.

…nge the default discovery behavior. Supported hints are 'db_name', 'db_target_prefix', 'db_target_suffix', …_db_domain' <true/false>. For example: db_name=PRODUCTS, discovery_timeout=15.

▽ Discovery Options

Discovery Hints

Search and Select: Targets - Oracle Enterprise Manager

Search and Select: Targets

Search

Target Type: All

Target Name:

On Host:

Configuration Search: <No configuration search selected>

Target Name	Target Type	On Host	Status
	Host		✖
	Cluster		⬆
	Host		⬆
	Cluster		⬆
	Host		⬆
	Host		✖
	Cluster		✖
	Host		✖
	Host		✖
	Host		✖
	Host		⬆

Mode: Single Select

If the host you select is a member of a cluster, then use this page to also indicate whether the Automated Storage Managers should be searched only on the selected host or on all hosts within the cluster.

After you click the search icon, a Select Targets dialog box appears. Specify the target name or select a target from the target table. You can filter the search to search for all down targets or up targets, by clicking the filter arrows in the Status column.

You can also configure your search by clicking the Configuration Search icon. After you select the CDB, click **Select.**

You get an option to choose if you want to search for ASM targets only on the current host that you selected, or on all the hosts in the cluster that the host target is a member of.

This option enables you to save time, if you want to search for ASM targets only on the current host selected.

The Discovery Options section enables you to specify discovery hints. You can specify additional discovery options to change the default discovery behavior. Supported hints are `db_name, db_target_prefix, db_target_suffix, discovery_timeout <in seconds per host>`, and `no_db_domain`.

For example,

`db_name=PRODUCTS, discovery timeout=15.`

Click **Next** to proceed with the discovery process.

Click **Next.**

5. The Database Discovery:Result page displays the targets discovered on the cluster ASM target.

   On the Database Discovery: Results page, in the Cluster ASM section, select the target that you want to promote.

   By default, selecting the cluster ASM target for promotion also selects all its associated discovered targets for promotion. If you want to add or remove a target from the ones selected for promotion, select the cluster ASM target, then click **Configure.** Select the **Instances** tab, then click **Add** or **Remove.** Click **Save.**

6. In the cluster ASM section, specify the monitoring credentials for the selected cluster ASM target, that is, the Monitor user name, Monitor password, and role. Also, if you want the selected target to be added to a group, specify a value for **Group.**

   If you specify Normal for **Role**, then the user name can be *dbsnmp* or a DB monitoring user. If you specify SYSDBA for **Role**, then you can provide any SYSDBA user.

7. Click **Test Connection** to test the connection made to the cluster ASM target using the specified monitoring credentials.

8. To specify global target properties for all the targets you have selected on the Promote Target: Results page, click **Set Global Target Properties,** specify the required properties, then click **OK.**

   To specify a common group for all the targets you have selected on the Promote Target: Results page, click **Specify Group for Targets,** select a group, then click **Select.**

9. If you have selected multiple targets and you want to set the same monitoring properties for all of them, select **Specify Common Monitoring Credentials.** Enter the monitoring credentials, monitoring password, and role. Click **Apply.**

10. Click **Next.**

11. Review the displayed information, then click **Submit.**

Adding Cluster ASM Targets By Using the Declarative Process

To add a cluster ASM target declaratively by specifying target monitoring properties, follow these steps:

> **Note:**
>
> A database system is automatically created on discovery of an Oracle Database. The system is built on the new target and association model that can be used to monitor the database's storage, connectivity, and high availability. This also enables you to monitor and manage applications that are dependent on the database. Database System topology can be used to view relationship between various entities within the database system as well as external dependencies.
>
> A database system contains a primary database and related targets such as Listener and Automatic Storage Management. It also includes standby databases and their related targets if the database is in a Data Guard configuration. However, you cannot create database systems for standby databases.

1. From the Setup menu, select **Add Target,** then select **Add Targets Manually.**
2. On the Add Targets Manually page, select **Add Target Declaratively.**
3. In the Add Target Declaratively dialog box, select **Cluster ASM.**
4. In the Monitoring Agent field, select the Management Agent monitoring the database.
5. Click **Add...**
6. On the Configure Cluster ASM: Properties page, specify a name for the Cluster ASM, the Oracle home path, username and password, role, cluster name, and service name.

   > **Note:**
   >
   > The Service Name is used to establish the cluster ASM connection. It should be one of the service names the cluster ASM registers with the listeners.

7. You can add instances, ASM IO server instances, and ASM proxy instances by clicking **Add** in the respective sections.

   Click **Test Connection** to test the connection made to the cluster ASM target using the specified monitoring credentials.

   OK.

Configuring a Target Database for Secure Monitoring

This section covers the following:

- About Secure Monitoring of Databases
- Configuring a Target Database for Secure Monitoring

About Secure Monitoring of Databases

The Oracle Database uses various encryption algorithms to secure the information moving across the network between the Oracle Database Server and a client. The Enterprise Manager agent communicates with the Database target over a TCP protocol to get real-time monitoring data. The Oracle Management Server (OMS) also communicates with the Database target in clear text over the network to manage the target. As TCP transfers data in clear text format over the network, anyone can access and modify the data. To secure the data exchanges between the OMS and managed target, Enterprise Manager can use the TCPS protocol to encrypt and protect the data.

The Secure Monitoring feature in Enterprise Manager allows you to monitor Oracle Database targets on secure channels using a TCP/IP with SSL (Secure Socket Layer) protocol. You can also discover and monitor the Listener processes running on TCPS ports using the Enterprise Manager console.

You can choose from the following options available for target monitoring:

- Target Monitoring over TCP

  In this case, Enterprise Manager connects to the target database for target monitoring using the TCP protocol. The data communication between the target database and the Enterprise Manager OMS happens in clear text form over the network.
- Target Monitoring over TCPS with Server Authentication using Trusted Certificate

  Enterprise Manager connects to the target database TCP over a secure socket layer (SSL) for target monitoring. In this case, the client authenticates the database server using a trusted certificate of the server. The data communicates over the network between Enterprise Manager and the target database in encrypted form.
- Target Monitoring over TCPS with Server Authentication using Kerberos

  Enterprise Manager connects to the target database TCP over a secure socket layer (SSL) for target monitoring. In this case, the client authenticates the database server using the Kerberos authentication protocol. The data communicates over the network between Enterprise Manager and the target database in encrypted form.

You can enable secure monitoring while discovering the target database, or you can change the secure monitoring settings for a target database on the Monitoring Configuration page.

Configuring a Target Database for Secure Monitoring

You can enable secure monitoring either while discovering a target database, or by changing the secure monitoring settings for a selected target database after discovery.

> **Note:**
>
> DB monitoring setup accepts case-sensitive user names. .Every Oracle database object has a name. In a SQL statement, you represent the name of an object with a quoted identifier or a non-quoted identifier.
>
> A quoted identifier begins and ends with double quotation marks (").
>
> - If you name a schema object using a quoted identifier, then you must use the double quotation marks whenever you refer to that object.
> - This is a case-sensitive identifier.
>
> A nonquoted identifier is not surrounded by any punctuation.
>
> - This is not a case sensitive name.
>
> When you set up the monitoring credentials you can provide either quoted (case-sensitive) or unquoted account names.

Follow these steps to change the settings of a discovered database target:

1. From the Enterprise Manager page, choose **Databases** from the **Targets** menu.

   Enterprise Manager displays the Databases page.

2. Choose the database for which you want to configure monitoring.

   The Database Home page for that database appears.

3. From the **Oracle Database** menu, select **Monitoring Configuration** from the **Target Setup** menu.

   Enterprise Manager displays the Configure Database Instance: Properties page.

4. On the Configure Database Instance: Properties page, scroll down to the Connection Protocol field and select from one of the following drop-down menu choices:

   - TCP -- Enterprise Manager connects to the target database for target monitoring using the TCP protocol. The data communication between the target database and the Enterprise Manager OMS happens in clear text form over the network.
   - TCPS -- Enterprise Manager connects to the target database TCP over a secure socket layer (SSL) for target monitoring. In this case, the client authenticates the database server using a trusted certificate of the server. The data communicates over the network between Enterprise Manager and the target database in encrypted form.
   - TCPS with Server Authentication using Kerberos -- Enterprise Manager connects to the target database TCP over a secure socket layer (SSL) for target monitoring. In this case, the client authenticates the database server using the Kerberos authentication protocol. The data communicates over the network between Enterprise Manager and the target database in encrypted form.

5. Optionally you can test the connection by clicking **Test Connection** to determine whether the change in Connection Protocol has impacted the connection to the database.
6. Click **Next** to move to the Configure Database Instance: Review Page.
7. Click **Submit** to apply your changes.

   The database target now connects using the protocol you selected.

Adding Connection Manager (CMAN) Targets By Using the Declarative Process

To add a Connection Manager target declaratively by specifying target monitoring properties, follow these steps:

1. From the Setup menu, select **Add Target,** then select **Add Targets Manually.**
2. On the Add Targets Manually page, select **Add Target Declaratively.**
3. In the Add Target Declaratively dialog box, enter the host and select **Connection Manager** as the target type.
4. Click **Add.**
5. On the Add: Connection Manager page, specify a name and the credentials.
6. Specify all the properties of the target, that is, **Connection Manager Name, Connection Protocol, Machine Name, Oracle Home, Port Number,** and **cman.ora Directory.**
7. Click **OK.**
8. Review the displayed information, and then click **Submit.**

Discovering and Adding Oracle Key Vault

Starting with Enterprise Manager 24ai, two new target types have been added to enable the monitoring of Oracle Key Vault (OKV):

- Oracle Key Vault Cluster (`oracle_kv_cluster`) - assists in the monitoring of OKV clusters.
- Oracle Key Vault Server (`oracle_kv_server`) - assists in the monitoring of individual OKV nodes (or a standalone server for non-cluster deployments).

Prerequisites:

> **Note:**
>
> Since OKV Server does not allow EM agent running on it, in this document, "EM Agent Host" refers to any other host where EM agent is installed and running.

- OKV Server to be monitored via EM must be at least version 21.10
- Ensure that Network Services and RESTful Services options are enabled in the OKV server console. For more information, see Getting Started with Oracle Key Vault RESTful Services.
- The OS specific "expect" package has to be installed on the host where the EM agent is installed.
- Monitoring User account to be used in OEM need to be created in OKV console first. And it must have `monitoring privilege` selected
- SNMP username and password need to be setup in OKV console.
- SNMP access allowed from must have OEM Agent Hosts ip address specified or All option selected.

> **Note:**
>
> If there is any firewall on the agent host, for the EM Agent to communicate with the OKV server being monitored, open the REST port and SNMP port.

To add a OKV Cluster or Server target, follow these steps:

1. From the **Setup** menu, select **Add Target**, then click **Add Targets Manually**.
2. On the **Add Targets Manually** dialog, select the **Agent Host**

   > **Note:**
   >
   > The agent needs to have HTTP/SNMP connectivity to the OKV server.

3. From the **Target Type** list, select `Oracle Key Vault Server` or `Oracle Key Vault Cluster`.
4. Click **Add**.

5. Specify all the properties of the target, as described in the OKV UI discovery input table, click **Next** and wait for the validation.

After validation completes, on the Review page displayed, select **Submit** to complete the target creation.

Input Field	Description
Target Name	Any name your OKV server or cluster as a target.
OKV Server IP Address	The IP address of the OKV console. OKV server must be version 21.10 onwards.
OKV Server Port	Port of the OKV console By default: `5695`
Monitoring Credentials: • Monitoring Username • Monitoring Password	Username and password created in OKV Console with a **monitoring privilege**
SNMP Credentials • SNMP Username • SNMP Password	SNMP User Name and Password created in the OKV console To create it, from the left navigation bar select the **System** tab, then **Settings** . In the **Monitoring and Alerts** area, click **SNMP**.
SNMP Port and SNMP Timeout	SNMP Port and Timeout found in the OKV console SNMP Port: `161` SNMP Timeout: `5 seconds`
OKV RESTful Services Utility Installation Location / OKV REST CLI Installation Location	Directory path accessible on the EM Agent host where OKV REST CLI / OKV REST Services Utilisation will be downloaded and configured automatically at the end of the discovery process.

For more information regarding the OKV Console, see Monitoring and Auditing Oracle Key Vault.

Preferred Connect Strings

A Preferred Connect String is a fully qualified connection descriptor that can override Enterprise Manager’s default connection configuration to the database with your preferred connection method. For example, you can create a preferred connect string using the host, port and service name (instead of SID), and specify that to be used for database connections. Preferred Connect Strings are supported for Single Instance, RAC and Pluggable databases.

There are two types of connections that support this configuration:

- From the OMS to a target database, used for access to database performance pages and for administration operations: To use a connect string for OMS-to-database connections, you specify an OMS Preferred Connect String (OMSPCS).
- Between an agent and a target database, used for metric collections: To use a connect string for agent-to-database connections, you specify an Agent Preferred Connect String (APCS).

The Preferred Connect Strings can be configured using EM console or EM CLI.

The OMS Preferred Connect String option is supported for all versions of Oracle Enterprise Manager 13c Release 5.

This Agent Preferred Connect String option is supported starting with Oracle Enterprise Manager 13c Release 5 Update 15 (13.5.0.15) and the EM agent version 13.5.0.15. The Agent Preferred Connect String is disabled by default. To enable it, run the following command, and restart the OMS:

```
emctl set property -name oracle.sysman.db.showapcs -value true -sysman_pwd
<sysman_password>
```

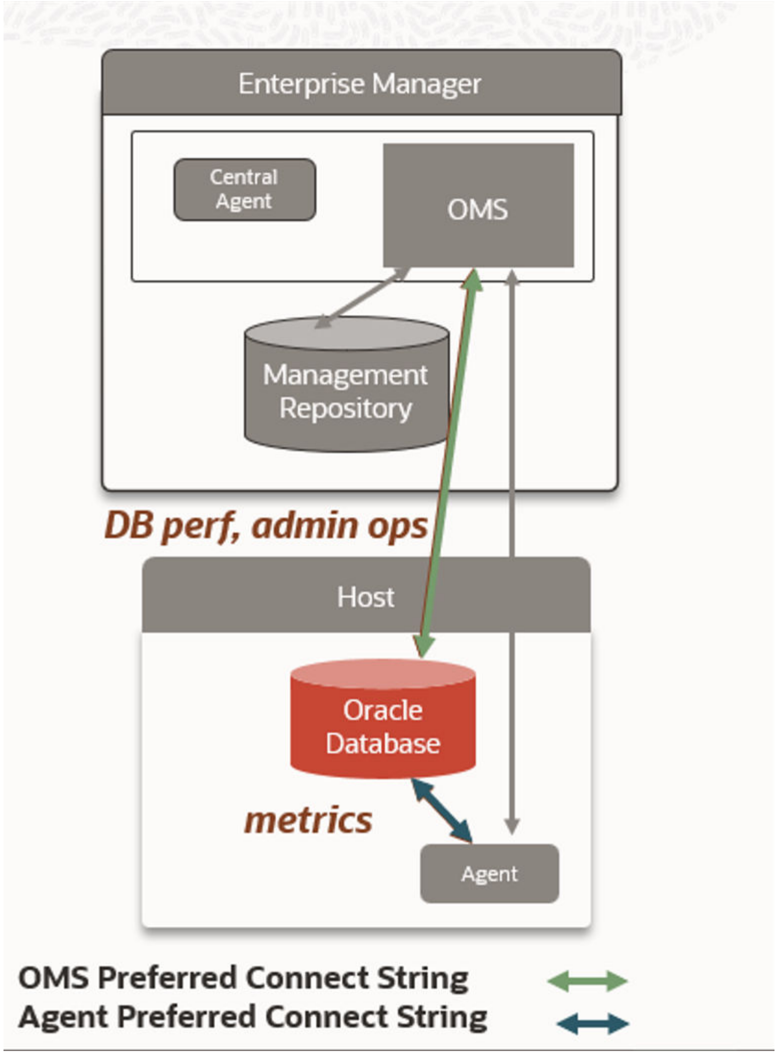

Preferred Connect String Configuration Steps

Choose one of the following methods to define the connect strings:

1. Using the EM Console: From the target's home page menu, go to **Target Setup**, then **Monitoring Configuration**.

   Here is an example for the RAC (Cluster Database) Instance:

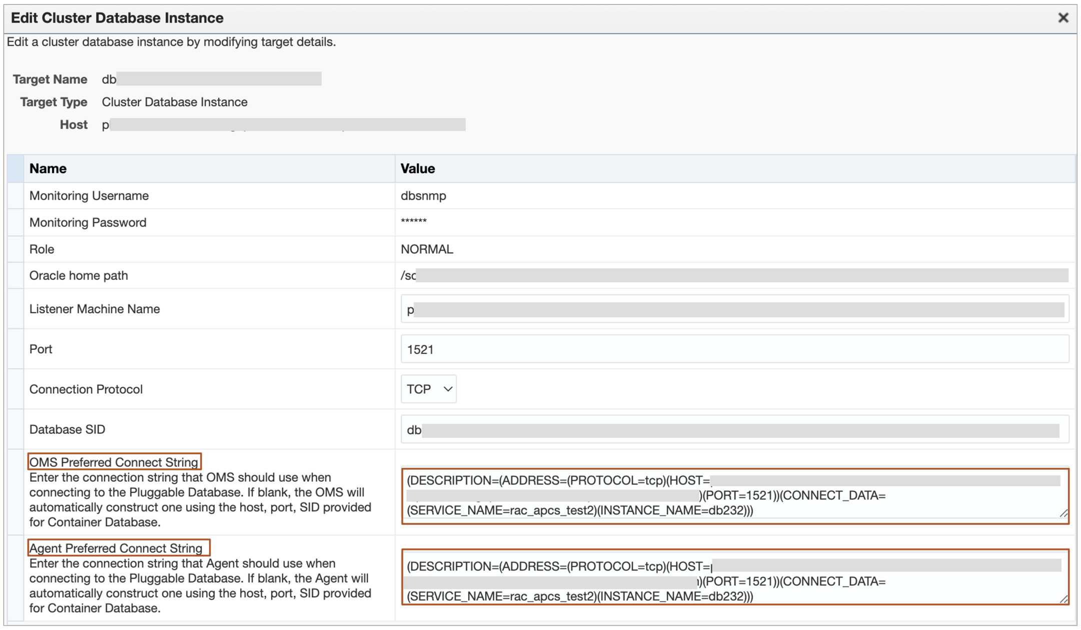

Edit Cluster Database Instance

Edit a cluster database instance by modifying target details.

Target Name	db[redacted]
Target Type	Cluster Database Instance
Host	p[redacted]

Name	Value
Monitoring Username	dbsnmp
Monitoring Password	******
Role	NORMAL
Oracle home path	/sc[redacted]
Listener Machine Name	p[redacted]
Port	1521
Connection Protocol	TCP
Database SID	db[redacted]
OMS Preferred Connect String Enter the connection string that OMS should use when connecting to the Pluggable Database. If blank, the OMS will automatically construct one using the host, port, SID provided for Container Database.	(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=[redacted])(PORT=1521))(CONNECT_DATA=(SERVICE_NAME=rac_apcs_test2)(INSTANCE_NAME=db232)))
Agent Preferred Connect String Enter the connection string that Agent should use when connecting to the Pluggable Database. If blank, the Agent will automatically construct one using the host, port, SID provided for Container Database.	(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=[redacted])(PORT=1521))(CONNECT_DATA=(SERVICE_NAME=rac_apcs_test2)(INSTANCE_NAME=db232)))

2. Using EMCLI: Specify the Preferred Connect String using EMCLI, using this format:

   - **OMS Preferred Connect String**

     ```
     emcli modify_target
         -name=“<target_name>"
         -type=“<target_type>"
         -properties="PreferredConnectString:<connect string>"
         -on_agent
     ```

     Example:

     ```
     emcli modify_target
         -name="database_CDB1_PDB1"
         -type="oracle_pdb"
         -properties="PreferredConnectString:
     (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=host1.example.com)(PORT=1521))
     (CONNECT_DATA=(SERVICE_NAME=myrac)(INSTANCE_NAME=db231)))"
         -on_agent
     ```

   - **Agent Preferred Connect String**

     ```
     emcli modify_target
         -name=“<target_name>"
         -type=“<target_type>"
         -properties="AgentPreferredConnectString:<connect string>"
         -on_agent
     ```

     Example:

     ```
     emcli modify_target
         -name="MyPDB"
         -type="oracle_pdb"
         -properties="AgentPreferredConnectString: (DESCRIPTION =
     ```

```
    (ADDRESS_LIST = (ADDRESS = (PROTOCOL = tcp)(HOST = host1.example.com)
    (PORT = 1522)))(CONNECT_DATA = (SERVICE_NAME = MYPDB)))"
        -on_agent
```

If you updated the OMS Preferred Connect String, log out and re-log in to the EM Console to update the OMS-database cached connections.

Examples of Using the Preferred Connect Strings

1. **Database Monitoring Using Service Name**

   By default, Enterprise Manager uses SID for connections to the database. You can use Preferred Connect Strings to have all EM database connections use the Service Names instead of SIDs. This includes OMS-to-database connections as well as Agent-to-database connections.

   For single instance databases or RAC instances, specify the connect string using this format:

   ```
   (DESCRIPTION=(ADDRESS=(PROTOCOL=<protocol,tcp/tcps)(HOST=<host name, e.g.
             myhost.example.com)(PORT=<port, e.g.
             1521>))(CONNECT_DATA=(SERVICE_NAME=<service name>)
   (INSTANCE_NAME=<instance
             name>)))
   ```

   For RAC databases, set the Preferred Connect String for each RAC instance, where the host should be set to the virtual IP of the RAC node.

2. **Monitoring PDBs in a Cluster Database (RAC)**

   You can use Preferred Connect Strings to monitor a PDB that is open only a preferred subset of RAC instances.

   For this configuration, follow the steps below:

   a. Create a service for this PDB.

   b. Construct a fully qualified connect string that uses the SCAN listener and service for the PDB:

      ```
      (DESCRIPTION = (ADDRESS_LIST = (ADDRESS = (PROTOCOL = tcp)(HOST =
      <hostname of the scan listener>)(PORT = 1521)))(CONNECT_DATA =
      (SERVICE_NAME = <service for the PDB)))
      ```

   c. Edit the Monitoring configuration UI for the PDB and specify your constructed connect string as the values for OMS Preferred Connect String and Agent Preferred Connect String

      The agent will connect via the SCAN listener to the open PDB on the preferred RAC instance.

3. **Monitoring Using TCPS-based connections with SSL Server DN Match**
   TCPS offers enhanced levels of security and it can be configured with SSL Domain Name.

   Configure the OMS Preferred Connect String with the SSL domain name as follows:

   a. Enable this setup in the OMS by running following command:

      ```
      emctl set property
          -sysman_pwd <password>
      ```

```
    -name oracle.sysman.db.ssl_server_dn_match
    -value true
```

b. Run the following command on the database host to get the host domain name:

```
orapki wallet display -wallet "<path_to_wallet>" -pwd <wallet_password>
```

Example output:

```
Oracle PKI Tool Release 21.0.0.0.0 - Production
Version 21.0.0.0.0
Copyright (c) 2004, 2020, Oracle and/or its affiliates. All rights
reserved.Requested Certificates:
User Certificates:
Subject:        CN=<server_cert_name> <---- This is the HOST DN
(CN=<server_cert_name>) you should use in Preferred Connect String
Trusted Certificates:
Subject:        CN=<client_cert_name>
Subject:        CN=<server_cert_name>
```

c. Construct the OMS Preferred Connect String using the certificate name as the SECURITY parameter, as shown:

```
(DESCRIPTION= (ADDRESS= (PROTOCOL=TCPS) (HOST=<DB_HOST>)
(PORT=<TCPS_PORT>)) (CONNECT_DATA= (SID=<DB_SID>))
(SECURITY=(SSL_SERVER_CERT_DN=CN=<server_cert_name>)))
```

d. Specify the OMS Preferred Connect String either using the EM Console or EM CLI.

4. **Monitoring Database Targets in Different Subnets**

If your OMS cannot have direct connections to the database targets because they are in separate subnets, or in separate data centers, you can use the OMS Preferred Connect String to connect with a proxy server:

```
(DESCRIPTION= (ADDRESS=(HTTPS_PROXY=<proxy name, e.g.sales-proxy)
(HTTPS_PROXY_PORT=8080) (PROTOCOL=TCPS) (HOST=<host name>) (PORT=443))
(CONNECT_DATA=(SERVICE_NAME=service1.example.com)))
```

Note:

This feature requires Oracle Enterprise Manager 13c Release 5 Update 15 (13.5.0.15) or higher.

3

Discovering and Adding Middleware Targets

This chapter describes how you can discover and add fusion middleware targets to be managed by Enterprise Manager. In particular, this chapter covers the following:

- Discovering and Adding WebLogic Domains
- Discovering New or Modified Domain Members
- Adding Standalone Oracle HTTP Servers
- Adding Exalytics Targets
- Removing Middleware Targets

Discovering and Adding WebLogic Domains

This section describes the different methods in which you can discover, promote, and add WebLogic domain targets in Enterprise Manager. In particular, this section covers the following:

- Discovering WebLogic Domains Using Autodiscovery
- Adding WebLogic Domains Using the Guided Discovery Process
- Adding Multiple WebLogic Domains Using EM CLI

Discovering WebLogic Domains Using Autodiscovery

To discover and promote WebLogic domains, follow these steps:

1. From the **Setup** menu, select **Add Target**, then select **Configure Auto Discovery**.
2. On the Configure Auto Discovery page, select the Oracle Fusion Middleware link in the table to configure auto discovery for Oracle Fusion Middleware or click the icon in the **Configure Host Discovery** column to configure that Oracle Fusion Middleware row.
3. Set the schedule at which the discovery job will be run, in days. This schedule will be applied to all selected hosts. By default the job will run every 24 hours.
4. Click **Add Host**. Select the host machines you want to include in the discovery.
5. Select a host in the table, and then click **Edit Parameters** to specify the Middleware Homes to search for targets. The Middleware Home is the top-level directory for all Oracle Fusion Middleware products, created when Oracle WebLogic Domain is installed.

   Enter ***** to search all Middleware Homes, or specify the path to one or more Middleware Homes on the host, each separated by a comma.
6. Click the **OK** button located at the right of the screen. At this point, automatic discovery has been enabled, and the discovery job will run at the scheduled frequency.
7. From the **Setup** menu, select **Add Target**, then select **Auto Discovery Results**.
8. Click the **Targets on Hosts** tab to view the discovered Oracle Fusion Middleware targets.
9. Select a target, then click **Promote**.

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

10. Supply or accept values for the following parameters:

    - Administration Server Host

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

    - Port

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

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

      You can access My Oracle Support at the following URL:

      `http://support.oracle.com/CSP/ui/flash.html`

    - Enter the WebLogic Administration Server user name and password.

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

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

    - Unique Domain Identifier.

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

    - Agent

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

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

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

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

      – To patch a WebLogic Domain, you need a local Management Agent on each WebLogic Server machine.

      – If you want to use Oracle Support Workbench for a WebLogic Domain target, then the target requires a local Management Agent.

  - Cloning requires a local Management Agent at the source administration server machine and a local Management Agent at all destination machines.

Advanced Parameters

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

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

`https://support.oracle.com/CSP/ui/flash.html`

- JMX Protocol

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

- Discover Down Servers

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

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

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

- JMX Service URL

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

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

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

- Discover Application Versions

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

- Enable Automatic Refresh

  This option refreshes the WebLogic domain every 24 hours.

- Use Host Name in Service URL

  You can use host name in service URL instead of JMX. It is recommended to use this option if you are using a private network and there are many hosts using the same IP address.

- Create Incident for Discovery Failure

  This option creates an OMS incident if discovery fails. You can view the incident from the Support workbench page.

- External Parameters

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

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

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

- Discovery Debug File Name

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

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

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

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

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

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

13. As a rare case, if you want to disable one or more target types for discovery, scroll down to the **Disable Target Types** section under **Advanced,** and move the target type from the **Available Target Types** list to the **Selected Target Types** list.

    Click **Refresh Targets** to view the refreshed **Targets and Agents Assignments** table.

    > **Note:**
    >
    > - If you disable a target type, it will remain disabled for future refresh operations.
    > - Child target types are disabled if you disable a parent target type.

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

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

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

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

> **Note:**
>
> After you discover a middleware target for the first time, it is recommended that you learn the best practices for monitoring and managing the discovered target. To navigate to the Target Management Best Practices page, from the target home page, select the **WebLogic Domain** menu, and then select **Target Management Best Practices.** The page lists the best practices items for a Fusion Middleware Domain.

Adding WebLogic Domains Using the Guided Discovery Process

Oracle WebLogic Domains and their respective components can be discovered using Enterprise Manager. A wizard guides you through the guided discovery process.

For any Java EE application deployed on any number of servers, only one Domain Application target will be discovered.

> **Note:**
>
> To discover a WebLogic Domain, the Administration Server must be up because the Management Agent must make a JMX connection to it. If the Administration Server is down, discovery cannot occur.
>
> Thereafter, to monitor the WebLogic Domain, the Administration Server need not be up.

To add a WebLogic domain using the guided process, follow these steps:

1. From the **Setup** menu, select **Add Target**, then select **Add Targets Manually**.
2. Select the **Add Using Guided Process** option.
3. In the Add Using Guided Process dialog box, select **Oracle Fusion Middleware/ WebLogic Domain** from the Guided Discovery column.

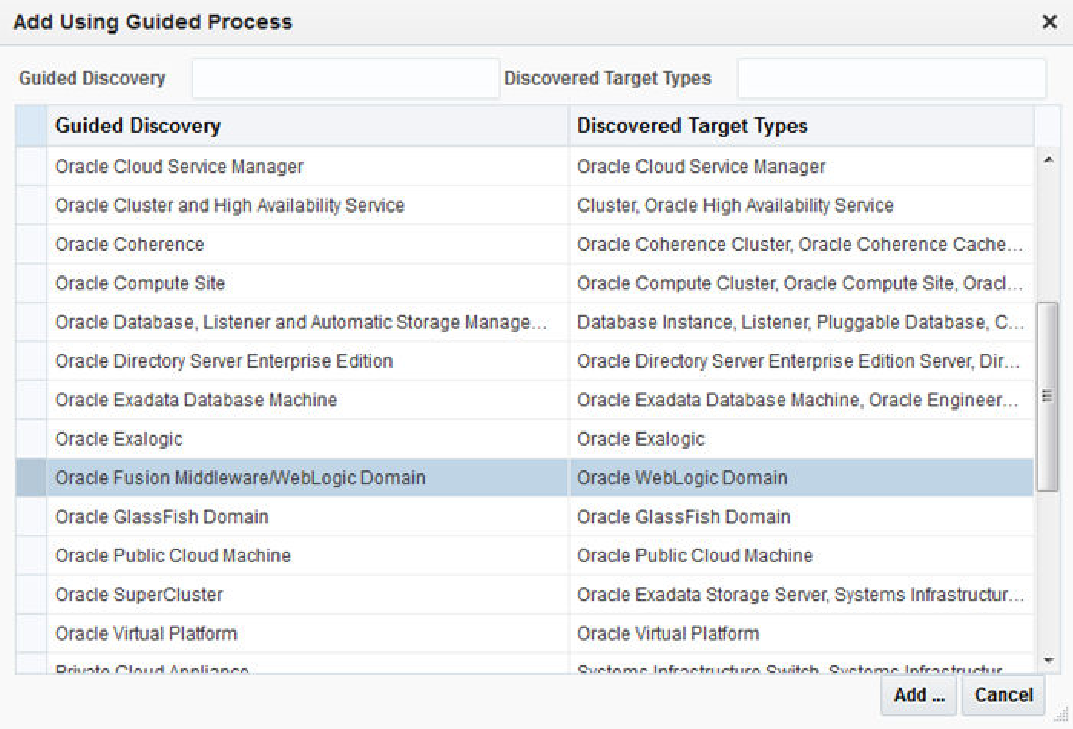

Add Using Guided Process

Guided Discovery | Discovered Target Types

Guided Discovery	Discovered Target Types
Oracle Cloud Service Manager	Oracle Cloud Service Manager
Oracle Cluster and High Availability Service	Cluster, Oracle High Availability Service
Oracle Coherence	Oracle Coherence Cluster, Oracle Coherence Cache...
Oracle Compute Site	Oracle Compute Cluster, Oracle Compute Site, Oracl...
Oracle Database, Listener and Automatic Storage Manage...	Database Instance, Listener, Pluggable Database, C...
Oracle Directory Server Enterprise Edition	Oracle Directory Server Enterprise Edition Server, Dir...
Oracle Exadata Database Machine	Oracle Exadata Database Machine, Oracle Engineer...
Oracle Exalogic	Oracle Exalogic
Oracle Fusion Middleware/WebLogic Domain	Oracle WebLogic Domain
Oracle GlassFish Domain	Oracle GlassFish Domain
Oracle Public Cloud Machine	Oracle Public Cloud Machine
Oracle SuperCluster	Oracle Exadata Storage Server, Systems Infrastructur...
Oracle Virtual Platform	Oracle Virtual Platform
Private Cloud Appliance	Systems Infrastructure Switch, Systems Infrastructur

4. Click **Add...**

5. Supply or accept values for the following parameters:

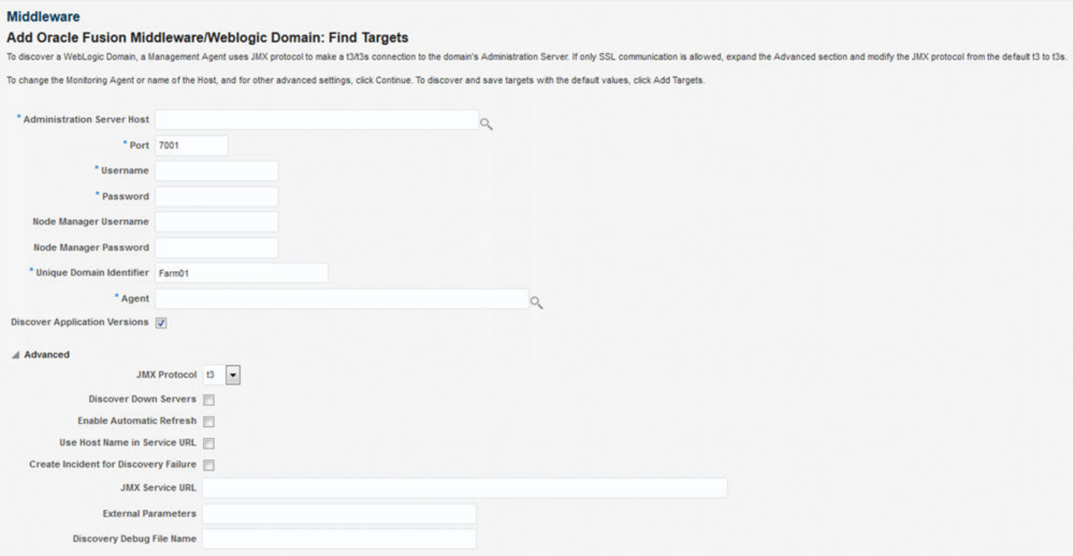

Middleware

Add Oracle Fusion Middleware/Weblogic Domain: Find Targets

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

To change the Monitoring Agent or name of the Host, and for other advanced settings, click Continue. To discover and save targets with the default values, click Add Targets.

* Administration Server Host	
* Port	7001
* Username	
* Password	
Node Manager Username	
Node Manager Password	
* Unique Domain Identifier	Farm01
* Agent	
Discover Application Versions	☑

Advanced

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

- Administration Server Host

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

- Port

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

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

You can access My Oracle Support at the following URL:

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

- WebLogic Administration Server user name and password

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

- Node Manager user name and password
- Unique Domain Identifier.

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

- Agent

  The host name for a Management Agent that will be used to discover the Fusion Middleware targets.If a Management Agent exists on the WebLogic Administration Server host, the host name for this Management Agent will be supplied by default. However, you can specify any Management Agent on any host that is managed by Enterprise Manager to perform the discovery.

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

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

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

Advanced Parameters

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

You can access My Oracle Support at the following URL:

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

- JMX Protocol

  Used to make a JMX connections to the Administration Server. For Secure domain JMX protocol - use t3s. If WebLogic domain is using a demo certificate, this certificate is automatically updated to monitoring and discovery Agent.

- Discover Down Servers

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

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

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

- Discover Application Versions

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

- JMX Service URL

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

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

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

- Discover Application Versions

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

- Enable Automatic Refresh

  This option refreshes the WebLogic domain every 24 hours.

- Use Host Name in Service URL

  You can use host name in service URL instead of JMX. It is recommended to use this option if you are using a private network and there are many hosts using the same IP address.

- Create Incident for Discovery Failure

  This option creates an OMS incident if discovery fails. You can view the incident from the Support workbench page.

- External Parameters

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

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

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

- Discovery Debug File Name

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

6. Click **Continue**. Enterprise Manager will discover all Fusion Middleware targets within the domain.
7. Click **Close** in the Finding Targets dialog to automatically assign Management Agents to the discovered targets.

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

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

   Note that you can also manually assign Management Agents to specific targets, if desired.
8. As a rare case, if you want to disable one or more target types for discovery, scroll down to the **Disable Target Types** section under **Advanced,** and move the target type from the **Available Target Types** list to the **Selected Target Types** list.

   Click **Refresh Targets** to view the refreshed **Targets and Agents Assignments** table.

   > **Note:**
   >
   > - If you disable a target type, it will remain disabled for future refresh operations.
   > - Child target types are disabled if you disable a parent target type.

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

   The Saving Target to Agent processing window appears, indicating how many total targets have been added and successfully saved. It will also indicate the number of targets that were unsuccessfully added.
10. Click **Close** in the processing window when finished. The Results page displays the targets and Agent assignments.
11. Click **OK** when finished. There may be a delay before these targets are visible and monitored. All the agents used for monitoring the targets must be up.

> **Note:**
>
> After you discover a middleware target for the first time, it is recommended that you learn the best practices for monitoring and managing the discovered target. To navigate to the Target Management Best Practices page, from the target home page, select the **WebLogic Domain** menu, and then select **Target Management Best Practices.** The page lists the best practices items for a Fusion Middleware Domain.

Adding Multiple WebLogic Domains Using EM CLI

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

The discover_wls verb can be used to discover the supported WebLogic domain versions. The verb reads a file named domain_discovery_file that contains the information required to discover each domain.

> **Note:**
>
> To know the supported WebLogic domain versions:
>
> 1. Log into **https://support.oracle.com/**.
> 2. On the My Oracle Support home page, select **Certifications** tab.
> 3. On the Certifications page, enter the product name as **Enterprise Manager Base Platform - OMS** in the Product field and select the relevant release number from the **Release list.**
> 4. Click Search.
> 5. In the Certification Results section, expand the **Application Servers** menu to view the supported versions.

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

Discovering New or Modified Domain Members

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

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

This section covers the following:

- Enabling Automatic Discovery of New Domain Members
- Manually Checking for New or Modified Domain Members

Enabling Automatic Discovery of New Domain Members

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

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

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

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

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

Manually Checking for New or Modified Domain Members

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

1. From the **Targets** menu, select **Middleware**.
2. Click the WebLogic Domain you want to (enable the job for in the Middleware home page) refresh.
3. From either the **Farm** or **WebLogic Domain** menu, select **Refresh WebLogic Domain**. The Refresh WebLogic Domain dialog opens.
4. Click **Add/Update Targets**. The Management Agent refreshes by connecting to the Administration Server. The Administration Server must be up for the refresh to occur. Click **Close** on the Confirmation page. Enterprise Manager will search the domain for new and modified targets.

   When any entity in a domain is removed from a WebLogic domain such as WebLogic j2eeaserver, j2eeapp, and the like, they are still displayed in Enterprise Manager. Click **Remove Targets** if you do not need the historical data of these targets. The obsolete targets which can be removed from the domain are then displayed.
5. Discovery Debug File Name option

   If problems occur while adding Middleware-related targets or refreshing domain membership, you can enable additional debugging information to quickly diagnose and resolve the issue. This file will be created in the discovery Agent's log directory.
6. The Assign Agents page displays the Fusion Middleware targets discovered and the Management Agent assigned to each. Click **Add Targets** to assign Management Agents as listed in the Assign Agents page.

   Agents are automatically assigned as follows:

   - If a local Agent can be found on the target host, that Agent will be assigned.

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

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

This page also provides the option of Selective Discovery. Using this option, you can disable the discovery of only new target types.

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

7. The Saving Targets to Agent processing window appears, indicating how many total targets have been added and successfully saved. It will also indicate the number of targets that were unsuccessfully added.
8. Click **Close** in the processing window when finished. The Results page displays the following options: Show Targets Details and Show WebLogic Domain Global Properties. The Show Targets Details page shows the targets and Agent assignments.

   Note: If there were targets that were removed, you can go back to the Refresh WebLogic Domain page and click **Remove Targets** to remove the targets and any historical information in the Management Repository. See Removing Middleware Targets.
9. Click **OK** when finished. There may be a delay before these targets are visible and monitored. All the agents used for monitoring the targets must be up.

   To do this, on the WebLogic domain homepage, from the WebLogic domain menu, select **Target setup**, and then click **Monitoring credentials.**

Adding Standalone Oracle HTTP Servers

To add standalone Oracle HTTP Servers, follow these steps:

- Meeting the Prerequisites
- Adding Standalone Oracle HTTP Servers Using the Guided Discovery Process

> **Note:**
>
> To view a visual demonstration on how to discover and manage standalone Oracle HTTP Servers, access the following URL and click **Begin Video.** The discovery described in this visual demonstration is based Enterprise Manager Cloud Control 12c Release 2 (12.1.0.3) and the Oracle Fusion Middleware Plug-in 12.1.0.5.
>
> `http://apex.oracle.com/pls/apex/f?p=44785:24:0::::P24_CONTENT_ID,P24_PREV_PAGE:8529,1`

Meeting the Prerequisites

Before you discover a standalone Oracle HTTP server, meet the following:

- Ensure that an Oracle Management Agent (Management Agent) is installed on the host where the standalone Oracle HTTP Server is running. For instructions to install a Management Agent, see the *Oracle Enterprise Manager Cloud Control Basic Installation Guide* available in the Enterprise Manager documentation library.

- Ensure that the standalone Oracle HTTP Server you are about to discover is of one of the following releases: 12.2.1.x, 12.1.3.x, 12.1.2.x, 11.1.1.9.x, 11.1.1.7.x, 11.1.1.6.x, 11.1.1.5.x, 11.1.1.4.x, 11.1.1.3.x, 11.1.1.2.x, 11.1.1.1.x, 10.1.2.x.

Adding Standalone Oracle HTTP Servers Using the Guided Discovery Process

To add a standalone Oracle HTTP server, follow these steps:

> **Note:**
>
> You can add only standalone Oracle HTTP Servers using this method. To add a managed Oracle HTTP Server, use the Add Oracle Fusion Middleware/Weblogic Domain wizard.

1. From the **Setup** menu, select **Add Target,** then select **Add Targets Manually.**
2. On the Add Targets Manually page, select **Add Using Guided Process.**
3. From the Target Types list, select **Standalone Oracle HTTP Server.**
4. Click **Add Using Guided Process.**
5. On the Add Standalone Oracle HTTP Server page, provide the following details, and click **Add Target.**

Element	Description
Oracle HTTP Server Host	Click the search icon to search and select the host where the standalone Oracle HTTP Server is running. In the Search and Select: Targets dialog, search the host, select it, and click **Select.** For example, `example.com`
Agent URL	URL of the Management Agent that is installed on the host where the standalone Oracle HTTP Server is running. Appears automatically by default when you select the standalone Oracle HTTP Server host. For example, `https://example.com:1838/emd/main/`
Oracle HTTP Server Port	Specify the Oracle HTTP Server port. The Oracle HTTP Server port is not validated. You need to specify the correct values.
Target Name	Specify the name with which you want to add and monitor the standalone Oracle HTTP Server target in Enterprise Manager. For example, `Standalone_OHS1` Once the standalone Oracle HTTP Server is discovered and added to Enterprise Manager , from the **Targets** menu, select **All Targets.** On the All Targets page, in the **Search Target Name** field, enter the target name with which you discovered and added the standalone Oracle HTTP Server target, and press **Enter.** Enterprise Manager displays the Home page of the standalone Oracle HTTP Server, with the target name you specified while discovering and adding it.
Oracle Home	Click the search icon to log in to the standalone Oracle HTTP Server host, and select the Oracle home where the standalone Oracle HTTP Server is running. For example, /u01/software/oracle/Standalone_OHS1

Element	Description
Configuration Path	Click the search icon to log in to the standalone Oracle HTTP Server host, and select the `httpd.conf` file. The value for this field must ideally be the absolute directory path to the `httpd.conf` file. For example, `/u01/software/oracle/Standalone_OHS1/Oracle_WT1/instances/instance1/config/OHS/ohs1/httpd.conf`
Version	Select the version of the target. If you selected a 11.x target, you need to specify the OHS Process Owner Credentials, or select from a list of available OHS Process Owner/ SSH Key Credentials.. If you selected a 12c target, you need to specify the Node Manager credentials.
OHS Process Owner Credentials	*(This field is only for discovering standalone Oracle HTTP Server 11.x targets)* Provide the OHS Process Owner Credentials. These credentials are optional and are required for collecting metrics for the target. These credentials are stored as monitoring credentials for the target and can be updated from the monitoring credentials page. To use SSH Key Credentials, create the SSH Host Credentials by accessing the Security tab, selecting Credentials, and then navigating to the Named Credentials page.
Node Manager User Name	*(This field is only for discovering standalone Oracle HTTP Server 12c target)* Specify the Node Manager user name.
Node Manager Password	*(This field is only for discovering standalone Oracle HTTP Server 12c target)* Specify the Node Manager password.

Adding Exalytics Targets

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

An Exalytics system consists of one or more Exalytics machines. The machine(s) can be physical, virtualized, or a mix of both.

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

Related targets running on the Exalytics machine such as OBIEE, Times Ten and WebLogic, need to be discovered following the respective flows for those components. If the Exalytics target has been discovered, the related targets will be associated with the Exalytics target when they are discovered. Else, the association will happen when the Exalytics target is discovered.

To add an Exalytics target, follow these steps:

- Meeting the Prerequisites
- Adding Exalytics System Targets Using the Guided Discovery Process

Meeting the Prerequisites

Before adding an Exalytics machine target, you must meet the following prerequisites:

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

  For discovering a virtual machine, make sure that at least one of the virtual Exalytics has a Management Agent running on it. The Management Agent should be on one of the VM Guest part of the virtual Exalytics Machine deployment. However, all virtual machines that contain components that need to be monitored require a Management Agent to be deployed on.
- To identify the Exalytics machine, ensure that you have the context info file which contains the Exalytics machine ID.
- To discover and monitor ILOM for a virtual Exalytics machine, you must install IMPITOOL on the VM guest. To install IMPITOOL, follow these steps:
  1. Download the latest Hardware Management Pack compatible with the operating system on your VM guest. Instructions for downloading the Hardware Management Pack are available here: `http://www.oracle.com/technetwork/server-storage/servermgmt/downloads/index.html`
  2. Extract the IMPITOOL package from the downloaded zip file and install it on the operating system on the VM guest.

Adding Exalytics System Targets Using the Guided Discovery Process

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

1. From the **Setup** menu, select **Add Target,** and then select **Add Targets Manually.**
2. On the Add Targets Manually page, select **Add Using Guided Process (Also Adds Related Targets).**
3. From the Target Types drop-down list, select **Exalytics System,** and then click **Add Using Guided Process...**

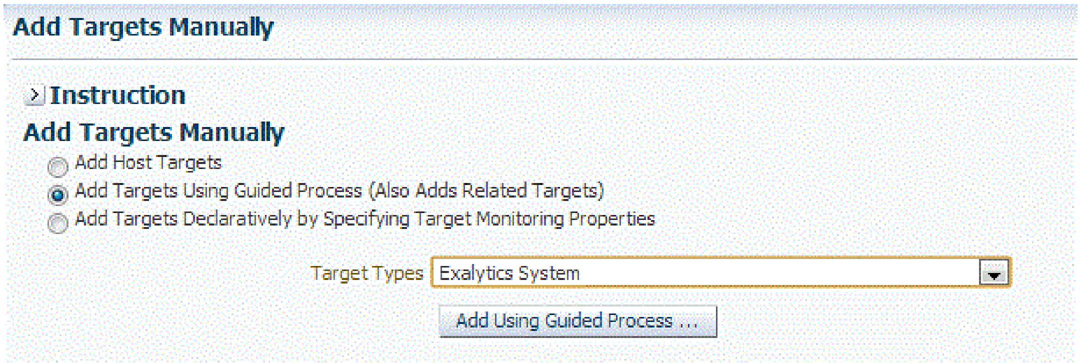

Add Targets Manually

› Instruction

Add Targets Manually

○ Add Host Targets
● Add Targets Using Guided Process (Also Adds Related Targets)
○ Add Targets Declaratively by Specifying Target Monitoring Properties

Target Types: Exalytics System

Add Using Guided Process ...

4. On the Discover Exalytics System page, provide a name for the Exalytics System, and then click **Add Machine.**

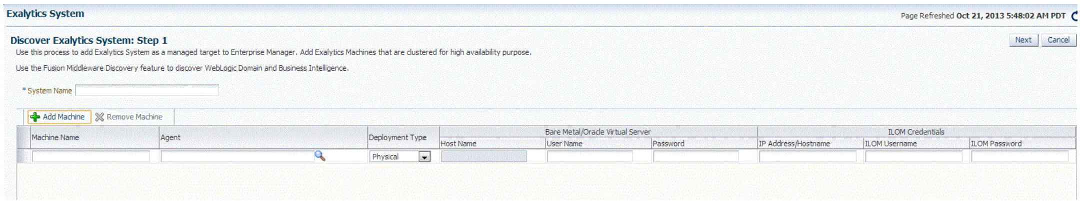

Discover Exalytics System: Step 1

Use this process to add Exalytics System as a managed target to Enterprise Manager. Add Exalytics Machines that are clustered for high availability purpose.

Use the Fusion Middleware Discovery feature to discover WebLogic Domain and Business Intelligence.

* System Name

Machine Name	Agent	Deployment Type	Bare Metal/Oracle Virtual Server			ILOM Credentials		
			Host Name	User Name	Password	IP Address/Hostname	ILOM Username	ILOM Password
		Physical						

5. Provide the following details:
   - **Machine Name**

     Provide a unique for the Oracle Exalytics machine target.
   - **Agent**

     Specify or select the Management Agent to use for the Oracle Exalytics Machine discovery process.
   - **Deployment Type**

     Select **Physical** or **Virtual** depending on if you want to discover a physical target or a virtual target.
   - **Host Name** (only for Virtual target)

     Enter the host name or IP Address where Oracle Virtual Server is running
   - **User Name and Password**

     For a physical Exalytics machine, provide credentials of the root user which has privileges to run the imageinfo command.

     For a virtual Exalytics machine, provide credentials to log in to Oracle Virtual Server (OVS). You should have privileges to run the imageinfo command.
   - **ILOM Credentials**

     Specify the ILOM IP address or hostname, ILOM username, and ILOM password of the root user.

     > **Note:**
     >
     > The ILOM Credentials are optional. If you do not add the ILOM Credentials, the ILOM target will not be discovered.
     >
     > You can add the ILOM target later by using the Refresh Exalytics System option.

   Click **Next.**
6. A confirmation box appears. Click **OK.**

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

Removing Middleware Targets

Removing Middleware targets from the Management Repository:

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

4

Discovering, Promoting, and Adding System Infrastructure Targets

This chapter describes how you can use Enterprise Manager to discover, promote, and add system infrastructure targets. This chapter covers the following:

- Discovering and Promoting Oracle MiniCluster
- About Discovering, Promoting, and Adding System Infrastructure Targets
- Discovering and Promoting Operating Systems
- Discovering and Promoting Oracle Solaris Zones
- Discovering and Promoting Oracle VM Server for SPARC
- Discovering and Promoting Servers
- Discovering and Promoting Oracle SuperCluster
- Configuring Snmp traps for Supercluster and Minicluster monitored hosts
- Discovering and Promoting PDUs
- Discovering and Promoting Oracle ZFS Storage
- Discovering Fabrics
- Related Resources for Discovering and Promoting System Infrastructure Targets

Discovering and Promoting Oracle MiniCluster

Before you begin the discovery process, there are several checks you should perform to ensure a smooth discovery.

Prerequisites

Enterprise Manager Agents have to be installed on Oracle Solaris global zones of both MiniCluster compute nodes. You can install EM Agents into Oracle Solaris Global and Non-Global Zones using Add Host Targets wizard.

Credentials Required for Oracle MiniCluster Discovery

The following are the credentials required for the discovery of Oracle MiniCluster.

Target Type	Credentials	Description
Systems Infrastructure Server	ILOM Monitoring Credentials, SNMP Credentials	Required to monitor MiniCluster compute nodes through their ILOMs

Oracle MiniCluster Discovery

To discover the Oracle MiniCluster system, perform the following steps:

Prerequisites

It is required to use Enterprise Manager Agents on both MiniCluster compute node Oracle Solaris global zones. Deploy the Enterprise Manager Agents to the Oracle Solaris global zones of the MiniCluster compute nodes prior to starting the Oracle MiniCluster discovery process.

> **Note:**
>
> To enable monitoring of Oracle VM for SPARC the non-privileged user used to install and run the Enterprise Manager agent must be granted the `solaris.ldoms.read` and `solaris.ldoms.ldmpower` authorizations and be assigned the LDoms Power Mgmt Observability rights profile.
>
> For example:
>
> ```
> /usr/sbin/usermod -A solaris.ldoms.read,solaris.ldoms.ldmpower
> oracle
> ```
>
> ```
> /usr/sbin/usermod -P 'LDoms Power Mgmt Observability' oracle
> ```

1. Log in to Enterprise Manager.
2. Under Setup, click **Add Target**, then click **Add Targets Manually**.
3. In the Overview section, click **Add using Guided Process**.
4. In the Add Using Guided Process screen, scroll down to Oracle MiniCluster and click **Add**. The Oracle MiniCluster Discovery wizard opens.
5. The Introduction wizard guides you through the steps required to discover Oracle MiniCluster in Oracle Enterprise Manager. Click **Next** to continue with the discovery process.

   > **Note:**
   >
   > The guided discovery process assumes that Oracle Enterprise Manager Agents are already deployed to Oracle Solaris Global zones of the MiniCluster Compute Nodes. The Enterprise Manager agent must be deployed in order to discover MiniCluster system, monitor disk shelves and a virtualization stack on compute nodes.
   >
   > If you want to monitor database instances running on a DB domain, you must install the Enterprise Manager Agents on all the Zones where the databases are installed.

6. In the Discovery Input screen provide the Agent EMD URL. Click the search icon and select an agent.

   Click **Next** to proceed to the next step. The Oracle MiniCluster discovery process is started. When the discovery is completed, a confirmation window displays the number of targets discovered. Click **Close** to close the window.

7. The Discovery Prerequisites screen opens. It displays basic information of the discovery. It also displays any warnings and errors found during the discovery process. Click **Next** to continue.

8. The Discovered Targets screen lists all the targets discovered in the Oracle MiniCluster system. All targets are selected by default. The Managed column indicates if the targets have already been discovered and managed by Enterprise Manager. Deselect the targets that you do not want to be discovered and monitored.

   a. Select monitoring and backup agents for MiniCluster compute node ILOMs.

   b. Provide ILOM Monitoring Credentials for the ILOMs of the MiniCluster system compute nodes.

   c. Select monitoring and backup agents for MiniCluster disk shelves.

   Click **Next** to proceed to the Review screen.

   > **Note:**
   >
   > Before you are taken to the Review screen, a validation is executed to find out if the correct agent for Disk Shelf Monitoring was selected. This agent has to be deployed to the host which Disk Shelf is directly connected using Serial Attached SCSI. If you selected an agent on another host where Disk Shelf is not directly connected, you will get a warning. You have to select the correct agent first otherwise you cannot proceed in MiniCluster discovery. After the Disk Shelf agent selection validation, another validation is executed. This validation checks if you correctly configured Solaris host so you can see hardware related incidents in Enterprise Manager. You can see result of this validation on the Review screen. If the Solaris Host is not configured properly to deliver incidents, you won't see FMA alerts with Disk Shelf failures as incidents in Enterprise Manager. This configuration is not mandatory to finish MiniCluster discovery but highly recommended. For information on how to configure incidents for Solaris Host, see Configuring Snmp traps for Supercluster and Minicluster monitored hosts.

9. In the System Review screen, review the system information.

   a. Click **Test Connection** to check whether the specified credentials are correct for all selected targets.

   b. Click **Promote Targets** to create and manage targets. This may take few minutes to complete. You are informed about any errors that occur during the process. In case there is at least one error, no targets are created. Fix all errors and rerun the discovery. Once the process completes without any errors, the targets are managed and you can view the Oracle MiniCluster system with all its targets.

About Discovering, Promoting, and Adding System Infrastructure Targets

Cloud Control enables you to monitor a variety of system infrastructure targets, including Oracle VM Server for SPARC, Oracle Solaris Zones, Oracle SuperClusters, servers, operating systems, storage, networks, racks, and power distribution units (PDUs). The steps to add System Infrastructure targets is the same for most targets. However, a few targets do require special procedures.

When fully managed, systems infrastructure targets provides monitoring and an enterprise-wide view of the bottom half of the stack, including Oracle Solaris and Linux operating systems, virtualized operating systems (zones) and virtual machines (logical domains), power distribution units (PDUs), servers, storage appliances, storage for a host, and network resources.

> **Note:**
>
> To monitor ILOM servers or virtualization targets, deploy the EM Agent with `sudo` privileges or manually install and run the `root.sh` script.
>
> Monitoring requires that the `root.sh` script is executed and installs the nmr binary. When you deploy an EM Agent with `sudo` privileges, the `root.sh` script is executed.

How long it takes to begin monitoring all discovered targets varies, depending on your environment, configuration, and the number of guests. For virtualization targets, the Summary dashlet on the Virtualization Platform includes the total number of guests being monitored. When guests are discovered and configured on the platform, but not yet monitored, an icon appears in the dashlet along with the number of guests that are pending monitoring. For example, in Figure 4-1 there are 63 guests configured on the host and 16 of those guests are not yet being monitored. The icon will not appear in the dashlet when all guests are monitored.

Figure 4-1 Summary Dashlet Showing the Number of Guests to be Monitored

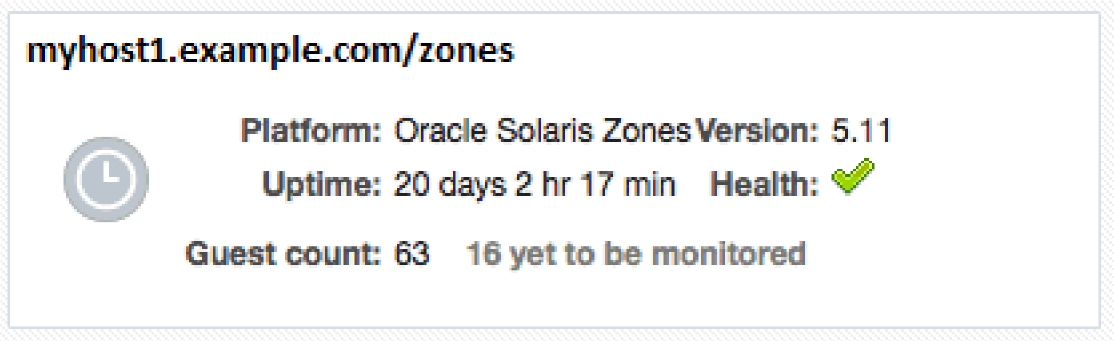

myhost1.example.com/zones

Platform: Oracle Solaris Zones Version: 5.11

Uptime: 20 days 2 hr 17 min Health: ✓

Guest count: 63 16 yet to be monitored

Discovering and Promoting Operating Systems

Oracle Solaris and Linux operating systems are discovered and promoted as part of the host discovery and promotion process. See Discovering and Adding Host and Non-Host Targets for more about discovery, promotion, and the steps to discover and promote a host target.

Discovering and Promoting Oracle Solaris Zones

Discovering and promoting Oracle Solaris Zones relies on the host discovery. When a host on a global zone is promoted, it triggers the discovery and monitoring of the zones.

> **Note:**
>
> To take advantage of all supported Oracle Solaris Zone monitoring, the Oracle Solaris release on the global zone must be Oracle Solaris 10 Update 11 or later.
>
> To discover and monitor zone targets, deploy the EM Agent with `sudo` privileges or manually install and run the `root.sh` script.

When a host is in a zone, you should discover the global zone before you add the virtual server targets. If you see a message stating that there is not a virtual platform associated with this virtual server, discover and add the virtual server's global zone as a managed target.

Oracle Solaris Zones are displayed with the following target display names:

- **Virtualization Platform:** <Oracle Solaris Global Zone OS Host Name> (Solaris Zones Virtual Platform)
- **Virtual Server:** <Oracle Solaris Zone Name>

You can change the display name of a discovered Oracle Solaris Zone using the CLI. For example, to change a Virtualization Platform display name:

```
emcli modify_target -name="Virtual Platform Target Name" -
type="oracle_si_virtual_platform_map" -display_name="New Display Name"
```

To change a Virtual Server display name:

```
emcli modify_target -name="Virtual Server Target Name" -
type="oracle_si_virtual_server_map" -display_name="New Display Name"
```

See Discovering and Adding Host and Non-Host Targets for more about discovery, promotion, and the steps to add a host.

Discovering and Promoting Oracle VM Server for SPARC

Discovering and promoting Oracle VM Server for SPARC relies on the host discovery. When a host on a control domain is promoted, it triggers the discovery and monitoring of the other domains.

> **Note:**
>
> To discover and monitor targets, deploy the EM Agent with `sudo` privileges or manually install and run the `root.sh` script.

When an agent is deployed to a control domain or logical domain OS, related targets including the ILOM server, virtual platform, and virtual server are automatically promoted.

When a host is in a logical domain, you should discover the control domain before you add the virtual server targets. If you see a message stating that a virtual platform is not associated with this virtual server, discover and add the virtual server's primary domain as a managed target.

An Oracle VM Server for SPARC must meet the following prerequisites to be discovered and added:

- The Oracle Solaris release on the control domain must be Oracle Solaris 11.1 or later.
- The Oracle VM Server for SPARC software must be version 3.1 or later.
- The non-privileged user used to install and run the Enterprise Manager agent must be granted the `solaris.ldoms.read` and `solaris.ldoms.ldmpower` authorizations and be assigned the LDoms Power Mgmt Observability rights profile. For example:

  ```
  /usr/sbin/usermod -A solaris.ldoms.read,solaris.ldoms.ldmpower oracle
  /usr/sbin/usermod -P 'LDoms Power Mgmt Observability' oracle
  ```

  > **Note:**
  >
  > You must run these commands only from the primary control domain because the Oracle VM Server for SPARC runs only on the primary control domain.

If other software is consuming a lot of resources on the system or control domain, it might take longer to discover and monitor the target.

Oracle VM Server for SPARC are displayed with the following target display names:

- **Virtualization Platform:** <Control Domain OS Host Name> (OVM SPARC Virtual Platform)
- **Virtual Server:** <Domain Name>

You can change the display name of a discovered Oracle VM Server for SPARC using the CLI. For example, to change a Virtualization Platform display name:

```
emcli modify_target -name="Virtual Platform Target Name" -
type="oracle_si_virtual_platform_map" -display_name="New Display Name"
```

To change a Virtual Server display name:

```
emcli modify_target -name="Virtual Server Target Name" -
type="oracle_si_virtual_server_map" -display_name="New Display Name"
```

See Discovering and Adding Host and Non-Host Targets for more about discovery, promotion, and the steps to promote a host.

Discovering and Promoting Servers

Discovering and adding servers can be performed separately or as part of the host discovery. When a host is promoted, it triggers the discovery and auto promotion of the server supporting the host. See Discovering and Adding Host and Non-Host Targets for more about discovery, promotion, and the steps to promote a host.

To discover and promote an ILOM server, see the following:

- Discover an ILOM Server Using ILOM-SSH Through the User Interface
- Discover an ILOM Server Using REST Through the User Interface
- Discover an ILOM Server Using the Command Line Interface

To change the display name after you have discovered an ILOM server, see Change the Display Name of a Discovered ILOM Server.

> **Note:**
>
> To auto-promote targets, deploy the EM Agent with `sudo` privileges or manually install and run the `root.sh` script.
>
> A Systems Infrastructure Server is displayed with the ILOM server name if the ILOM is discovered before the host, and if the HMP package is installed on the host. Otherwise, the Systems Infrastructure Server is displayed with the name "<*host name*>/server".

For some servers, a minimum firmware version is recommended to improve performance.

> **Note:**
>
> For SPARC M6-32 servers, system firmware 9.4.2.E or higher is recommended.

Discover an ILOM Server Using ILOM-SSH Through the User Interface

To discover an ILOM server using ILOM-SSH, perform the following steps:

1. From the **Setup** menu, select **Add Target**.
2. Click **Add Targets Manually**.
3. Click **Add Using Guided Process** listed under Add Non-Host Targets Using Guided Process.

   The Add Using Guided Process window is displayed with the list of Guided Discovery and Discovered Target Types.
4. Select **Systems Infrastructure Server ILOM** from the list in the Add Using Guided Process window.
5. Click **Add**.
6. Select **ILOM SSH Credentials** which is the default option.
7. Click **Discover Server Target**.
8. For Target, enter the following details:
   a. Target Name: Enter the name of the ILOM server.
   b. Server ILOM DNS Name or IP Address: Enter the Server DNS Name or IP Address.
9. For Monitoring Agents, enter the following details:
   a. Enter the Monitoring Agent EMD URL.
   b. (Optional): Enter the Backup Agent EMD URL.
10. For Monitoring Credentials, enter the following SSH monitoring credentials:
    a. Select the Credential type as **ILOM SSH Credentials**.
    b. Enter the root user name in the Username field.
    c. Enter the root password in the Password and Confirm Password fields.
11. For SNMP v1 and v2 configurations, enter the following credential parameters:

   a. Select the Credential type as **SNMP V1/V2 Credentials**.
   b. Enter your community string in the Community String and Confirm Community String fields.

12. For SNMP v3 configurations, enter the following parameters:
    a. User name
    b. Authorization password
    c. Confirm authorization password
    d. Authorization protocol, either MD5 or SHA
    e. Privacy password
    f. Confirm privacy password
    g. Privacy Protocol
13. Click **Add**.

A confirmation window appears when the target is successfully added.

> **Note:**
>
> You must at least provide an SNMP V1/V2 community string even if you want to explicitly disable the SNMP configuration to occur on the ILOM server.
>
> For disabling the SNMP Monitoring Configuration, type `true` in the **Skip SNMP Subscription** field.

Discover an ILOM Server Using REST Through the User Interface

> **Note:**
>
> - Ensure that System Infrastructure Server is discovered using ILOM-SSH credentials for full monitoring.
> - REST monitoring is available only for ILOM *version 5.0.1 or later*.
> - REST Access Point is not supported for SPARC-based machines.

To discover an ILOM server using REST, perform the following steps:

1. From the **Setup** menu, select **Add Target**.
2. Click **Add Targets Manually**.
3. Click **Add Using Guided Process** listed under **Add Non-Host Targets Using Guided Process**.

   The **Add Using Guided Process** window is displayed with the list of guided discovery and discovered target types.
4. Select **Systems Infrastructure Server ILOM** from the list.
5. Click **Add**.

6. Select **HTTPS Credentials** option.
7. Default **HTTPS Port** is set to `443`. If the server uses a HTTPS Port other than default one, enter it in the **HTTPS Port** text box.
8. Click **Discover Server Target**.
9. For **Target**, enter the following details:
    a. **Target Name**: Enter the name of the ILOM server.
    b. **DNS Name or IP Address**: Enter the Server DNS Name or IP Address.
10. For **Monitoring Agents**, enter the following details:
    a. Enter the **Monitoring Agent EMD URL**.
    b. (Optional): Enter the **Backup Agent EMD URL**.
11. For **Monitoring Credentials**, enter the following details:
    a. Enter the **Alias** and **Password** in the respective fields.
    b. Enter the password again in the **Confirm Password** field.

Discover an ILOM Server Using the Command Line Interface

You can discover a server using the *`emcli`* command line tool. You must configure the command line interface before you can issue commands. For more information, click the **Setup** menu, then click **Command Line Interface**. Follow the Download and Deploy instructions.

To discover a server using the `emcli`, perform the following steps:

1. Open your command line on the host where OMS is running.
2. Login to emcli using command `emcli login -username=`<*Your user name*>.
3. Type the password when prompted.
4. Execute command `emcli sync`.
5. Discover a new server using the `emcli add_target` command and define the following options:
    - `-name=`Name of the server to be displayed within Enterprise Manager
    - `-type=`**`oracle_si_server_map`**
    - `-host=`Host name from which you discover the server target for monitoring
    - `-access point name=`Access point name to be displayed within EM
    - `-access_point_type=`**`oracle_si_server_ilom`** for ILOM-SSH Access Point

      `-access_point_type=`**`oracle_si_server_http`** for REST Access Point
    - `-properties=`key:ILOM IP address that will be used in discovery. For example,
        – `-properties='dispatch.url=ilom-ssh://<`*`ILOM_IP_ADDRESS`*`>'` for ILOM-SSH Access Point
        – `-properties='dispatch.url=https://<`*`ILOM_IP_ADDRESS`*`>'` for REST Access Point
    - `-subseparator=`properties with sub separator string, For example, you can use `=` as a separator between key value pairs. Alternatively, you can use the `-separator` option when there are multiple properties (key value pairs) that need to be separate.

- `-monitoring_cred=`The credentials of the server ILOM to be discovered

Example 4-1 SNMP Credential-based Configuration

The user will have to provide "snmpv1v2_v3" monitoring credentials, specifying the SNMPV1Creds and the desired SNMP Community string as shown in the example below:

```
emcli add_target -name=<target name> -type=oracle_si_server_map -host=<EM Agent>  -
access_point_name=<ILOM Access Point Name> -access_point_type=oracle_si_server_ilom -
subseparator=properties== -properties=dispatch.url=ilom-ssh://<ILOM server host name> '-
monitoring_cred=ilom_creds_set;oracle_si_server_ilom;ilom_creds;username:<ILOM user
name>;password:<ILOM user password>' '-
monitoring_cred=snmp_v1v2_v3;oracle_si_server_ilom;SNMPV1Creds;COMMUNITY:<SNMP Community
string>'
```

Example 4-2 Legacy Community String Property-based Configuration

Despite the fact that this not encouraged, it is still possible to perform the ILOM discovery using a SNMP community string as a property as in the example below:

```
emcli add_target -name=<target name> -type=oracle_si_server_map -host=<EM Agent> -
access_point_name=<ILOM server Access Point name>  -
access_point_type=oracle_si_server_ilom -subseparator=properties== '-
properties=SNMPCommunity=<SNMP Community String>;dispatch.url=ilom-ssh://<ILOM Server
host name>' '-
monitoring_cred=ilom_creds_set;oracle_si_server_ilom;ilom_creds;username:<ILOM user
name>;password:<ILOM User password>'
```

Example 4-3 Discovery without SNMP subscription

As with the UI discovery flow, the user can disable the automatic SNMP rules configuration on the ILOM server by setting the "SkipSnmpSubscription" property to "true" as shown below:

```
emcli add_target -name=<target name> -type=oracle_si_server_map -host=<EM Agent> -
access_point_name=<ILOM server Access Point name>  -
access_point_type=oracle_si_server_ilom -subseparator=properties== '-
properties=SkipSnmpSubscription=true;dispatch.url=ilom-ssh://<ILOM Server host name>' '-
monitoring_cred=ilom_creds_set;oracle_si_server_ilom;ilom_creds;username:<ILOM user
name>;password:<ILOM User password>'
```

Example 4-4 SNMP V3 Monitoring configuration

The user needs to provide the "snmp_v1v2_v3" monitoring credentials configuration to perform SNMP V3 subscription on the ILOM server. The following parameters will be configured:

- authUser
- authPwd
- authProtocol, either MD5or SHA
- privProtocol, which is optional
- privPwd, which is optional

```
emcli add_target -name=<target name> -type=oracle_si_server_map -host=<EM
Agent> -access_point_name=<ILOM server Access Point name> -
access_point_type=oracle_si_server_ilom -subseparator=properties='=' -
properties='dispatch.url=ilom-ssh://<ILOM Server host name>' -
monitoring_cred='ilom_creds_set;oracle_si_server_ilom;ilom_creds;username:<ILO
M user name>;password:<ILOM User password>' -
monitoring_cred='snmp_v1v2_v3;oracle_si_server_ilom;SNMPV3Creds;authUser:<SNMP
 Authorization User>;authPwd:<SNMP Authorization User
```

```
password>;authProtocol:<auth protocol>;privPwd:<SNMP privacy
password>;privProtocol:<privacy protocol>'
```

Example 4-5 REST Credential Based Discovery

We can discover REST Access Point for Systems Infrastructure Server using emcli as in the example below:

```
emcli add_target -name=<target name> -type=oracle_si_server_map -host=<EM
Agent> -access_point_name=<REST Access Point name> -
access_point_type=oracle_si_server_http -subseparator=properties='=' -
properties='dispatch.url=https://<ILOM Server host name>' -
monitoring_cred='ServerHttpCredentialSet;oracle_si_server_http;AliasCredential
;Alias:<REST user name>;Password:<REST user password>'
```

Change the Display Name of a Discovered ILOM Server

You can change the display name of a discovered server using the CLI by running the **emcli** command to modify the target name. For example:

```
emcli modify_target -name="Server Target Name" -type="oracle_si_server_map" -
display_name="New Display Name"
```

Discovering and Promoting Oracle SuperCluster

Before you begin the discovery process, there are several checks you should perform to ensure a smooth discovery.

Prerequisites

A discovery precheck script is available to automatically verify many of the common problem areas prior to discovery.

Some of Oracle SuperCluster discoveries on Oracle Enterprise Manager 13c may run into issues due to configuration mismatches in the software setup. The discovery pre-check script helps you resolve most common configuration problems. Run the script before the Oracle SuperCluster discovery and examine the output before proceeding with the discovery in Oracle Enterprise Manager.

For the best possible discovery result, run the script and resolve all issues even though some of them might be ignored to successfully finish the process.

The script is part of Enterprise Manager Agent bundle.

Obtain the Discovery Precheck Script

You can obtain the script in one of the following ways:

1. Access the script as part of Systems Infrastructure plug-in 13.2.2.0.0 after the plug-in is deployed to the agent:

   ```
   <agent installation directory>/plugins/
   oracle.sysman.si.discovery.plugin_13.2.2.0.0/discover/sscDiscoveryPreCheck.pl
   ```

2. Check My Oracle Support for the latest version of the prerequisites script. If you run the script downloaded from My Oracle Support, make sure you run it on a compute node (CDOM) where the Enterprise Manager Agent bundle is deployed.

Run the Discovery Precheck Script

Type the following commands to run the script:

- `$ cd <agent installation directory>/plugins/oracle.sysman.si.discovery.plugin_13.5.1.0.0/discover`
- When using the CLI, set the values of the following environment variables:

  `AGENT_DEP_SSHD_CP` to the location of `sshd` libraries in agent

  `AGENT_DEP_XMLPARSER_CP` to the location of `xmlparser2` libraries in agent

  For example:

  ```
  export AGENT_DEP_SSHD_CP=<sshd_jar_path_in_agent>/sshd-core-2.5.0.jar
  ```

  ```
  export AGENT_DEP_XMLPARSER_CP=<xmlparserv2_jar_path_in_agent>/xmlparserv2.jar
  ```
- `$ perl ./sscDiscoveryPrecheck.pl`

Type the following command to check the help for optional parameters:

- $ perl ./sscDiscoveryPrecheck.pl --help

As the script runs, you are prompted for various inputs. The script executes all the built-in checks and displays important messages on standard output. Detailed information is stored in a log file and can be used to debug execution errors.

The discovery pre-check script performs the following checks:

- Execution environment and network
- Network configuration (IP, host name validity, ping)
- Hardware monitoring credentials (optional)
- Detailed hardware or software checks (may require credentials)
- Correct ILOM versions of the monitored targetsExadata Cell management and cell server version and status
- PDU firmware version, Trap and NMS table availability
- IPMI tool version

Credentials Required for Oracle SuperCluster Discovery

The following are the credentials required for the discovery of Oracle SuperCluster.

Table 4-1 Credentials for SSC Discovery

Target Type	Credentials
Systems Infrastructure Server	ILOM Monitoring Credentials, SNMP Credentials
Systems Infrastructure InfiniBand Switch	ILOM Monitoring Credentials, SNMP Credentials
Systems Infrastructure CISCO Switch	Cisco Switch IOS Credentials, SNMP Credentials

Table 4-1 (Cont.) Credentials for SSC Discovery

Target Type	Credentials
Systems Infrastructure ZFS Storage Appliance Controller	ZFS SA Storage Controller SSH Credentials
Systems Infrastructure PDU	HTTP Monitoring Credentials, SNMP Credentials
Oracle Exadata Storage Server	Exadata Privileged Credentials, SNMP Community String. Allows SNMP subscription and can be used once only. Use other lower privileged credentials for monitoring.
Host	Agent Host User Credentials. This credential is required to setup passwordless access to the storage cell.

Manual Prerequisite Verification

Ensure that the host names of compute nodes and Exadata cells in each individual Oracle SuperCluster system have a unique prefix.

Oracle SuperCluster Discovery

It is recommended to use multiple Enterprise Manager Agents to manage targets. Deploy the agents prior to the Oracle Supercluster discovery process. The Enterprise Manager Agent has to be deployed to each of the virtualization platforms that is planned to be managed.

Figure 4-2 is a pictorial representation of Oracle SuperCluster discovery workflow.

Figure 4-2 Oracle SuperCluster Discovery Workflow

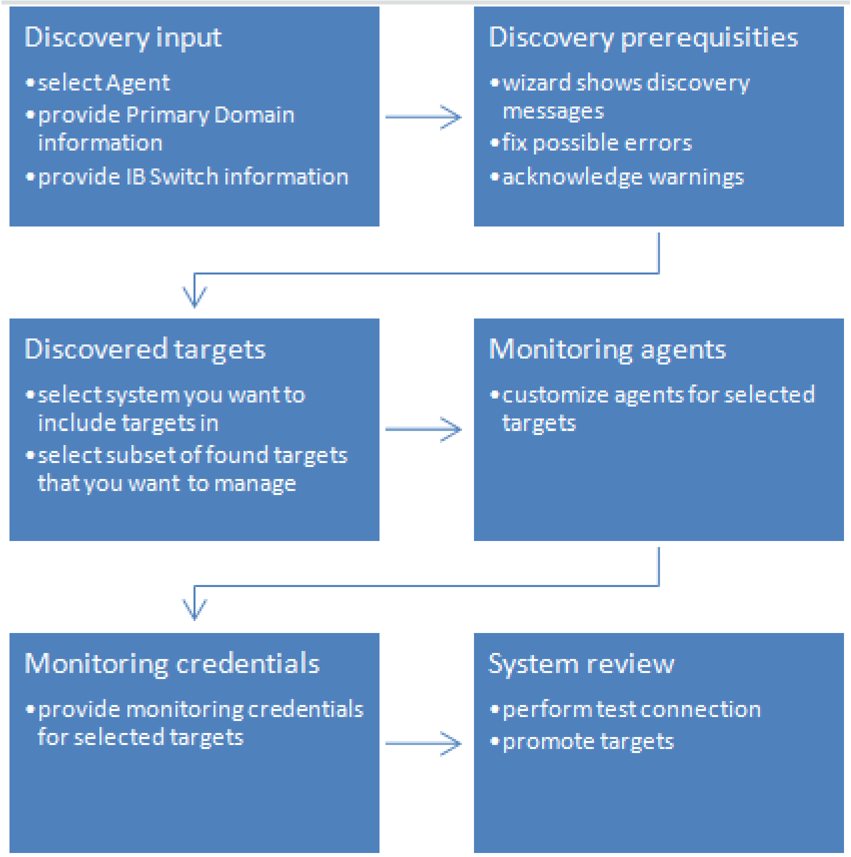
Discovery input
•select Agent
•provide Primary Domain information
•provide IB Switch information
Discovery prerequisities
•wizard shows discovery messages
•fix possible errors
•acknowledge warnings
Discovered targets
•select system you want to include targets in
•select subset of found targets that you want to manage
Monitoring agents
•customize agents for selected targets
Monitoring credentials
•provide monitoring credentials for selected targets
System review
•perform test connection
•promote targets

To discover the Oracle SuperCluster system, perform the following steps:

1. Log in to Enterprise Manager.
2. Under Setup, click **Add Target**, then click **Add Targets Manually**.
3. In the Overview section, click **Add Targets using Guided Process**.
4. In the Add Using Guided Process screen, scroll down to Oracle SuperCluster and click **Add**. The Oracle SuperCluster Discovery wizard opens.
5. The Introduction wizard guides you through the steps required to discover Oracle SuperCluster in Oracle Enterprise Manager. Click **Next** to continue with the discovery process.

> **Note:**
>
> The guided discovery process assumes that Oracle Enterprise Manager Agents are already deployed to Compute Nodes. The Enterprise Manager agent must be deployed in order to monitor a virtualization stack on compute nodes.

> **Note:**
>
> If you want to monitor database instances running on a DB domain, you must install the Enterprise Manager Agents on all the Zones and Guest LDOMs where the databases are installed.

6. In the Discovery Input screen, enter the required information.
   a. In the Agent EMD URL field, click the search icon and select an agent.

      > **Note:**
      >
      > The agent must be on the same network where one of the Oracle SuperCluster InfiniBand switches is located.

   b. In the Primary Domain Host name field, specify the host name of the primary domain that was used for OneCommand execution during setup of the Oracle SuperCluster system.

      > **Note:**
      >
      > The domain is used to discover system targets. This is a one-time operation. Discovery requires privileged user (root) credentials to discover details of the system and presence of catalog.xml or databasemachine.xml file in OneCommand directory located in /opt/oracle.SupportTools/onecommand. It is suggested to verify that configuration files mentioned above are up-to-date.

   c. In the Credential field, select **New** to create credentials.
      - Click **Named** if you already have existing credentials.
   d. In the User Name field, enter the name of the user with root privileges.
   e. In the Password field, enter the password of the compute node.
   f. Click **Save As** to save the credentials for later use.
   g. In the InfiniBand Switch Host name field, enter the DNS name or IP address of the InfiniBand switch.
   h. In the Credential field, select **New** to create credentials.
      - Click **Named** if you already have existing credentials.
   i. In the Username field, enter the user name to connect to the InfiniBand switch ILOM.
   j. In the Password field, enter the password for the InfiniBand switch ILOM.
   k. In the Run Privilege field, select **None**.
      - Click **Save As** to save the credentials for later use.
   l. (Optional) Click **Test Connection** to test your connection.
7. Click **Next** to proceed to the next step. The Oracle SuperCluster discovery process is started. When the discovery is completed, a confirmation window displays the number of targets discovered. Click **Close** to close the window.

8. The Discovery Prerequisites screen opens. It displays basic information of the discovery. It also displays any warnings and errors found during the discovery process. Click **Next** to continue.

   > **Note:**
   >
   > If you receive errors, you must fix the errors before you can proceed to the next step. To continue with warnings, read the warnings carefully and then acknowledge them.

9. The **Discovered Targets** screen lists all the targets discovered in the Oracle SuperCluster system. All targets are selected by default. The Managed column indicates if the targets have already been discovered and managed by Enterprise Manager. Deselect the targets that you do not want to have discovered and monitored.

   To associate discovered targets to an existing SuperCluster system target, user can provide name or use the target selector to pick an existing system. If newly discovered hardware belongs to a rack that is already member of the system, then the racks will be updated.

10. Click **Next** to proceed.

11. The **Monitoring Agents** screen allows you to assign monitoring and backup agents to targets selected in the previous step. The discovery wizard assigns available agents automatically to achieve best possible performance and reliability. For targets that are not already managed by Enterprise Manager you can manually select other agents than the default ones if necessary. You can assign the same primary or backup monitoring agent to all targets in a group of targets (like Compute nodes, etc.) using buttons next to first drop down selection with agents in every target group. You can reset automatically assigned agents using Reset button at the upper top corner of the page. Click **Next** when finished with the agent selection.

    > **Note:**
    >
    > The agent must be on the same network and be able to reach the target. Selected agents for a target should be on different compute nodes to prevent failure in case one compute node goes down.

12. The Monitoring Credentials screen lists the credentials that are used to monitor the targets. Click **Edit** to change or provide monitoring credentials for targets that do not have the information. Click **Next.**

    > **Note:**
    >
    > Use the root username and password to discover an InfiniBand network switch.

> **Note:**
>
> For the Monitoring Agent Hosts section, enter the user name and password of the user under which the EM agent is running on a given host.

> **Note:**
>
> Select the **Use for all** option to use the same credentials for all targets of the same target type.

> **Note:**
>
> Before you are taken to the Review Screen, a validation is executed to find out if the Solaris host was correctly configured so you can see hardware related incidents in Enterprise Manager. You can see result of this validation on the Review screen. If the Solaris Hosts was not configured properly to deliver incidents, you won't see FMA alerts generated on in Solaris OS with disk, memory and other hardware failures as incidents in Enterprise Manager. This configuration is not mandatory to finish SuperCluster discovery but highly recommended. To see how to configure incidents for Solaris Host, see "Enabling Incidents on Solaris Host".

13. In the System Review screen, review the system information. Click **Test Connection** to check whether the specified credentials are correct for all selected targets.
14. Click **Promote Targets** to create and manage targets. This may take few minutes to complete. You are informed about any errors that occur during the process. In case there is at least one error, no targets are created. Fix all errors and rerun the discovery. Once the process completes without any errors, the targets are managed and you can view the Oracle SuperCluster system with all its targets.

> **Note:**
>
> In the Oracle SuperCluster discovery wizard, if any hardware target is not discovered, you can return to any screen, correct the input data or deselect the problematic hardware, and then rerun the discovery. You can rerun the discovery later and discover any missing targets of the system.
>
> At this point, the Oracle SuperCluster is discovered and monitored by Enterprise Manager. If you want to discover and monitor the database cluster running on the system, you need to proceed with Oracle Exadata Database Machine discovery.

Configuring Snmp traps for Supercluster and Minicluster monitored hosts

After the EM agents are deployed to the engineered system, it is recommended to configure the hosts to send the snmp traps to the agents. Perform the following steps to configure the snmp:

1. Find the EM agent numeric IP address and port number.

   Login to the host where the EM agent is deployed. From the agent directory `<AGENT_HOME>/bin` run the following command to obtain the port number of the agent:

   ```
   $ emctl status agent | grep 'Agent URL'
   Agent URL: https://hostname.domain:3863/emd/main/
   ```

   The port number can be seen after the fully qualified domain name of the agent, it is 3863 in the example above.

   To get the numeric IP address, use ping or `nslookup` command.

2. Configure the snmpd.

   Add the`trap2sink` entry information into the `snmpd.conf` configuration file.

   The configuration file `snmpd.conf` is in different location for Solaris 10 and Solaris 11.

   For Solaris 11 the location is

   ```
   vi /etc/net-snmp/snmp/snmpd.conf
   ```

   For Solaris 10 the location is

   ```
   vi /etc/snmp/snmpd.conf
   ```

   Add the following line:

   ```
   trap2sink <numericip> public <transport>
   ```

   where,

   - The <transport> is the port number obtained in step (1). It is the port number of the EM agent.
   - The <numericip> is the IP address of the EM agent host, obtained in step (1).

   For example

   ```
   trap2sink 10.133.249.68 public 3863
   ```

3. Restart the SNMP services.

   For Solaris11:

   ```
   svcadm restart net-snmp
   svcadm restart snmp-notify
   ```

For Solaris10:

```
svcadm restart sma
```

If the SNMP service was not enabled before, ensure that the services are enabled.

For Solaris11:

```
svcadm enable net-snmp
svcadm enable snmp-notify
```

For Solaris10:

```
svcadm enable sma
```

4. Optionally, verify the snmpd configuration by making sure the status for Telemetry Status is **ON** in the host target, in the **Telemetry Alert Config** metric.

   After the configuration of the snmp through steps (1) – (3), navigate to **Host** home page and perform the following steps:

   a. Select **Host** and click **Configuration.**
   b. Click **Last collected.**
   c. Select **Telemetry Alert Config** metric.
   d. Click **Refresh,** wait for few minutes until the metrics get refreshed.

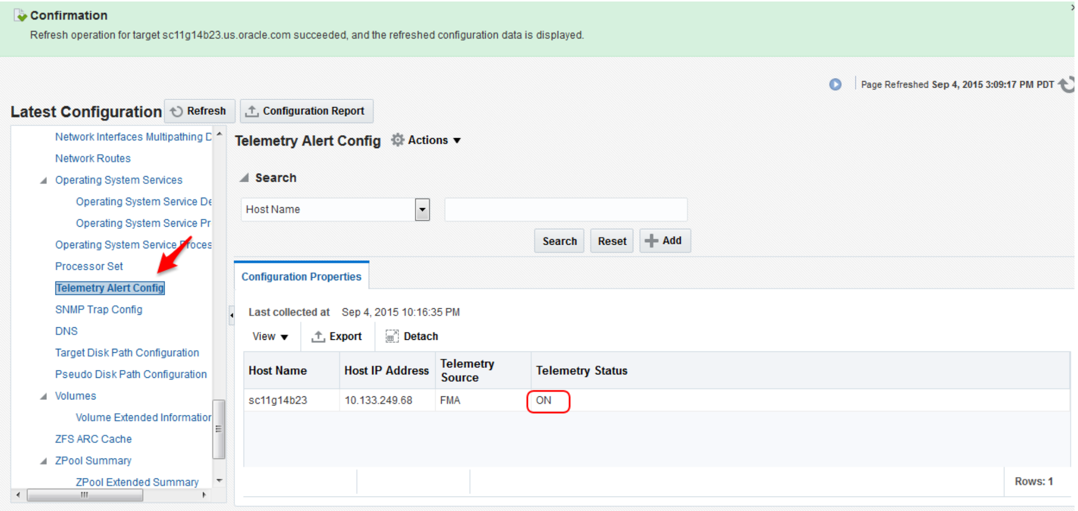

Confirmation

Refresh operation for target sc11g14b23.us.oracle.com succeeded, and the refreshed configuration data is displayed.

Page Refreshed Sep 4, 2015 3:09:17 PM PDT

Last collected at Sep 4, 2015 10:16:35 PM

Host Name	Host IP Address	Telemetry Source	Telemetry Status
sc11g14b23	10.133.249.68	FMA	ON

Rows: 1

If the Telemetry Status is **ON** in **Telemetry Alert Config** metric, it indicates the the configuration of FMA SNMP traps is done correctly.

Discovering and Promoting PDUs

Before you run the PDU discovery, verify that the NMS table and Trap Hosts Setup table of the PDU have an empty row for monitoring agent and an empty slot for backup agent, if you plan to use backup monitoring agent.

Verify PDU v1 NMS Table and Trap Hosts Setup Table

To verify and modify the NMS table and Trap Hosts setup table on a PDU v1 (see, PDU Version Identification), perform the following steps in the PDU user interface:

1. Open the PDU Management Interface in the web browser.
2. In the PDU User Interface, click **Net Configuration**.
3. Login with your user name and password.
4. Locate the NMS table and make sure there are enough empty slots for IP addresses of EM agents that will monitor the PDU.
    - Empty slot contains 0.0.0.0 value in the IP address field and empty string in the Community string field.
5. Enter 0.0.0.0 value in the IP address field and empty Community string field if there are not enough empty slots.
6. Click **Submit**.
7. Locate the Trap Hosts Setup table on the same page and make sure there are enough empty slots for IP addresses of EM agents that will monitor the PDU.
    - Empty slot contains 0.0.0.0 value in the IP address field and empty string in the Community string field.
8. Enter 0.0.0.0 value in the IP address field and empty Community string field if there are not enough empty slots.
9. Click **Submit**.
10. Logout from the PDU interface.

Verify PDU v2 NMS Table, SNMPv3 Access Table, and Trap Hosts Setup Table

To verify and modify the NMS table and Trap Hosts setup table use SNMPv1 credentials. If you’re using SNMPv3 credentials, use the SNMPv3 Access Table and Trap Hosts Setup table on a PDU v2. Perform the following steps in the PDU user interface:

1. Open the PDU Management Interface in the web browser.
2. In the PDU User Interface, click **Net Configuration**.
3. Login with your user name and password.
4. Click **SNMP Access** tab.
5. Make sure there are enough empty slots in NMS (SNMP v1/v2) table for IP addresses of EM agents that will monitor the PDU.
    - Empty slot is a slot where the Enable check box is deselected.
6. Deselect the Enable check boxes if there are not enough empty slots.
7. Click **Submit**.
8. Make sure there is an empty slot in SNMP v3 table for user that will be used to monitor PDU.
    - Empty slot is a slot where the Enable check box is deselected.
9. Deselect the Enable check boxes if there are not enough empty slots.

10. Click **Submit**.
11. Click **SNMP Traps** tab.
12. Make sure there are enough empty slots in Trap Remote Host Setup table for IP addresses of EM agents that will monitor the PDU.
    - Empty slot is a slot where the Enable check box is deselected.
13. Deselect Enable check boxes if there are not enough empty slots.
14. Click **Submit**.
15. Logout from the PDU interface.

PDU Discovery in the Enterprise Manager

To discover the PDU, perform the following steps:

1. Log in to Enterprise Manager.
2. Under Setup, click **Add Target**, then click **Add Targets Manually**.
3. In the Overview section, click **Add Targets using Guided Process**.
4. In the Add Using Guided Process screen, scroll down to Systems Infrastructure PDU, then click **Add**.
5. In the Systems Infrastructure PDU Discovery screen, enter the required information.
    a. In the Target Name field, enter a name for the target.
    b. In the PDU DNS Name or IP Address field, enter the IP address or DNS name of the PDU.
    c. In the Monitoring Agent EMD URL field, click the search icon and select a monitoring agent. The agent must be able to communicate with the PDU over the network. (Optional) You can select backup monitoring agent for the PDU. In the Backup Agent EMD URL field, click the search icon and select a monitoring agent.

    > **Note:**
    >
    > Backup agent is used to monitor a target and collect its metrics when the primary agent is not reachable or is in maintenance state.

    d. Enter credentials for the PDU Management Interface in the HTTP Monitoring Credentials section.
    - In the Credential Type field, choose **SNMPv1 or SNMPv3 Creds** .

      > **Note:**
      >
      > Currently PDU monitoring supports SNMPv1 credentials only for SNMPv3 only no privacy and authentication options. Privacy password is not supported. Please do not enter a privacy password..

    - Enter the SNMP Community String or SNMPv3 user and authentication password and method to be used to communicate with the PDU using SNMP protocol.

The agent IP address and community string is automatically added to the NMS table and the Trap Hosts Setup table for SNMPv1 or user is added to SNMPv3 Access table of the PDU when the PDU is discovered.

e. (Optional) The Properties section is populated by default. Enter the SNMP Port and Timeout value if you want to change the port and timeout setting.

f. (Optional) If you discover PDU v2 you can specify in properties section in property SNMP MIB version whether PDU should be monitored using original or enhanced SNMP MIB variant. Please enter word “Original”, “Enhanced” or leave property empty.

   If you select enhanced or original MIB here, this MIB version will be enforced and configured in PDU web management interface.

   If you leave property empty, no MIB version will be enforced in PDU and Enterprise Manager will select appropriate monitoring method automatically. Enhanced MIB has advantage over the Original one that Enterprise Manager monitors PDU completely using SNMP which performance is higher than when PDU is monitored using SNMP only partially.

   Be careful when selecting Enhanced MIB. Some obsolete PDU monitoring tools which depends on Original MIB may loose connection to PDU.

6. It is highly recommended to test all entered values and PDU reachability and configuration correctness using the Test Connection button in the right corner of the screen. Once the test connection is successfully completed, you can continue to add the target.

7. Click **Add** in the top right corner of the screen. Once the job is successfully run, the PDU is discovered and the Add Targets Manually screen is displayed. It might take a few minutes for the data to load before you open the PDU target landing page to view.

Discovering a PDU Using Command Line Interface

You can discover a power distribution unit using the emcli command line tool.

To discover PDU using the emcli, perform the following steps:

1. Open the command line on the host where OMS is running.
2. Login to emcli using command *emcli login –username=<Your user name>*.
3. Type the password when prompted.
4. Execute command *emcli sync*.
5. Discover a new PDU using the following command if you want to use SNMPv1 credentials:

```
emcli add_target \

-name='Name of your PDU' \

-type=oracle_si_pdu \

-host=’Host on which the deployed agent is used to monitor the PDU’ \

-subseparator=properties='=' \

-separator=properties=';' \

-properties='dispatch.url=http://PDU IP or DNS name' \

-monitoring_cred='http;oracle_si_pdu;http;username:PDU admin user
username;password:PDU admin user password' \

-monitoring_cred='snmp_v1v2_v3;oracle_si_pdu;SNMPV1Creds;COMMUNITY:SNMP
community string'
```

6. Discover a new PDU using the following command if you want to use SNMPv3 credentials:

   ```
   emcli add_target \

   -name='Name of your PDU' \

   -type=oracle_si_pdu \

   -host='Host on which the deployed agent is used to monitor the PDU' \

   -subseparator=properties='=' \

   -separator=properties=';' \

   -properties='dispatch.url=http://PDU IP or DNS name' \

   -monitoring_cred='http;oracle_si_pdu;http;username:PDU admin user
   username;password:PDU admin user password' \

   -monitoring_cred='snmp_v1v2_v3;oracle_si_pdu;SNMPV3Creds;authUser:SNMPv3 user
   name;authProtocol:SHA;authPwd:SNMPv3 user password'
   ```

Set the following in the emcli add_target command:

- Replace *Name of your PDU* with the name of your PDU.
- To host, set the Monitoring Agent EMD URL without port, that is, the host name of host where agent is deployed, the agent must be able to communicate with the PDU over the network.
- Replace *PDU IP or DNS name* with the IP address or DNS name of the PDU.
- Replace *PDU admin user username* with the user name of the admin user that can manage the PDU using the PDU Web Management Interface.
- Replace *PDU admin user password* with the password of the admin user that can manage the PDU using the PDU Web Management Interface.
- If you're using SNMPv1, replace *SNMP community string* with the SNMP Community String to be used to communicate with the PDU using SNMP protocol.
- In you're using SNMPv3, replace SNMPv3 user name and SNMPv3 user password with the user name and password of user which will be used to communicate with the PDU using SNMPv3 protocol. (Optional) You can replace SHA authentication protocol with MD5.

The following is a sample command to create a generic PDU:

```
emcli add_target \

-name=pdu.example.com \

-type=oracle_si_pdu \

-host=host.example.com \

-subseparator=properties='=' \

-separator=properties=';' \

-properties='dispatch.url=http://pdu.example.com'\

-monitoring_cred='http;oracle_si_pdu;http;username:admin;password:password123' \

-monitoring_cred='snmp_v1v2_v3;oracle_si_pdu;SNMPV1Creds;COMMUNITY:public'
```

Discovering and Promoting Oracle ZFS Storage

This section describes the manual storage discovery procedure for the Oracle ZFS Storage Appliance target and Oracle ZFS Storage Appliance Cluster.

The ZFS Storage Server target is a multi access point (MAP) target that supports two access points, AKCLI and REST/WebSvc access points. If both the REST/WebSvc and AKCLI access points are available, then all the metrics are uploaded using the ReST/WebSvc access point. The following command is used to add the REST/WebSvc access point:

```
emcli add_target -name=<name> -type=oracle_si_zfssa_storage_server -host=<emagent name>
-access_point_name=<any name> -access_point_type='oracle_si_zfssa_storage_server_websvc'
-properties='dispatch.url=https://<applianceURL>:215/api/' -subseparator=properties='='
-
monitoring_cred='ZfssaHttpCredentialSet;oracle_si_zfssa_storage_server_websvc;AliasCreden
tial;Alias:root;Password:<password>'
```

Oracle recommends that you use ReST/Websvc mode of discovery for monitoring ZFS targets. This mode reduces the load on the storage server and provides better response from the server.

This flowchart illustrates the procedure for discovering an Oracle ZFS Storage Appliance.

Figure 4-3 Discovering an Oracle ZFS Storage Appliance

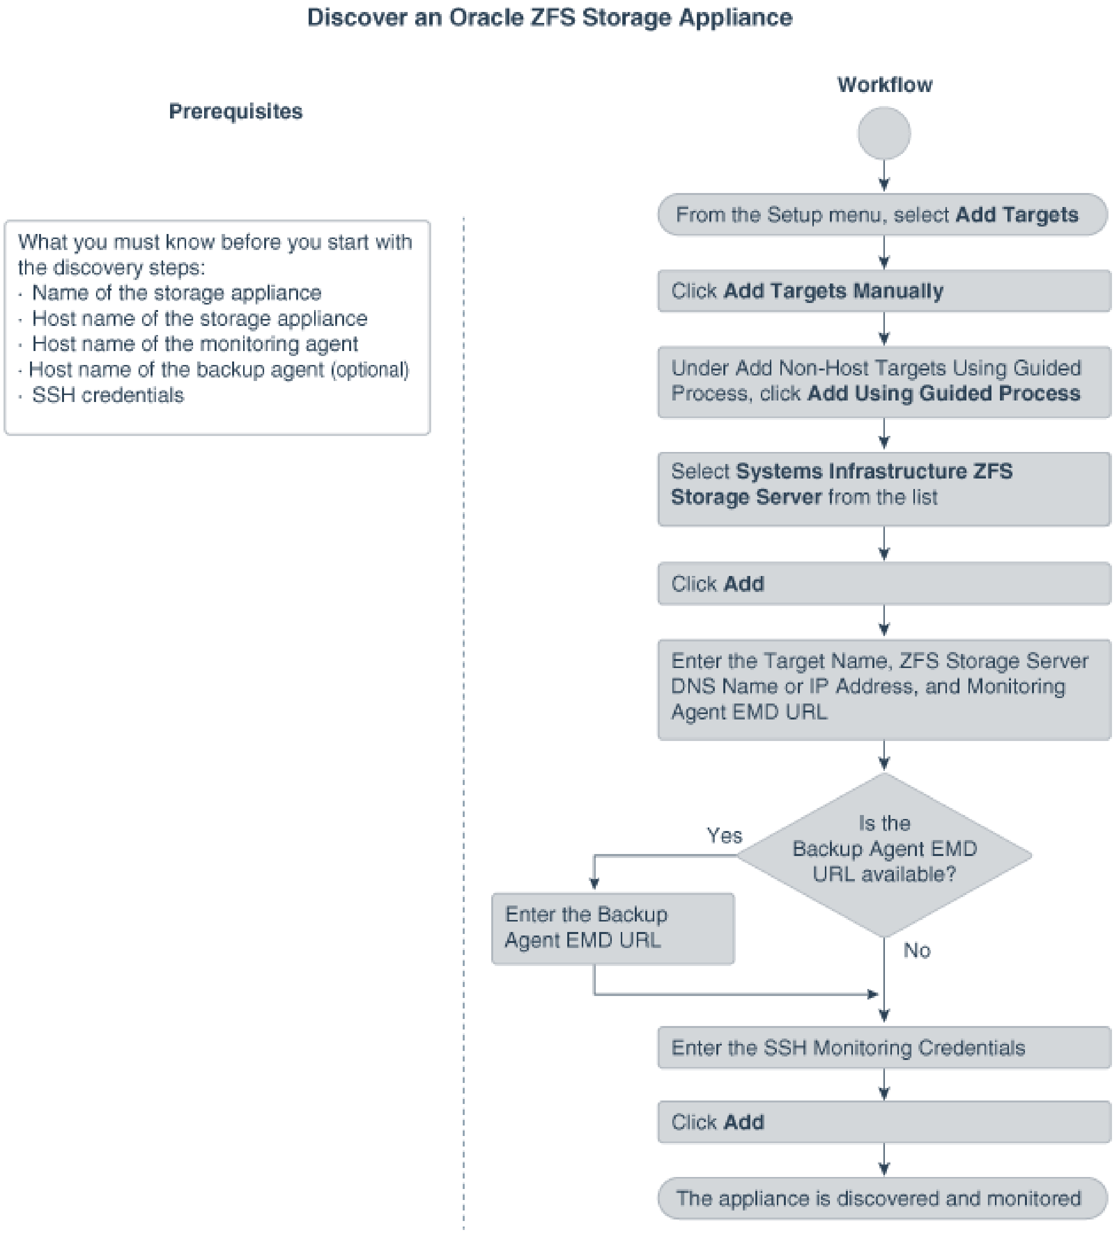
Discover an Oracle ZFS Storage Appliance
Prerequisites
What you must know before you start with the discovery steps:
· Name of the storage appliance
· Host name of the storage appliance
· Host name of the monitoring agent
· Host name of the backup agent (optional)
· SSH credentials
Workflow
From the Setup menu, select Add Targets
Click Add Targets Manually
Under Add Non-Host Targets Using Guided Process, click Add Using Guided Process
Select Systems Infrastructure ZFS Storage Server from the list
Click Add
Enter the Target Name, ZFS Storage Server DNS Name or IP Address, and Monitoring Agent EMD URL
Is the Backup Agent EMD URL available?
Yes
Enter the Backup Agent EMD URL
No
Enter the SSH Monitoring Credentials
Click Add
The appliance is discovered and monitored

Discovering an Oracle ZFS Storage Appliance using AKCLI

To discover an Oracle ZFS Storage Appliance using AKCLI, perform the following steps:

1. From the **Setup** menu, select **Add Target**.
2. Click **Add Targets Manually**.
3. Click **Add Using Guided Process** listed under Add Non-Host Targets Using Guided Process.

   The Add Using Guided Process window is displayed with the list of Guided Discovery and Discovered Target Types.

4. Select **Systems Infrastructure ZFS Storage Server** from the list in the Add Using Guided Process window.
5. Click **Add**.
6. Select **SSH Credentials** option.
7. Click **Discover ZFS Target**.
8. For Target, enter the following details:
    1. Enter the Target Name.
    2. Enter the ZFS Storage Server DNS Name or IP Address.
9. For Monitoring Agents, enter the following details:
    1. Enter the Monitoring Agent EMD URL.
    2. (Optional): Enter the Backup Agent EMD URL.
10. For Monitoring Credentials, enter the following details:
    1. Select the Credential type as **SSH Credentials**.
    2. Enter the User name and Password in the respective fields.
    3. Retype the password in the Confirm Password field.
    4. (Optional): You can enter the Role Name, Role Password and retype the password in the Confirm Role Password field.
11. Click **Add**.

Target Members of an Oracle ZFS Storage Appliance

The following target members are automatically promoted when you add an Oracle Systems Infrastructure ZFS Storage Server target through the discovery wizard:

- ZFS Storage Server

  Target Type: Oracle ZFS Storage Server
- Diskshelf

  Target Type: ZFS Diskshelf Storage

Target Members of an Oracle ZFS Storage Appliance Cluster

When you add two ZFS Storage Appliance nodes, which are setup as cluster nodes, then the ZFS Storage Appliance Cluster is auto-discovered.

The following targets are added when two ZFS Storage Servers have cluster configuration setup:

- ZFS Storage Server

  Target Type: Oracle ZFS Storage Server
- Diskshelf

  Target Type: ZFS Diskshelf Storage
- ZFS Storage Appliance Cluster

  Target Type: Oracle ZFS Storage Server Cluster

Discovering an Oracle ZFS Storage Appliance using WebSvc

To discover an Oracle ZFS Storage Appliance using WebSvc, perform the following steps:

1. From the **Setup** menu, select **Add Target**.
2. Click **Add Targets Manually**.
3. Click **Add Using Guided Process** listed under Add Non-Host Targets Using Guided Process.

   The Add Using Guided Process window is displayed with the list of Guided Discovery and Discovered Target Types.
4. Select **Systems Infrastructure ZFS Storage Server** from the list in the Add Using Guided Process window.
5. Click **Add**.
6. Select **HTTP Credentials** option.
7. Click **Discover ZFS Target**.
8. For Target, enter the following details:
   a. Enter the Target Name.
   b. Enter the ZFS Storage Server DNS Name or IP Address.
9. For Monitoring Agents, enter the following details:
   a. Enter the Monitoring Agent EMD URL.
   b. (Optional): Enter the Backup Agent EMD URL.
10. For Monitoring Credentials, enter the following details:
    a. Enter the Alias and Password in the respective fields.
    b. Retype the password in the Confirm Password field.
11. Click **Add**.

Discovering Fabrics

Enterprise Manager discovers and manages Ethernet fabrics and InfiniBand fabrics. A network switch contributes its ports, datalinks, and networks to a fabric.

- Discover an InfiniBand Network Switch
- Discover an Ethernet Network Switch
- Use the Command Line To Discover a Switch

After discovery, you can manage the components of these fabrics and view their attributes and metrics, according to the procedures in Managing Networks .

Discover an InfiniBand Network Switch

1. In the Setup menu, select **Add Target**.
2. Click **Add Targets Manually**.
3. In the **Add Non-Host Target Using Guided Process** section, click **Add Using Guided Process**.

4. Scroll to **Systems Infrastructure Oracle InfiniBand Switch** and then click the **Add...** button.

Figure 4-4   InfiniBand Switch Target Type

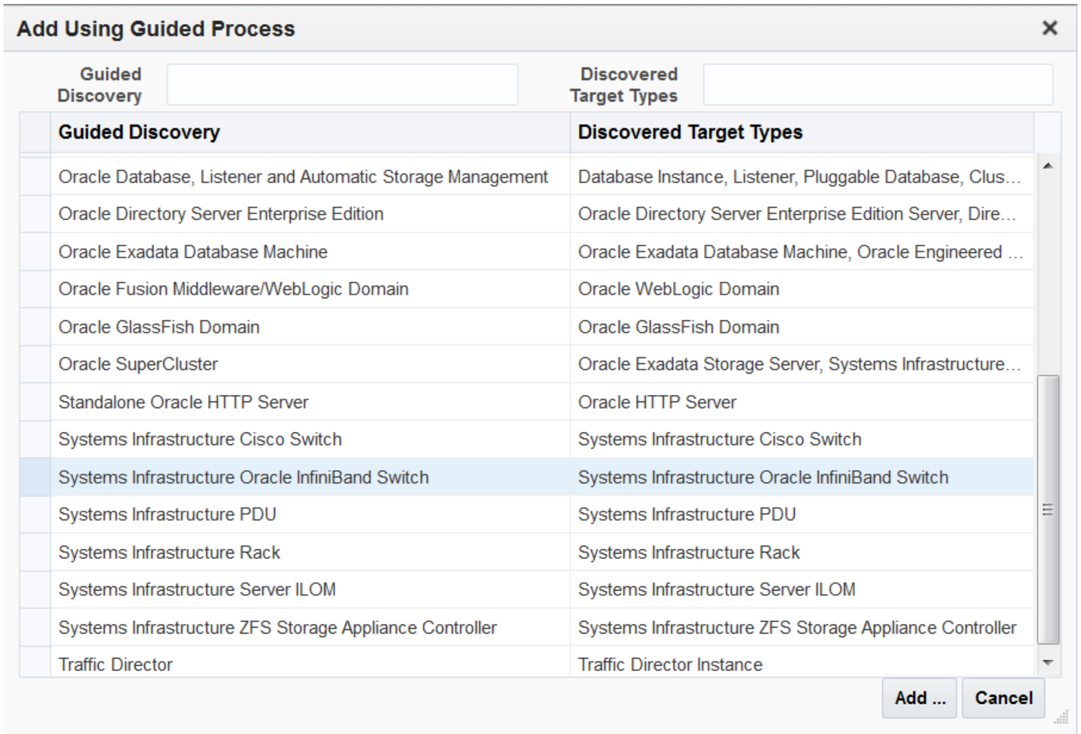

Guided Discovery	Discovered Target Types
Oracle Database, Listener and Automatic Storage Management	Database Instance, Listener, Pluggable Database, Clus…
Oracle Directory Server Enterprise Edition	Oracle Directory Server Enterprise Edition Server, Dire…
Oracle Exadata Database Machine	Oracle Exadata Database Machine, Oracle Engineered …
Oracle Fusion Middleware/WebLogic Domain	Oracle WebLogic Domain
Oracle GlassFish Domain	Oracle GlassFish Domain
Oracle SuperCluster	Oracle Exadata Storage Server, Systems Infrastructure…
Standalone Oracle HTTP Server	Oracle HTTP Server
Systems Infrastructure Cisco Switch	Systems Infrastructure Cisco Switch
Systems Infrastructure Oracle InfiniBand Switch	Systems Infrastructure Oracle InfiniBand Switch
Systems Infrastructure PDU	Systems Infrastructure PDU
Systems Infrastructure Rack	Systems Infrastructure Rack
Systems Infrastructure Server ILOM	Systems Infrastructure Server ILOM
Systems Infrastructure ZFS Storage Appliance Controller	Systems Infrastructure ZFS Storage Appliance Controller
Traffic Director	Traffic Director Instance

5. In the Target section, enter a name for the new target and its DNS name or IP address.

Figure 4-5 Discovery Specifications for InfiniBand Switches

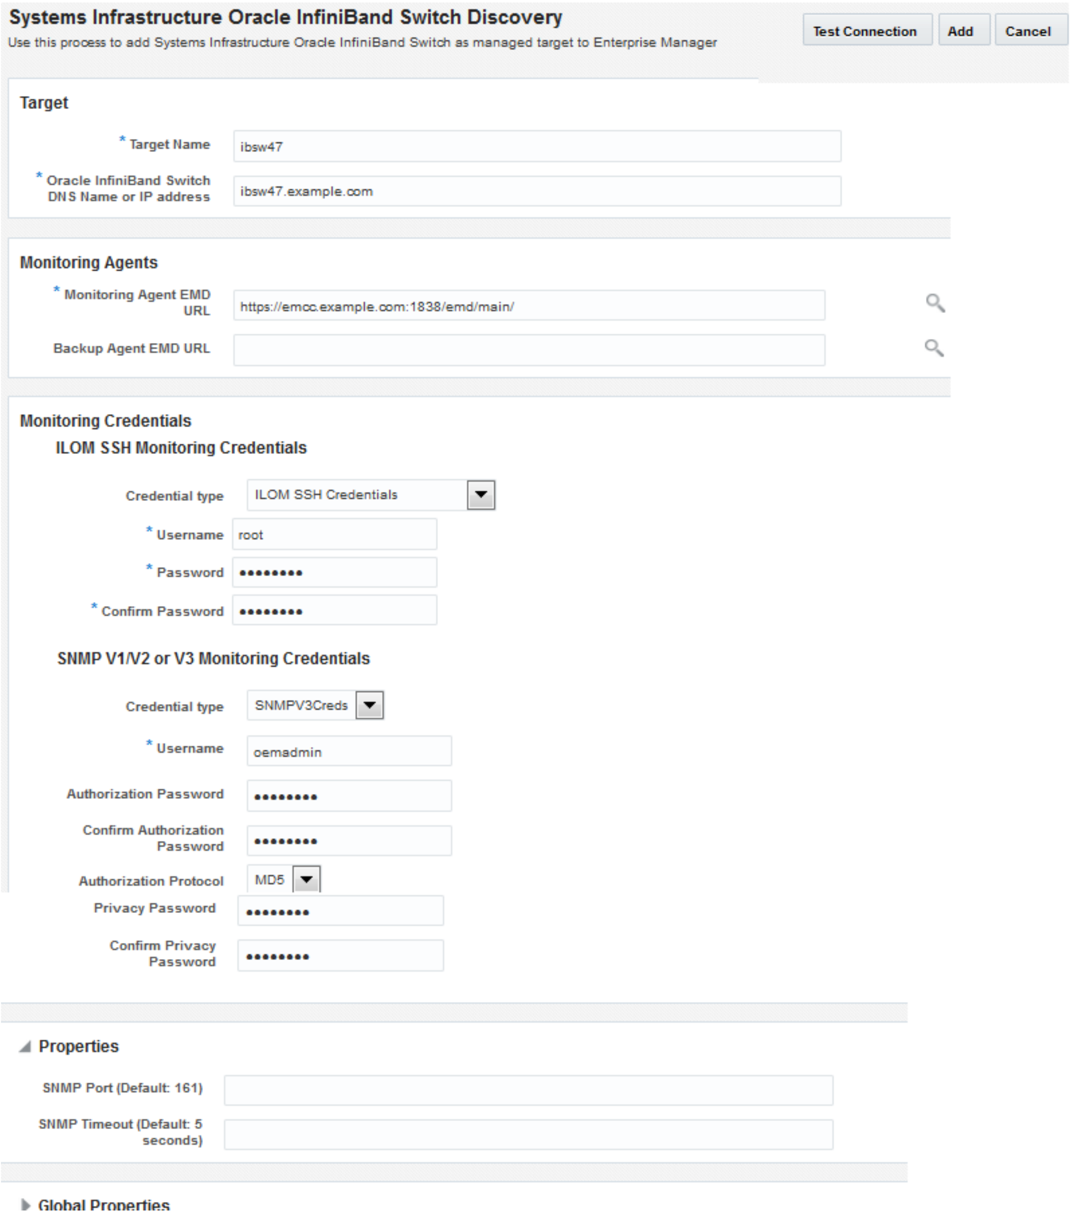

Systems Infrastructure Oracle InfiniBand Switch Discovery

Use this process to add Systems Infrastructure Oracle InfiniBand Switch as managed target to Enterprise Manager

Target

* Target Name	ibsw47
* Oracle InfiniBand Switch DNS Name or IP address	ibsw47.example.com

Monitoring Agents

* Monitoring Agent EMD URL	https://emcc.example.com:1838/emd/main/
Backup Agent EMD URL	

Monitoring Credentials

ILOM SSH Monitoring Credentials

Credential type	ILOM SSH Credentials
* Username	root
* Password	••••••••
* Confirm Password	••••••••

SNMP V1/V2 or V3 Monitoring Credentials

Credential type	SNMPV3Creds
* Username	oemadmin
Authorization Password	••••••••
Confirm Authorization Password	••••••••
Authorization Protocol	MD5
Privacy Password	••••••••
Confirm Privacy Password	••••••••

Properties

SNMP Port (Default: 161)	
SNMP Timeout (Default: 5 seconds)	

6. In the Monitoring Agents section, enter the URL of the monitoring system and the URL of a backup system.

7. In the ILOM SSH Monitoring Credentials section, select the type of credentials and enter the username and the password. The credentials enable Enterprise Manager to monitor the switch's service processor.

> **Note:**
>
> Use the `root` username and password to discover an InfiniBand network switch.

8. In the SNMP Monitoring Credentials section, specify the SNMP version number that Enterprise Manager uses to monitor the hardware components. Version 2c is the default version and requires only a community string. Version 3 requires a username, a password

of exactly eight characters, and the encryption type. Version 1 is supported but not recommended, as described in About Performance of Fabrics.

9. In the Properties section, you can accept the default port number and timeout interval or change them.
10. In the Global Properties section, you can specify site-specific values, such as the Oracle Customer Support Identifier (CSI).
11. Click **Add**, located at the top of the window.

Discover an Ethernet Network Switch

1. In the Setup menu, select **Add Target**.
2. Click **Add Targets Manually**.
3. In the **Add Non-Host Target Using Guided Process** section, click **Add Using Guided Process**.
4. Scroll to **Systems Infrastructure Cisco Switch** or **Systems Infrastructure Juniper Switch** and then click the **Add...** button.

Figure 4-6 Juniper Ethernet Switch Target Type

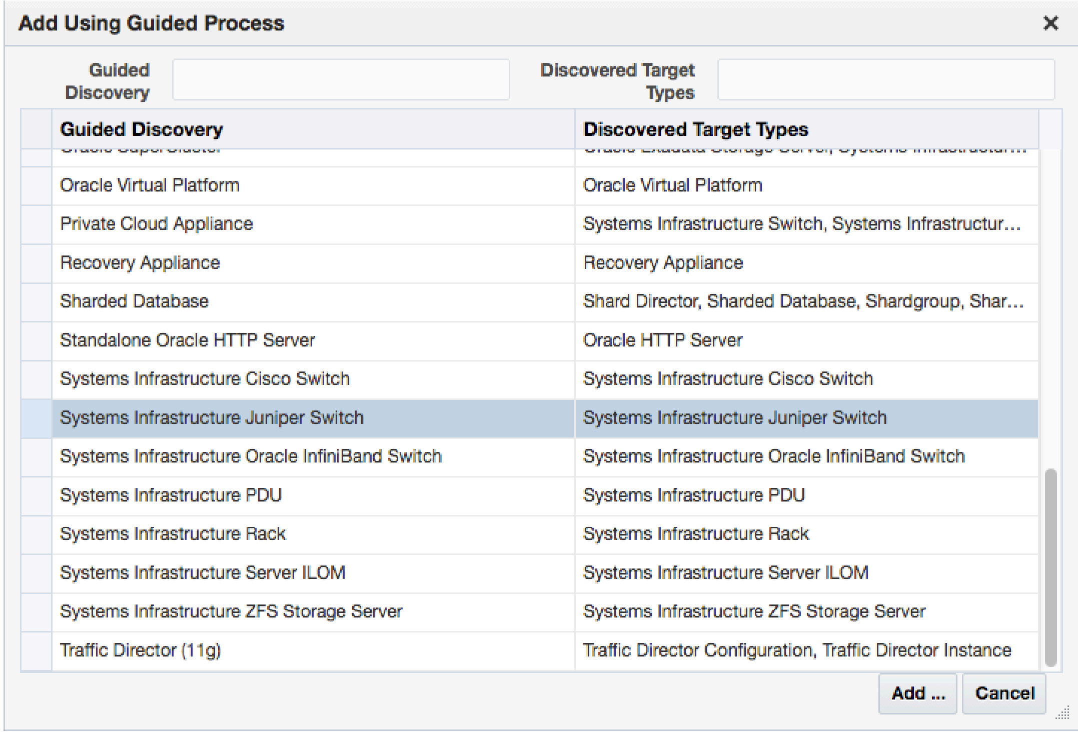

Guided Discovery	Discovered Target Types
Oracle SuperCluster	Oracle Exadata Storage Server, Systems Infrastructur…
Oracle Virtual Platform	Oracle Virtual Platform
Private Cloud Appliance	Systems Infrastructure Switch, Systems Infrastructur…
Recovery Appliance	Recovery Appliance
Sharded Database	Shard Director, Sharded Database, Shardgroup, Shar…
Standalone Oracle HTTP Server	Oracle HTTP Server
Systems Infrastructure Cisco Switch	Systems Infrastructure Cisco Switch
Systems Infrastructure Juniper Switch	Systems Infrastructure Juniper Switch
Systems Infrastructure Oracle InfiniBand Switch	Systems Infrastructure Oracle InfiniBand Switch
Systems Infrastructure PDU	Systems Infrastructure PDU
Systems Infrastructure Rack	Systems Infrastructure Rack
Systems Infrastructure Server ILOM	Systems Infrastructure Server ILOM
Systems Infrastructure ZFS Storage Server	Systems Infrastructure ZFS Storage Server
Traffic Director (11g)	Traffic Director Configuration, Traffic Director Instance

5. In the Target section, enter a name for the new target and provide its DNS name or IP address.

Figure 4-7 Discovery Specifications for Juniper Ethernet Switches

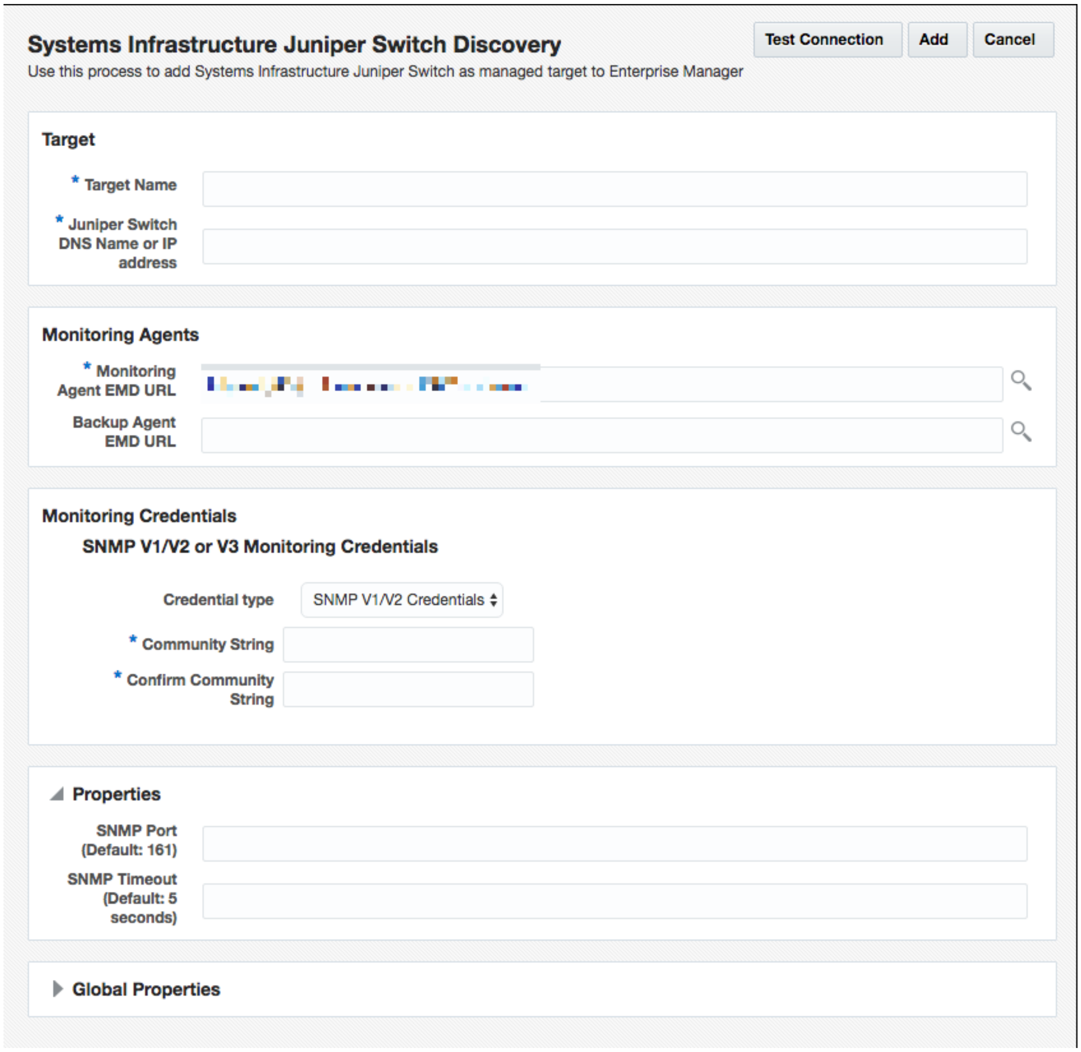

Systems Infrastructure Juniper Switch Discovery

Use this process to add Systems Infrastructure Juniper Switch as managed target to Enterprise Manager

Target

* Target Name

* Juniper Switch DNS Name or IP address

Monitoring Agents

* Monitoring Agent EMD URL [redacted]

Backup Agent EMD URL

Monitoring Credentials

SNMP V1/V2 or V3 Monitoring Credentials

Credential type: SNMP V1/V2 Credentials

* Community String

* Confirm Community String

Properties

SNMP Port (Default: 161)

SNMP Timeout (Default: 5 seconds)

Global Properties

6. In the Monitoring Agents section, enter the URL of the monitoring system and the URL of a backup system.
7. In the Cisco Switch IOS Monitoring Credential section, specify the type of credentials and enter the username and password. Also, enter the Cisco EXEC password. The credentials enable Enterprise Manager to monitor the switch's service processor.
8. In the SNMP Monitoring Credential section, specify the SNMP version number that Enterprise Manager uses to monitor the hardware components. Version 3 requires a username, a password of exactly eight characters, and the encryption type. Version 1 is supported but not recommended, as described in About Performance of Fabrics.
9. In the Properties section, you can change the default configuration details such as port number and timeout interval.
10. In the Global Properties section, you can specify site-specific values, such as the Oracle Customer Support Identifier (CSI).
11. Click **Add**, located at the top of the window.

Example 4-6   Example. Command to Discover a Juniper Ethernet Switch Using SNMP Version 3

The following command discovers and manages a Juniper Ethernet switch named ***MY_SWITCH*** and sets the SNMP version to Version 3.

```
emcli add_target -name=MY_SWITCH -type=oracle_si_netswitch -host=AGENT_HOST -
access_point_name=MY_SWITCH -access_point_type=oracle_si_switch_juniper_junos
-properties='dispatch.url=snmp://MY_SWITCH_IP_ADDR' -
subseparator=properties='='  -
monitoring_cred='snmp_v1v2_v3;oracle_si_switch_cisco_ios;SNMPV3Creds;authUser:
PRINCIPAL;authPwd:AUTHCRED;authProtocol:MD5;privPwd:PRIV_CREDS'
```

Use the Command Line To Discover a Switch

These examples use the Command Line Interface to discover and manage a network switch. You must configure the command line interface before you can issue commands. For more information, click the **Setup** menu, then click **Command Line Interface**. Follow the Download and Deploy instructions.

Example 4-7   Command to Discover an InfiniBand Switch Using SNMP Version 3

The following command discovers and manages an InfiniBand switch named `MY_SWITCH` and sets the SNMP version to Version 3.

```
emcli add_target -name=MY_SWITCH -type=oracle_si_netswitch -host=AGENT_HOST -
access_point_name=MY_SWITCH -access_point_type=oracle_si_switch_oracle_ib -
properties='dispatch.url=ilom-ssh://MY_SWITCH_IP_ADDR' -subseparator=properties='=' -
monitoring_cred='ilom_creds_set;oracle_si_switch_oracle_ib;ilom_creds;username:PRIV_USER;
password:PASSWORD' -
monitoring_cred='snmp_v1v2_v3;oracle_si_switch_oracle_ib;SNMPV3Creds;authUser:PRINCIPAL;a
uthPwd:AUTHCRED;authProtocol:MD5;privPwd:PRIV_CREDS'
```

Example 4-8   Command to Discover a Cisco Ethernet Switch Using SNMP Version 3

The following command discovers and manages a Cisco Ethernet switch named `MY_SWITCH` and sets the SNMP version to Version 3.

```
emcli add_target -name=MY_SWITCH -type=oracle_si_netswitch -host=AGENT_HOST -
access_point_name=MY_SWITCH -access_point_type=oracle_si_switch_cisco_ios -
properties='dispatch.url=ios-ssh://MY_SWITCH_IP_ADDR' -subseparator=properties='=' -
monitoring_cred='cisco_creds_set;oracle_si_switch_cisco_ios;cisco_creds;username:PRIV_USE
R;userpass:USER_PASSWORD;privpass:PRIV_PASSWORD' -
monitoring_cred='snmp_v1v2_v3;oracle_si_switch_cisco_ios;SNMPV3Creds;authUser:PRINCIPAL;a
uthPwd:AUTHCRED;authProtocol:MD5;privPwd:PRIV_CREDS'
```

Example 4-9   Command to Discover an InfiniBand Switch Using SNMP Version 1

The following command discovers and manages an InfiniBand switch named `MY_SWITCH`, sets the SNMP version to Version 1, and changes the time interval for the SNMP property.

```
 emcli add_target -name=MY_SWITCH -type=oracle_si_netswitch -host=AGENT_HOST -
access_point_name=MY_SWITCH -access_point_type=oracle_si_switch_oracle_ib -
properties='dispatch.url=ilom-ssh://MY_SWITCH_IP_ADDR;SNMPTimeout=180' -
subseparator=properties='='-
monitoring_cred='ilom_creds_set;oracle_si_switch_oracle_ib;ilom_creds;username:PRIV_USER;
password:PASSWORD' -
monitoring_cred='snmp_v1v2_v3;oracle_si_switch_oracle_ib;SNMPV1Creds;COMMUNITY:<COMMUNITY
>'
```

Related Resources for Discovering and Promoting System Infrastructure Targets

See the following for more information:

- Discovering and Adding Host and Non-Host Targets
- Working with Systems Infrastructure Targets
- Managing Storage
- Managing Networks

Part III

Monitoring and Managing Targets

This section contains the following chapters:

5 Using Incident Management

Incident management allows you to monitor and resolve service disruptions quickly and efficiently by allowing you to focus on what is important from a broader management perspective (incidents) rather than isolated, discrete events that may point to the same underlying issue.

In this chapter:	You will learn:
Advanced Topics	How to perform specialized incident management operations. • Defining Custom Incident Statuses • Clearing Stateless Alerts for Metric Alert Event Types • User-reported Events • Additional Rule Applications • Event Prioritization • Root Cause Analysis (RCA) and Target Down Events
Moving from Enterprise Manager 10/11g to 12c and Greater	Migrating notification rules to incident rules.

Management Concepts

Enterprise Manager exposes three levels of management granularity that, when combined, provide complete monitoring/management coverage of your environment. These management levels are:

- Event Management
- Incident Management
- Problem Management
- Rule Sets
- Incident Manager

Event Management

Intuitively, you monitor for specific events in your monitored environment. An event is a significant occurrence on a managed target that typically indicates something has occurred outside normal operating conditions--they provide a uniform way to indicate that something of interest has occurred in an environment managed by Enterprise Manager. Examples of events are:

- Metric Alerts
- Compliance Violations
- Job Events
- Availability Alerts

Existing Enterprise Manager customers may be familiar with metric alerts and metric collection errors. For Enterprise Manager 12*c*, metric alerts are a type of event, one of many different event types. The notion of an event unifies the different exception conditions that are detected by Enterprise Manager, such as monitoring issues or compliance issues, into a common concept. It is backed by a consistent and uniform set of event management capabilities that can indicate something of interest has occurred in a datacenter managed by Enterprise Manager.

All events have the following attributes:

Table 5-1 Event Attributes

Attribute	Description
Type	Type of event that is being reported. All events of a specific type share the same set of attributes that describe the exact nature of the problem. For example, Metric Alert, Compliance Standard Score Violation, or Job Status Change.
Severity	Event severity. For example, Fatal, Warning, or Critical.
Internal Name	An internal name that describes the nature of the event and can be used to search for events. For example, you can search for all *tablespacePctUsed* events.
Entity on which the event is raised.	An event can be raised on a target, a non-target source object (such as a job) or be related to a target and a non-target source object. Note: This attribute is important when determining what privileges are required to manage the event.
Message	Informational text associated with the event.
Reported Date	Time the event was reported.
Category	Functional or operational classification for an event. Available Categories: • Availability • Business • Capacity • Configuration • Diagnostics • Error • Fault • Jobs • Load • Performance • Security
Causal Analysis Update	Used for Root Cause Analysis of target down events. Possible Values: Root Cause or Symptom

Event Types

The *type* of an event defines the structure and payload of an event and provides the details of the condition it is describing. For example, a metric alert raised by threshold violation has a specific payload whereas a job state change has a different structure. As shown in the following table, the range of events types greatly expands Enterprise Manager's monitoring flexibility.

Event Type	Description
Target Availability	The Target Availability Event represents a target's availability status (Example: Up, Down, Agent Unreachable, or Blackout).
Metric Alert	A metric alert event is generated when an alert occurs for a metric on a specific target (Example: CPU utilization for a host target) or metric on a target and object combination Example: Space usage on a specific tablespace of a database target.

Event Type	Description
Metric Evaluation Error	A metric evaluation error is generated when the collection for a specific metric group fails for a target.
Job Status Change	All changes to the status of an Enterprise Manager job are treated as events, and these events are made available via the Job Status Change event class. **Note**: A prerequisite to creating Incident Rules, is to enable the relevant job status and add required targets to job event generation criteria. To change this criteria, from the **Setup** menu, select **Incidents**, and then **Job Events**.
Compliance Standard Rule Violation	Events are generated for compliance standard rule violations. Each event corresponds to a violation of a compliance rule on a specific target.
Compliance Standard Score Violation	Events are generated for compliance standard score violations. An event is generated when the compliance score for a compliance standard on a specific target falls below predefined thresholds.
High Availability	High Availability events are generated for database availability operations (shutdown and startup), database backups and Data Guard operations (switchover, failover, and other state changes).
Service Level Agreement Alert	These events are generated when a service level or service level objective is violated for a service. occurs for a Service Level Agreement or a Service Level Objective.
User-reported	These events are created by end-users.
Application Performance Management KPI Alert	An Application Performance Management (APM) Key Performance Indicator (KPI) alert event is generated when a KPI violation alert occurs for a metric on an APM managed entity associated with a Business Application target.
JVM Diagnostics Threshold Violation	A JVMD Diagnostics event is raised when a JVMD metric exceeds its threshold value on a Java Virtual Machine target.
Application Dependency and Performance Alert	Alerts are raised by ADP monitoring when metrics related to a J2EE application or component have crossed some thresholds.
Blackout Infrastructure Alert	The Blackout Event represents blackout infrastructure events such as execution failure of a blackout operation as the agent is not reachable. The blackout operations include blackout start, blackout stop, and blackout edit.
Service Infrastructure Alert	These alerts are generated when there is a problem in the service infrastructure.

Event Type	Description
Target Monitoring Disruption	This event is generated when normal monitoring of the target cannot proceed. This is due to issues such as too many hung threads or higher than expected resource consumption during metric collection. These issues might result in limited target monitoring such as monitoring of target Status only, delayed collection evaluation, or no monitoring of the target.

Event Severity

The severity of an event indicates the criticality of a specific issue. The following table shows the various event severity levels along with the associated icon.

Icon	Severity	Description
⛔	Fatal	Corresponding service is no longer available. For example, a monitored target is down (target down event). A Fatal severity is the highest level severity and only applies to the Target Availability event type.
⊗	Critical	Immediate action is required in a particular area. The area is either not functional or indicative of imminent problems.
⚠	Warning	Attention is required in a particular area, but the area is still functional.
⚠	Advisory	While the particular area does not require immediate attention, caution is recommended regarding the area's current state. This severity can be used, for example, to report Oracle best practice violations.
✓	Clear	Conditions that raised the event have been resolved.
ℹ	Informational	A specific condition has just occurred but does not require any remedial action. Events with an informational severity: • do not appear in the incident management UI. • cannot create incidents. • are not stored within Enterprise Manager.

Incident Management

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). Of all events raised within your managed environment, there is likely only a subset that you need to act on because they impact your business applications (such as a target down event). However, managing by incident also allows you to address more complex situations where the subset of events you are interested in are related and may indicate a higher level issue needs to be addressed as a single issue and not as individual events: A cluster of events by themselves may indicate a minor administrative issue, but when viewed together may signify a larger problem that can potentially consist of events from multiple domains/layers of your monitored infrastructure.

For example, you are monitoring a host. If you want to monitor 'load' being placed on one or more hosts you might be interested in events such as CPU utilization, memory utilization, and swap utilization exceeding acceptable metric thresholds. Individually, these events may or may not indicate an issue with the host, but together, these events form an incident indicating extreme load is being placed on a monitored host.

Incidents represent the larger service disruptions that may impact your business instead of discrete events. Managing by incidents, therefore, allows you to monitor for complex operational issues that may affect multiple domains that may impact your business. These incidents typically need to be tracked, assigned to appropriate personnel, and resolved as quickly as possible. You can effectively implement a centralized monitoring that consolidates monitoring information and more effectively allocate resource across your ecosystem to resolve or prevent issues from occurring. The end result is better implementation of your business processes that in turn lead to better performance of your IT resources.

While events indicate issues requiring attention in your managed environment, it is more efficient to work on a collective subset of related events as a single unit of work-- you can work on different events representing the same issue or you can work on one incident containing multiple space-related events. For example, you have multiple space events from various targets that indicate you are running low on space. Instead of managing numerous discrete events, you can more efficiently manage a smaller set of incidents.

An incident is a significant event or set of related significant events that need to be managed because it can potentially impact your business applications. These incidents typically need to be tracked, assigned to appropriate personnel, and resolved as quickly as possible. You perform these incident management operations through Incident Manager, an intuitive UI within Enterprise Manager.

Incident Manger provides you with a central 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 for more information about this UI.

Working with Incidents

When an incident is created, Enterprise Manager makes available a rich set of incident management workflow features that let you to manage and track the incident through its complete lifecycle.

- Assign incident ownership.
- Track the incident resolution status.
- Set incident priority.
- Set incident escalation level.

- Ability to provide a manual summary.
- Ability to add user comments.
- Ability to suppress/unsuppress
- Ability to manually clear the incident.
- Ability to create a ticket manually.

All incident management/tracking operations are carried out from Incident Manager. Creation of incidents for events, assignment of incidents to administrators, setting priority, sending notifications and other actions can be automated using (incident) rules.

Incident Status

The lifecycle of an incident within an organization is typically determined by two pieces of information: The current resolution state of the incident (Incident Status) and how important it is to resolve the incident relative to other incidents (Priority). As key incident attributes, the following options are available:

- New
- Work in Progress
- Closed
- Resolved

You can define additional statuses if the default options are not adequate. In addition, you can change labels using the Enterprise Manager Command Line Interface (EM CLI). See **Advanced Topics** for more information.

Priority

By changing the priority, you can escalate the incident and perform operations such as assigning it to a specific IT operator or notifying upper-management. The following priority options are available:

- None
- Low
- Medium
- High
- Very High
- Urgent

Priority is often based on simple business rules determined by the business impact and the urgency of resolution.

Incident Attributes

Every incident possesses attributes that provide information as identification, status for tracking, and ownership. The following table lists available incident attributes.

Incident Attribute	Definition
Escalated	An escalation level signifying a escalation to raise the level of attention on the incident from your organization's IT or management hierarchy. Available escalation levels: • None (Not escalated) • Level 1 through Level 5

Incident Attribute	Definition
Category	Operational or organizational classification for an incident. Incidents (and events) can have multiple categories. Categories for all events within an incident are aggregated. Available Categories: • Availability • Business • Capacity • Configuration • Diagnostics • Error • Fault • Jobs • Load • Performance • Security
Summary	An intuitive message indicating what the incident is about. By default, the incident summary is pulled from the message of the last event of the incident, however, this message can be changed to a fixed summary by any administrator working on the incident.
Incident Created	Date and time the incident was created.
Last Updated	Date and time the incident was last updated or when the incident was closed.
Severity	Severity is based on the worst severity of the events in the incident. For example, Fatal, Warning, or Critical.
Source	Source entities of the incident.
Priority	Priority Values • None (Default) • Low • Medium • High • Very High • Urgent

Incident Attribute	Definition
Status	Incident Status. • New (Default) • Work in Progress • Closed (Terminal state when the incident is closed. See below for more information.) • Resolved You can define additional statuses if the default options are not adequate. In addition, you can change labels using the Enterprise Manager Command Line Interface (EM CLI). Closed Status: Enterprise Manager automatically sets the status to closed when an incident severity is cleared--administrators do not manually select the Closed status. The incident severity is set to Clear when all of the events contained within the incident have been cleared. Typically the Agent sets the Clear severity, as would be the case when a metric alert value falls below a severity threshold. If an event or incident supports manual clearing, then the Clear option will be shown in the Incident Manager UI. Once an incident has been cleared by an administrator or by Enterprise Manager, only then will Enterprise Manager set the status to Closed.If you do not see the option to clear the incident in the UI, this means Enterprise Manager will automatically set the status to Clear if it detects the monitored condition no longer holds true. For example, you want to indicate that an incident has been fixed. You can set the status to Resolved and Enterprise Manager will set the status to Closed when it clears the severity.
Comment	Annotations added by an administrator to communicate analysis information or actions taken to resolve the incident.
Owner	Administrator/user currently working on the incident.
Acknowledged	Indicates that a user has accepted ownership of an incident or problem. Available options: Yes or No. When an incident is acknowledged, it will be implicitly assigned to the user who acknowledged it. When a user assigns an incident to himself, it is considered 'acknowledged'. Once acknowledged, an incident cannot be unacknowledged, but can be assigned to another user. Acknowledging an incident stops any repeat notifications for that incident.
Causal Analysis Update	Used for Root Cause Analysis of target down incidents. Possible Values: Root Cause or Symptom

Incident Composed of a Single Event

The simplest incident is composed of a single event. In the following example, you are concerned whenever any production target is down. You can create an incident for the target down event which is raised by Enterprise Manager if it detects the monitored target is down. Once the incident is created, you will have all incident management functionality required to track and manage its resolution.

Figure 5-1 Incident with a Single Event

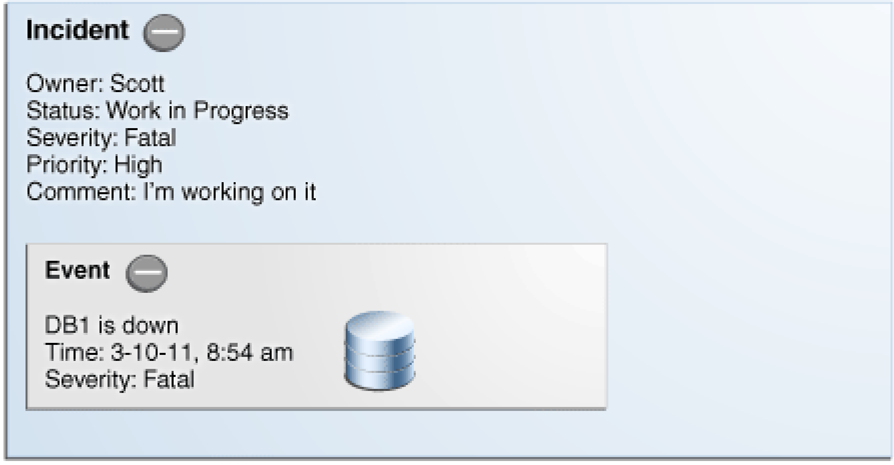

Incident

Owner: Scott
Status: Work in Progress
Severity: Fatal
Priority: High
Comment: I'm working on it

Event

DB1 is down
Time: 3-10-11, 8:54 am
Severity: Fatal

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. When the event is newly generated, there is no ownership or status. An incident is opened that can be updated manually or by automated rules to set owners, status, as well as other attributes. In the example, 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.

Incident Composed of Multiple Events

Situations of interest may involve more than a single event. It is an incident's ability to contain multiple events that allows you to monitor and manage complex and more meaningful issues. These multi-event incidents can be automatically generated through event compression enabled at rule set level (see Rule-based Event Compression) or through global event compression policies enabled (see Event Compression Policies).

For example, if a monitored system is running out of space, separate multiple events such as *tablespace full* and *filesystem full* may be raised. Both, however, are related to running out of space. Another machine resource monitoring example might be the simultaneous raising of CPU utilization, memory utilization, and swap utilization events. Together, these events form an incident indicating extreme load is being placed on a monitored host. The following figure illustrates this example.

Figure 5-2 Incident with Multiple Events

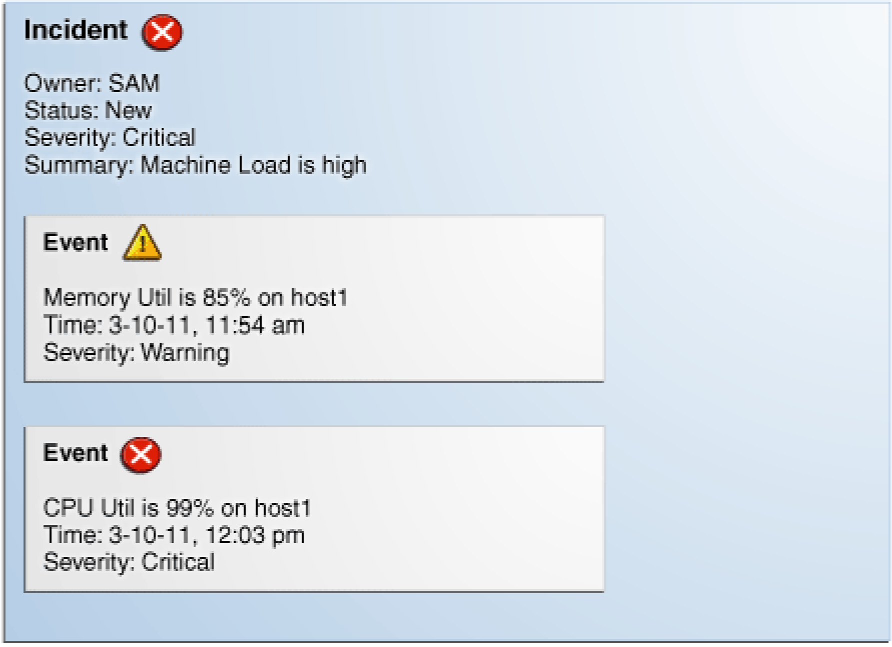

Incident ☒

Owner: SAM
Status: New
Severity: Critical
Summary: Machine Load is high

Event ⚠

Memory Util is 85% on host1
Time: 3-10-11, 11:54 am
Severity: Warning

Event ☒

CPU Util is 99% on host1
Time: 3-10-11, 12:03 pm
Severity: Critical

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 last event of the incident, however, this message can be changed by any administrator working on the incident.

Because administrators are interested in overall machine load, administrator Sam has an incident created for these two metric events because they are related—together these events 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.

How are Incidents Created?

Incidents are most commonly created automatically through rules and rule sets (user-defined instructions that tell Incident Manager how to handle specific events when they occur). As shown in the preceding examples, incidents can also be created manually. Once an incident is raised, its severity is inherited from the worst severity of all events within the incident. The latest event *Message*, by default, becomes the *Incident Summary*. Beginning with Enterprise Manager 13*c*, you can also define customized messages for grouped incidents. Incidents can also be created manually. See "Creating an Incident Manually" for more information.

Problem Management

Problem management involves the functionality that helps track the underlying root causes of incidents. Once the immediate service disruptions represented by incidents are resolved, you can then progress to understanding and resolving the underlying root cause of the issue.

For Enterprise Manager 12*c*, problems focus on the diagnostic incidents and problem diagnostic incidents/problems stored in Advanced Diagnostic Repository (ADR), which are automatically raised by Oracle software when it encounters critical errors in the software. A problem, therefore, represents the root cause of all the Oracle software incidents. For these diagnostic incidents, in order to address root cause, a problem is created that represents the root cause of these diagnostic incidents. A problem is identified by a *problem key* which uniquely identifies the particular error in software. Each occurrence of this error results in a diagnostic incident which is then associated with the problem object.

When a problem is raised for Oracle software, Oracle has determined that the recommended recourse is to open a service request (SR), send support the diagnostic logs, and eventually provide a solution from Oracle. 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 Incident Manager, or in context of a target/group home page, you can easily determine what issues are actually affecting your monitored target.

To manage problems, you can use Support Workbench to package the diagnostic details gathered in ADR and open SR. Users should then manage the problems in Incident Manager. Access to Support Workbench functionality is available through Incident Manager (**Guided Resolution** area) in context of the problem.

Rule Sets

Incident rules and rule sets automate actions related to events, incidents and problems. They can automate the creation of incidents based on important events, perform notification actions such as sending email or opening helpdesk tickets, or perform operations to manage the incident workflow lifecycle such as changing incident ownership, priority, or escalation level.

With previous versions of Enterprise Manager, you used notification rules to choose the individual targets and conditions for which you want to perform actions or receive notifications (send email, page, open a helpdesk ticket) from Enterprise Manager. For Enterprise Manager 13c, the concept and function of notification rules has been replaced with incident rules and rule sets.

- **Rules**: A rule instructs Enterprise Manager to take specific actions when incidents, events, or problems occur, such as performing notifications. Beyond notifications, rules can also instruct Enterprise Manager to perform specific actions, such as creating incidents, updating incidents and problems. The actions can also be conditional in nature. For example, a rule action can be defined to page a user when an incident severity is critical or just send email if it is warning.
- **Rule Set**: An incident rule set is a collection of rules that apply to a common set of objects such as targets (hosts, databases, groups), jobs, metric extensions, or self updates and take appropriate actions to automate the business processes underlying event, incident and problem management.

Operationally, individual rules within a rule set are executed in a specified order as are the rule sets themselves. Rule sets are executed in a specified order. By default, the execution order for both rules and rule sets is the order in which they are created, but they can be reordered from the Incident Rules UI.

The following figure shows typical rule set structure and how the individual rules are applied to a heterogeneous group of targets.

Figure 5-3 Rule Set Application

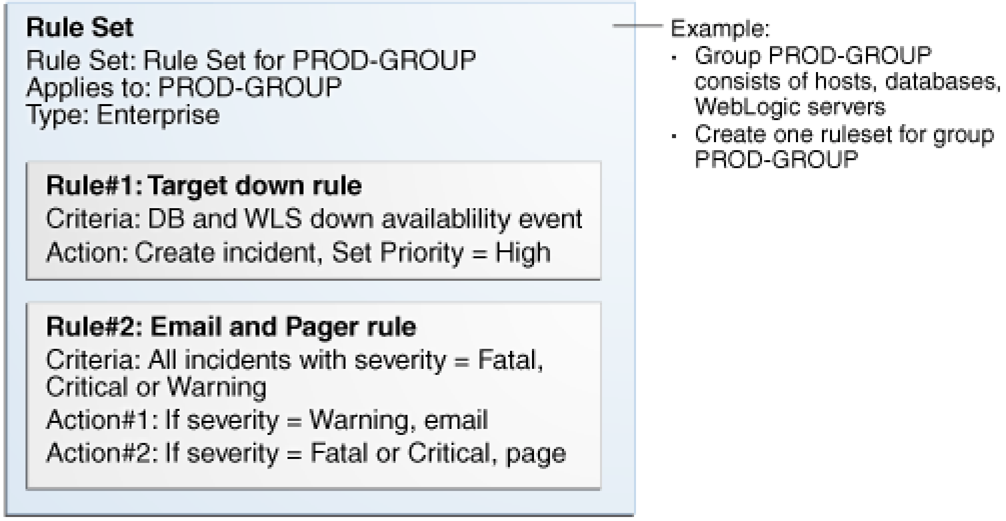
Rule Set
Rule Set: Rule Set for PROD-GROUP
Applies to: PROD-GROUP
Type: Enterprise
Rule#1: Target down rule
Criteria: DB and WLS down availablility event
Action: Create incident, Set Priority = High
Rule#2: Email and Pager rule
Criteria: All incidents with severity = Fatal, Critical or Warning
Action#1: If severity = Warning, email
Action#2: If severity = Fatal or Critical, page
Example:
· Group PROD-GROUP consists of hosts, databases, WebLogic servers
· Create one ruleset for group PROD-GROUP

The graphic illustrates a situation where all rules pertaining to a group of targets can be put into a single rule set (this is also a best practice). In the above example, a group named *PROD-GROUP* consists of hosts, databases, and WebLogic servers exists as part of a company's managed environment. A single rule set is created to manage the group.

In addition to the actual rules contained within a rule set, a rule set possesses the following attributes:

- **Name**: A descriptive name for the rule set.
- **Description**: Brief description stating the purpose of the rule set.
- **Applies To**: Object to which all rules in the rule set apply: Valid rule set objects are targets, jobs, metric extensions, and self update.
- **Owner**: The Enterprise Manager user who created the rule set. Rule set owners have the ability to update or delete the rule set and the rules in the rule set.
- **Enabled**: Whether or not the rule set is actively being applied.
- **Type**: Enterprise or Private. See "Rule Set Types"

Out-of-Box Rule Sets

Enterprise Manager provides out-of-box rule sets for incident creation and event clearing based on typical scenarios. Out-of-box rule sets cannot be edited or deleted, however, they can be disabled. As a best practice, you should create your own copies of out-of-box rule sets and then subscribe to the rule set copies rather than subscribing directly to the out-of-box rule sets. Effectively, you are making a copy of the rule set and changing the target criteria to fit your enterprise needs by selecting an appropriate group of targets (preferably an administration group).

Note that out-of-box rule set definitions and actions they perform can be changed by Oracle at any time and will be applied during patching or software upgrade.

Regular Enterprise Manager administrators are allowed to perform the following operations on rule sets:

- Subscribe for email notifications

- Unsubscribe from email notifications
- Enable
- Disable

> **Note:**
>
> Even though administrators can subscribe to a rule set, they will only receive notification from the targets for which they have at least the View Target privilege.

Enterprise Manager Super Administrators have the added ability to reorder the rule sets.

Enterprise rule sets are evaluated sequentially and may go through multiple passes as needed. When there is a change to the entity being processed - such as an incident being created for an event or an incident priority changing due to a rule - we rerun through all the rules from the beginning again until there are no matches. Any rule that is matched in a prior pass will not match again (to prevent infinite loops).

For example, when a new event, incident, or problem arises, the first rule set in the list is checked to see if any of its member rules apply and appropriate actions specified in those rules are taken. The second rule is then checked to see if its rules apply and so on. Private rule sets are only evaluated once all enterprise rule set evaluations are complete and in no particular order.

> **Note:**
>
> Use caution when reordering rule sets as their order defines the event, incident, and problem handling workflow. Reordering rule sets without fully understanding the impact on your system can result in unintended actions being taken on incoming events, incidents, and problems.

Rule Set Types

There are two types of 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. However, if the rule set owner loses the *Create Enterprise Rule Set* system privilege at some future time, he can still edit or delete the rule set. Super Administrators can edit or delete any rule set. 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 by adding co-authors. Enterprise rule sets are visible to all administrators.
- **Private**: Used when an administrator wants to be notified about something he is monitoring but not as a standard business practice. The only action a private rule set can perform is to send email to the rule set owner. Any administrator can create a private rule set regardless of whether they have been granted the *Create Enterprise Rule Set* resource privilege. Oracle recommends that private rule sets be used only in rare or exceptional situations.

When a rule set performs actions, the privileges of the rule set creator are used. For example, a rule set owner/creator must have at least View Target privilege in order to receive notifications and at least Manage Target Events privilege in order to update the incident. The exception is when a rule set sends a notification. In this case, the privileges of the user it is sent to is used.

Rules

Rules are instructions within a rule set that automate actions on incoming events or incidents or problems. 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 the rule applies.
- **Action**(s): The ordered set of one or more operations on the specified events, incidents, or problems. Each action can be executed based on additional conditions.

The following table shows how rule criteria and actions determine rule application. In this rule operation example there are three rules which take actions on selected events and incidents. Within a rule set, rules are executed in a specified order. The rule execution order can be changed at any time. By default, rules are executed in the order they are created.

Table 5-2 Rule Operation

Rule Name	Execution Order	Criteria	Condition	Actions
Rule 1	First	CPU Util(%), Tablespace Used(%) metric alert events of warning or critical severity	_	Create incident.
Rule 2	Second	Incidents of warning or critical severity	If severity = critical If severity =warning	Notify by page Notify by email
Rule 3	Third	Incidents are unacknowledged for more than six hours	_	Set escalation level to 1

In the rule operation example, *Rule 1* 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.

When the incident severity level (the incident severity is inherited from the worst event severity) reaches Warning, *Rule 2* is applied according to its first condition and Enterprise Manager sends an email to the administrator. If the incident severity level reaches Critical, *Rule 2*'s second condition is applied and Enterprise Manager sends a page to the administrator.

If the incident remains open for more than six hours, *Rule 3* applies and the incident escalation level is increased from None to Level 1. At this point, Enterprise Manager runs through all the rule sets and their rules from the beginning again.

Rule Application

Each rule within a rule set applies to an event, incident OR problem. For each of these, you can choose rule application criteria such as:

- Apply the rule to incoming events or updated events only

- Apply the rule to critical events only.

Rules are applied to events, incidents, and problems according to criteria selected at the time of rule creation (or update). The following situations illustrate the methodology used to apply rules.

- If one of the rules creates a new incident in response to an incoming event, Enterprise Manager finishes matching the event to any further rules/rule sets. Once completed, Enterprise Manager then matches the newly created incident to all the rule sets from the beginning to see if any incident-specific rules match.
- If an incoming event is already associated with an incident (for example, a *Warning* event creates an incident and then a *Critical* event is generated for the same issue), Enterprise Manager applies all the matching rules to the event and then matches all rules to the incident.
- If, while applying a rule to an incident, changes are made to the incident (change priority. for example), Enterprise Manager stops rule application at that point and then re-applies the rules to the incident from the beginning. The conditional action that updated the incident will not be matched again in the same rule application cycle.

Rule Criteria

The following tables list selectable criteria for each type.

Table 5-3 Rule Criteria: Events

Criteria	Description
Type	Rule applies to a specific event type.
Severity	Rule applies to a specific event severity.
Category	Rule applies to a specific event category.
Target type	Rule applies to a specific target type.
Target Lifecycle Status	Rule applies to a specific lifecycle status for a target. Lifecycle status is a target property that specifies a target's operational status.
Associated with incident	Typically, events are associated with incidents through rules. Specify Yes or No.
Event name	Rule applies to events with a specific name. The specified name can either be an exact match or a pattern match.
Causal analysis update	Upon completion of Root Cause Analysis (RCA) event, the rule applies to the event that is marked either as root cause or symptom. Alternatively, the rule can act on an RCA event when it is no longer a symptom.
Associated incident acknowledged	Rule applies to an event that is associated with a specific incident when that incident is acknowledged by an administrator. Specify Yes or No.
Total occurrence count	For duplicated events, the rule is applies when the total number of event occurrences reaches a specified number.
Comment added	Rule applies to events where an administrator adds a comment.

For incidents, a rule can apply to all new and/or updated incidents, or newly created incidents that match specific criteria shown in the following table.

Table 5-4 Rule Criteria: Incidents

Criteria	Description
Rules that created the incident	Rule applies to incidents raised by a specific rule.
Category	Rule applies to a specific incident category.
Target Type	Rule applies to a specific target type.
Target Lifecycle Status	Rule applies to a specific lifecycle status for a target. Lifecycle status is a target property that specifies a target's operational status.
Severity	Rule applies to a specific incident severity.
Acknowledged	Rule applies if the incident has been acknowledged by an administrator. Specify Yes or No.
Owner	Rule applies for a specified incident owner.
Priority	Rule applies when incident priority matches a selected priority.
Status	Rule applies when the incident status matches a selected incident status.
Escalation Level	Rule applies when the incident escalation level matches the selected level. Available escalation levels: None, Level 1, Level 2, Level 3, Level 4, Level 5
Associated with Ticket	Rule applies when the incident is associated with a helpdesk ticket. Specify Yes or No.
Associated with Service Request	Rule applies when the incident is associated with a service request. Specify Yes or No.
Diagnostic Incident	Rule applies when the incident is a diagnostic incident. Specify Yes or No.
Unassigned	Rule applies if the newly raised incident does not have an owner.
Comment Added	Rule applies if an administrator adds a comment to the incident.

For problems, a rule can apply to all new and/or updated problems, or newly created problems that match specific criteria shown in the following table.

Table 5-5 Rule Criteria: Problems

Criteria	Description
Problem key	Each problem has a problem key, which is a text string that describes the problem. It includes an error code (such as ORA 600) and in some cases, one or more error parameters. Rule can apply to a specific problem key or a key matching a specific pattern (using a wildcard character).
Category	Rule applies to a specific problem category.
Target Type	Rule applies to a specific target type.
Target Lifecycle Status	Rule applies to a specific lifecycle status for a target. Lifecycle status is a target property that specifies a target's operational status.
Acknowledged	Rule applies when the problem is acknowledged.
Owner	Rule applies for a specified problem owner.
Priority	Rule applies when problem priority matches a selected priority.
Status	Rule applies when the problems matches a specific status.

Table 5-5 (Cont.) Rule Criteria: Problems

Criteria	Description
Escalation Level	Rule applies when the problem escalation level matches the selected level. Available escalation levels: None, Level 1, Level 2, Level 3, Level 4, Level 5
Incident Count	Rule applies when the number of incidents related to the problem reaches the specified count limit. The problem owner and the Operations manager are notified via email.
Associated with Service Request	Rule applies if the incoming problem is has an associated Service Request. Specify Yes or No.
Associated with Bug	Rule applies if the incoming problem is has an associated bug. Specify Yes or No.
Unassigned	Rule applies if the newly raised incident does not have an owner.
Comment Added	Rule applies if an administrator adds a comment to the problem.

Rule Actions

For each rule, Enterprise Manager allows you to define specific actions.

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

- Create an incident based on an event.
- Perform notification actions such as sending an email or generating a helpdesk ticket.
- Perform actions to manage incident workflow notification via email/PL/SQL methods/ SNMP traps. For example, if a target down event occurs, create an incident and email administrator Joe about the incident. If the incident is still open after two days, set the escalation level to one and email Joe's manager.

The following table summarizes available actions for each rule application.

Table 5-6 Available Rule Actions

Action	Event	Incident	Problem
Email	Yes	Yes	Yes
Page	Yes	Yes	Yes
Advanced Notifications			
Send SNMP Trap	Yes	No	No
Run OS Command	Yes	Yes	Yes
Run PL/SQL Procedure	Yes	Yes	Yes
Create an Incident	Yes	No	No
Set Workflow Attributes	Yes Note: Within an event rule, the workflow attributes of the associated incident can also be updated.	Yes	Yes

Table 5-6 (Cont.) Available Rule Actions

Action	Event	Incident	Problem
Create a Helpdesk Ticket	Yes Note: Action performed indirectly by first creating an incident and then creating a ticket for the incident.	Yes	No

Note:

you can test rule actions against targets without actually performing the actions using Enterprise Manager's event rule simulation feature. For more information, see "Testing Rule Sets".

Incident Manager

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:

- Filter incidents, problems, and events by using custom views
- Search for specific incidents by properties such as target name, summary, status, or target lifecycle status
- Respond and work on an incident
- Manage incident lifecycle including assigning, acknowledging, tracking its status, prioritization, and escalation
- Access (in context) My Oracle Support knowledge base articles and other Oracle documentation to help resolve the incident.
- Access direct in-context diagnostic/action links to relevant Enterprise Manager functionality allowing you to quickly diagnose or resolve the incident.

For example, you have an open incident. You can use Incident Manager to track its ownership, its resolution status, set the priority and, 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 issues quickly. By drilling down on an open incident, you can access this information and modify it accordingly.

Incident Manager Overview - YouTube Video

Displaying Target Information in the Context of an Incident

You can directly view information about a target for which an incident or event has been raised. The type of information shown varies depending on the target type.

To display in-context target information:

1. From the **Enterprise** menu, select **Monitoring** and then **Incident Manager**.

2. From the Incident Manager UI, choose an incident. Information pertaining to the incident displays.
3. From the Incident Details area of the General tab, click on the information icon "i" next to the *target*. Target information as it pertains to the incident displays.

Being able to display target information in this way provides you with more operational context about the targets on which the events and incidents are raised. This in turn helps you manage the lifecycle of the incident more efficiently.

Views

Views let you work efficiently with incidents by allowing you to categorize and focus on only those incidents of interest. A view is a set of search criteria for filtering incidents and problems in the system. Incident Manager provides a set of predefined *standard* views that cover the most common event, incident, and problem search scenarios. In addition, Incident Manager also allows you to create your own custom views. Custom views can be shared with other users. For instructions on creating custom views, see "Setting Up Custom Views". For instructions on sharing a custom view, see "Sharing/Unsharing Custom Views".

Summing Up

- **Event**: A significant occurrence of interest on a target that has been detected by Enterprise Manager.

  Goal: Ensure that your environment is monitored.
- **Incident**: A set of significant events or combination of related events that pertain to the same issue.

  Goal: Ensure that service disruptions are either avoided or resolved quickly.
- **Problems**: The underlying root cause of incidents. Currently, this represents critical errors in Oracle software that represents the underlying root cause of diagnostic incidents.

  Goal: Ensure underlying root causes of issues are resolved to avoid future occurrence of issues.

Events, incidents, and problems work in concert to allow you to manage your complete IT ecosystem both effectively and efficiently. The following illustration summarizes how they work within your managed environment.

Figure 5-4 Event/Incident/Problem Flow

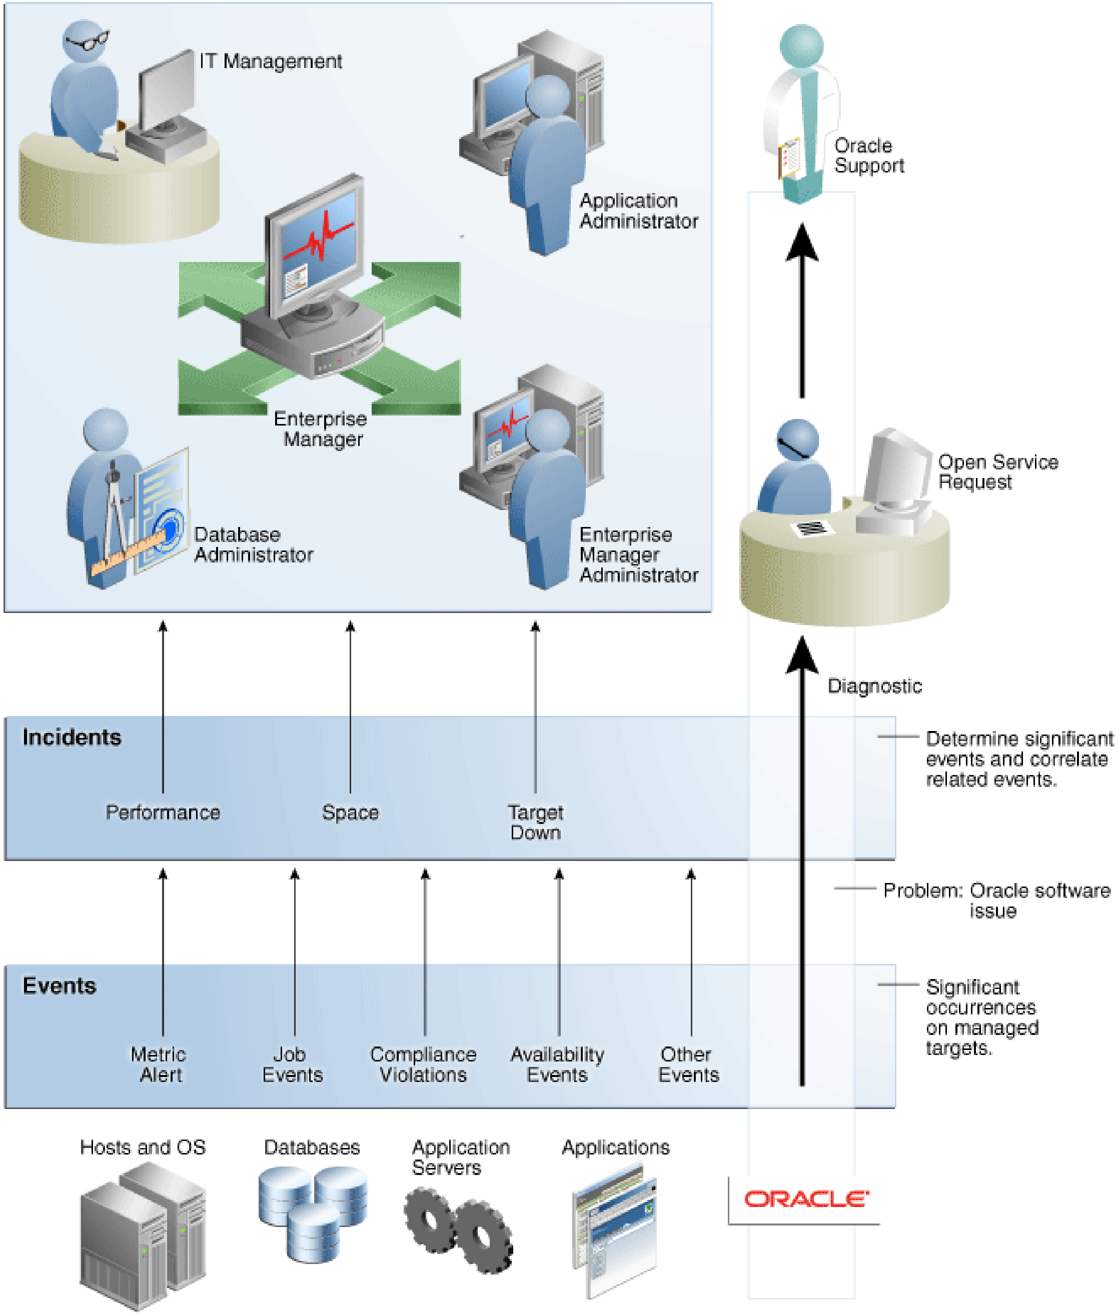

The following sections delve into events, incidents, and problems in more detail.

Setting Up Your Incident Management Environment

Before you can monitor and manage your environment using incidents, you must ensure that your monitoring environment is properly configured. Proper configuration consists of the following:

- Setting Up Your Monitoring Infrastructure

- Setting Up Administrators and Privileges
- Monitoring Privileges
- Setting Up Rule Sets

Setting Up Your Monitoring Infrastructure

The first step in setting up your monitoring infrastructure is to determine which conditions need to be monitored and hence are the source of events. To prevent an inordinate number of extraneous events from being generated, thus reducing system and administrator overhead, you need to determine what is of interest to you and enable monitoring based on your requirements. You can leverage Enterprise Manager features such as Administrations Groups to automatically apply management settings such as monitoring settings or compliance standards when new targets are added to your monitored environment. This greatly simplifies the task of ensuring that events are raised only for those conditions in which you are interested. For more information, see Using Administration Groups .

Example: You want to ensure that the database containing your human resource information is available round the clock. One condition you are monitoring for is whether that database target is up or down. If it goes down, you want the appropriate person to be notified and have them resolve the problem as quickly as possible. Other conditions that you may want to monitor include performance threshold violations, any changes in application configuration files, or job failures. Working with events, you are monitoring and managing individual targets and issues directly related to those targets. For example, you monitor for individual database availability, individual host threshold violations such as CPU and I/O load, or perhaps the performance of a Web service.

In general, if you are primarily interested in availability and some key performance related metrics, you should use default monitoring templates and other template features to ensure the only those specific metrics are collected and events are raised only for those metrics.

Job Events: The status of a job can change throughout its lifecycle - from the time it is submitted to the time it has executed. For each of these job statuses, events can be raised to notify administrators of the status of the job.

As a general rule, events should be generated only for job status values that require administration attention. These job status values include Action Required and Problem status values such as Failed or Stopped. However, in order to avoid overloading the system with unnecessary events, job events are not enabled for any target by default. Hence, if you would like to generate events for jobs, you must:

1. Set the appropriate job status. You can use the default settings or modify them as required.
2. Specify the set of targets for which you would like job-related events to be generated.

   You can perform these operations from the *Job Event Generation Criteria* page. From the **Setup** menu, choose **Incidents** and then **Job Events**.

Rule Set Development

Before creating incident rules/rule sets, the first step is to strategically determine when incidents should be created based on the business requirements of your organization. Important questions to consider are:

1. What events should create incidents? Which service disruptions need to be tracked and resolved by IT administrators?
2. Which administrators should be notified for incoming events or incidents?

3. Are any of the events or incidents being forwarded to external systems (such as a helpdesk ticketing system)?

Example 5-1   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 incidents
  3. Rule(s) that escalate incidents based on some condition. For example, the length of time an incident is open.

Example 5-2   Example Rule Set in Greater Detail

- Rule Set for Production Group G
  - Target: Production Group G
  - Rule 1: Create an incident for all *target down* events.
  - Rule 2: Create an incident for specific database, host, and WebLogic Server metric alert event of critical or warning severity.
  - Rule 3: Create an incident for any problem job events.
  - Rule 4: For all critical incidents, sent a page. For all warning incidents, send email.
  - Rule 5: If a Fatal incident is open for more than 12 hours, set the escalation level to 1 and email a manager.

Once the exact business requirements are understood, you translate those into enterprise rule sets. Adhering to the following guidelines 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 targets into a smaller number of monitoring entities for the rule set. Groups should be composed of targets that have similar monitoring requirements including incident management and response.
- All the rules that apply to the same groups of targets should be consolidated into one rule set. You can create multiple rules that apply to the targets in the rule set. You can create rules for events specific to an event class, rules that apply to events of a specific event class and target type, or rules that apply to incidents on these targets.
- Leverage the execution order of rules within the rule set. Rule sets and rules within a rule set are executed in sequential order. Therefore, ensure that rules and rule sets are sequenced with that in mind.

When creating a new rule, you are given a choice as to what object the rule will apply—events, incidents or problems. Use the following rule usage guidelines to help guide your selection.

Table 5-7 Rule Usage Guidelines

Rule Usage	Application
Rules on Event	To create incidents for the events managed in Enterprise Manager. To send notifications on events. To create tickets for incidents managed by helpdesk analysts, you want to create an incident for an event, then create a ticket for the incident. Send events to third-party management systems.
Rules on Incidents	Automate management of incident workflow operations (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.
Rules on Problems	Automate management of problem workflow operations (assign owner, set priority, escalation levels..) and send notifications

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 set. In this example, the rule set applies to a group (Production Group G) that consists of the following targets:

- DB1 (database)
- Host1 (host)
- WLS1 (WebLogic Server)

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

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, execution efficiency is achieved. As shown in this example, by using conditional actions that take different actions for the same set of events based on severity, it is easier to change the event selection criteria in the future without having to change multiple rules. **Note**: This assumes that the action requirements for all incidents (from rules 1 - 3) are the same.

The following table illustrates explicit rule set operation for this example. All targets are within Production Group G.

Table 5-8 Example Rule Set for Production Group G

Rule Name	Execution Order	Criteria	Triggering Condition	Actions
Rule 1	First	DB1 goes down . Host1 goes down. WLS1 goes down.	N/A	Create incident.
Rule 2	Second	**DB1** Tablespace Full (%) Note: The warning and critical thresholds are defined in Metric and Policy settings, not from the rules UI. **Host1** CPU Utilization (%) **WLS1** Heap Usage (%)	If severity=Warning If severity=Critical	Create incident.
Rule 3	Third	Event generated for problem job status changes for DB1, Host1, and WLS1.	N/A	Create incident.

Rule 4	Fourth	All incidents for Production Group G	Severity=Warning
			Severity=Critical
				Send email
				Send page

Rule 5	Fifth	Incident remains open for more than 12 days.	Status=Fatal	Increase escalation level to 1.

Before Using Rules

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

- User's Enterprise Manager account has notification preferences (email and schedule). This is required not just for the administrator who is creating/editing a rule, but also for any user who is being notified as a result of the rule action.
- If you decide to use connectors, tickets, or advanced notifications, you need to configure them before using them in the actions page.
- Ensure that the SMTP gateway has been properly configured to send email notifications.
- User's Enterprise Manager account has been granted the appropriate privileges to manage incidents from his managed system.

Setting Up Notifications

After determining which events should be raised for your monitoring environment, you need to establish a comprehensive notification infrastructure for your enterprise by configuring Enterprise Manager to send out email and or pages, setting up email addresses for administrators and tagging them as email/paging. In addition, depending on the needs of your organization, notification setup may involve configuring advanced notification methods such as OS scripts, PL/SQL procedures, SNMP traps, Webhooks or Slack. For detailed information and setup instructions for Enterprise Manager notifications, see Using Notifications .

Setting Up Administrators and Privileges

This step involves defining the appropriate administrators (which includes assigning the proper privileges for security) and then setting up notification assignments based on their defined roles and domain ownership within your organization.

To perform user account administration, click **Setup** on the Enterprise Manager home page, select **Security**, then select **Administrators** to access the Administrators page.

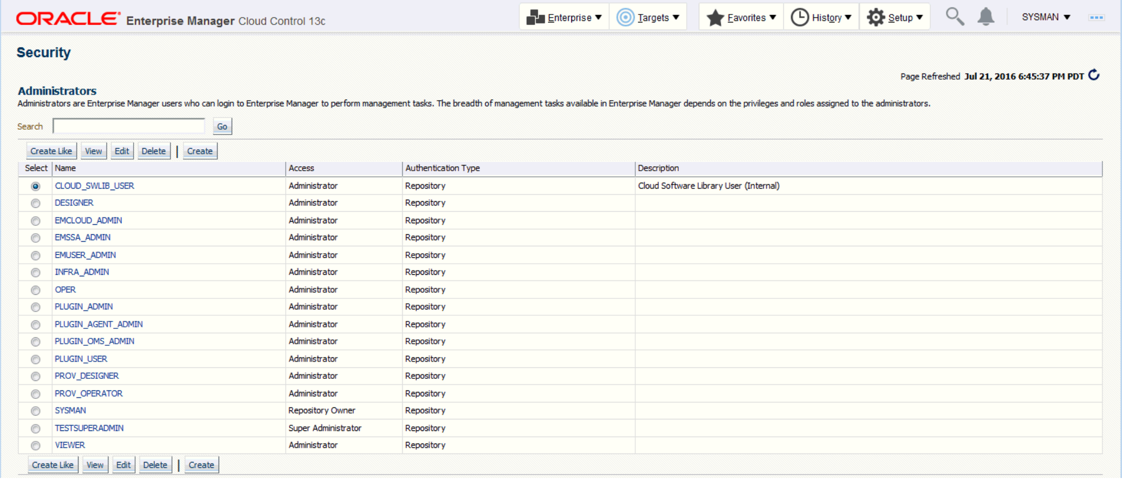
ORACLE Enterprise Manager Cloud Control 13c
Enterprise
Targets
Favorites
History
Setup
SYSMAN

Security

Page Refreshed Jul 21, 2016 6:45:37 PM PDT

Administrators

Administrators are Enterprise Manager users who can login to Enterprise Manager to perform management tasks. The breadth of management tasks available in Enterprise Manager depends on the privileges and roles assigned to the administrators.

Search
Go
Create Like | View | Edit | Delete | Create

Select	Name	Access	Authentication Type	Description
●	CLOUD_SWLIB_USER	Administrator	Repository	Cloud Software Library User (Internal)
○	DESIGNER	Administrator	Repository	
○	EMCLOUD_ADMIN	Administrator	Repository	
○	EMSSA_ADMIN	Administrator	Repository	
○	EMUSER_ADMIN	Administrator	Repository	
○	INFRA_ADMIN	Administrator	Repository	
○	OPER	Administrator	Repository	
○	PLUGIN_ADMIN	Administrator	Repository	
○	PLUGIN_AGENT_ADMIN	Administrator	Repository	
○	PLUGIN_OMS_ADMIN	Administrator	Repository	
○	PLUGIN_USER	Administrator	Repository	
○	PROV_DESIGNER	Administrator	Repository	
○	PROV_OPERATOR	Administrator	Repository	
○	SYSMAN	Repository Owner	Repository	
○	TESTSUPERADMIN	Super Administrator	Repository	
○	VIEWER	Administrator	Repository	

There are two types of administrators typically involved in incident management.

- *Business Rules Architect/Analyst*: Administrator who has a deep understanding of how the business works and translates this knowledge to operational rules. Once these rules have been deployed, the business architect uses their knowledge of the dynamic organization to keep these rules up-to-date.

  In order to create or edit an enterprise rule set, the business architect/analyst must have been granted the *Create Enterprise Rule Set* privilege on the *Enterprise Rule Set* resource. The architect/analyst can share ownership of the rule sets with other administrators who may or may not have the *Create Enterprise Rule Set* privilege but are responsible for managing a specific rule set.

- *IT Operator/Manager*: The IT manager is responsible for day-to-day management of incident assignment. The IT operator is assigned the incidents and is responsible for their resolution.

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

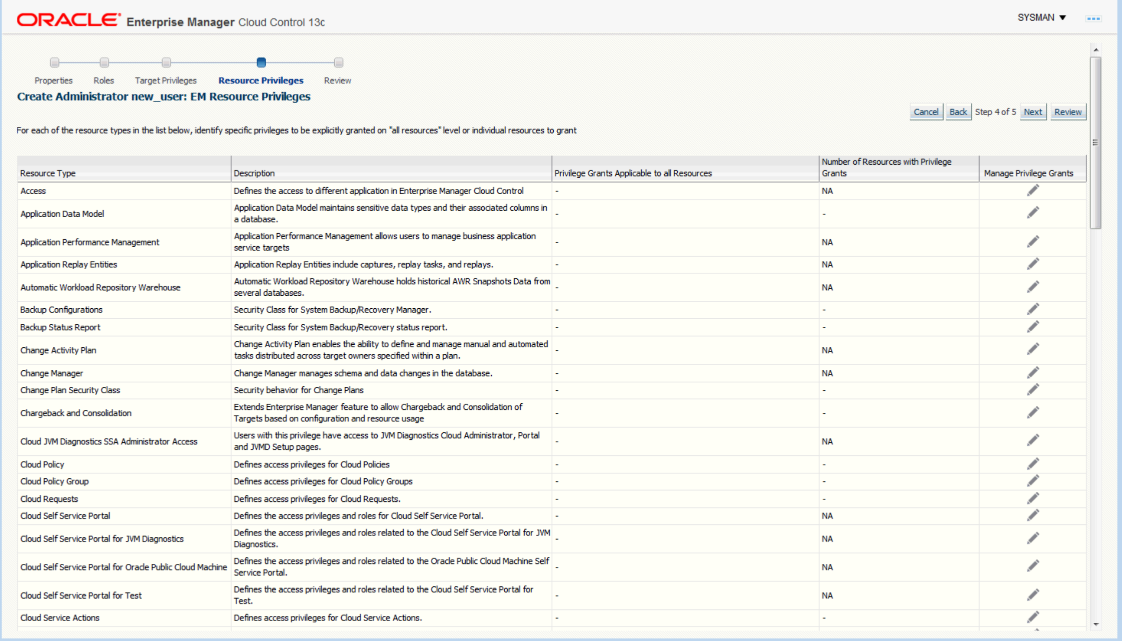

Create Administrator new_user: EM Resource Privileges

For each of the resource types in the list below, identify specific privileges to be explicitly granted on "all resources" level or individual resources to grant

Resource Type	Description	Privilege Grants Applicable to all Resources	Number of Resources with Privilege Grants	Manage Privilege Grants
Access	Defines the access to different application in Enterprise Manager Cloud Control	-	NA	✎
Application Data Model	Application Data Model maintains sensitive data types and their associated columns in a database.	-	-	✎
Application Performance Management	Application Performance Management allows users to manage business application service targets	-	NA	✎
Application Replay Entities	Application Replay Entities include captures, replay tasks, and replays.	-	NA	✎
Automatic Workload Repository Warehouse	Automatic Workload Repository Warehouse holds historical AWR Snapshots Data from several databases.	-	NA	✎
Backup Configurations	Security Class for System Backup/Recovery Manager.	-	-	✎
Backup Status Report	Security Class for System Backup/Recovery status report.	-	-	✎
Change Activity Plan	Change Activity Plan enables the ability to define and manage manual and automated tasks distributed across target owners specified within a plan.	-	NA	✎
Change Manager	Change Manager manages schema and data changes in the database.	-	NA	✎
Change Plan Security Class	Security behavior for Change Plans	-	-	✎
Chargeback and Consolidation	Extends Enterprise Manager feature to allow Chargeback and Consolidation of Targets based on configuration and resource usage	-	-	✎
Cloud JVM Diagnostics SSA Administrator Access	Users with this privilege have access to JVM Diagnostics Cloud Administrator, Portal and JVMD Setup pages.	-	NA	✎
Cloud Policy	Defines access privileges for Cloud Policies	-	-	✎
Cloud Policy Group	Defines access privileges for Cloud Policy Groups	-	-	✎
Cloud Requests	Defines access privileges for Cloud Requests.	-	-	✎
Cloud Self Service Portal	Defines the access privileges and roles for Cloud Self Service Portal.	-	NA	✎
Cloud Self Service Portal for JVM Diagnostics	Defines the access privileges and roles related to the Cloud Self Service Portal for JVM Diagnostics.	-	NA	✎
Cloud Self Service Portal for Oracle Public Cloud Machine	Defines the access privileges and roles related to the Oracle Public Cloud Machine Self Service Portal.	-	NA	✎
Cloud Self Service Portal for Test	Defines the access privileges and roles related to the Cloud Self Service Portal for Test.	-	NA	✎
Cloud Service Actions	Defines access privileges for Cloud Service Actions.	-	-	✎

Granting User Privileges for Events, Incidents and Problems

In order to work with incidents, all relevant Enterprise Manager administrator accounts must be granted the appropriate privileges to manage incidents. Privileges for events, incidents, and problems are determined according to the following rules:

- 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 the events in the incident.
- Similarly, problem privileges are calculated based on privileges on underlying incidents.

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

For events, two privileges are defined in the system:

- 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 an event, creating an incident for an event, and creating a ticket for an event. You can also associate an event with an incident.

> **Note:**
>
> Incident privilege is inherited from the underlying events.

If an event is raised on a target alone (the majority of event types are raised on targets such as metric alerts, availability events or service level agreement), you will need the following privileges:

- *View* on target to view the event.
- *Manage Target Events* to manage the event.

  Note: This is a sub-privilege of Operator.

If an event is raised on both a target and a job, you will need the following privileges:

- *View on target* and *View* on the job to view the event.
- *View on target* and *Full* on the job to manage the event.

If the event is raised on a job alone, you will need the following privileges:

- *View* on the job to view the event.
- *Full* on the job to manage the event.

If an event is raised on a metric extension, you will need *View* privilege on the metric extension to view the event. Because events raised on metric extensions are informational (and do not appear in Incident Manager) event management privileges do not apply in this situation.

If an event is raised on a Self-update, only system privilege is required. Self-update events are strictly informational.

For incidents, two privileges are defined in the system:

- 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 events, and creating tickets for incidents.

If an incident consists of a single event, you can view the incident if you can view the event and manage the incident if you can manage the event.

If an incident consists of more than one event, you can view the incident if you can view at least one event and manage incident if you can manage at least one of the events.

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, and manually closing the problem.

In Enterprise Manager 12*c*, problems are always related to a single target. So the View Problem privilege, if an administrator has View privilege on the target, and the Manage Problem privilege, if an administrator has *manage_target_events* privilege on the target, implicitly grants management privileges on the associated event. This, in turn, grants management privileges on the incident within the problem.

Monitoring Privileges

The monitoring functions that an administrator can perform within the Enterprise Manager environment depend on privileges that have been granted to that user. To maintain the integrity and security of a monitored infrastructure, only the required privileges for a specific role should be granted. The following guidelines can be used to grant proper privilege levels based on user roles.

Administrators who set up monitoring

Create a role with privileges and grant it to administrators:

- Recommend using individual user accounts instead of shared account
- If using super administrator, do not use sysman
- If privilege is based on targets, create privilege-propagating group containing the targets (or use administration group if it meets requirements) and grant privilege on the group to the role

Administrators who respond to events / incidents

- Create a role and grant it to administrators
- Create privilege-propagating group (or use administration group if it meets requirements) containing relevant targets and grant appropriate privilege on the group to the role

Example: You create the role *DB_Admins* and grant *Manage Target Events* on a the privilege-propagating group named *DB-group* containing relevant databases. You then grant role DB_Admins to the DBAs.

Monitoring Actions and Required Privileges

Enterprise Manager supports fine-grained privileges to enable more granular control over actions performed in Enterprise Manager.

The table below shows a (non-exhaustive) list of various job responsibilities and the corresponding privilege in Enterprise Manager required to support these

The following tables summarize the privilege levels required to perform specific monitoring responsibilities.

Table 5-9 Monitoring Operations and Required Privileges

Monitoring Operation	Required Privilege(s)
Monitoring Setup	_
Configure SMTP gateway (email)	Super Administrator
Create Advanced Notification Methods (e.g. SNMP traps)	Super Administrator
Configure event or ticketing connector	Super Administrator
Creating Roles	Super Administrator
Create Administration Group Hierarchy	Full Any Target Create Privilege Propagating Group
Edit Administration Group Hierarchy	Full Any Target Create Privilege Propagating Group (if adding new target property values as group criteria within a level of the administration group hierarchy)

Table 5-9 (Cont.) Monitoring Operations and Required Privileges

Monitoring Operation	Required Privilege(s)
Delete Administration Group Hierarchy	Full Any Target
View entire Administration Group hierarchy in Group Administration pages	View Any Target Note: Administrators who have privileges to only a subset of the groups can view these groups in the Groups list page accessible via Targets-->Groups
Use Monitoring Templates	No privileges required to create new monitoring templates. However if the monitoring template contains a corrective action, then Create on Job System privilege is required View on specific monitoring template to use the template created by another user (e.g. to add the monitoring template to a Template Collection
Use Template Collections	Create Template Collection (to create new Template Collections)View Template Collection on specific Template Collection to view/associate the Template Collection created by another userView Any Template Collection to view/associate any Template CollectionFull Template Collection on specific Template Collection to edit/delete the Template Collection created by another user
Associate a Template Collection with an Administration Group	Manage Template Collection Operations on the group (this includes Manage Target Compliance and Manage Target Metrics privileges) View Template Collection on the Template Collection
Operations on the Administration Group	_
Manage privileges on the group (for example, grant to other users)	Group Administration on the group
Add a target to an Administration Group by setting its target properties	Configure Target (on the target to be added to the Administration Group)
Perform a manual sync of the group with the associated Template Collection	Manage Template Collection Operations on the group
Operations on the members of the Administration Group	_
Delete the target from Enterprise Manager	Full on the target (Full also contains the privileges enumerated below

Table 5-9 (Cont.) Monitoring Operations and Required Privileges

Monitoring Operation	Required Privilege(s)
Set blackout for planned downtime Change monitoring settings Change monitoring configuration Manage events and incidents on the target View target, receive notifications for events or incidents	Operator on the target also contains the following privileges: • Blackout Target on the target • Manage Target Metrics on the target • Configure Target on the target • Manage Target Events on the target • View on the target
Create Incident Rule Sets	Create Enterprise Rule Set Manage Target Events on target if rule is creating incidents for the target
Granting privileges on administration group to roles	No extra privilege required if creator of the administration group
Set a target's property values	Configure Target
Edit Monitoring Template that is part of Template Collection	Full on the Monitoring Template Manage Target Metrics on administration group
Change monitoring settings on specific target	Manage Target Metrics
Receive email for events, incidents	View on Target and/or View on source object (for example, view on job for job events)
Create incident for event	Manage Target Events
Incident management actions (for example, acknowledge, assign incident, prioritize, set escalation level)	Manage Target Events

Note:

SYSMAN is a system account intended for Enterprise Manager infrastructure installation and maintenance. It should never be used for administrator access to Enterprise Manager as a Super Administrator.

Setting Up Rule Sets

Rule sets automate actions in response to incoming events, incidents and problems or updates to them. This section covers the most common tasks and examples.

- Creating a Rule Set
- Creating a Rule to Create an Incident
- Creating a Rule to Manage Escalation of Incidents
- Creating a Rule to Escalate a Problem
- Testing Rule Sets

- Subscribing to Receive Email from a Rule
- Receiving Email for Private Rules

Creating a Rule Set

In general, to create a rule set, perform the following steps:

1. From the **Setup** menu, select **Incidents** then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, edit the existing rule set or create a new rule set. For new rule sets, you will need to first select the targets to which the rules apply. Rules are created in the context of a 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.

Narrowing Rule Set Scope Based on Target Lifecycle Status

When creating a new rule set, you can choose to have the rule set apply to a narrower set of targets based on the target's *Lifecycle Status* value. For example, you can create one rule set that only applies only to targets that have a *Lifecycle Status* of *Staging* and *Production*. As shown in the following graphic, you determine rule set scope by setting the *Lifecycle Status* filter.

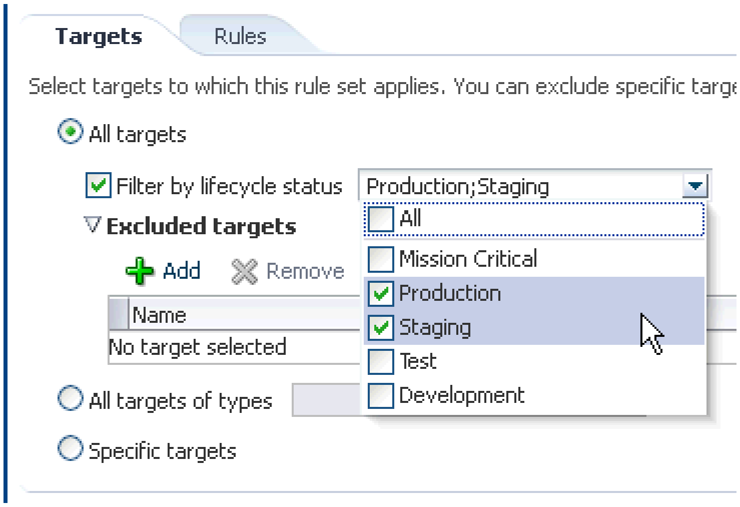

Targets | Rules

Select targets to which this rule set applies. You can exclude specific targe

● All targets

☑ Filter by lifecycle status: Production;Staging

□ All
□ Mission Critical
☑ Production
☑ Staging
□ Test
□ Development

▽ Excluded targets

No target selected

○ All targets of types

○ Specific targets

Using this filter allows you to create rules for targets based on their *Lifecycle Status* without having to first create a group containing only such targets.

Narrowing Rule Set Scope by Excluding Targets

You can also choose to narrow the scope of the rule set by excluding specific targets. The Exclude targets option lets you add one or more targets to be omitted. You will need to query the Enterprise Manager Repository for the specific names of targets.

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 pop-up dialog. Click **Continue**.

4. In the Create New Rule wizard, provide the required information.
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.

Creating a Rule to Create an Incident

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

1. From the **Setup** menu, 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 rule/rule set name, description, target name, or target type for the target and the associated rule set. You can search by target name or the group target name to which this target belongs to locate the rule sets that manage the targets.

   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,
   a. Click **Create ...**
   b. Select "Incoming events and updates to events"
   c. 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.) **Note**: Only one event type can be selected in a single rule and, once selected, it cannot be changed when editing a rule.

      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.

Creating a Rule to Designate a Dynamic Runbook for an Incident

1. From the **Setup** menu, 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 rule/rule set name, description, target name, or target type for the target and the associated rule set. You can search by target name or the group target name to which this target belongs to locate the rule sets that manage the targets.

   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,

   a. Click **Create ...**

   b. Select "Incoming events and updates to events"

   c. 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.) **Note**: Only one event type can be selected in a single rule and, once selected, it cannot be changed when editing a rule.

      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.

Creating a Rule to Manage Escalation of Incidents

To create a rule to manage incident escalation, perform the following steps:

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.

2. Determine whether there is an existing rule set that contains a rule that manages the incident. You can add it to any of your existing rule sets on incidents.

   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:

   a. Click **Create ...**
   b. Select "Newly created incidents or updates to incidents"
   c. 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 where a fatal incident is open for 48 hours. Similarly, the DBA director is paged if the incident has been escalated to level 2, the severity is fatal and it has been open for 72 hours. If the fatal incident is still open after 96 hours, then it is escalated to level 3 and the operations VP is notified.

   Provide the rule details using the Create New Rule wizard.

   a. To set up the rule to apply to all newly created incidents or when the incident is updated with *fatal* severity, select the **Specific Incidents** option and add the condition *Severity is Fatal* .
   b. In the **Conditions for Actions** region located on the Add Actions page, select **Only execute the actions if specified conditions match**.

      Select **Incident has been open for some time and is in a particular state (select time and optional expressions)**.

      Select the time to be 48 hours and Status is not resolved or closed.
   c. In the **Notification** region, type the name of the administrator to be notified by email or page. Click **Continue** to save the current set of conditions and actions.
   d. Repeat steps b and c to page the DBA director (Time in this state is 72 hours, Status is Not Resolved or Closed). If open for more than 96 hours, set escalation level to 3, page Operations VP.
   e. After reviewing added actions sets, click **Next**. Click **Next** to go to the Summary screen. Review the summary information and click **Continue** to save the rule.

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

   In Edit Rule Set page, click on the **desired rule from the Rules** table and select **Reorder Rules** from the **Actions** menu to reorder rules within the rule set, then click **Save** to save the rule sequence changes.

Example Scenario

To facilitate the incident escalation 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, email addresses and schedule).

Creating a Rule to Escalate a Problem

In an organization, whenever an unresolved problem has more than 20 occurrences of associated incidents, the problem should be auto-assigned to the appropriate administrator based on target type of the target on which the problem has been raised.

Accordingly, a problem rule is created to observe the count of incidents attached to the problem and notify the appropriate administrator handling that specific target type.

The problem owner and the Operations manager are notified by email.

To create a rule to escalate a problem, perform the following steps:

1. Navigate to the Incident Rules page.

   From the **Setup** menu, 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.

   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. In the Rules section of the Edit Rule Set page, select **Create...**

4. From the **Select Type of Rule to Create** dialog, select **Newly created problems or updates to problems** and click **Continue**.

5. 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**.

6. 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 email to the owner of the problem and to the Operations Manager.
   - In the Update Problem region, enter the email address of the appropriate administrator in the **Assign to** field.

   Click **Continue**.

7. Review the rules summary. Make corrections as needed. Click **Continue** to return to Edit Rule Set page and then click **Save** to save the rule set.

Testing Rule Sets

When developing a rule set, it can be difficult to develop rule criteria to match all possible event conditions. Previously, the only way to test rules was to trigger an event within your monitored environment and seeing which rules match the event and what actions the rules perform. Beginning with Enterprise Manager Release 12.1.0.4, you can simulate existing events, thus allowing you to test rule actions during the rule set development phase and not waiting for specific event conditions to occur. The rule simulation feature lets you see how the rules will perform given a specific event. You immediately see which rules match for a given event and then see what actions are taken.

> **Note:**
>
> The simulate rule feature can only be used with event rules. Incident rules cannot be tested with this feature.

To simulate rules:

This procedure assumes you have already created rule sets. See "Creating a Rule Set" for instructions on creating a rule set. Ensure that the rule type is *Incoming events and updates to events*.

1. From the Setup menu, select Incidents, and then Incident Rules. The *Incident Rules - All Enterprise Rules* page displays.
2. Click **Simulate Rules**. The Simulate Rules dialog displays.

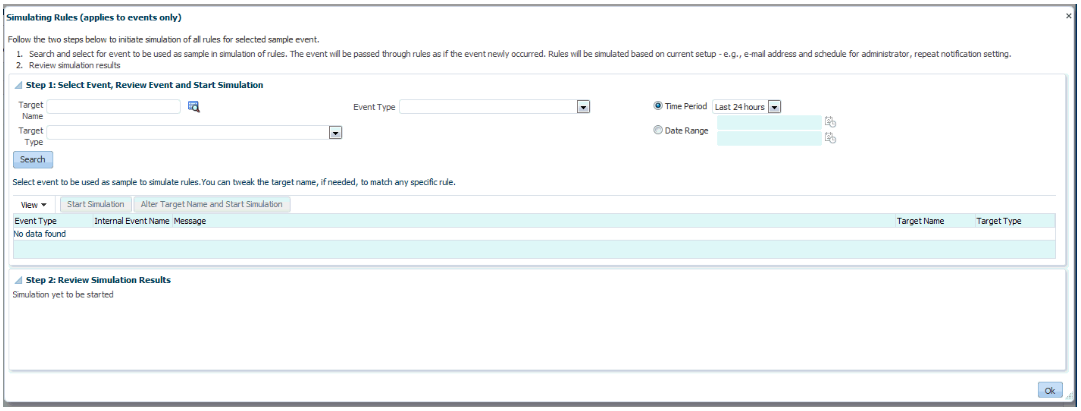

Simulating Rules (applies to events only)

Follow the two steps below to initiate simulation of all rules for selected sample event.

1. Search and select for event to be used as sample in simulation of rules. The event will be passed through rules as if the event newly occurred. Rules will be simulated based on current setup - e.g., e-mail address and schedule for administrator, repeat notification setting.
2. Review simulation results

Step 1: Select Event, Review Event and Start Simulation

Target Name:
Target Type:
Event Type:
● Time Period: Last 24 hours
○ Date Range

Select event to be used as sample to simulate rules.You can tweak the target name, if needed, to match any specific rule.

Event Type	Internal Event Name	Message	Target Name	Target Type
No data found				

Step 2: Review Simulation Results

Simulation yet to be started

3. Enter the requisite search parameters to find matching events and click **Search**.
4. Select an event from the list of results.

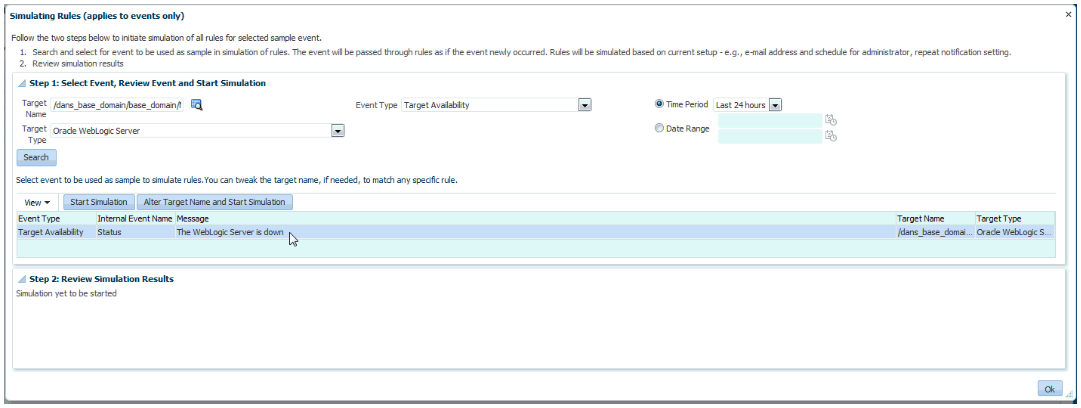

Simulating Rules (applies to events only)

Follow the two steps below to initiate simulation of all rules for selected sample event.

1. Search and select for event to be used as sample in simulation of rules. The event will be passed through rules as if the event newly occurred. Rules will be simulated based on current setup - e.g., e-mail address and schedule for administrator, repeat notification setting.
2. Review simulation results

Step 1: Select Event, Review Event and Start Simulation

Target Name: /dans_base_domain/base_domain/
Target Type: Oracle WebLogic Server
Event Type: Target Availability
● Time Period: Last 24 hours
○ Date Range

Select event to be used as sample to simulate rules.You can tweak the target name, if needed, to match any specific rule.

Event Type	Internal Event Name	Message	Target Name	Target Type
Target Availability	Status	The WebLogic Server is down	/dans_base_domai...	Oracle WebLogic S...

Step 2: Review Simulation Results

Simulation yet to be started

5. Click **Start Simulation**. The event will be passed through the rules as if the event had newly occurred. Rules will be simulated based on the current notification configuration (such as email address, schedule for the assigned administrator, or repeat notification setting).

   Changing the Target Name: Under certain circumstances, an event matching rule criteria may occur on a target that is not a rule target. For testing purposes, you are only interested in the event. To use the alternate target for the simulation, click **Alter Target Name and Start Simulation.**

Results are displayed.

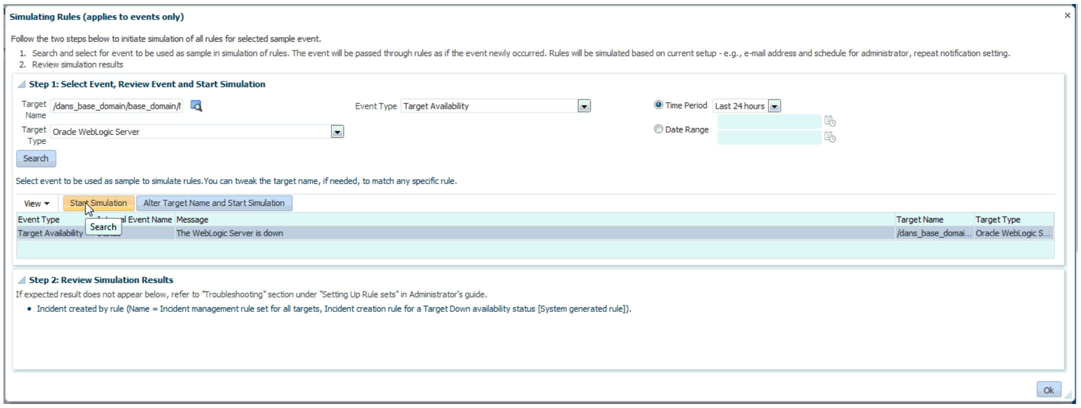

Simulating Rules (applies to events only)

Follow the two steps below to initiate simulation of all rules for selected sample event.

1. Search and select for event to be used as sample in simulation of rules. The event will be passed through rules as if the event newly occurred. Rules will be simulated based on current setup - e.g., e-mail address and schedule for administrator, repeat notification setting.
2. Review simulation results

Step 1: Select Event, Review Event and Start Simulation

Target Name: /dans_base_domain/base_domain/
Event Type: Target Availability
● Time Period: Last 24 hours
○ Date Range
Target Type: Oracle WebLogic Server

Select event to be used as sample to simulate rules.You can tweak the target name, if needed, to match any specific rule.

Event Type	Internal Event Name	Message	Target Name	Target Type
Target Availability	Status	The WebLogic Server is down	/dans_base_domai...	Oracle WebLogic S...

Step 2: Review Simulation Results

If expected result does not appear below, refer to "Troubleshooting" section under "Setting Up Rule sets" in Administrator's guide.

- Incident created by rule (Name = Incident management rule set for all targets, Incident creation rule for a Target Down availability status [System generated rule]).

Testing Event Rules on a Production Target: Although you can generate an event on a test target, you may want to check the actions on a production target for final verification. You can safely test event rules on production targets without performing rule actions (sending email, SNMP traps, opening trouble tickets). To test your event rule on a production target, change the Target Name to a production target. When you run the simulation, you will see a list of actions to be performed by Enterprise Manager. None of these actions, however, will actually be performed on the production target.

6. If the rule actions are not what you intended, edit the rules and repeat the rule simulation process until the rules perform the desired actions. The following guidelines can help ensure predictable/expected rule simulation results.

   If you do not see a rule action for email:

   - Make sure there is a rule that includes that event and has an action to send email.
   - If the specified email recipient is an Enterprise Manager administrator, make sure that administrator has an email address and notification schedule set up.
   - Make sure the email recipient has at least View privileges on the target of the event.
   - Check the SMTP gateway setup and make sure that the administrator has performed a Test Email.

   If you do not see other rule actions such as creating an incident or opening a ticket:

   - Make sure there is a rule that includes the event and corresponding action (create incident, for example).
   - Make sure the target is included in the rule set.
   - Make sure the rule set owner has at least *Manage Events* target privilege on the target of the event.
   - For notifications such as Open Ticket, Send SNMP trap, or Call Event Connector, make sure these are specified as actions in the event rule.

Subscribing to Receive Email from a Rule

A DBA is aware that incidents owned by him will be escalated when not resolved in 48 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 there exists a rule that escalates High Priority Incidents for databases that have not been resolved in 48 hours

Perform the following steps:

1. From the **Setup** menu, select **Incidents**, and then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, click on the rule set containing incident escalation rule in question and click **Edit..**. Rules are created in the context of a 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. In the Rules section of the Edit Rule Set page, highlight the escalation rule and click **Edit...**.
4. Navigate to the Add Actions page.
5. Select the action that escalates the incident and click **Edit...**
6. In the Notifications section, add the DBA to the **email cc** list.
7. Click **Continue** and then navigate back to the **Edit Rule Set** page and click **Save**.

As a result of the edit to the enterprise rule, when an incident stays unresolved for 48 hours, the rule marks it to escalation level 1. An email is sent out to the DBA notifying him about the escalation of the incident.

Alternate Rule Set Subscription Method: From the Incident Rules - All Enterprise Rules page, select the rule in incident rules table. From the **Actions** menu, select **email** and then **Subscribe me** (or **Subscribe administrator....**).

Receiving Email for Private Rules

A DBA has setup a backup job on the database that he is administering. As part of the job, the DBA has subscribed to email notification for "completed" job status. Before you create the rule, ensure that the DBA has the requisite privileges to create jobs. See Utilizing the Job System and Corrective Actions for job privilege requirements.

Perform the following steps:

1. Navigate to the Rules page.

   From the **Setup** menu, 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 **Incoming events and updates to events**. 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 email Me for notifications.
6. Review the rules summary. Make corrections as needed. Click **Save**.

7. Create a database backup job and subscribe for email 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 is considered 'matching' against the incoming job events and email will be sent to the DBA.

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

Working with Incidents

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:

- Send notifications to the appropriate administrators.
- Create incidents and rules.
- 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.
- Accessing My Oracle Support knowledge articles.
- Opening Oracle Service Requests to request assistance with issues with Oracle software (Problems).

You can update resolution information for an incident by performing the following:

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

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

Working with incidents involves the following stages:

1. Finding What Needs to be Worked On
2. Searching for Incidents
3. Setting Up Custom Views
4. Responding and Working on a Simple Incident
5. Responding to and Managing Multiple Incidents, Events and Problems in Bulk
6. Managing Workload Distribution of Incidents
7. Creating an Incident Manually

Finding What Needs to be Worked On

Enterprise Manager provides multiple access points that allow you to find out what needs to be worked on. The primary focal point for incident management is the Incident Manager console, however Enterprise Manager also provides other methods of notification. The most common way to be notified that you have an issue that needs to be addressed is by email. However, incident information can also be found in the following areas:

Custom Views (See "Setting Up Custom Views")

Group or System Homepages (See Managing Groups)

Target Homepages

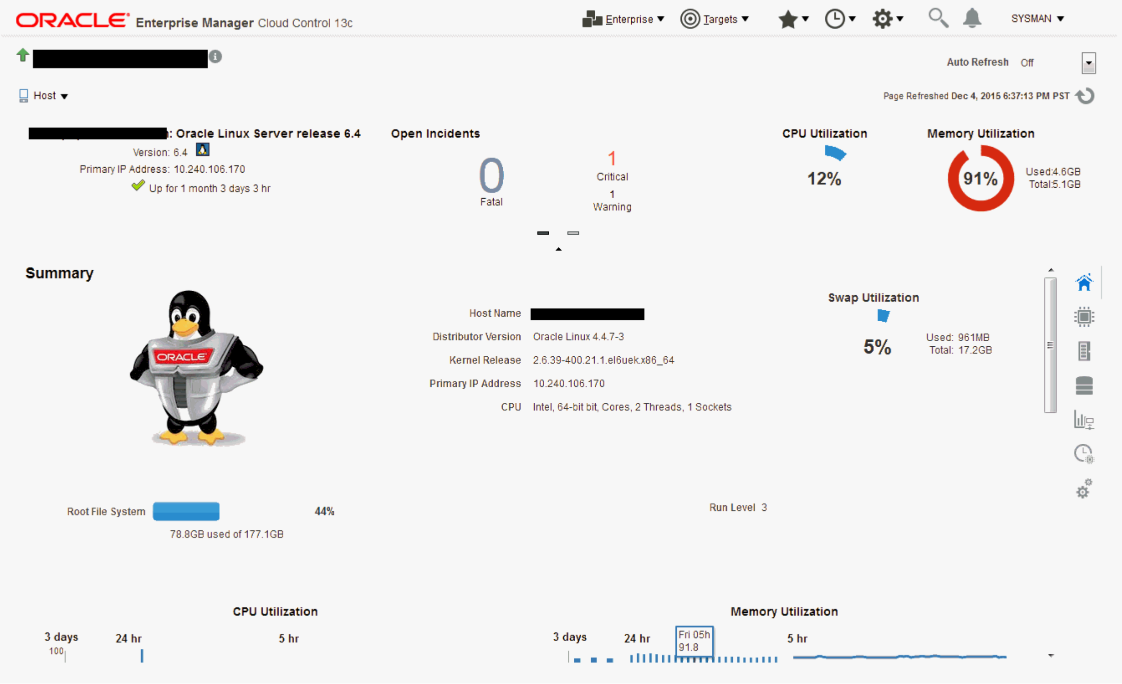

[redacted]

Page Refreshed Dec 4, 2015 6:37:13 PM PST

[redacted]: Oracle Linux Server release 6.4
Version: 6.4
Primary IP Address: 10.240.106.170
Up for 1 month 3 days 3 hr

0 Fatal
1 Critical
1 Warning

CPU Utilization 12%

Memory Utilization 91% Used:4.6GB Total:5.1GB

Host Name [redacted]
Distributor Version Oracle Linux 4.4.7-3
Kernel Release 2.6.39-400.21.1.el6uek.x86_64
Primary IP Address 10.240.106.170
CPU Intel, 64-bit bit, Cores, 2 Threads, 1 Sockets

Swap Utilization 5% Used: 961MB Total: 17.2GB

Root File System 44%
78.8GB used of 177.1GB

Run Level 3

Incident Manager (in context of a system or target)

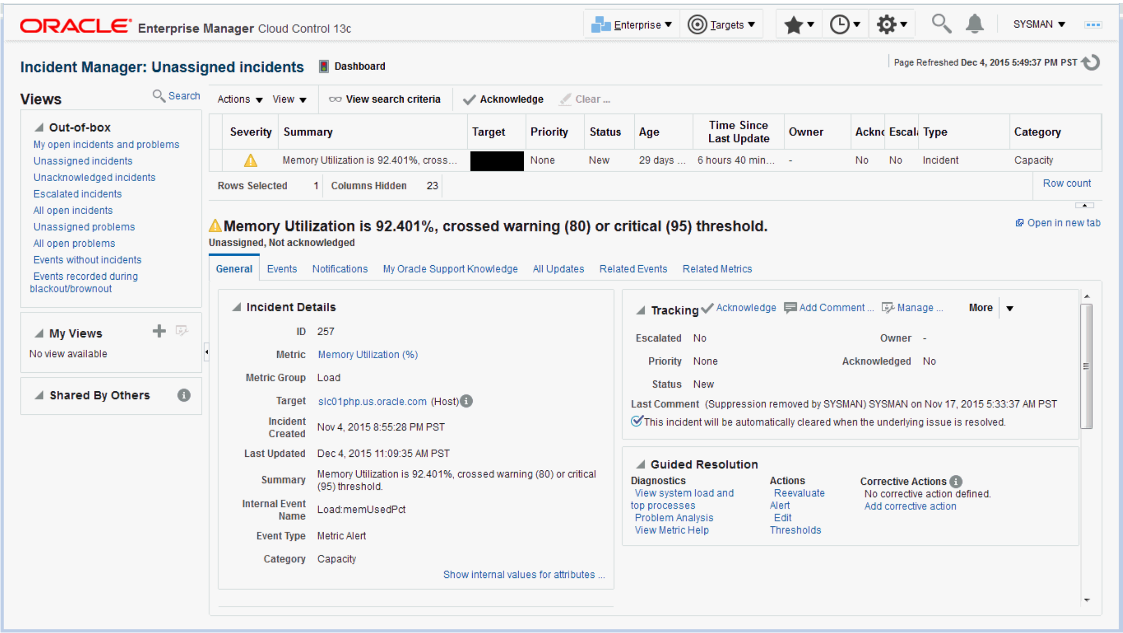

Incident Manager: Unassigned incidents

Page Refreshed Dec 4, 2015 5:49:37 PM PST

No view available

Severity	Summary	Target	Priority	Status	Age	Time Since Last Update	Owner	Ackno	Escal	Type	Category
⚠	Memory Utilization is 92.401%, cross...	[redacted]	None	New	29 days ...	6 hours 40 min...	-	No	No	Incident	Capacity

Rows Selected 1 Columns Hidden 23

⚠ Memory Utilization is 92.401%, crossed warning (80) or critical (95) threshold.
Unassigned, Not acknowledged

Incident Details

ID 257
Metric Memory Utilization (%)
Metric Group Load
Target slc01php.us.oracle.com (Host)
Incident Created Nov 4, 2015 8:55:28 PM PST
Last Updated Dec 4, 2015 11:09:35 AM PST
Summary Memory Utilization is 92.401%, crossed warning (80) or critical (95) threshold.
Internal Event Name Load:memUsedPct
Event Type Metric Alert
Category Capacity

Tracking

Escalated No
Owner -
Priority None
Acknowledged No
Status New
Last Comment (Suppression removed by SYSMAN) SYSMAN on Nov 17, 2015 5:33:37 AM PST
✓ This incident will be automatically cleared when the underlying issue is resolved.

Guided Resolution

Diagnostics
View system load and top processes
Problem Analysis
View Metric Help

Actions
Reevaluate Alert
Edit Thresholds

Corrective Actions
No corrective action defined.
Add corrective action

Enterprise Manager Console

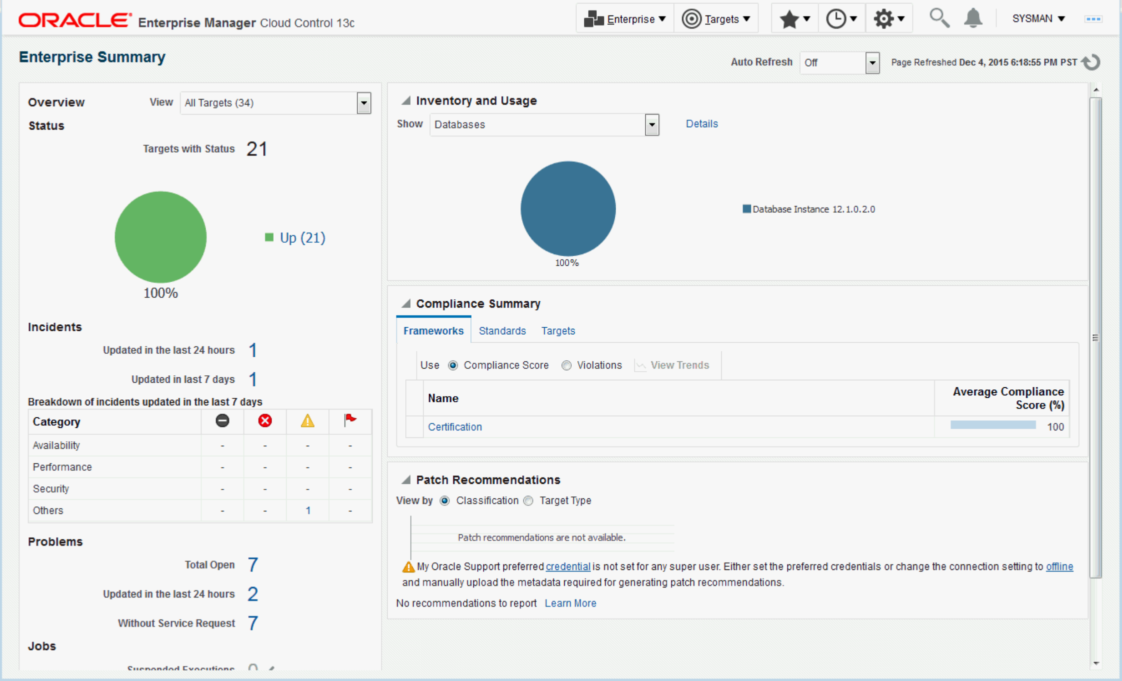
ORACLE Enterprise Manager Cloud Control 13c
Enterprise ▾
Targets ▾
SYSMAN ▾
100%
Database Instance 12.1.0.2.0
100%

Enterprise Summary

Auto Refresh: Off

Page Refreshed Dec 4, 2015 6:18:55 PM PST

Overview

View: All Targets (34)

Status

Targets with Status 21

Up (21)

Incidents

Updated in the last 24 hours 1

Updated in last 7 days 1

Breakdown of incidents updated in the last 7 days

Category	⊖	⊗	⚠	⚑
Availability	-	-	-	-
Performance	-	-	-	-
Security	-	-	-	-
Others	-	-	1	-

Problems

Total Open 7

Updated in the last 24 hours 2

Without Service Request 7

Jobs

Suspended Executions 0

Inventory and Usage

Show: Databases   Details

Compliance Summary

Frameworks   Standards   Targets

Use ● Compliance Score ○ Violations   View Trends

Name	Average Compliance Score (%)
Certification	100

Patch Recommendations

View by ● Classification ○ Target Type

Patch recommendations are not available.

⚠ My Oracle Support preferred credential is not set for any super user. Either set the preferred credentials or change the connection setting to offline and manually upload the metadata required for generating patch recommendations.

No recommendations to report   Learn More

Searching for Incidents

You can search for incidents based on a variety of incident attributes such as the time incidents were last updated, target name, target type, or incident status.

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**.

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

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

Searching for Incidents by Target Lifecycle Status

In addition to searching for incidents using high-level incident attributes, you can also perform more granular searches based on individual target lifecycle status. Briefly, lifecycle status is a target property that specifies a target's operational status. Status options for which you can search are:

- All

- Mission Critical
- Production
- Staging
- Test
- Development

For more discussion on lifecycle status, see Event Prioritization.

To search for incidents by target lifecycle status:

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**.
3. In the **Search** region, click **Add Fields**. A pop-up menu appears showing the available lifecycle statuses.
4. Choose on one or more of the lifecycle status options.
5. Enter any additional search criteria.
6. Click **Get Results**.

Setting Up Custom Views

Incident Manager also allows you to define custom 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 critical database incidents that you own. By specifying and saving view preferences to display only those incident attributes that you are interested in Enterprise Manager will show only the list of matching incidents.

You can then search the incidents for only the ones with specific attributes, such as priority 1. The view allows easy access to pertinent incidents for daily triage. Accordingly, you can save the search criteria as a filter named "All priority 1 incidents for my targets". The view becomes available in the UI for immediate use and will be available anytime you log in to access the specific incidents. The last view you used will be the default view used on your next login.

How to create a view in Incident Manager - YouTube Video

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 **MyViews** region located on the left, click the create "+" icon.
   1. In the **Search** region, search for Incidents using the **Type** list and select **Incidents**.
   2. 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.
   3. After you have provided the appropriate criteria, click **Get Results**.

      Validate that the list of incidents match what you are looking for. If not, change the search criteria as needed.
   4. To view all the columns associated with this table, in the **View** menu, select **Columns**, then select **Show All**.

To select a subset of columns to display and also the order in which to display them, from the **View** menu, select **Columns**, then **Manage Columns**. A dialog displays showing a list of columns available to be added in the table.

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

f. Enter the view name. If you want other administrators to use this view, check the **Share** option.

g. Click **OK** to save the view.

> **Note:**
>
> From the View creation dialog, you can also mark the view as shared. See Sharing/Unsharing Custom Views for more information.

Incident Dashboard

An incident dashboard allows you to track and monitor the state of different aspects of incident management, such as getting a sense of how the incidents are distributed. Specifically, the incident dashboard presents another way of looking at a scoped set of incidents using custom views (one dashboard per view). In addition to providing an intuitive way to interpret incidents and incident distribution patterns, the dashboard provides you with quick and easy access to incident lifecycle actions such as acknowledging, assigning, and adding comments.

An incident dashboard is shown in the following figure.

Figure 5-5 Incident Dashboard

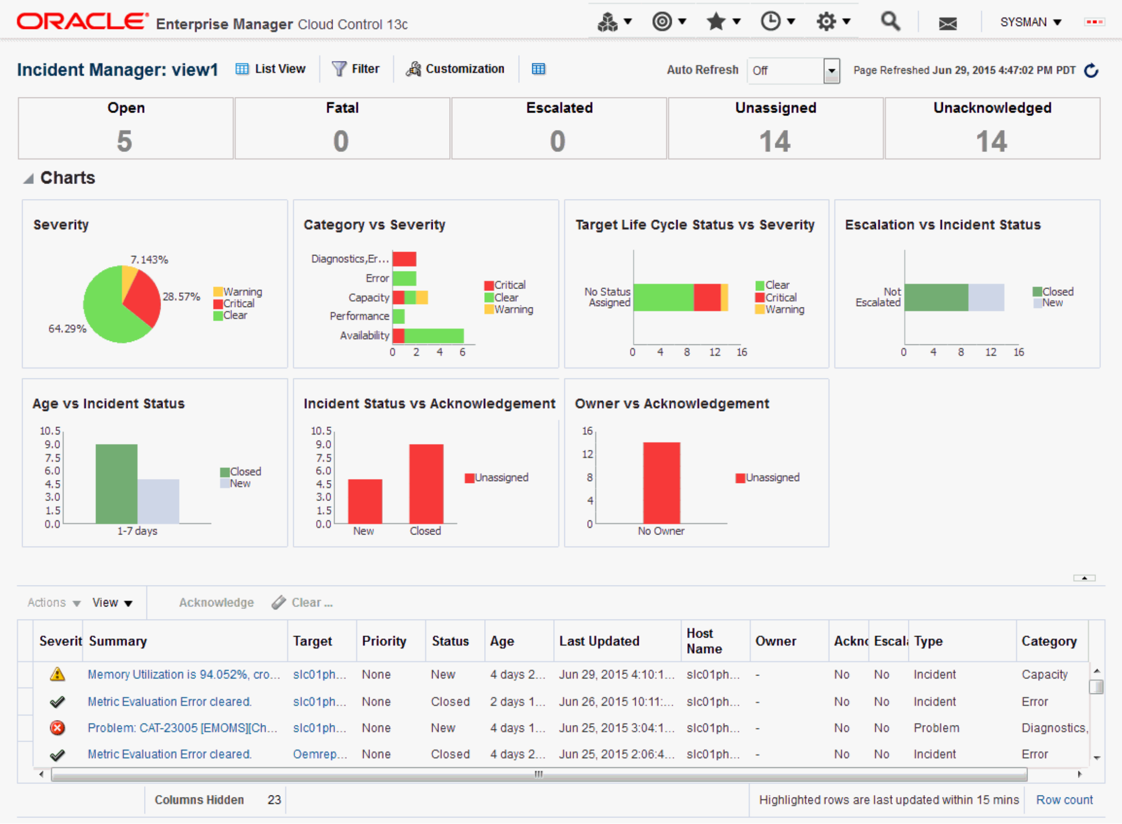

Incident Manager: view1

Auto Refresh: Off

Page Refreshed Jun 29, 2015 4:47:02 PM PDT

Open	Fatal	Escalated	Unassigned	Unacknowledged
5	0	0	14	14

Severit	Summary	Target	Priority	Status	Age	Last Updated	Host Name	Owner	Ackno	Escal	Type	Category
⚠	Memory Utilization is 94.052%, cro...	slc01ph...	None	New	4 days 2...	Jun 29, 2015 4:10:1...	slc01ph...	-	No	No	Incident	Capacity
✓	Metric Evaluation Error cleared.	slc01ph...	None	Closed	2 days 1...	Jun 26, 2015 10:11:...	slc01ph...	-	No	No	Incident	Error
⊗	Problem: CAT-23005 [EMOMS][Ch...	slc01ph...	None	New	4 days 1...	Jun 25, 2015 3:04:1...	slc01ph...	-	No	No	Problem	Diagnostics,
✓	Metric Evaluation Error cleared.	Oemrep...	None	Closed	4 days 2...	Jun 25, 2015 2:06:4...	slc01ph...	-	No	No	Incident	Error

Columns Hidden 23

Highlighted rows are last updated within 15 mins

Row count

Incident Dashboard Areas

The content of the incident dashboard is divided into three sections:

> **Note:**
>
> By default, data is automatically refreshed every 30 seconds. You can increase or decrease the refresh interval as required.

- **Summary** area displays the number of:
  - **Open** incidents including those created in the last hour.
  - **Fatal** incidents including those created or updated to Fatal in the last hour.
  - **Escalated** incidents including those escalated in the last hour.
  - **Unassigned** incidents.
  - **Unacknowledged** incidents.

  Incident dashboard elements that are highlighted in red require immediate attention. Clicking on the summary numbers allows you to view only incidents pertaining to that incident area. Data displayed in the charts and incident list are modified accordingly.
- **Charts** provide you with an easy-to-understand look at the current incident distribution and management status for each incident. You can click on slices of the charts to filter the data displayed in the incident dashboard only to those incidents.
- **Incident List** that shows the open incidents listed in reverse chronological order by last updated time stamp. From this list, you can perform requisite incident lifecycle actions such as escalating, prioritizing, acknowledging, assigning owners, adding comments to incident.

Creating an Incident Dashboard

To create an incident dashboard for all open incidents, navigate to Incident Manager and click the **Dashboard** button located at the upper-left side of the page above the incident table.

While an open incident dashboard can be useful, a more typical scenario involves creating a custom view so that the incident dashboard only displays data for incidents that are of interest to you. See "Setting Up Custom Views" for information on creating a custom view. To view an incident dashboard for a specific view, select the desired view from the **My Views** list in the Incident Manager UI and then click **Dashboard**.

Customizing the Incident Dashboard

If the default dashboard does not meet your requirements, you can modify the dashboard in a variety of ways, such as removing an out-of-box chart, adding one of the predefined charts, or altering an existing chart, such as changing the dimensions of the chart or changing from a pie chart to a bar chart, for example.

Click **Customization** located above the *Summary* section in order to customize the Incident Dashboard. Once changes are made, click **Save** to save changes.

> **Note:**
>
> Only the view owner and Super Administrators can create or edit view customizations. Without view ownership, users can only change the chart auto-refresh frequency and chart type. These changes, however, are not permanent.

Sharing/Unsharing Custom Views

When you create your own views, they are private (only you can see them). Beginning with Enterprise Manager Release 12.1.0.4, you can share your private views with other administrators. When you share a view, all Enterprise Manager users will be able to use the view.

As mentioned previously, you are given the opportunity to share a view during the view creation process. If you have already created custom views, you can share them at any time.

1. Navigate to Incident Manager.

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

2. From the My Views region, click the Manage icon.

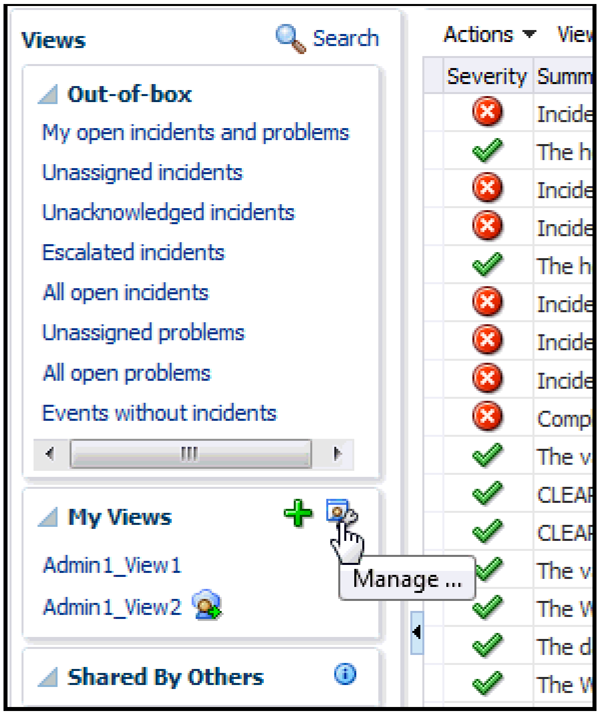

3. From the Manage Custom Views dialog, choose a custom view.
4. Click **Share** (or **Unshare** if the view is already shared and you want to unshare it.)
5. Click **Yes** to confirm the share/unshare operation.

Responding and Working on a Simple Incident

The following steps take you through one possible incident management scenario.

1. Navigate to Incident Manager.

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

2. Use a view to filter the list of incidents. For example, you should use **My Open Incidents and Problems** view to see incidents and problems assigned to you. You can then sort the list by priority.

3. To work on an incident, select the incident. In the **General** tab, click **Acknowledge** to indicate that you are working on this incident, and to stop receiving repeat notifications for the incident.

   In addition to the acknowledging the incident, you can perform other incident management operations such as:

   - Adding a comment.
   - Managing the incident. See Responding to and Managing Multiple Incidents, Events and Problems in Bulk for more information on incident management options.
   - Editing the summary.
   - Manually creating a ticket.
   - Suppressing/unsuppressing the incident.
   - Clearing the incident.

   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.

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 incident.

Alternatively, you can work with incidents for a specific target from that target's home page. From the *target* menu, select **Monitoring** and then select **Incident Manager** to access incidents for that target (or group).

Responding to and Managing Multiple Incidents, Events and Problems in Bulk

There may be situations where you want to respond to multiple incidents in the same way. For example, you find that a cluster of incidents that are assigned to you are due to insufficient tablespace issues on several production databases. Your manager suggests that these tablespaces be transferred to a storage system being procured by another administrator. In this situation, you want to set all of the tablespace incidents to a customized resolution state "Waiting for Hardware." You also want to assign the incidents to the other administrator and add a comment to explain the scenario. In this situation, you want to update all of these incidents in bulk rather than individually.

To respond to incidents in bulk:

1. Navigate to Incident Manager.

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

2. Use a view to filter the list of incidents to the subset of incidents you want to work on. For example, you can use **My Open Incidents and Problems** view to see incidents and problems assigned to you. You can then sort the list by priority.

3. Select the incidents to which you want to respond. You can select multiple incidents by holding down the Control key and selecting individual incidents or you can hold down the Shift key and select the first and last incidents to select a contiguous block of incidents.

4. From the **Action** menu, choose the desired response action.

   - **Acknowledge**: Indicate that you have viewed the incidents. This option also stops any repeat notifications sent out for the incidents. This sets the *Acknowledged* flag to *Yes* and also makes you the owner of the incident

   - **Manage**: Allows you to perform a multi-action response to the incidents.

     - *Acknowledge*: If an incident is acknowledged, it will be implicitly assigned to the user who acknowledged it. When a user assigns an incident to himself, it is considered acknowledged. Once acknowledged, an incident cannot be unacknowledged. Acknowledgement also stops any repeat notifications for that incident

     - *Assign to*: Assign the incident(s) to the administrator who will take ownership of the incident.

     - *Prioritization*: The priority level of an incident can be set by selecting one of the out-of-the-box priority values: None, Urgent, Very High, High, Medium, Low

     - *Incident Status*: The resolution state for the incident can be set by selecting either Work in Progress or Resolved or to any custom status defined.

     - *Escalation Level*: Administrators can update incidents to set an escalation level: Level 1 through 5, in addition to the default value of None. An escalated issue can be de-escalated by setting the escalation to None. The appropriate *Escalation Level* depends on the IT procedures you have in place.

     - *Comment*: You can enter comments such as those you want to pass to the owner of the incident.

   - **Suppress**: Suppressing an incident stops corresponding notifications, and removes it from out-of-the-box views and default totals (such as those presented in the summary region). Suppression is typically performed when you want to defer action on the incident until a future time and in the meantime want to visually hide them from appearing in the console. Administrators can see suppressed incidents by explicitly searching for them such as performing a search on incidents where the search criteria includes the *Suppressed* search field

     Incidents can be suppressed until any of the following conditions are met:

     - Until the suppression is manually removed

     - Until specified date in the future

     - Until the severity state changes (incidents only)

     - Until it is closed

   - **Clear**: Administrators can clear incidents or problems manually. For incidents, this applies only to incidents containing incidents that can be manually cleared.

- **Add Comment**: Users can add comments on incidents and events. Comments may be used for sharing information with other users or to provide tracking information on any actions being taken. Comments can be added even on closed issues.

> **Note:**
>
> The single action **Acknowledge** and **Clear** buttons are enabled for open incidents and can be used for multiple incident selection.

If any of the above actions applies only to a subset of selected incidents (for example, if an administrator tries to acknowledge multiple incidents, of which some are already acknowledged), the action will be performed only where applicable. The administrator will be informed of the success or failure of the action. When an administrator selects any of these actions, a corresponding annotation is added to the incident for future reference.

5. Click **OK**. Enterprise Manager displays a process summary and confirmation dialogs.
6. Continue working with the incidents as required.

Searching My Oracle Support Knowledge

To access My Oracle Support Knowledge base entries from within Incident Manager, 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. 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**.

   If your My Oracle Support (MOS) login credentials have been saved as MOS Preferred Credentials, you do not need to log in manually. If not, you will need to sign in to My Oracle Support. To save your MOS login information as Preferred Credentials.

   Setting MOS Preferred Credentials: From the **Setup** menu, select **Security** and then **Preferred Credentials**. From the My Oracle Support Preferred Credentials region, click **Set MOS Credentials**.
4. On the My Oracle Support page, click the **Knowledge** tab to browse the knowledge base.

   From this page, in addition to accessing formal Oracle documentation, you can also change the search string in to look for additional knowledge base entries.

Submitting an Open Service Request (Problems-only)

There are times when you may need assistance from Oracle Support to resolve a *problem*. This procedure is not relevant for incidents or events.

To submit a service request (SR), 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. Use one of the views to find the problem or search for it or use one of your custom views. Select the appropriate problem from table.
3. Click on the **Support Workbench: Package Diagnostic** link.
4. Complete the workflow for opening an SR. Upon completing the workflow, a draft SR will have been created.
5. Sign in to My Oracle Support if you are not already signed in.
6. On the My Oracle Support page, click the **Service Requests** tab.
7. Click **Create SR** button.

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 need to defer work on the incident until a future date (for example, until maintenance window). In order to avoid having it appear in the UI, you want to temporarily hide or suppress the incident until a future date. In order to find a suppressed incident, you must explicitly search for the incident using either the *Show all* or the *Only show suppressed* search option. In order to unhide a suppressed incident or problem, it must be manually unsuppressed.

To suppress an incident or problem:

1. Navigate to Incident Manager.

   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.
5. Click **OK**.

To unsuppress an incident or problem:

1. Navigate to Incident Manager.

   From the **Enterprise** menu on the Enterprise Manager home page, select **Monitoring**, then select **Incident Manager**.
2. Click **Search**.
3. From the **Suppressed** menu, select **Only show suppressed**.
4. Click **Get Results**.

   The suppressed incidents are displayed. Choose the appropriate incident or problem.
5. Click the **General** tab.
6. In the resulting details region, click **More**, and then select **Unsuppress**.

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, reassign 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 reassigns 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.

Reviewing Events on a Periodic Basis

Oracle recommends managing via incidents in order to focus on important events or groups of related events. Due to the variety and sheer number of events that can be generated, it is possible that not all important events will be covered by incidents. To help you find these important yet untreated events, Enterprise Manager provides the **Events without incidents** standard view.

Perform the following steps:

1. From the **Enterprise** menu, select **Monitoring**, then select **Incident Manager**.
2. In the Views region, click **Events without incidents**.
3. Select the desired event in the table. The event details display.
4. In the details area, choose **More** and then either **Create Incident** or **Add Event to Incident**.

Example Scenario

During the initial phase of Enterprise Manager uptake, every day the DBA manager reviews the events for 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.

Creating an Incident Manually

If an event of interest occurs that is not covered by any rule and you want to convert that event to an incident, perform the following:

1. Using an available view, find the event of interest.
2. Select the event in the table.
3. From the **More...** drop-down menu, choose **Create Incident...**

4. Enter the incident details and click **OK**.
5. Should you decide to work on the incident, 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 email 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 email. 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".

Advanced Topics

The following sections discuss incident/event management features relating advanced applications or operational areas.

Automatic Diagnostic Repository (ADR): Incident Flood Control

ADR is a file-based repository that stores database diagnostic data such as traces, dumps, the alert log, and health monitor reports. ADR's unified directory structure and a unified set of tools enable customers and Oracle Support to correlate and analyze diagnostic data across multiple instances and Oracle products.

Like Enterprise Manager, ADR creates and tracks incidents and problems to allow you to resolve issues.

- A *problem* is a critical error in the database. Critical errors manifest as internal errors, such as ORA-00600, or other severe errors, such as ORA-07445 (operating system exception) or ORA-04031 (out of memory in the shared pool).
- An *incident* is a single occurrence of a problem. When a problem (critical error) occurs multiple times, an incident is created for each occurrence. Incidents are timestamped and tracked in ADR. When an incident occurs, ADR sends a diagnostic incident alert to Enterprise Manager.

Working with ADR Diagnostic Incidents Using Incident Manager

Each diagnostic incident recorded in the ADR is also recorded as an incident in Enterprise Manager, thus providing you with a unified view of ADR/Enterprise Manager incidents and problems from within Incident Manager. For the ADR diagnostic incidents, you can access Enterprise Manager Support Workbench to take further action, such as packaging a problem or raising a service request with Oracle Support.

Incident Flood Control

Prior to Enterprise Manager Release 12.1.0.4, there was no limit to the number of diagnostic incidents recorded for a single problem in Incident Manager. It is conceivable that a problem could generate dozens or perhaps hundreds of incidents in a short period of time. While incidents generated during the early stages of a problem may be useful, after a certain point the excess diagnostic data would provide little value and possibly slow down your efforts to diagnose and resolve the problem. Because diagnostic problems typically tend to be longlived, a significant number of incidents could be generated over time. Also, depending on the size of your monitored environment, the diagnostic data may consume considerable system resources.

For these reasons, the Enterprise Manager applies flood control limits on the number of diagnostic incidents that can be raised for a given problem in Incident Manager. Flood-controlled incidents provide a way of informing you that a critical error is ongoing, without overloading the system with diagnostic data.

Beginning with Enterprise Manager Release 12.1.0.4, two limits are placed on the number of diagnostic incidents that can be raised for a given problem in Incident Manager. A problem is identified by a unique problem signature called a problem key and is associated with a single target.

Enterprise Manager Limits on Diagnostic Incidents

Enterprise Manager enforces two limits for diagnostic incidents:

- For any given *hour*, Enterprise Manager only records up to five (default value) diagnostic incidents for a given target and problem key combination.
- On any given *day*, Enterprise Manager only records up to 25 (default value) diagnostic incidents for a given problem key and target combination.

When either of these limits is reached, any diagnostic incidents for the same target/problem key combination will not be recorded until the corresponding hour or day is over. Diagnostic incident recording will commence once a new hour or day begins.

> **Note:**
>
> Hour and day calculations are based on UTC (or GMT).

These diagnostic incident limits only apply to Incident Manager and not to the underlying ADR. All incidents continue to be recorded in the ADR repository. Using Enterprise Manager Support Workbench, users can view all the incidents for a given problem at any time and take appropriate actions.

Enterprise Manager diagnostic incident limits are configurable. As mentioned earlier, the defaults for these two limits are set to 5 incidents per hour and 25 incidents per day. These defaults should not be changed unless there is a clear business reason to track all diagnostic incidents.

Changing Enterprise Manager Diagnostic Incident Limits

To update the diagnostic limits, execute the following SQL against the Enterprise Manager repository as the SYSMAN user using the appropriate limit values as shown in the following example.

The PL/SQL shown in the following example prints out the current limits.

> **Note:**
>
> The Enterprise Manager incident limits are **in addition to** any diagnostic incident limits imposed by underlying applications such as Oracle database, Middleware and Fusion Applications. These limits are specific to each application. See the respective application documentation for more information.

Example 5-3 SQL Used to Change Diagnostic Incident Limits

```
exec  EM_EVENT_UTIL.SET_ADR_INC_LIMITS(5,25);
```

Example 5-4 SQL Used to Print Out Current Diagnostic Incident Limits

```
DECLARE
  l_adr_hour_limit NUMBER;
  l_adr_day_limit NUMBER;
BEGIN
      em_event_util. GET_ADR_INC_LIMITS
              (p_hourly_limit => l_adr_hour_limit,
               p_daily_limit => l_adr_day_limit);
      dbms_output.put_line(l_adr_hour_limit || '-' || l_adr_day_limit);
END;
```

Defining Custom Incident Statuses

As discussed in "Working with Incidents", one of the primary incident workflow attributes is *status*. For most conditions, these predefined status attributes will suffice. However, the uniqueness of your monitoring and management environment may require an incident workflow requiring specialized incident states. To address this need, you can define custom states using the *create_resolution_state* EM CLI verb.

Creating a New Resolution State

```
emcli create_resolution_state
     -label="Label for display"
     -position="Display position"
     [-applies_to="INC|PBLM"]
```

This verb creates a new resolution state for describing the state of incidents or problems.

> **Note:**
>
> This command can only be executed by Enterprise Manager Super Administrators.

The new state is always added between the *New* and *Closed* states. You must specify the exact position of this state in the overall list of states by using the `-position` option. The position can be between 2 and 98.

By default, the new state is applicable to both incidents and problems. The `-applies_to` option can be used to indicate that the state is applicable only to incidents or problems.

A success message is reported if the command is successful. An error message is reported if the change fails.

Examples

The following example adds a resolution state that applies to both incidents and problems at position 25.

```
emcli create_resolution_state -label="Waiting for Ticket" -position=25
```

The following example adds a resolution state that applies to problems only at position 35.

```
emcli create_resolution_state -label="Waiting for SR" -position=35 -
applies_to=PBLM
```

Modifying an Existing Resolution State

You can chance the both the display label and the position of an existing state by using the *modify_resolution_state* verb.

```
emcli modify_resolution_state
        -label="old label of the state to be changed"
        -new_label="New label for display"
        -position="New display position"
        [-applies_to=BOTH]
```

This verb modifies an existing resolution state that describes the state of incidents or problems. As with the create_resolution_state verb, this command can only be executed by Super Administrators.

You can optionally indicate that the state should apply to both incidents and problems using the `-applies_to` option.

Examples

The following example updates the resolution state with old label "Waiting for TT" with a new label "Waiting for Ticket" and if necessary, changes the position to 25.

```
emcli modify_resolution_state -label="Waiting for TT" -new_label="Waiting for
Ticket" -position=25
```

The following example updates the resolution state with the old label "SR Waiting" with a new label "Waiting for SR" and if necessary, changes the position to 35. It also makes the state applicable to incidents and problems.

```
emcli modify_resolution_state -label="SR Waiting" -new_label="Waiting for SR" -
position=35 -applies_to=BOTH
```

Clearing Stateless Alerts for Metric Alert Event Types

For *metric alert* event types, an event (metric alert) is raised based on the metric threshold values. These metric alert events are called *stateful* alerts. For those metric alert events that are not tied to the state of a monitored system (for example, *snapshot too old*, or *resumable session suspended*), these alerts are called stateless alerts. Because stateless alerts are not cleared automatically, they need to be cleared manually. You can perform a bulk purge of stateless alerts using the *clear_stateless_alerts* EM CLI verb.

> **Note:**
>
> For large numbers of incidents, you can manually clear incidents in bulk. See "Responding to and Managing Multiple Incidents, Events and Problems in Bulk".

clear_stateless_alerts clears the stateless alerts associated with the specified target. The clearing must be manually performed as the Management Agent does not automatically clear stateless alerts. To find the metric internal name associated with a stateless alert, use the EM CLI *get_metrics_for_stateless_alerts* verb.

Format

```
emcli clear_stateless_alerts -older_than=number_in_days -target_type=target_type -
target_name=target_name [-include_members][-
metric_internal_name=target_type_metric:metric_name:metric_column] [-unacknowledged_only]
[-ignore_notifications] [-preview][ ] indicates that the parameter is optional
```

Options

- **older_than**

  Specify the age of the alert in days. (Specify 0 for currently open stateless alerts.)
- **target_type**

  Internal target type identifier, such as host, oracle_database, and emrep.
- **target_name**

  Name of the target.
- **include_members**

  Applicable for composite targets to examine alerts belonging to members as well.
- **metric_internal_name**

  Metric to be cleaned up. Use the get_metrics_for_stateless_alerts verb to see a complete list of supported metrics for a given target type.
- **unacknowledged_only**

  Only clear alerts if they are not acknowledged.
- **ignore_notifications**

  Use this option if you do not want to send notifications for the cleared alerts. This may reduce the notification sub-system load.
- **ignore_notifications**

  Use this option if you do not want to send notifications for the cleared alerts. This may reduce the notification sub-system load.
- **preview**

  Shows the number of alerts to be cleared on the target(s).

Example

The following example clears alerts generated from the database alert log over a week old. In this example, no notifications are sent when the alerts are cleared.

```
emcli clear_stateless_alerts -older_than=7 -target_type=oracle_database -tar
get_name=database -metric_internal_name=oracle_database:alertLog:genericErrStack -
ignore_notifications
```

Automatically Clearing "Manually Clearable" Events

There are those events that clear automatically, such as CPU Utilization and those events that must be manually cleared, either through the Incident Manager UI or automatically via rule (such as Job Failure, or Log Metric events). Auto-clear events, as the term implies, are cleared automatically by Enterprise Manager once the underlying issue is resolved. In the case of CPU Utilization, the event CPU Utilization clears automatically once the percent utilization falls below the warning threshold. However, for those events that must be cleared manually, a user must intervene and clear the event using Incident Manager either by selecting the incident/ event and clicking **Clear**, or creating an event rule to do the job (recommended method).

As mentioned previously, an event rule automates the clearing of *manually clearable* events. Enterprise Manager provides a limited number of out-of-box rules that automatically clear *manually clearable* events, such as job failures or ADP events that remain open for seven days. However, to more accurately meet the needs of your monitoring environment, Oracle recommends creating your own event rules to automatically clear those *manually clearable* events that are most prevalent in your environment.

During the rule creation process, you can specify that an event be automatically cleared by selecting the **Clear Event** option while you are adding conditional actions.

Getting Notified when the Event Clears

The event clearing action is an asynchronous operation, which means that when the rule action (clear) is initiated, the manually clearable event will be enqueued for clearing, but not actually cleared. Hence, an email notification sent upon rule execution will indicate that the event has not been cleared. Asynchronous clearing is by design as it reduces overall rule engine processing load and processing time. Subscribing to this event clearing rule with the intent to be notified when the event clears will be of little value. If you want to be notified when the event clears, you must create a new event rule and explicitly specify a *Clear* severity. In doing so, you will be notified once the event is actually cleared.

User-reported Events

Users may create (publish) events manually using the EM CLI verb *publsh_event*. A User-reported event is published as an event of the "User-reported event" class. Only users with Manage Target privilege can publish these events for a target. An error message is reported if the publish fails.

After an event is published with a severity other than CLEAR (see below), end-users with appropriate privileges can manually clear the event from the UI, or they can publish a new event using a severity level of CLEAR and the same details to report clearing of the underlying situation.

Format

```
emcli publish_event
        -target_name="Target name"
        -target_type="Target type internal name"
        -message="Message for the event"
        -severity="Severity level"
        -name="event name"
        [-key="sub component name"
         -context="name1=value1;name2=value2;.."
```

```
    -separator=context="alt. pair separator"
    -subseparator=context="alt. name-value separator"]

[ ] indicates that the parameter is optional
```

Options

- **target_name**

  Target name.
- target_type

  Target type name.
- message

  Message to associate for the event. The message cannot exceed 4000 characters.
- **severity**

  Numeric severity level to associate for the event. The supported values for severity level are as follows:

  "CLEAR"
  "MINOR_WARNING"
  "WARNING"
  "CRITICAL"
  "FATAL"
- **name**

  Name of the event to publish. The event name cannot exceed 128 characters.

  This is indicative of the nature of the event. Examples include "Disk Used Percentage," "Process Down," "Number of Queues," and so on. The name must be repeated and identical when reporting different severities for the same sequence of events. This should not have any identifying information about a specific event; for example, "Process xyz is down." To identify any specific components within a target that the event is about, see the key option below.
- **key**

  Name of the sub-component within a target this event is related to. Examples include a disk name on a host, name of a tablespace, and so forth. The key cannot exceed 256 characters.
- **context**

  Additional context that can be published for a given event. This is a series of strings of format name:value separated by a semi-colon. For example, it might be useful to report the percentage size of a disk when reporting space issues on the disk. You can override the default separator ":" by using the sub-separator option, and the pair separator ";" by using the separator option.

  The context names cannot exceed 256 characters, and the values cannot exceed 4000 characters.
- **separator**

  Set to override the default ";" separator. You typically use this option when the name or the value contains ";". Using "=" is not supported for this option.
- **subseparator**

Set to override the default ":" separator between the name-value pairs. You typically use this option when the name or value contains ":". Using "=" is not supported for this option.

Examples

Example 1

The following example publishes a warning event for "my acme target" indicating that a HDD restore failed, and the failure related to a component called the "Finance DB machine" on this target.

```
emcli publish_event  -target_name="my acme target" -target_type="oracle_acme"
-name="HDD restore failed" -key="Finance DB machine" -message="HDD restoration
failed due to corrupt disk" -severity=WARNING
```

Example 2

The following example publishes a minor warning event for "my acme target" indicating that a HDD restore failed, and the failure related to a component called the "Finance DB machine" on this target. It specifies additional context indicating the related disk size and name using the default separators. Note the escaping of the \ in the disk name using an additional "\".

```
emcli publish_event  -target_name="my acme target" -target_type="oracle_acme"
-name="HDD restore failed" -key="Finance DB machine" -message="HDD restoration
failed due to corrupt disk" -severity=MINOR_WARNING -context="disk size":800GB\;"disk
name":\\uddo0111245
```

Example 3

The following example publishes a critical event for "my acme target" indicating that a HDD restore failed, and the failure related to a component called the "Finance DB machine" on this target. It specifies additional context indicating the related disk size and name. It uses alternate separators, because the name of the disk includes the ":" default separator.

```
emcli publish_event  -target_name="my acme target" -target_type="oracle_acme"
-name="HDD restore failed" -key="Finance DB machine" -message="HDD restoration
failed due to corrupt disk" -severity=CRITICAL -context="disk size"^800GB\;"disk name"^\
\sdd1245:2 -subseparator=context=^
```

Additional Rule Applications

Rules can be set up to perform more complicated tasks beyond straightforward notifications. The following tasks illustrate additional rule capabilities.

- Setting Up a Rule to Send Different Notifications for Different Severity States of an Event
- Creating a Rule to Notify Different Administrators Based on the Event Type
- Creating a Rule to Create a Ticket for Incidents
- Creating a Rule to Send SNMP Traps for Events to Third Party Systems

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

Before you perform this task, ensure the DBA has set appropriate thresholds for the metric so that a critical metric alert is generated as expected.

Consider the following example:

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 email 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.

Perform the following tasks to set appropriate thresholds:

1. From the **Setup** menu, 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. Click **Create...**
4. From the Select Type of Rule to Create dialog, choose **Incoming events and updates to events**. Click **Continue**.
5. Provide the rule details as follows:
    a. For Type, select **Metric Alerts** as the Type.
    b. In the criteria section, select **Severity**. From the drop-down list, check and **Critical** and **Warning** as the selected values. Click **Next**.
    c. On the Add Actions page, click **+Add**.

       In the Create Incident section, check the **Create Incident** option. Click **Continue**. The Add Action page displays with the new rule. Click **Next**.
    d. Specify a name for the rule and a description. Click **Next**.
    e. On the Review page, ensure your settings are correct and click **Continue**. A message appears informing you that the rule has been successfully created. Click **OK** to dismiss the message.

       Next, you need to create a rule to perform the notification actions.
6. From the Rules section on the Edit Rules page, click **Create**.
7. Select **Newly created incidents or updates to incidents** as the rule type and click **Continue**.
8. Check **Specific Incidents.**
9. Check **Severity** and from the drop-down option selector, check **Critical** and **Warning**. Click **Next**.
10. On the Add Actions page, click **Add**. The Conditional Actions page displays.
11. In the **Conditions for actions** section, choose **Only execute the actions if specified conditions match**.
12. From the **Incident matches the following criteria** list, choose **Severity** and then **Critical** from the drop-down option selector.
13. In the **Notifications** section, enter the DBA in the **Page** field. Click **Continue**. The Add Actions page displays.
14. Click **Add** to create a new action for the Warning severity.
15. In the **Conditions for actions** section, choose **Only execute the actions if specified conditions match**.
16. From the **Incident matches the following criteria** list, choose **Severity** and then **Warning** from the drop-down option selector.

17. In the **Notifications** section, enter the DBA in the **Email to** field. Click **Continue**. The Add Actions page displays with the two conditional actions. Click **Next**.
18. Specify a rule name and description. Click **Next**.
19. On the Review page, ensure your rules have been defined correctly and click **Continue**. The Edit Rule Set page displays.
20. Click **Save** to save your newly defined rules.

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

As per operations policy for production databases, the incidents that relate to application issues should go to the application DBAs and the incidents 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 email.

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.
- Rule has been setup to create incident for all such events.
- Respective notification setup is complete, for example, global SMTP gateway, email address, and schedule for individual DBAs.

Perform the following steps:

1. Navigate to the Incident Rules page.

   From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. Search the list of enterprise rules matching the events from the production database.
3. 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.
4. From the Edit Rule Set page (Rules tab), select the rule which creates the incidents for the metric alert events for the database. Click **Edit**
5. From the Select Events page, click **Next**.
6. From the Add Actions page, click **+Add**. The Add Conditional Actions page displays.
7. In the Notifications area, enter the email address of the DBA you want to be notified for this specific event type and click **Continue** to add the action. Enterprise Manager returns you to the Add Actions page. Click **Next**.
8. On the Specify Name and Description page, enter an intuitive rule name and a brief description.
9. Click **Next**.
10. On the Review page, review the **Applies to**, **Actions** and **General** information for correctness .
11. Click **Continue** to create the rule.
12. Create/Edit additional rules to handle alternate additional administrator notifications according to event type.
13. Review the rules summary and make corrections as needed. Click **Save** to save your rule set changes.

Creating a Rule to Create a Ticket for Incidents

If your IT process requires a helpdesk ticket be created to resolve incidents, then you can use the ticketing connector to associate the incident with a helpdesk ticket and have Enterprise Manager automatically open a ticket when the incident is created. Communication between Incident Manager and your ticketing system is bidirectional, thus allowing you to check the changing status of the ticket from within Incident Manager. Enterprise Manager also allows you to link out to a Web-based third-party console directly from the ticket so that you can launch the console in context directly from the ticket.

For example, according to the operations policy of an organization, all critical incidents from a production database should be tracked by way of ServiceNow tickets. A rule is set up to create a ServiceNow ticket when a critical incident occurs for the database. When such an incident occurs, the ticket is generated by the 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 ServiceNow to get the details about the ticket.

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

- Monitoring support has been set up.
- ServiceNow ticketing connector has been configured.

Perform the following steps:

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, select the appropriate rule set and select **Edit...**. (Rules are created in the context of a rule set. If there is no applicable rule set , create a new rule set.)
3. Select the appropriate rule that covers the incident conditions for which tickets should be generated and select **Edit...**
4. Select **Next** to proceed to the **Add Actions** page.
5. Select **+Add** to access the **Add Conditional Actions** page.
   a. Specify that a ticket should be generated for incidents covered by the rule.
   b. Specify the ticket template to be used.
6. Click **Continue** to return to the Add actions page.
7. On the Add Actions page, click **Next**.
8. On the Review page, click **Continue**.
9. On the Specify Name and Description page, click **Next**.
10. On the Review page, click **Continue**. A message displays indicating that the rule has been successfully modified. Click **OK** to close the message.
11. Repeat steps 3 through 10 until all appropriate rules have been edited.
12. Click **Save** to save your changes to the rule set.

Creating a Rule to Send SNMP Traps for Events to Third Party Systems

As mentioned in Using Notifications , Enterprise Manager supports integration with third-party management tools through the SNMP. Sending SNMP traps to third party systems is a two-step process:

Step 1: Create an advanced notification method based on an SNMP trap.

Step 2: Create an incident rule that invokes the SNMP trap notification method.

The following procedure assumes you have already created the SNMP trap notification method. For instruction on creating a notification method based on an SNMP trap, see "Sending SNMP Traps to Third Party Systems".

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, click **Create Rule Set...**
3. Enter the rule set **Name**, a brief **Description**, and select the type of source object the rule **Applies to** (Targets).
4. Click on the **Rules** tab and then click **Create...**
5. On the Select Type of Rule to Create dialog, select **Incoming events and updates to events** and then click **Continue**.
6. On the Create New Rule : Select Events page, specify the criteria for the events for which you want to send SNMP traps and then click **Next**.

   > **Note:**
   >
   > You must create one rule per event type. For example, if you want to send SNMP traps for Target Availability events and Metric Alert events, you must specify two rules.

7. On the Create New Rule : Add Actions page, click **Add**. The Add Conditional Actions page displays.
8. In the Notifications section, under Advanced Notifications, select an existing SNMP trap notification method.

   For information on creating SNMP trap notification methods, see "Sending SNMP Traps to Third Party Systems".

9. Click **Continue** to return to the Create New Rule : Add Actions page.
10. Click **Next** to go to the Create New Rule : Specify Name and Description page.
11. Specify a rule name and a concise description and then click **Next**.
12. Review the rule definition and then click **Continue** add the rule to the rule set. A message displays indicating the rule has been added to the rule set but has not yet been saved. Click **OK** to close the message.
13. Click **Save** to save the rule set. A confirmation is displayed. Click **OK** to close the message.

Creating a Rule to Send SNMP Traps for Incidents to Third Party Systems

As mentioned in Using Notifications , Enterprise Manager supports integration with third-party management tools through the SNMP. Sending SNMP traps for incidents to third party systems is a two-step process:

Step 1: Create an advanced notification method based on an SNMP trap.

Step 2: Create an incident rule that invokes the SNMP trap notification method.

The following procedure assumes you have already created the SNMP trap notification method. For instruction on creating a notification method based on an SNMP trap, see "Sending SNMP Traps to Third Party Systems".

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. On the **Incident Rules - All Enterprise Rules** page, click **Create Rule Set...**
3. Enter the rule set **Name**, a brief **Description**, and select the type of source object the rule applies to (Targets).
4. Click on the **Rules** tab and then click **Create**.
5. On the **Select Type of Rule to Create** dialog, select **Newly created incidents or updates** to incidents and then click **Continue**.
6. On the **Create New Rule: Select Incidents** page, specify the criteria for the incidents for which you want to send SNMP traps and then click **Next**.
7. On the **Create New Rule: Add Actions** page, click **Add**. The **Add Conditional Actions** page displays.
8. In the **Notifications** section, under **Advanced Notifications**, select an existing SNMP trap notification method.
   For information on creating SNMP trap notification methods, see "Sending SNMP Traps to Third Party Systems".
9. Click **Continue** to return to the **Create New Rule: Add Actions** page.
10. Click **Next** to go to the **Create New Rule: Specify Name and Description** page.
11. Specify a rule name and a concise description and then click **Next**.
12. Review the rule definition and then click **Continue** add the rule to the rule set. A message displays indicating the rule has been added to the rule set but has not yet been saved. Click **OK** to close the message.
13. Click **Save** to save the rule set. A confirmation is displayed. Click **OK** to close the message.

Exporting and Importing Incident Rules

You invest a great deal of time and effort carefully designing and testing the incident rule sets that automate Enterprise Manager incident management practices within your organization. Typically, the design and test phase of rule set creation is carried out in a separate Enterprise Manager test environment. Incident Manager's rule set import/export functionality simplifies moving rule sets from your development environment to your production environment.

In addition to moving rule sets from a test environment to a production environment, the import/export functionality also allows you to back up incident rule sets so they can be safely archived in case of disaster. More importantly, the import/export functionality makes it easy to standardize incident management automation processes across your Enterprise Manager environments.

Exporting Rule Sets using the Enterprise Manager Console

To export an incident rule set:

1. From the **Setup** menu, select **Incidents** then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, select the desired rule set you wish to export.

> **Note:**
>
> You cannot export Oracle-supplied out-of-box rule sets.

3. Click **Export**. Your browser's file dialog appears prompting you to save or open the file. Save the file to your local disk. By default the file name will be the name of your rule set with a.xml extension.

> **Note:**
>
> You should not edit the generated rule set XML files.

Importing Rule Sets using the Enterprise Manager Console

In order to import an incident rule set, administrators must have the *Create Enterprise Rule Set* privilege.

When an incident rule set is first imported, it will be disabled by default. You will need to edit the imported rule set in order to specify environment-specific parameters such as target names for specific target selection or user names for email notification. You will then need to enable the rule set.

To import an incident rule set:

1. From the **Setup** menu, select **Incidents** then select **Incident Rules**.
2. Click **Import**. The Import Rule Set dialog displays.
3. From the Import Rule Set dialog, click **Choose File**. The File Upload dialog displays.
4. Select the incident rule set XML file and click **Open**.
5. Click **OK**.

   If there is a naming conflict for the name, you will be asked to select one of the following:

   - Override rule set with same name
   - Create rule set with different name

Importing Rule Sets Using EM CLI

Using EM CLI, you can write scripts to import/export large numbers of rule sets. The *Create Enterprise Rule Set* privilege is required in order to run the import operation from the command line or script.

You can import a rule set from list of enterprise rule set(s) except for predefined (out-of-box) rule sets supplied by Oracle.

```
emcli import_incident_rule_set
      -import_file=<XML file name along with the file path for the exported rule set
earlier>
      [-alt_rule_set_name=<rule set name>]
```

Options

- import_file=<XML file name along with the file path for the exported rule set earlier>

- alt_rule_set_name=<rule set name>

  Optionally, you can specify the name of an enterprise rule set to use in case rule set already exists.

Example

```
emcli import_incident_rule_set -import_file="/tmp/TEST_RULESET.xml" -
alt_rule_set_name=COPY_OF_TEST_RULESET
```

This command imports the rule set and names it as 'COPY_OF_TEST_RULESET' from rule set XML specified 'TEST_RULESET.xml'

Exporting Rule Sets Using EM CLI

You can export a rule set from list of enterprise rule set(s) except for predefined (out-of-box) rule sets supplied by Oracle. Any user can run the export operation. No special privileges are required.

```
emcli export_incident_rule_set
      -rule_set_name=<rule set name>
      [-rule_set_owner=<ruleset owner>]
      -export_file=<XML file name along with the file path for the exported rule set>
```

Options

- rule_set_name=<rule set name>

  Name of an enterprise rule set.
- rule_set_owner=<ruleset owner>

  Optionally, you can specify the owner of the rule set.
- export_file=<XML file name along with the file path for the exported rule set>

  If the filename is specified as directory, it will create a file with rule set name in that directory.

Examples:

```
emcli export_incident_rule_set -rule_set_name=TEST_RULESET -rule_set_owner=sysman
-export_file="/tmp/"
```

This command exports the ruleset named 'TEST_RULESET' from rule set(s) and saves at '/tmp/TEST_RULESET.xml'

Creating Corrective Actions for Events

Prior to Enterprise Manager release 13.1, corrective actions could only be associated with metric alerts. Enterprise Manager release 13.1 now allows script-based corrective actions to fire on an event by associating them with event rules. This greatly increases the number of situations where corrective actions can be used, such as compliance standard violations, metric errors, or target availability. By associating corrective actions with event rules, you can have the corrective action performed automatically.

You can also initiate the corrective action manually through the event details *Guided Resolutions* area of Incident Manager. For a detailed discussion about corrective actions, see "Creating Corrective Actions.".

Corrective Actions in Event Rules

When you create an event rule to be triggered when a matching event occurs, you can select an appropriate predefined corrective action from the *Corrective Actions Library*. The corrective actions available for selection will depend on the event type and target type selected for the rule.

When an event rule set is exported or imported, the associated corrective actions will be exported/imported as well. For more information about importing/exporting event rules, see "Exporting and Importing Incident Rules."

Create the Corrective Action

In order to associate a corrective action with an event rule, you must first add it to the Corrective Action Library. After a corrective action is in the library, you can reuse the corrective action definition whenever you define a corrective action for an event rule.

1. From the Enterprise menu, select **Monitoring**, and then **Corrective Actions**. The Corrective Action Library page appears.
2. Select a job type from the **Create Library Corrective Action** drop-down. For events, you must create an *OS Command* job type so that a script can be executed. Select **OS Command**, specify a name and then click **Go**. The Create OS Command Corrective Action page displays.

   Specify a corrective action **Name** and a brief **Description** or event type.
3. From the Target Type drop-down menu, choose a target type. Click on the **Parameters** tab.
4. From the Command Type drop-down menu, choose **Script**.
5. Enter the OS script text.

   All target and event Properties that can be used in the script are listed in the table to the right.

   Tip: When accessing an Event Details page from Incident Manager, you can click **Show Internal Values for Attributes** to display the internal name and values for the event attributes. You can use this to determine what information you can access when writing the script for the corrective action. Just copy and paste the information from the dialog into a text editor and refer to this list of attributes when creating your script

   > **Note:**
   >
   > If you are using an event context parameter, it must be prefixed with EVTCTX.
6. Specify an interpreter. For example, `%perlbin%/perl`
7. Once you have finished, click **Save to Library.** The Corrective Actions Library page displays and your corrective action appears in the library list.

   At this point, the corrective action will be in *draft* status. At this stage, you can test and revise the corrective action. However, only you, as owner, can test the CA by running the CA manually from Incident Manager.

   To test the corrective action, you must trigger an event that matches the event rule with the associated corrective action to see if the actions are what you expect. Once you are satisfied and are ready for other administrators to use the corrective action, proceed to the next step.

Note: The Access tab on the "Create 'OS Command' Corrective Action" page displays administrators and roles that have access to this corrective action. You can change access to this corrective action from this tab, if required.

8. Navigate to the Corrective Actions Library page and select the Corrective Action and then click **Publish**. A confirmation message displays. Click **Yes** to confirm publication.
9. Set the Preferred Credentials. From the **Setup** menu, select **Security** and then **Preferred Credentials**. The Preferred Credentials page displays. Note that the preferred credential of the rule set owner will be used by the corrective action linked to the rule.

> **Note:**
>
> The corrective action will use these credentials to access the system and carry out the actions (in this case, running the script). For example, set credential for host if your corrective action is going to perform corrective actions on a specific host.

10. If not already set, select the **Target Type** to be accessed by the corrective action and click **Manage Preferred Credentials**. You need to define the Default Preferred Credentials for the specific target type that the CA is going to perform the actions on. The target type's Preferred Credentials page displays. On the My preferences tab, navigate to the Default Preferred Credentials region and select the applicable credential. Click **Set**.

> **Note:**
>
> Preferred credentials must be set or the corrective action will fail.

Associate the Corrective Action with an Event Rule

Once you have created the corrective action to be associated with an event, you are now ready to create an event rule that uses the corrective action. You can only associate one corrective action per conditional action of the rule.

1. From the Setup menu, select **Incidents** and then **Incident Rules**. The Incident Rules - All Enterprise Rules page displays.
2. Click **Create Incident Rule Set.** The Create Rule Set page displays.
3. Enter a rule set **Name** and **Description**.
4. Select the appropriate **Targets**.
5. Scroll down to the Rules section and click **Create...** The Select Type of Rule dialog displays. Choose **Incoming events and updates to events** and click **Continue**. The Create Rule Set wizard appears.
6. From the Type drop-down menu, select the event **Type**. By default, **Metric Alert** is selected. Choose one of the event types, Compliance Standard Rule Violation, for example. Expand the **Advanced Selection Options** and set any event parameters to which the event rule should apply.
7. Click **Next** to proceed to the Add Actions page.
8. On the Add Actions page, click **Add**. The Add Conditional Actions page displays.
9. Scroll down to the Submit Corrective Action section and click **Select Corrective Action**. The corrective action selection dialog displays.

10. Choose the corrective action to be attached and click **OK**.

> **Note:**
>
> You are not prompted for credentials because the rules are run in the background and the rule set owner's preferred credentials are used to execute the corrective action.

11. Click **Continue**. You are returned to the main Add Actions page. Continue to add more actions, if necessary.
12. Complete the rule set definition and ensure that it appears in the list of incident rule sets on the Incident Rules - All Enterprise Rules page.

You will need to recreate the particular rule violation in order to test the CA.

Running the Corrective Action Manually

If you are aware that there exists a corrective action in the Corrective Action Library that can resolve the current event, you can run the corrective action manually from the library. In the Guided Resolution section of an Event Details page, the Corrective Actions area displays the **Submit from Library** link.

Click **Submit from Library** to display the Corrective Action Library dialog. This dialog lists ONLY those corrective actions that apply to the current event conditions. Select a corrective action from the list. The credential settings are displayed. By default, the preferred credentials are shown. You have the option of using alternate credentials.

Once set, click **Submit**. The *Corrective action <CA name> submitted successfully* dialog displays. Click the link **Click here to view the execution details**."to go to the job execution page. Here, you can view the job status and output.

Compressing Multiple Events into a Single Incident

An incident is created for an event when there is a corresponding event rule that has an action to create an incident for the event. In this situation, multiple events will generate multiple incidents. However, if the events relate to the same issue, instead of generating multiple incidents, it is better from a manageability standpoint to generate a single incident composed of multiple related events, i.e. an incident with compressed events. This is especially true if these related events are to be managed by the same administrator.

There are two types of Event Compression available in Enterprise Manager.

Event Compression Types

- Event Compression Policies
- Rule-based Event Compression

Event Compression Policies

Event Compression Policies, introduced in Enterprise Manager 13c Release 5 Update 8, work with your incident rule sets to decide if sets of related events can compress into a single incident.
Event Compression Policies serve as a noise reduction mechanism by identifying the condition under which multiple correlated events are grouped together, or compressed, in one incident. These are global policies that apply to all incident-creating rules and are enabled by default.

With Event Compression, you can reduce the overall volume of incidents to a more meaningful and manageable set.

Event Compression Policies

There are several out-of-box policies that collectively represent recommended ways in which related events should be compressed together into a single incident. These Event Compression Policies work with your incident rule sets to compress related events into a single incident. An event rule specifies a set of events and an action that should be taken in case any of the events occur, such as *create an incident for the event*. Without event compression, when the event occurs, an incident will be created for each event, potentially resulting in a larger number of incidents. The event rule now has an option, as part of the create incident action, to use Event Compression Policies. With this option enabled, before an incident is created, Enterprise Manager will locate an Event Compression Policy that is applicable to the set of events in the event rule and then compress the events into a single incident based on that policy. New rules that have actions to create incidents will automatically be configured to use Event Compression Policies.

To access these policies and the controls that enable/disable them, from the Enterprise Manager console, click **Setup**, select **Incidents**, and then **Event Compression Policies**. The Event Compression Policies page displays as shown below.

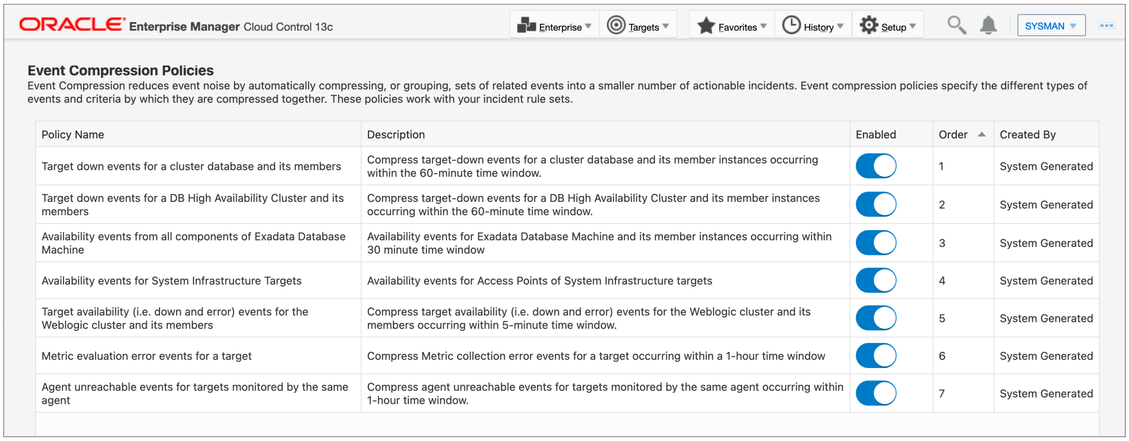

Event Compression Policies

Event Compression reduces event noise by automatically compressing, or grouping, sets of related events into a smaller number of actionable incidents. Event compression policies specify the different types of events and criteria by which they are compressed together. These policies work with your incident rule sets.

Policy Name	Description	Enabled	Order ▲	Created By
Target down events for a cluster database and its members	Compress target-down events for a cluster database and its member instances occurring within the 60-minute time window.	●	1	System Generated
Target down events for a DB High Availability Cluster and its members	Compress target-down events for a DB High Availability Cluster and its member instances occurring within the 60-minute time window.	●	2	System Generated
Availability events from all components of Exadata Database Machine	Availability events for Exadata Database Machine and its member instances occurring within 30 minute time window	●	3	System Generated
Availability events for System Infrastructure Targets	Availability events for Access Points of System Infrastructure targets	●	4	System Generated
Target availability (i.e. down and error) events for the Weblogic cluster and its members	Compress target availability (i.e. down and error) events for the Weblogic cluster and its members occurring within 5-minute time window.	●	5	System Generated
Metric evaluation error events for a target	Compress Metric collection error events for a target occurring within a 1-hour time window	●	6	System Generated
Agent unreachable events for targets monitored by the same agent	Compress agent unreachable events for targets monitored by the same agent occurring within 1-hour time window.	●	7	System Generated

Alternatively, you can also access the Event Compression Policies page directly from the Incident Rules page.

1. From the Enterprise Manager console, click **Setup**, then select **Incidents** and then **Incident Rules**.
2. In the banner area at the top, click **View Event Compression Policies**.

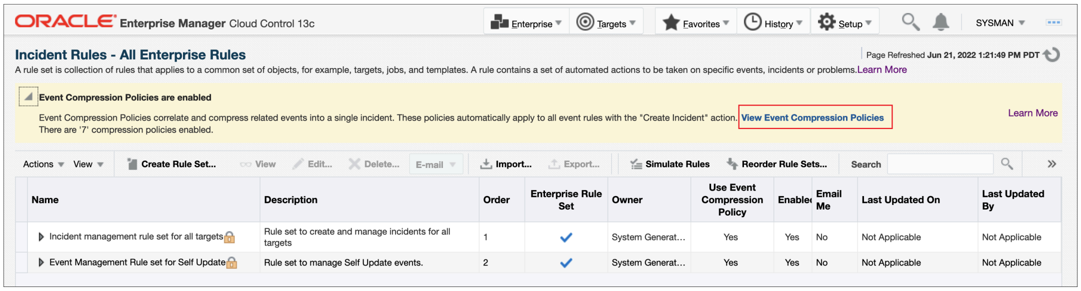

Incident Rules - All Enterprise Rules

Page Refreshed **Jun 21, 2022 1:21:49 PM PDT**

A rule set is collection of rules that applies to a common set of objects, for example, targets, jobs, and templates. A rule contains a set of automated actions to be taken on specific events, incidents or problems.Learn More

Event Compression Policies are enabled

Event Compression Policies correlate and compress related events into a single incident. These policies automatically apply to all event rules with the "Create Incident" action. **View Event Compression Policies**
There are '7' compression policies enabled.

Learn More

Name	Description	Order	Enterprise Rule Set	Owner	Use Event Compression Policy	Enablec	Email Me	Last Updated On	Last Updated By
Incident management rule set for all targets	Rule set to create and manage incidents for all targets	1	✓	System Generat…	Yes	Yes	No	Not Applicable	Not Applicable
Event Management Rule set for Self Update	Rule set to manage Self Update events.	2	✓	System Generat…	Yes	Yes	No	Not Applicable	Not Applicable

These compression policies provide for common event scenarios and contain logic that defines the following:

- The events that should be compressed based on the target types and/or particular attributes of the events, e.g., event type or severity
- The conditions by which they are compressed/grouped
- The time window within which the events should have occurred in order to be compressed
- The format of the message that the compressed incident uses when it is in a non-clear state and when it is in a clear state

Event Compression policies are evaluated from top down, with the top having higher priority (*Order*). The evaluation order is the sequence in which policies will be executed to match to a rule.

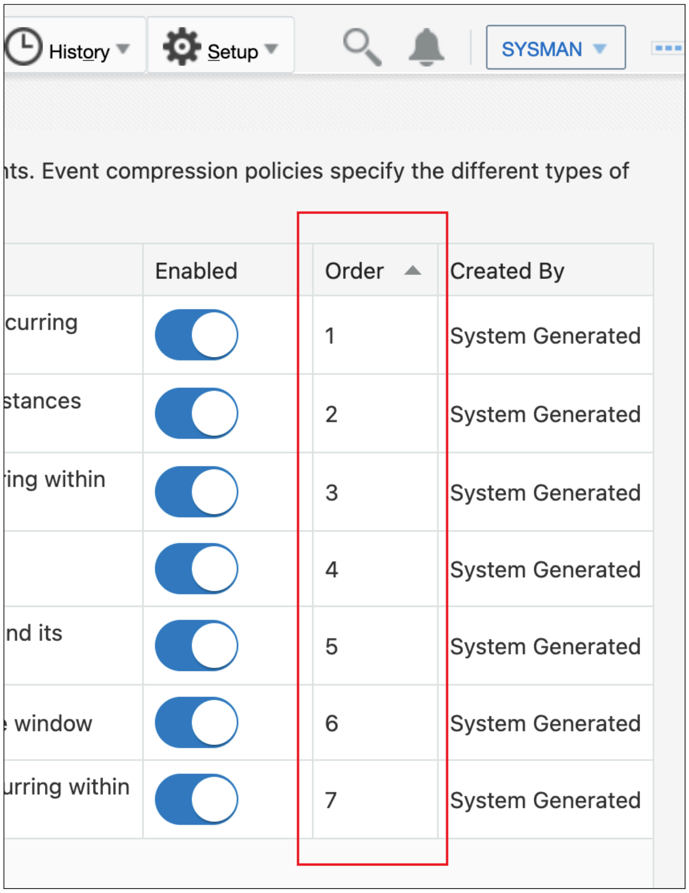

ıts. Event compression policies specify the different types of

	Enabled	Order	Created By
curring	●	1	System Generated
stances	●	2	System Generated
ing within	●	3	System Generated
	●	4	System Generated
nd its	●	5	System Generated
window	●	6	System Generated
ırring within	●	7	System Generated

You can view a policy's compression logic and incident messaging by clicking on the desired policy to display the policy's detail page. In the following example, the policy details page shows that target-down events for a cluster database and its members, that have a severity of *fatal* and that have occurred within a 60-minute window, will be compressed into a single incident. The format of the incident message will be as follows:

- **Non-clear State**: *There are %EVENT_COUNT% target down events on members of Cluster Database: %PARENT_TARGET_NAME%*
- **Clear State**: *Target down events on members of Cluster Database: %PARENT_TARGET_NAME% are cleared*

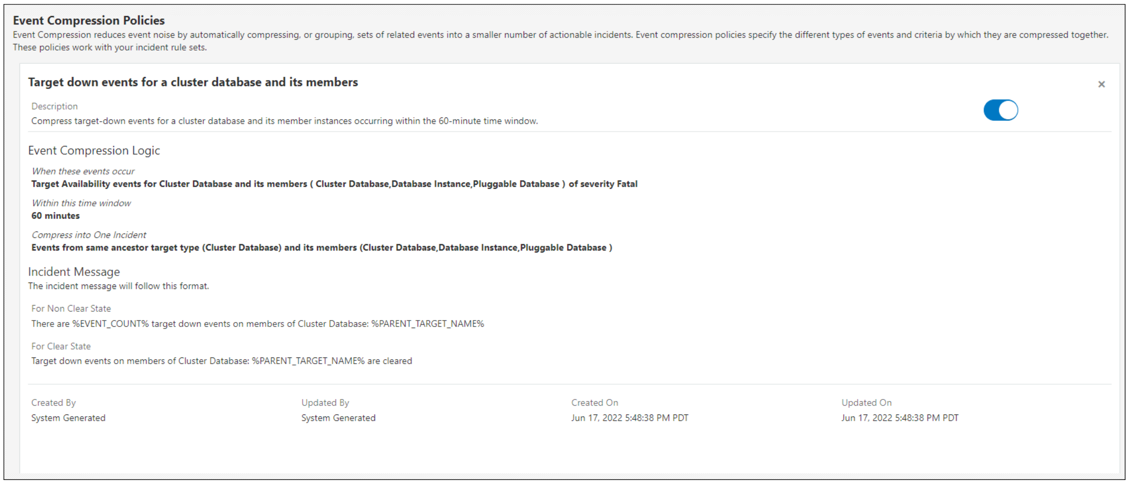

Event Compression Policies

Event Compression reduces event noise by automatically compressing, or grouping, sets of related events into a smaller number of actionable incidents. Event compression policies specify the different types of events and criteria by which they are compressed together. These policies work with your incident rule sets.

Target down events for a cluster database and its members

Description
Compress target-down events for a cluster database and its member instances occurring within the 60-minute time window.

Event Compression Logic

When these events occur
Target Availability events for Cluster Database and its members (Cluster Database,Database Instance,Pluggable Database) of severity Fatal

Within this time window
60 minutes

Compress into One Incident
Events from same ancestor target type (Cluster Database) and its members (Cluster Database,Database Instance,Pluggable Database)

Incident Message
The incident message will follow this format.

For Non Clear State
There are %EVENT_COUNT% target down events on members of Cluster Database: %PARENT_TARGET_NAME%

For Clear State
Target down events on members of Cluster Database: %PARENT_TARGET_NAME% are cleared

Created By	Updated By	Created On	Updated On
System Generated	System Generated	Jun 17, 2022 5:48:38 PM PDT	Jun 17, 2022 5:48:38 PM PDT

User-defined Event Compression Policies

Available in Enterprise Manager 13*c* Release 5 Update 13 and later, there is now the option to create a *user-defined event compression policy*. This means you can now also author your own policies that will apply to all incident-creating rules.

> **Note:**
>
> To create, update, or delete a user-defined event compression policy, you must have Create Business Rule privilege.

1. Click **Create New Policy**.

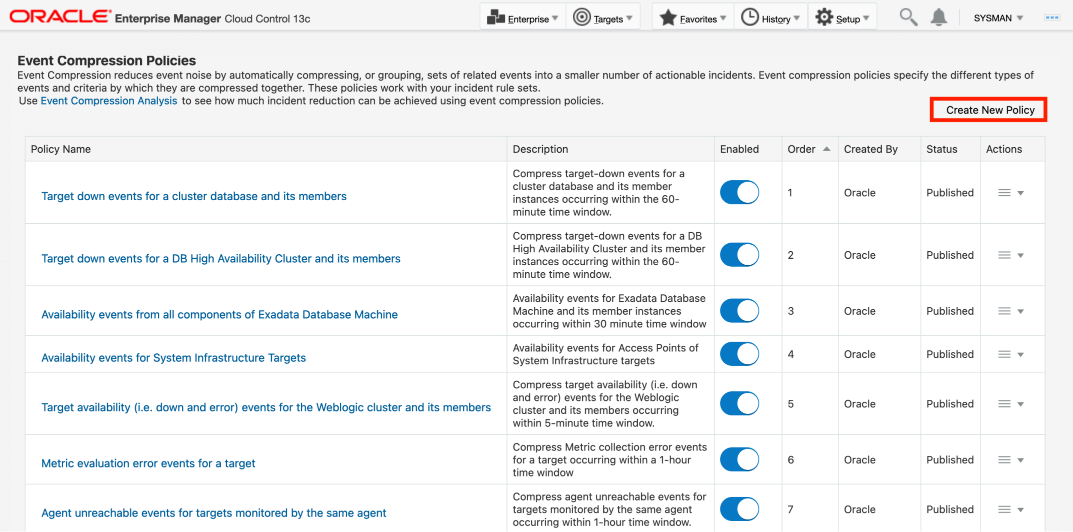

Event Compression Policies

Event Compression reduces event noise by automatically compressing, or grouping, sets of related events into a smaller number of actionable incidents. Event compression policies specify the different types of events and criteria by which they are compressed together. These policies work with your incident rule sets.
Use Event Compression Analysis to see how much incident reduction can be achieved using event compression policies.

Policy Name	Description	Enabled	Order	Created By	Status	Actions
Target down events for a cluster database and its members	Compress target-down events for a cluster database and its member instances occurring within the 60-minute time window.	●	1	Oracle	Published	≡ ▾
Target down events for a DB High Availability Cluster and its members	Compress target-down events for a DB High Availability Cluster and its member instances occurring within the 60-minute time window.	●	2	Oracle	Published	≡ ▾
Availability events from all components of Exadata Database Machine	Availability events for Exadata Database Machine and its member instances occurring within 30 minute time window	●	3	Oracle	Published	≡ ▾
Availability events for System Infrastructure Targets	Availability events for Access Points of System Infrastructure targets	●	4	Oracle	Published	≡ ▾
Target availability (i.e. down and error) events for the Weblogic cluster and its members	Compress target availability (i.e. down and error) events for the Weblogic cluster and its members occurring within 5-minute time window.	●	5	Oracle	Published	≡ ▾
Metric evaluation error events for a target	Compress Metric collection error events for a target occurring within a 1-hour time window	●	6	Oracle	Published	≡ ▾
Agent unreachable events for targets monitored by the same agent	Compress agent unreachable events for targets monitored by the same agent occurring within 1-hour time window.	●	7	Oracle	Published	≡ ▾

The *Create Compression Policy* region displays.

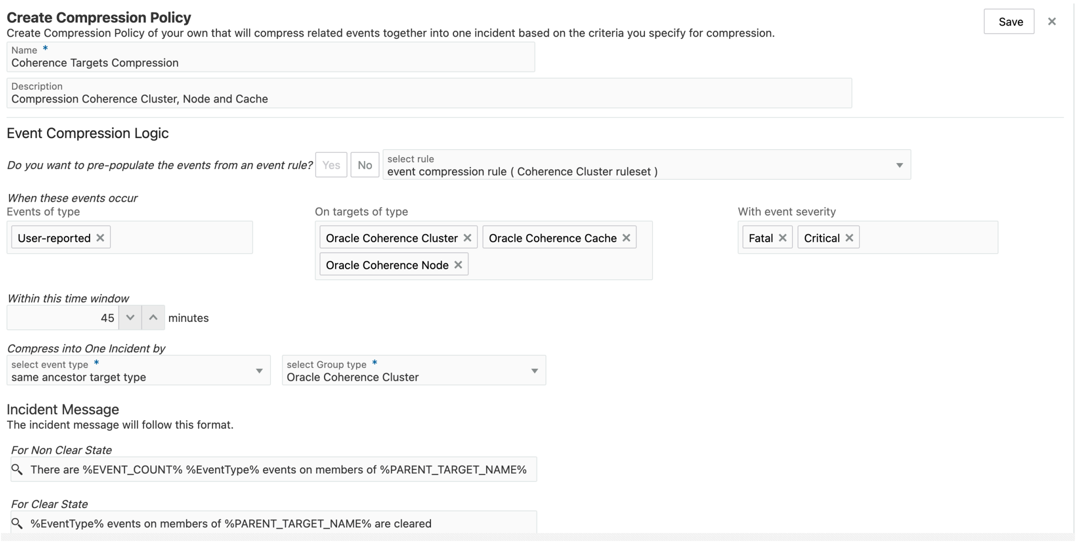

Create Compression Policy

Create Compression Policy of your own that will compress related events together into one incident based on the criteria you specify for compression.

Name *
Coherence Targets Compression

Description
Compression Coherence Cluster, Node and Cache

Event Compression Logic

Do you want to pre-populate the events from an event rule?

select rule
event compression rule (Coherence Cluster ruleset)

When these events occur

Events of type: User-reported ×

On targets of type: Oracle Coherence Cluster × Oracle Coherence Cache × Oracle Coherence Node ×

With event severity: Fatal × Critical ×

Within this time window
45 minutes

Compress into One Incident by

select event type *
same ancestor target type

select Group type *
Oracle Coherence Cluster

Incident Message

The incident message will follow this format.

For Non Clear State
There are %EVENT_COUNT% %EventType% events on members of %PARENT_TARGET_NAME%

For Clear State
%EventType% events on members of %PARENT_TARGET_NAME% are cleared

2. Enter a **Name** and **Description**. Define a name and description that is meaningful to you and describes the policy you will be creating.

> **Note:**
>
> The policy name must be unique (not used by any out-of-box or existing user-defined compression policy).

Next, you define the compression logic.

3. Decide if you want to pre-populate your compression policy with conditions that have already been specified in an existing event rule for incident creation. If you click **Yes**, a drop-down menu appears where you can select an existing event rule. Once you select an event rule, it will populate the ***When these events occur*** region where it is applicable.
If you click **No** to pre-populate from an event rule, then you can manually select the fields (event type, target type, event severity) for this section.

   To define the compression logic, you can add or delete *event types*, *target types*, and *event severities* whether or not you chose to pre-populate this region via an existing event rule.

> **Note:**
>
> In Enterprise Manager 13.5 RU18, you may now define multiple event types in your compression logic. For example, a database down may cause connectivity-related metric alerts on the dependent applications. In EM, this would translate to Target Availability events and Metric Alert events being generated. Using Multi-Event Type compression, you may group these different but related events into one incident instead.

4. Define the time range for compression to apply. For example, a time window of 45 minutes means that if multiple related events occurred within 45 minutes, then they will be compressed into one incident.
Next, you need to choose how you want the events to be compressed, which basically determines the effectiveness of your compression policy.

5. Under *Compress into One Incident by*, select the **event type**.

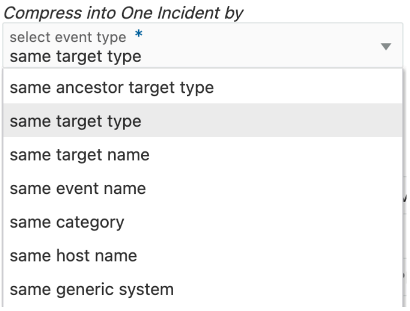

Compress into One Incident by

select event type *
same target type

> **Note:**
>
> When you select **same ancestor target type**, all events for all target types chosen in Step 3 will be compressed into a single incident. For that reason, you must also select a **Group type** to limit the event compression to a specific target type.

6. Define incident *Clear* and *Non-Clear* messages you want to display when the incident occurs. By default, this field will be pre-populated based on what is entered for the *Compress into One Incident by* section. You may modify the incident message as needed. The format of the incident message are as follows:

   - **Non-clear** State: `There are %EVENT_COUNT% target down events on members of Cluster Database: %PARENT_TARGET_NAME%`
   - **Clear** State: `Target down events on members of Cluster Database: %PARENT_TARGET_NAME% are cleared`
     There are additional variables that can be used in an incident message. These variables will be replaced with its corresponding attribute when the message is generated. For example, if the variable `%TARGET_NAME%` was used in the event compression logic, when an incident is triggered it will be replaced with the name of the target that the incident was created on.

   > **Note:**
   >
   > Only certain variables make sense for each *Compress into One Incident by* selection. It is best to refer to the default message that gets populated with applicable variables during the selection.

   Variables that can be used in the incident message are:

Variable	Definition
%EventType%	The type of event that triggered the incident. Examples include: Metric Alert, Target Availability, High Availability, etc.
%EVENT_NAME%	The internal event name describing the nature of the events
%TARGET_NAME%	The name of the target that the incident was created on

Variable	Definition
%TARGET_TYPE%	The type of target that the incident was created on. Examples include: Database Instance, WebLogic, etc.
%PARENT_TARGET_NAME%	Name of the parent target that an incident was created on. Example: Cluster Database (parent target) has members instances: Database Instance and Pluggable Database (children targets)
%HOST_TARGET%	Target where targets that triggered the incident are hosted
%GENERIC_SYSTEM%	Group together events from targets that are members of the same generic system. Example: Users create their own systems for their applications and want to group all events related to a specific system in order to manage it as a unit.
%CATEGORY_NAME%	The name of the category. Functional or operational classification for an event.Available Categories: – Availability – Business – Capacity – Configuration – Diagnostics – Error – Fault – Jobs – Load – Performance – Security
%EVENT_COUNT%	The number of events associated with this incident

7. After completing the policy definition, click **Save**. The newly created user-defined event compression policy appears in the policy list in *Draft* status where you can continue to edit and update the user-defined event compression policy.

> **Note:**
>
> While in *Draft* status, you can test the effectiveness of this compression policy and any changes you make to it by running the *Event Compression Policy Analyzer*. See Assessing the Benefits of Using Event Compression Policies for more information.

Publishing a User-defined Event Compression Policy

Once the policy is ready, publish it so that it can be used by other administrators. Select **Publish** from the action menu.

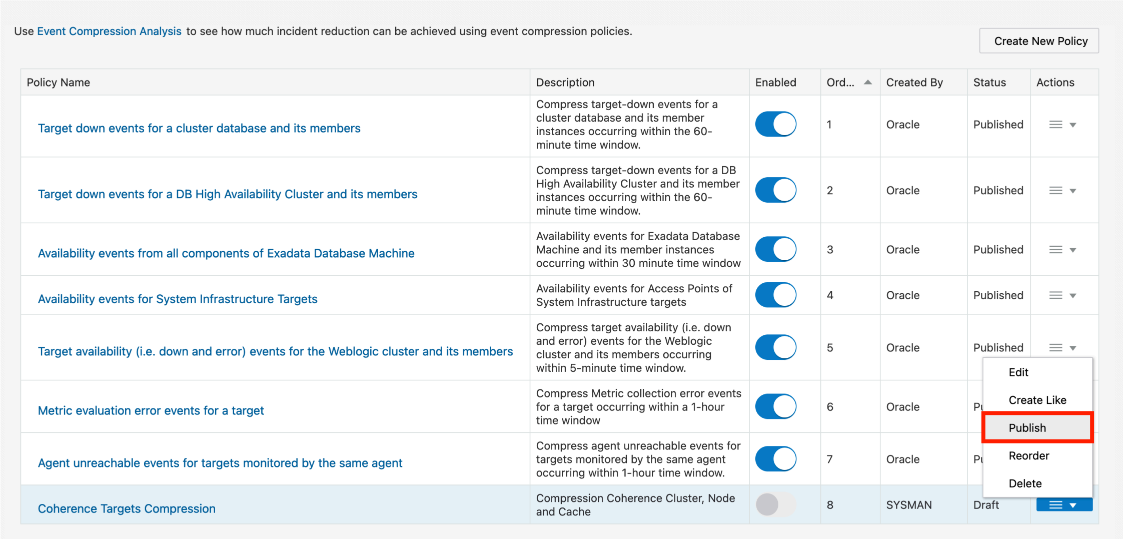

Use Event Compression Analysis to see how much incident reduction can be achieved using event compression policies.

Policy Name	Description	Enabled	Ord... ▲	Created By	Status	Actions
Target down events for a cluster database and its members	Compress target-down events for a cluster database and its member instances occurring within the 60-minute time window.	●	1	Oracle	Published	≡ ▾
Target down events for a DB High Availability Cluster and its members	Compress target-down events for a DB High Availability Cluster and its member instances occurring within the 60-minute time window.	●	2	Oracle	Published	≡ ▾
Availability events from all components of Exadata Database Machine	Availability events for Exadata Database Machine and its member instances occurring within 30 minute time window	●	3	Oracle	Published	≡ ▾
Availability events for System Infrastructure Targets	Availability events for Access Points of System Infrastructure targets	●	4	Oracle	Published	≡ ▾
Target availability (i.e. down and error) events for the Weblogic cluster and its members	Compress target availability (i.e. down and error) events for the Weblogic cluster and its members occurring within 5-minute time window.	●	5	Oracle	Published	≡ ▾
Metric evaluation error events for a target	Compress Metric collection error events for a target occurring within a 1-hour time window	●	6	Oracle	Pu	
Agent unreachable events for targets monitored by the same agent	Compress agent unreachable events for targets monitored by the same agent occurring within 1-hour time window.	●	7	Oracle	Pu	
Coherence Targets Compression	Compression Coherence Cluster, Node and Cache	○	8	SYSMAN	Draft	≡ ▾

Once published, you’ll be able to enable the user-defined event compression policy.

Editing a Published User-defined Event Compression Policy

You can edit a published user-defined event compression policy. You must first disable the published compression policy before you can make edits. If this compression policy had been previously enabled, it will remain inactive until it is re enabled.

Evaluation Order and Placement

The policies are evaluated from top down, with the top having higher priority (Order). The evaluation order is the sequence in which policies will be executed to match to a rule. For user-defined event compression policies, they will always appear at the bottom of the policy list when created. However, you may change their order in the execution list. They can be placed anywhere above or below the out-of-box policies, but cannot be placed in between the out-of-box policies.

Assessing the Benefits of Using Event Compression Policies

If you’re trying to decide whether you want to use Event Compression Policies, you can view the beneficial impact of using them by running an Event Compression Analysis (available with Enterprise Manager 13c Release 5 Update 11 and later). This analysis uses historical monitoring data from your environment to analyze the impact Event Compression Policies would have had on events and incidents over a selected period in the past. The results of the analysis will show what effect Event Compression Policies would have had on events that generated incidents for the specified time range.

Running an Analysis

To initiate an Event Compression Analysis:

1. Navigate to the Event Compression Policies page (Setup->Incidents->Event Compression Policies).
2. Enable the desired Event Compression Policies that are appropriate for the event rules you currently have defined.
3. Click **Event Compression Analysis** in the introductory text at the top of the page.

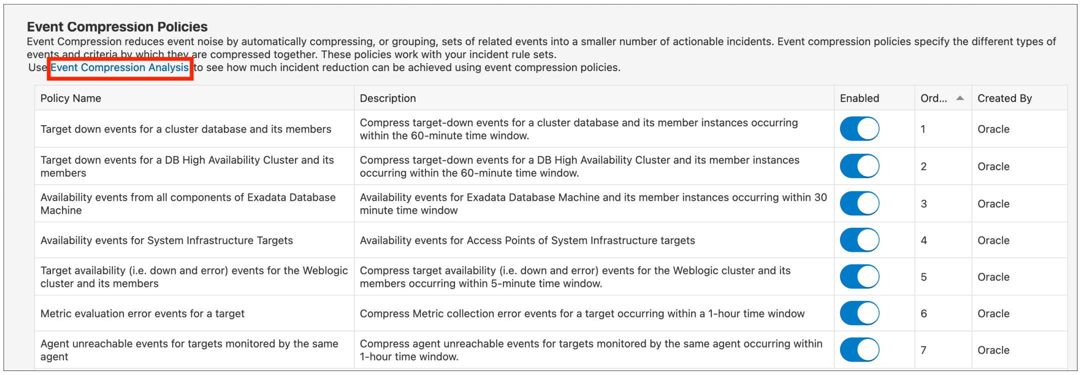

Event Compression Policies

Event Compression reduces event noise by automatically compressing, or grouping, sets of related events into a smaller number of actionable incidents. Event compression policies specify the different types of events and criteria by which they are compressed together. These policies work with your incident rule sets.

Use Event Compression Analysis to see how much incident reduction can be achieved using event compression policies.

Policy Name	Description	Enabled	Ord...	Created By
Target down events for a cluster database and its members	Compress target-down events for a cluster database and its member instances occurring within the 60-minute time window.	●	1	Oracle
Target down events for a DB High Availability Cluster and its members	Compress target-down events for a DB High Availability Cluster and its member instances occurring within the 60-minute time window.	●	2	Oracle
Availability events from all components of Exadata Database Machine	Availability events for Exadata Database Machine and its member instances occurring within 30 minute time window	●	3	Oracle
Availability events for System Infrastructure Targets	Availability events for Access Points of System Infrastructure targets	●	4	Oracle
Target availability (i.e. down and error) events for the Weblogic cluster and its members	Compress target availability (i.e. down and error) events for the Weblogic cluster and its members occurring within 5-minute time window.	●	5	Oracle
Metric evaluation error events for a target	Compress Metric collection error events for a target occurring within a 1-hour time window	●	6	Oracle
Agent unreachable events for targets monitored by the same agent	Compress agent unreachable events for targets monitored by the same agent occurring within 1-hour time window.	●	7	Oracle

The Event Compression Policy Analyzer page displays.

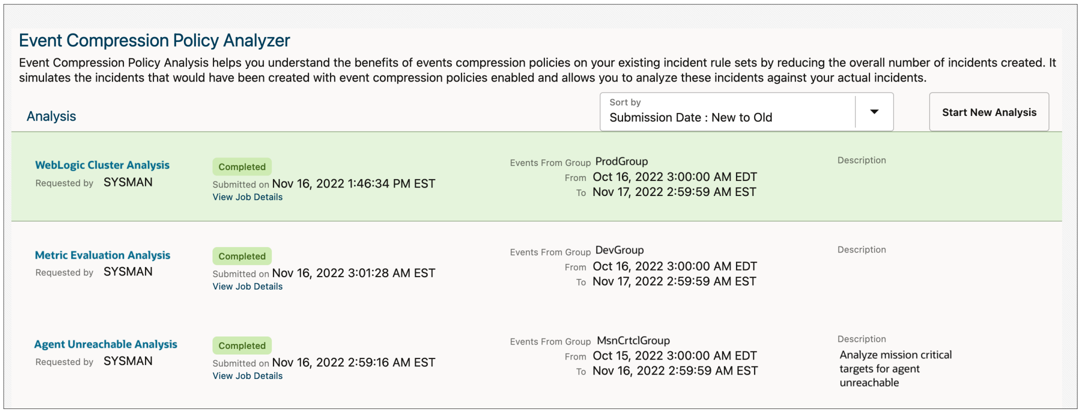

Event Compression Policy Analyzer

Event Compression Policy Analysis helps you understand the benefits of events compression policies on your existing incident rule sets by reducing the overall number of incidents created. It simulates the incidents that would have been created with event compression policies enabled and allows you to analyze these incidents against your actual incidents.

WebLogic Cluster Analysis
Requested by SYSMAN
Completed
Submitted on Nov 16, 2022 1:46:34 PM EST
View Job Details
Events From Group ProdGroup
From Oct 16, 2022 3:00:00 AM EDT
To Nov 17, 2022 2:59:59 AM EST
Description

Metric Evaluation Analysis
Requested by SYSMAN
Completed
Submitted on Nov 16, 2022 3:01:28 AM EST
View Job Details
Events From Group DevGroup
From Oct 16, 2022 3:00:00 AM EDT
To Nov 17, 2022 2:59:59 AM EST
Description

Agent Unreachable Analysis
Requested by SYSMAN
Completed
Submitted on Nov 16, 2022 2:59:16 AM EST
View Job Details
Events From Group MsnCrtclGroup
From Oct 15, 2022 3:00:00 AM EDT
To Nov 16, 2022 2:59:59 AM EST
Description
Analyze mission critical targets for agent unreachable

4. Click **Start New Analysis** to display the *Compression Policy Analysis* definition dialog.

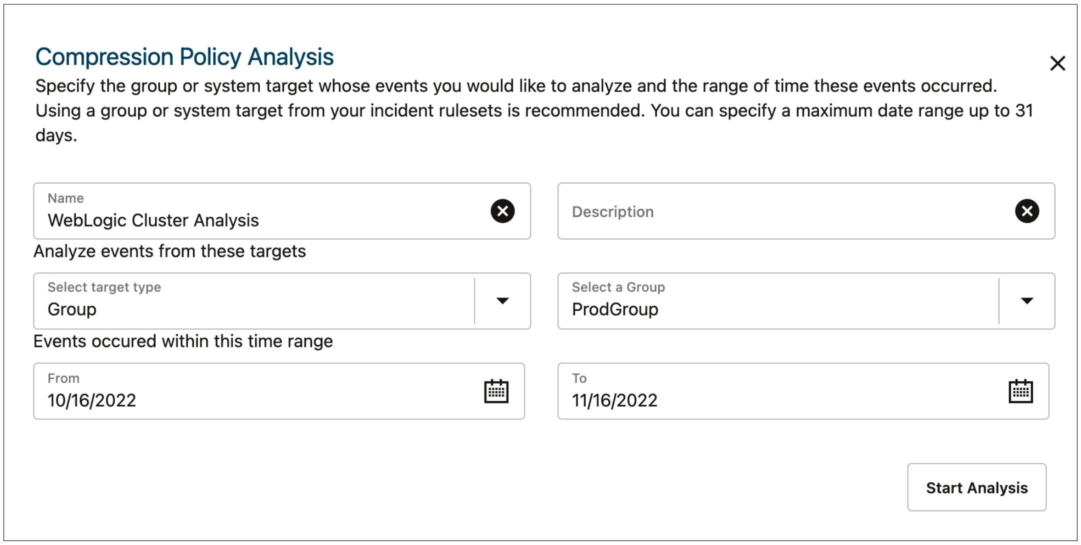

Compression Policy Analysis

Specify the group or system target whose events you would like to analyze and the range of time these events occurred. Using a group or system target from your incident rulesets is recommended. You can specify a maximum date range up to 31 days.

Name: WebLogic Cluster Analysis
Description

Analyze events from these targets

Select target type: Group
Select a Group: ProdGroup

Events occured within this time range

From: 10/16/2022
To: 11/16/2022

- **Name** and **Description:** Define a name and description that is meaningful to you and any individuals with whom you'll share your analysis.
- **Select Target Type** and **Select** *targets*: Select type of target generating events on which you want to run the compression analysis against. Once you've selected the Target Type, you can then choose available targets from the drop-down menu.

> **Note:**
>
> Only group or system target types are supported as you want to run compression analysis on a set of related targets generating events.

- **Time Range:** Define the duration of the compression analysis. By default, historical event data for the last 30 days is used.

5. Click **Start Analysis** once you have specified all criteria for the compression analysis. A job is then submitted to start the analysis. You can view the Job Progress by selecting **View Job Progress** on the analysis page .

Interpreting Analysis Results

When the analysis is complete, click on the analysis name. The Event Compression Policy Analyzer page displays. From here, you can view how enabling Event Compression policies would have affected the number of incidents created for the selected time period.

For example, in the image below, 52 events occurred during the last 30 days. Without compression, 52 incidents would have been created. With compression, the number of incidents created was reduced to 47, a 9% reduction in the number of incidents. The Compression Ratio indicates the average number of events that were compressed for each incident.

In the graph below, the *analysis summary* shows the number of incidents created over time for the selected analysis period and the incident reduction percentage. The bars are color coded to specify the number of incidents when Event Compression Policies are used versus the number of incidents when compression policies aren't used.

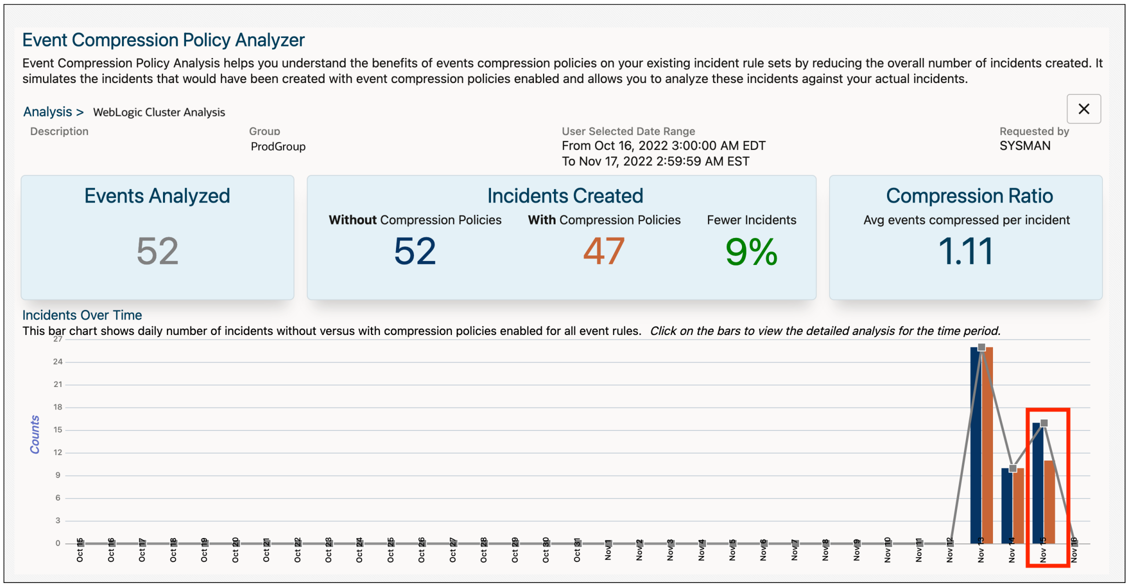

Event Compression Policy Analyzer

Event Compression Policy Analysis helps you understand the benefits of events compression policies on your existing incident rule sets by reducing the overall number of incidents created. It simulates the incidents that would have been created with event compression policies enabled and allows you to analyze these incidents against your actual incidents.

Analysis > WebLogic Cluster Analysis

Description

Group: ProdGroup

User Selected Date Range: From Oct 16, 2022 3:00:00 AM EDT To Nov 17, 2022 2:59:59 AM EST

Requested by: SYSMAN

Events Analyzed: 52

Incidents Created — **Without** Compression Policies: 52 — **With** Compression Policies: 47 — Fewer Incidents: 9%

Compression Ratio — Avg events compressed per incident: 1.11

Incidents Over Time

This bar chart shows daily number of incidents without versus with compression policies enabled for all event rules. *Click on the bars to view the detailed analysis for the time period.*

To further analyze how events and incidents are compressed on a specific date, click on the date's corresponding bars to see a visual representation of the compression. In this visualization, you will see how events are mapped from incidents without compression policies to incidents with compression policies.

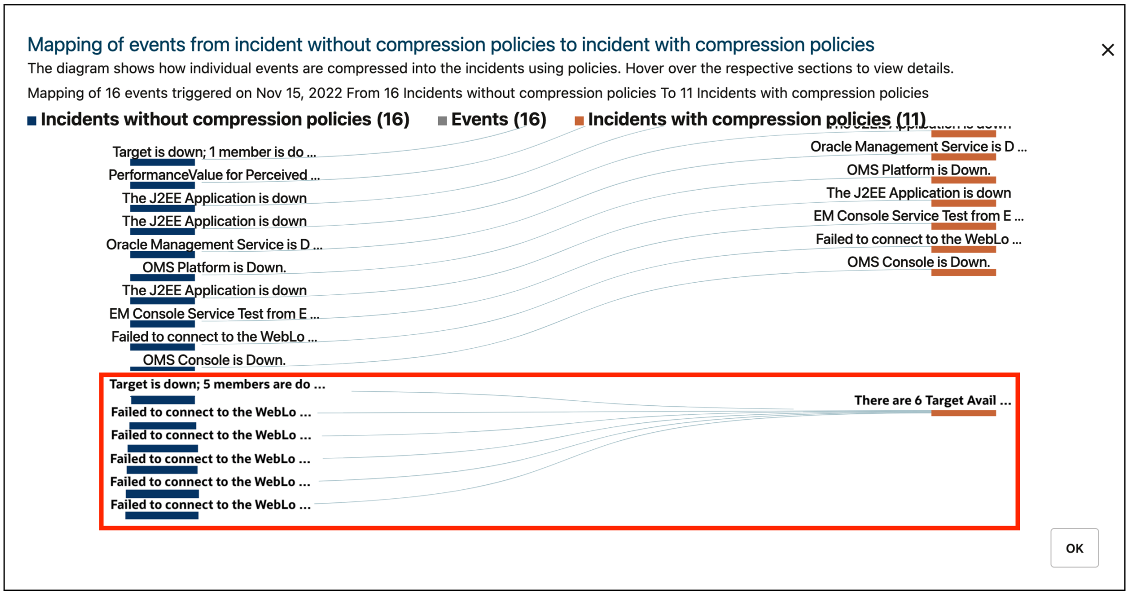

Mapping of events from incident without compression policies to incident with compression policies

The diagram shows how individual events are compressed into the incidents using policies. Hover over the respective sections to view details.

Mapping of 16 events triggered on Nov 15, 2022 From 16 Incidents without compression policies To 11 Incidents with compression policies

■ Incidents without compression policies (16)
■ Events (16)
■ Incidents with compression policies (11)
×
Target is down; 1 member is do ...
PerformanceValue for Perceived ...
The J2EE Application is down
The J2EE Application is down
Oracle Management Service is D ...
OMS Platform is Down.
The J2EE Application is down
EM Console Service Test from E ...
Failed to connect to the WebLo ...
OMS Console is Down.
Target is down; 5 members are do ...
Failed to connect to the WebLo ...
Failed to connect to the WebLo ...
Failed to connect to the WebLo ...
Failed to connect to the WebLo ...
Failed to connect to the WebLo ...
Oracle Management Service is D ...
OMS Platform is Down.
The J2EE Application is down
EM Console Service Test from E ...
Failed to connect to the WebLo ...
OMS Console is Down.
There are 6 Target Avail ...
OK

You can also hover and click the respective areas (i.e., events, incidents with/without compression policies) to see an additional pop-up with more details.

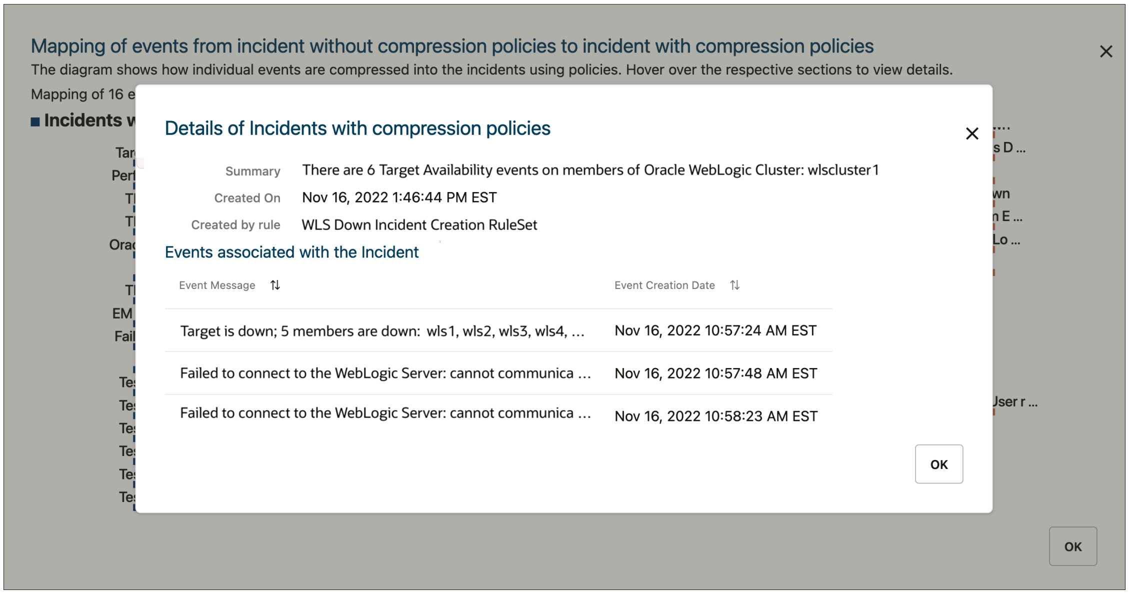

Details of Incidents with compression policies

Summary	There are 6 Target Availability events on members of Oracle WebLogic Cluster: wlscluster1
Created On	Nov 16, 2022 1:46:44 PM EST
Created by rule	WLS Down Incident Creation RuleSet

Events associated with the Incident

Event Message ⇅	Event Creation Date ⇅
Target is down; 5 members are down: wls1, wls2, wls3, wls4, ...	Nov 16, 2022 10:57:24 AM EST
Failed to connect to the WebLogic Server: cannot communica ...	Nov 16, 2022 10:57:48 AM EST
Failed to connect to the WebLogic Server: cannot communica ...	Nov 16, 2022 10:58:23 AM EST

Using Event Compression Policies

1. Rule Level Compression

   As shown below, in your event rules, you can enable the use of Event Compression Policies when creating an incident in the **Add Actions** page.

   To keep using compression for the same event type, maintain the default selection of **Allow policies to compress events within the same rule (Rule Level Compression)** in the main ruleset page.

> **Note:**
>
> The use of Event Compression Policies for rules that create incidents will be automatically enabled for all new event rules created after Enterprise Manager 13c Release 5 Update 8 or later. All event rules that have pre-existed prior to Enterprise Manager 13c Release 5 Update 8 (13.5.0.8) will remain as is and will not have Event Compression Policies automatically enabled. You can choose to enable Event Compression Policies for these event rules using the methods described below.

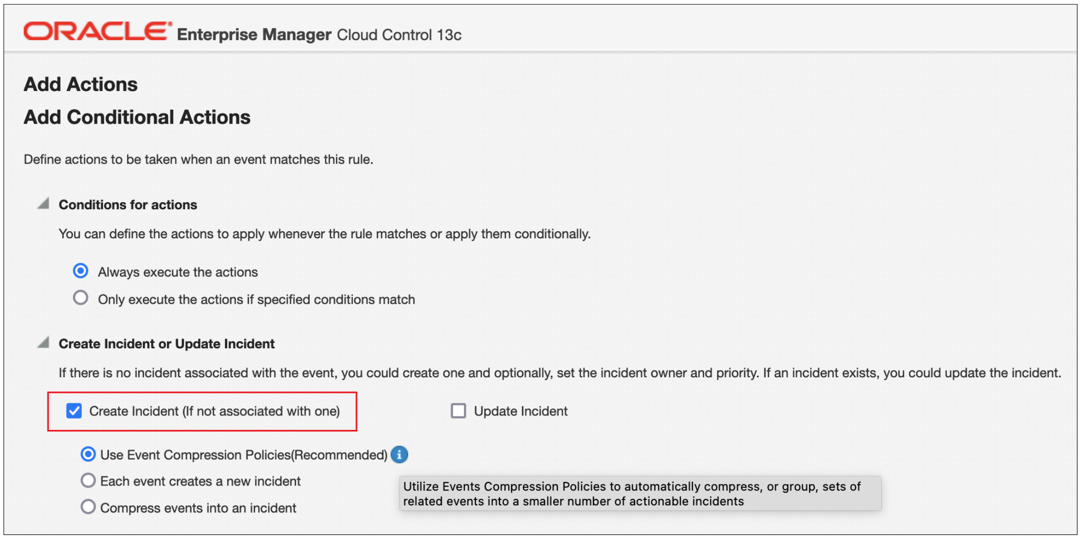

ORACLE Enterprise Manager Cloud Control 13c

Add Actions

Add Conditional Actions

Define actions to be taken when an event matches this rule.

Conditions for actions

You can define the actions to apply whenever the rule matches or apply them conditionally.

● Always execute the actions
○ Only execute the actions if specified conditions match

Create Incident or Update Incident

If there is no incident associated with the event, you could create one and optionally, set the incident owner and priority. If an incident exists, you could update the incident.

☑ Create Incident (If not associated with one)   □ Update Incident

● Use Event Compression Policies(Recommended)
○ Each event creates a new incident
○ Compress events into an incident

Utilize Events Compression Policies to automatically compress, or group, sets of related events into a smaller number of actionable incidents

2. Ruleset Level Compression

   Ruleset Level Compression works across multiple rules in a ruleset, allowing compression in a single incident of multiple event types that follow different rules from the ruleset.

> **Note:**
>
> The rules forming the ruleset must be creating incidents that use event compression policies.

To enable the policies to compress events across multiple event types from multiple rules in the ruleset, select **Allow policies to compress events across event types from multiple rules in the rule set (Ruleset Level Compression)**.

Example:

Let's say you have two event rules defined. The first rule is to create an incident, using Event Compression Policies, for target availability events for a database instance. The second rule is to create an incident, using Event Compression Policies, for metric evaluation error events for a database instance. In order to compress both the target availability and metric evaluation error events across these two event rules into a single incident, select option **Allow policies to compress events across event types from multiple rules in the rule set (Ruleset Level Compression)**. This will work with the corresponding Event Compression Policy to compress **both** target availability events and metric evaluation events for a database instance into an incident.

- Compress Target Availability Events for Database Instance
- Compress Metric Evaluation Error Events for Database Instance

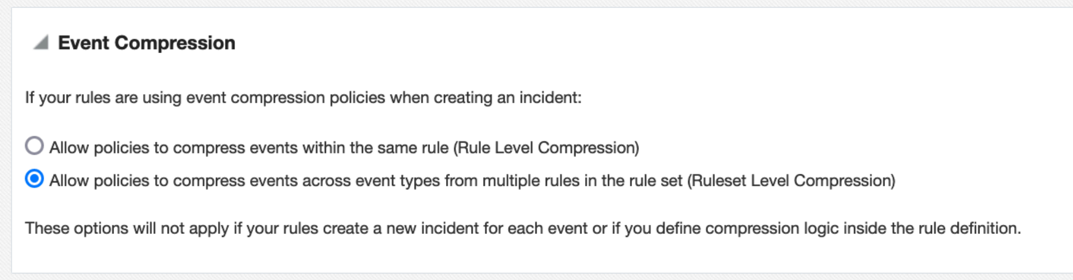

Event Compression

If your rules are using event compression policies when creating an incident:

○ Allow policies to compress events within the same rule (Rule Level Compression)

● Allow policies to compress events across event types from multiple rules in the rule set (Ruleset Level Compression)

These options will not apply if your rules create a new incident for each event or if you define compression logic inside the rule definition.

Enabling Event Compression Policies from Incident Rules Page

You can enable/disable Event Compression policies for specific rules in different Rule Sets from the *Incident Rules-All Enterprise Rules* page.

From the Enterprise Manager console, click **Setup**, then **Incidents**, and then select **Incident Rules**. The Incident Rules - All Enterprise Rules page displays.

1. Select a rule set and click on an individual rule in the table.
2. From the **Actions** menu, select **Enable Event Compression** or, if disabling a compression policy, select **Disable Event Compression**.

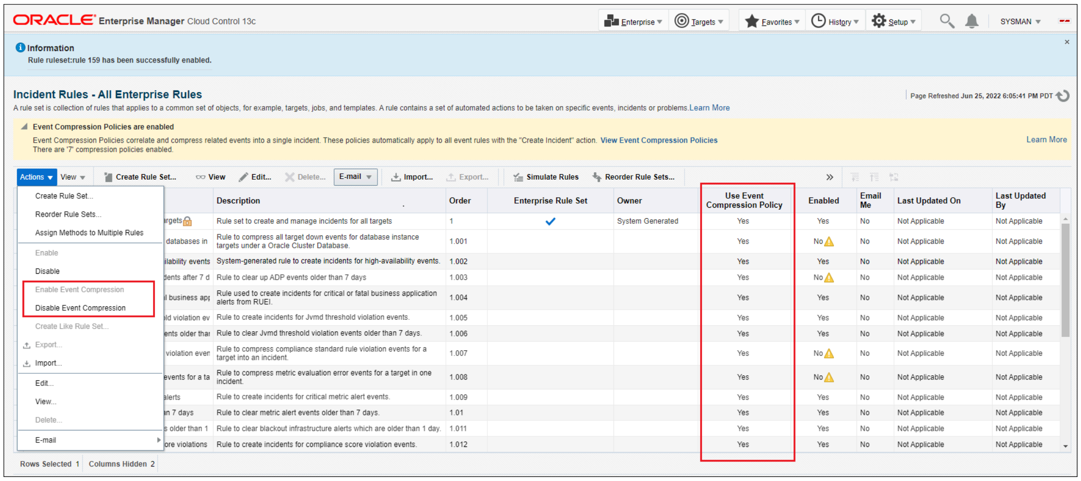

Information
Rule ruleset:rule 159 has been successfully enabled.

Incident Rules - All Enterprise Rules

A rule set is collection of rules that applies to a common set of objects, for example, targets, jobs, and templates. A rule contains a set of automated actions to be taken on specific events, incidents or problems.Learn More

Page Refreshed Jun 25, 2022 6:05:41 PM PDT

Event Compression Policies are enabled

Event Compression Policies correlate and compress related events into a single incident. These policies automatically apply to all event rules with the "Create Incident" action. View Event Compression Policies

There are '7' compression policies enabled.

Learn More

	Description	Order	Enterprise Rule Set	Owner	Use Event Compression Policy	Enabled	Email Me	Last Updated On	Last Updated By
argets	Rule set to create and manage incidents for all targets	1	✓	System Generated	Yes	Yes	No	Not Applicable	Not Applicable
databases in	Rule to compress all target down events for database instance targets under a Oracle Cluster Database.	1.001			Yes	No ⚠	No	Not Applicable	Not Applicable
ilability events	System-generated rule to create incidents for high-availability events.	1.002			Yes	Yes	No	Not Applicable	Not Applicable
dents after 7 d	Rule to clear up ADP events older than 7 days	1.003			Yes	No ⚠	No	Not Applicable	Not Applicable
l business apr	Rule used to create incidents for critical or fatal business application alerts from RUEI.	1.004			Yes	Yes	No	Not Applicable	Not Applicable
ld violation ev	Rule to create incidents for Jvmd threshold violation events.	1.005			Yes	Yes	No	Not Applicable	Not Applicable
ents older thar	Rule to clear Jvmd threshold violation events older than 7 days.	1.006			Yes	Yes	No	Not Applicable	Not Applicable
violation even	Rule to compress compliance standard rule violation events for a target into an incident.	1.007			Yes	No ⚠	No	Not Applicable	Not Applicable
events for a ta	Rule to compress metric evaluation error events for a target in one incident.	1.008			Yes	No ⚠	No	Not Applicable	Not Applicable
alerts	Rule to create incidents for critical metric alert events.	1.009			Yes	Yes	No	Not Applicable	Not Applicable
n 7 days	Rule to clear metric alert events older than 7 days.	1.01			Yes	Yes	No	Not Applicable	Not Applicable
s older than 1	Rule to clear blackout infrastructure alerts which are older than 1 day.	1.011			Yes	Yes	No	Not Applicable	Not Applicable
ore violations	Rule to create incidents for compliance score violation events.	1.012			Yes	Yes	No	Not Applicable	Not Applicable

Rows Selected 1 Columns Hidden 2

Enabling Event Compression Policies for Rules in an Individual Rule Set

You can also enable/disable Event Compression Policies for rules within a Rule Set.

From the Enterprise Manager console, click **Setup**, then **Incidents**, and then select **Incident Rules**. The *Incident Rules - All Enterprise Rules* page displays.

1. Click on an individual rule or rule set in the table.
2. Click **Edit** in the menu bar. The *Edit Rule Set* page displays.
3. Select a rule from the *Rules* region.

4. From the **Actions** menu, select **Enable Event Compression** or, if disabling a compression policy, select **Disable Event Compression**.

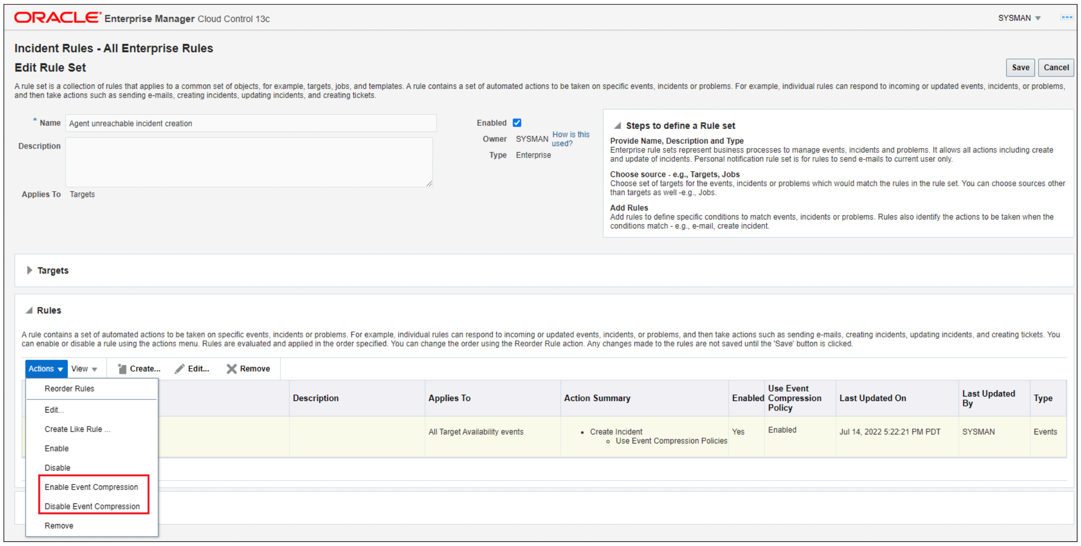

Incident Rules - All Enterprise Rules

Edit Rule Set

A rule set is a collection of rules that applies to a common set of objects, for example, targets, jobs, and templates. A rule contains a set of automated actions to be taken on specific events, incidents or problems. For example, individual rules can respond to incoming or updated events, incidents, or problems, and then take actions such as sending e-mails, creating incidents, updating incidents, and creating tickets.

* Name: Agent unreachable incident creation
Description:
Applies To: Targets
Enabled: ☑
Owner: SYSMAN
Type: Enterprise

Steps to define a Rule set

Provide Name, Description and Type
Enterprise rule sets represent business processes to manage events, incidents and problems. It allows all actions including create and update of incidents. Personal notification rule set is for rules to send e-mails to current user only.

Choose source - e.g., Targets, Jobs
Choose set of targets for the events, incidents or problems which would match the rules in the rule set. You can choose sources other than targets as well -e.g., Jobs.

Add Rules
Add rules to define specific conditions to match events, incidents or problems. Rules also identify the actions to be taken when the conditions match - e.g., e-mail, create incident.

Rules

A rule contains a set of automated actions to be taken on specific events, incidents or problems. For example, individual rules can respond to incoming or updated events, incidents, or problems, and then take actions such as sending e-mails, creating incidents, updating incidents, and creating tickets. You can enable or disable a rule using the actions menu. Rules are evaluated and applied in the order specified. You can change the order using the Reorder Rule action. Any changes made to the rules are not saved until the 'Save' button is clicked.

Description	Applies To	Action Summary	Enabled	Use Event Compression Policy	Last Updated On	Last Updated By	Type
	All Target Availability events	• Create Incident ∘ Use Event Compression Policies	Yes	Enabled	Jul 14, 2022 5:22:21 PM PDT	SYSMAN	Events

Showing Incidents with Compressed Events in Incident Manager

The Incident Manager UI provides direct access to Event Compression details for a given incident.

1. From the Enterprise menu, select **Monitoring** and then **Incident Manager**. The Incident Manager dashboard displays.
2. Select an incident with multiple events. *Incident Details* are displayed showing the number of events that have been compressed.

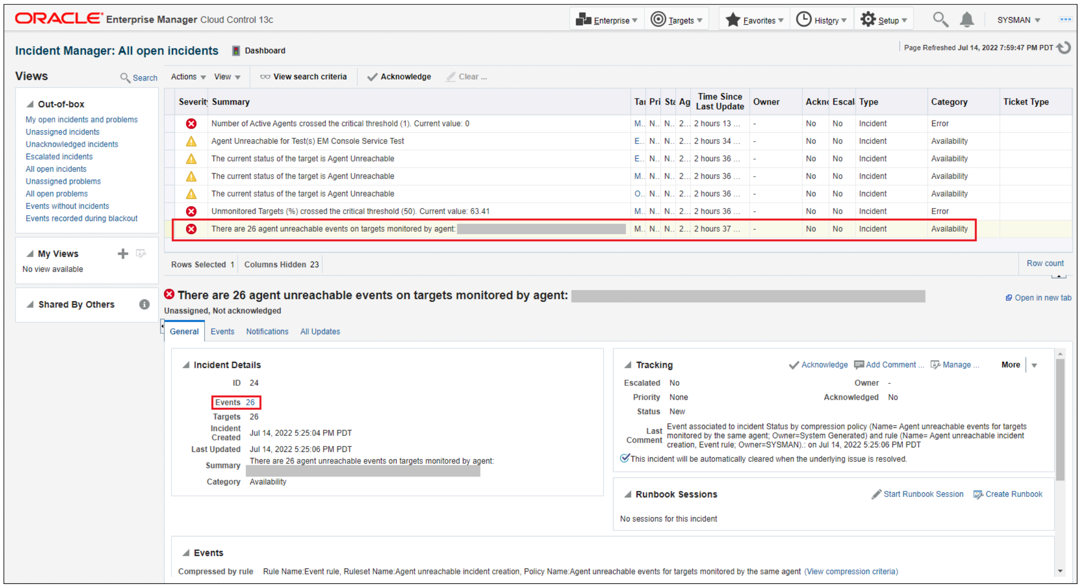

Incident Manager: All open incidents

Page Refreshed Jul 14, 2022 7:59:47 PM PDT

Severit	Summary	Tai	Pri	Sta	Ag	Time Since Last Update	Owner	Ackn	Escal	Type	Category	Ticket Type
✖	Number of Active Agents crossed the critical threshold (1). Current value: 0	M..	N..	N..	2...	2 hours 13 ...	-	No	No	Incident	Error	
⚠	Agent Unreachable for Test(s) EM Console Service Test	E..	N..	N..	2...	2 hours 34 ...	-	No	No	Incident	Availability	
⚠	The current status of the target is Agent Unreachable	E..	N..	N..	2...	2 hours 36 ...	-	No	No	Incident	Availability	
⚠	The current status of the target is Agent Unreachable	M..	N..	N..	2...	2 hours 36 ...	-	No	No	Incident	Availability	
⚠	The current status of the target is Agent Unreachable	O..	N..	N..	2...	2 hours 36 ...	-	No	No	Incident	Availability	
✖	Unmonitored Targets (%) crossed the critical threshold (50). Current value: 63.41	M..	N..	N..	2...	2 hours 36 ...	-	No	No	Incident	Error	
✖	There are 26 agent unreachable events on targets monitored by agent: [redacted]	M..	N..	N..	2...	2 hours 37 ...	-	No	No	Incident	Availability	

Rows Selected 1 Columns Hidden 23

✖ There are 26 agent unreachable events on targets monitored by agent: [redacted]

Unassigned, Not acknowledged

Incident Details

ID: 24
Events: 26
Targets: 26
Incident Created: Jul 14, 2022 5:25:04 PM PDT
Last Updated: Jul 14, 2022 5:25:06 PM PDT
Summary: There are 26 agent unreachable events on targets monitored by agent: [redacted]
Category: Availability

Tracking

Escalated: No
Priority: None
Status: New
Owner: -
Acknowledged: No
Last Comment: Event associated to incident Status by compression policy (Name= Agent unreachable events for targets monitored by the same agent; Owner=System Generated) and rule (Name= Agent unreachable incident creation, Event rule; Owner=SYSMAN).: on Jul 14, 2022 5:25:06 PM PDT

✓This incident will be automatically cleared when the underlying issue is resolved.

Runbook Sessions

No sessions for this incident

Events

Compressed by rule: Rule Name:Event rule, Ruleset Name:Agent unreachable incident creation, Policy Name:Agent unreachable events for targets monitored by the same agent (View compression criteria)

3. In the General tab, scroll down to the **Events** region. This region displays the triggered events within the incident.

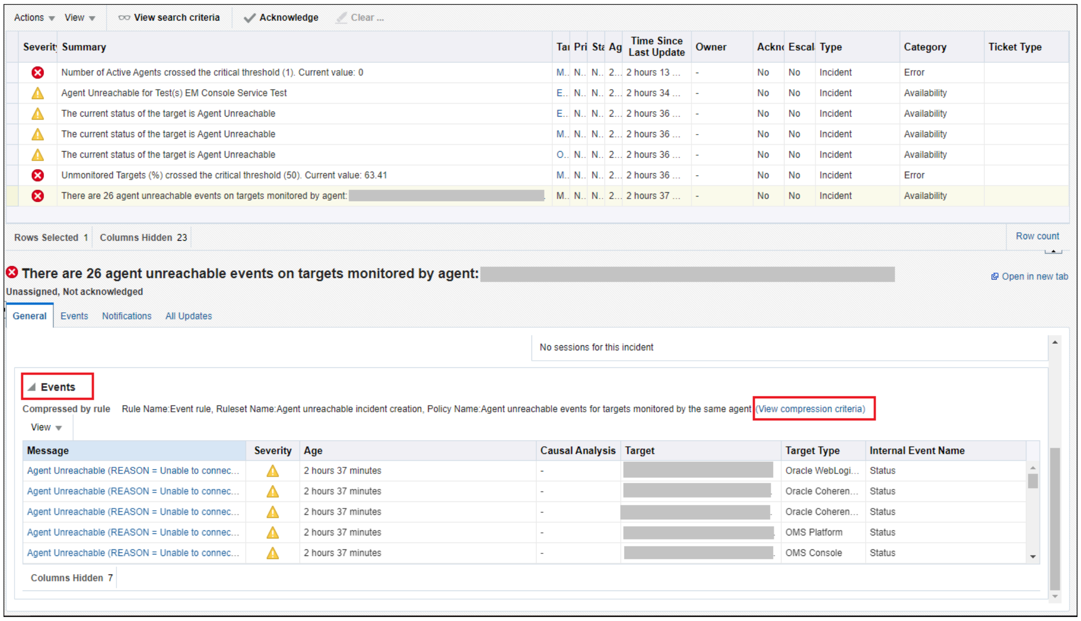

Severit	Summary	Tai	Pri	Sta	Ag	Time Since Last Update	Owner	Ackno	Escal	Type	Category	Ticket Type
⊗	Number of Active Agents crossed the critical threshold (1). Current value: 0	M..	N..	N..	2...	2 hours 13 ...	-	No	No	Incident	Error	
⚠	Agent Unreachable for Test(s) EM Console Service Test	E..	N..	N..	2...	2 hours 34 ...	-	No	No	Incident	Availability	
⚠	The current status of the target is Agent Unreachable	E..	N..	N..	2...	2 hours 36 ...	-	No	No	Incident	Availability	
⚠	The current status of the target is Agent Unreachable	M..	N..	N..	2...	2 hours 36 ...	-	No	No	Incident	Availability	
⚠	The current status of the target is Agent Unreachable	O..	N..	N..	2...	2 hours 36 ...	-	No	No	Incident	Availability	
⊗	Unmonitored Targets (%) crossed the critical threshold (50). Current value: 63.41	M..	N..	N..	2...	2 hours 36 ...	-	No	No	Incident	Error	
⊗	There are 26 agent unreachable events on targets monitored by agent: [redacted]	M..	N..	N..	2...	2 hours 37 ...	-	No	No	Incident	Availability	

Rows Selected 1 Columns Hidden 23

There are 26 agent unreachable events on targets monitored by agent: [redacted]

Unassigned, Not acknowledged

No sessions for this incident

Events

Compressed by rule Rule Name:Event rule, Ruleset Name:Agent unreachable incident creation, Policy Name:Agent unreachable events for targets monitored by the same agent (View compression criteria)

Message	Severity	Age	Causal Analysis	Target	Target Type	Internal Event Name
Agent Unreachable (REASON = Unable to connec...	⚠	2 hours 37 minutes	-	[redacted]	Oracle WebLogi...	Status
Agent Unreachable (REASON = Unable to connec...	⚠	2 hours 37 minutes	-	[redacted]	Oracle Coheren...	Status
Agent Unreachable (REASON = Unable to connec...	⚠	2 hours 37 minutes	-	[redacted]	Oracle Coheren...	Status
Agent Unreachable (REASON = Unable to connec...	⚠	2 hours 37 minutes	-	[redacted]	OMS Platform	Status
Agent Unreachable (REASON = Unable to connec...	⚠	2 hours 37 minutes	-	[redacted]	OMS Console	Status

Columns Hidden 7

4. Click **View compression criteria**. The *Compression Criteria* dialog is displayed showing the compression criteria used to group the events into a single incident.

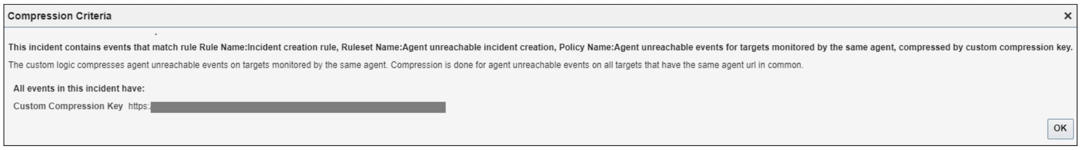

This incident contains events that match rule Rule Name:Incident creation rule, Ruleset Name:Agent unreachable incident creation, Policy Name:Agent unreachable events for targets monitored by the same agent, compressed by custom compression key.

The custom logic compresses agent unreachable events on targets monitored by the same agent. Compression is done for agent unreachable events on all targets that have the same agent url in common.

All events in this incident have:

Custom Compression Key https:[redacted]

5. Click **OK** to close the dialog.
6. Click on the Events tab to view a more detailed view of the compressed events for the incident.

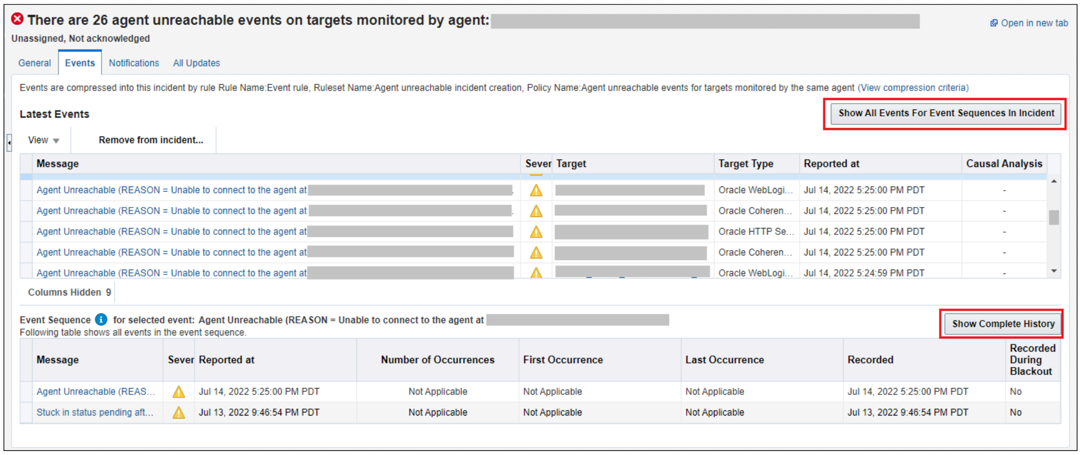

There are 26 agent unreachable events on targets monitored by agent: [redacted]

Unassigned, Not acknowledged

Events are compressed into this incident by rule Rule Name:Event rule, Ruleset Name:Agent unreachable incident creation, Policy Name:Agent unreachable events for targets monitored by the same agent (View compression criteria)

Latest Events

Message	Sever	Target	Target Type	Reported at	Causal Analysis
Agent Unreachable (REASON = Unable to connect to the agent at [redacted]	⚠	[redacted]	Oracle WebLogi...	Jul 14, 2022 5:25:00 PM PDT	-
Agent Unreachable (REASON = Unable to connect to the agent at [redacted]	⚠	[redacted]	Oracle Coheren...	Jul 14, 2022 5:25:00 PM PDT	-
Agent Unreachable (REASON = Unable to connect to the agent at [redacted]	⚠	[redacted]	Oracle HTTP Se...	Jul 14, 2022 5:25:00 PM PDT	-
Agent Unreachable (REASON = Unable to connect to the agent at [redacted]	⚠	[redacted]	Oracle Coheren...	Jul 14, 2022 5:25:00 PM PDT	-
Agent Unreachable (REASON = Unable to connect to the agent at [redacted]	⚠	[redacted]	Oracle WebLogi...	Jul 14, 2022 5:24:59 PM PDT	-

Columns Hidden 9

Event Sequence ⓘ for selected event: Agent Unreachable (REASON = Unable to connect to the agent at [redacted]
Following table shows all events in the event sequence.

Message	Sever	Reported at	Number of Occurrences	First Occurrence	Last Occurrence	Recorded	Recorded During Blackout
Agent Unreachable (REAS...	⚠	Jul 14, 2022 5:25:00 PM PDT	Not Applicable	Not Applicable	Not Applicable	Jul 14, 2022 5:25:00 PM PDT	No
Stuck in status pending aft...	⚠	Jul 13, 2022 9:46:54 PM PDT	Not Applicable	Not Applicable	Not Applicable	Jul 13, 2022 9:46:54 PM PDT	No

You can:

- Show all events for event sequences in the incident. You can choose to show events in the current sequence or show the complete history of the event sequence.
- Show only the latest events.

Importing or Exporting Event Compression Policies

Import/Export feature in user defined compression policies enables users to export compression policies which are created by users where the policy details will be exported as a `json` file, similarly user can also create a compression policy by importing compression policies where the importing file must be in `json` format.

To import a compression policy from a `json` file:

1. Click the **Import Policy** button at the top right of the table.

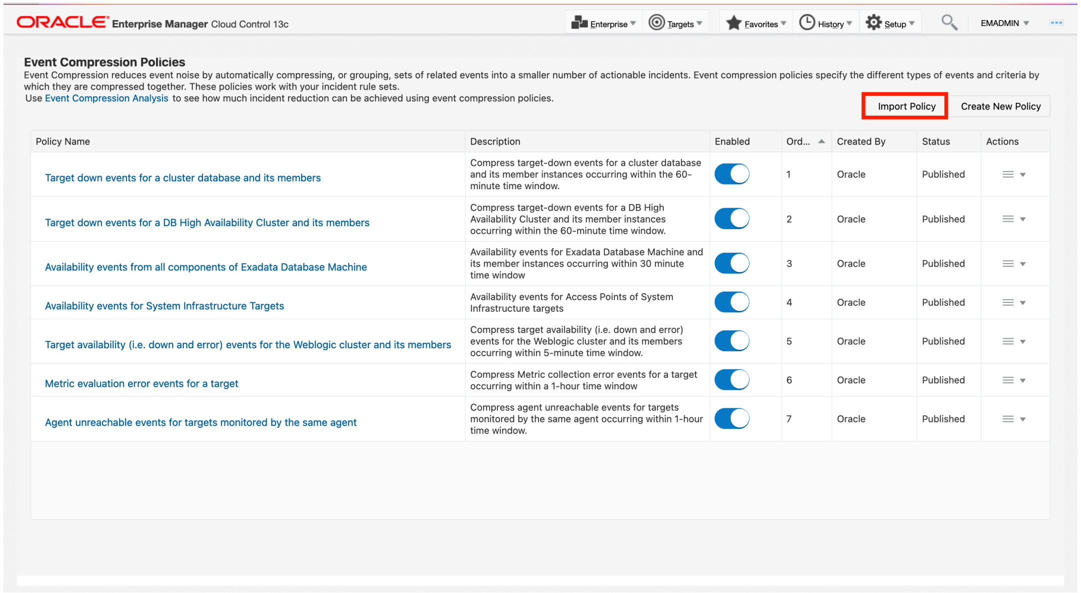

Event Compression Policies

Event Compression reduces event noise by automatically compressing, or grouping, sets of related events into a smaller number of actionable incidents. Event compression policies specify the different types of events and criteria by which they are compressed together. These policies work with your incident rule sets.

Use Event Compression Analysis to see how much incident reduction can be achieved using event compression policies.

Policy Name	Description	Enabled	Ord...	Created By	Status	Actions
Target down events for a cluster database and its members	Compress target-down events for a cluster database and its member instances occurring within the 60-minute time window.	●	1	Oracle	Published	≡
Target down events for a DB High Availability Cluster and its members	Compress target-down events for a DB High Availability Cluster and its member instances occurring within the 60-minute time window.	●	2	Oracle	Published	≡
Availability events from all components of Exadata Database Machine	Availability events for Exadata Database Machine and its member instances occurring within 30 minute time window	●	3	Oracle	Published	≡
Availability events for System Infrastructure Targets	Availability events for Access Points of System Infrastructure targets	●	4	Oracle	Published	≡
Target availability (i.e. down and error) events for the Weblogic cluster and its members	Compress target availability (i.e. down and error) events for the Weblogic cluster and its members occurring within 5-minute time window.	●	5	Oracle	Published	≡
Metric evaluation error events for a target	Compress Metric collection error events for a target occurring within a 1-hour time window	●	6	Oracle	Published	≡
Agent unreachable events for targets monitored by the same agent	Compress agent unreachable events for targets monitored by the same agent occurring within 1-hour time window.	●	7	Oracle	Published	≡

2. Select a file or drop one in the dialog box.

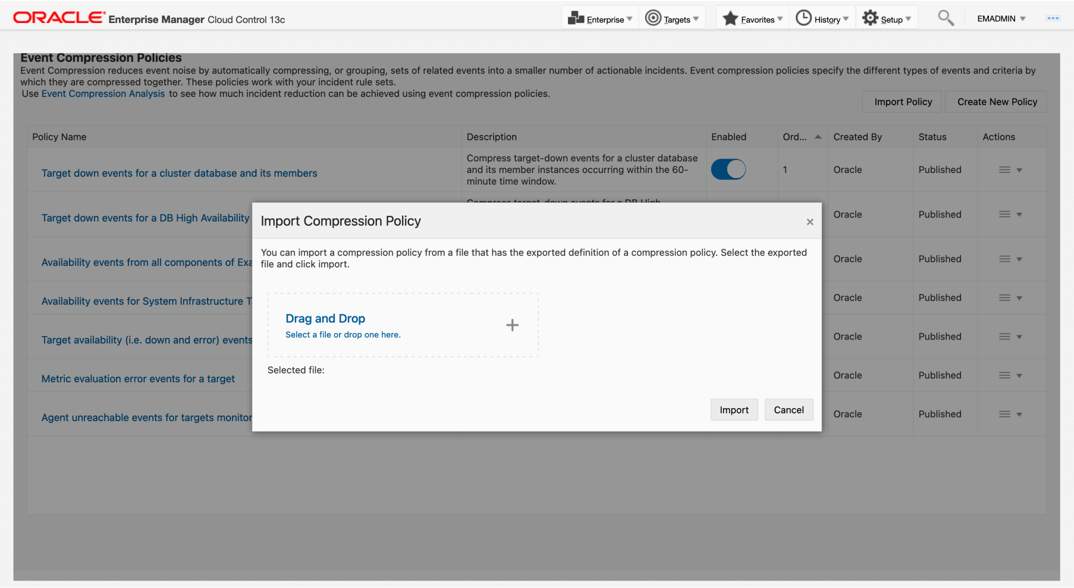

Import Compression Policy

You can import a compression policy from a file that has the exported definition of a compression policy. Select the exported file and click import.

Select a file or drop one here.

Selected file:

3. Check the import status.

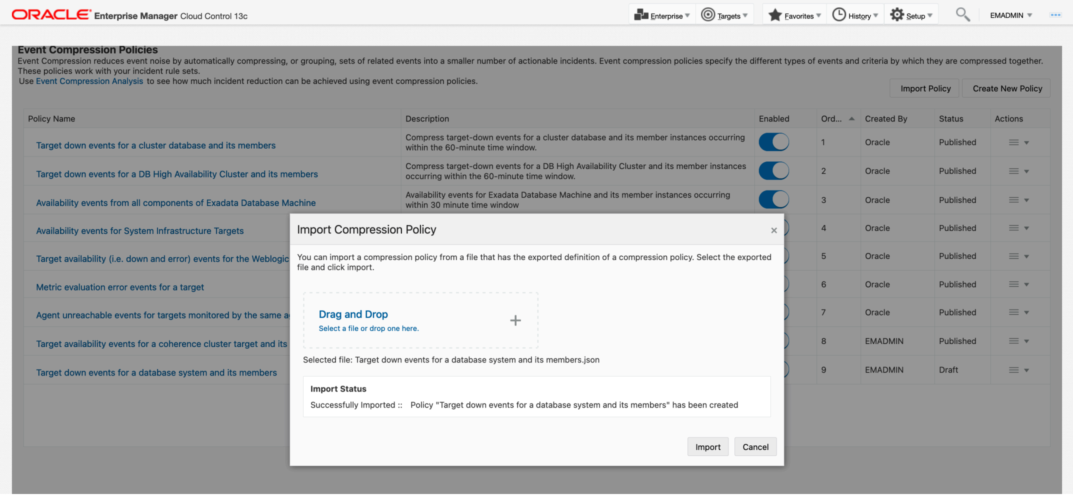

Event Compression Policies

Event Compression reduces event noise by automatically compressing, or grouping, sets of related events into a smaller number of actionable incidents. Event compression policies specify the different types of events and criteria by which they are compressed together. These policies work with your incident rule sets.
Use Event Compression Analysis to see how much incident reduction can be achieved using event compression policies.

Policy Name	Description	Enabled	Ord...	Created By	Status	Actions
Target down events for a cluster database and its members	Compress target-down events for a cluster database and its member instances occurring within the 60-minute time window.		1	Oracle	Published	
Target down events for a DB High Availability Cluster and its members	Compress target-down events for a DB High Availability Cluster and its member instances occurring within the 60-minute time window.		2	Oracle	Published	
Availability events from all components of Exadata Database Machine	Availability events for Exadata Database Machine and its member instances occurring within 30 minute time window		3	Oracle	Published	
Availability events for System Infrastructure Targets			4	Oracle	Published	
Target availability (i.e. down and error) events for the Weblogic			5	Oracle	Published	
Metric evaluation error events for a target			6	Oracle	Published	
Agent unreachable events for targets monitored by the same a			7	Oracle	Published	
Target availability events for a coherence cluster target and its			8	EMADMIN	Published	
Target down events for a database system and its members			9	EMADMIN	Draft	

Import Compression Policy

You can import a compression policy from a file that has the exported definition of a compression policy. Select the exported file and click import.

Drag and Drop
Select a file or drop one here.

Selected file: Target down events for a database system and its members.json

Import Status
Successfully Imported :: Policy "Target down events for a database system and its members" has been created

4. If the **Import Status** is *Successfully Imported*, click **Import**

To export a compression policy to a `json` file:

1. Open the drop-down **Actions** menu of the policy you want to export and click **Export**.

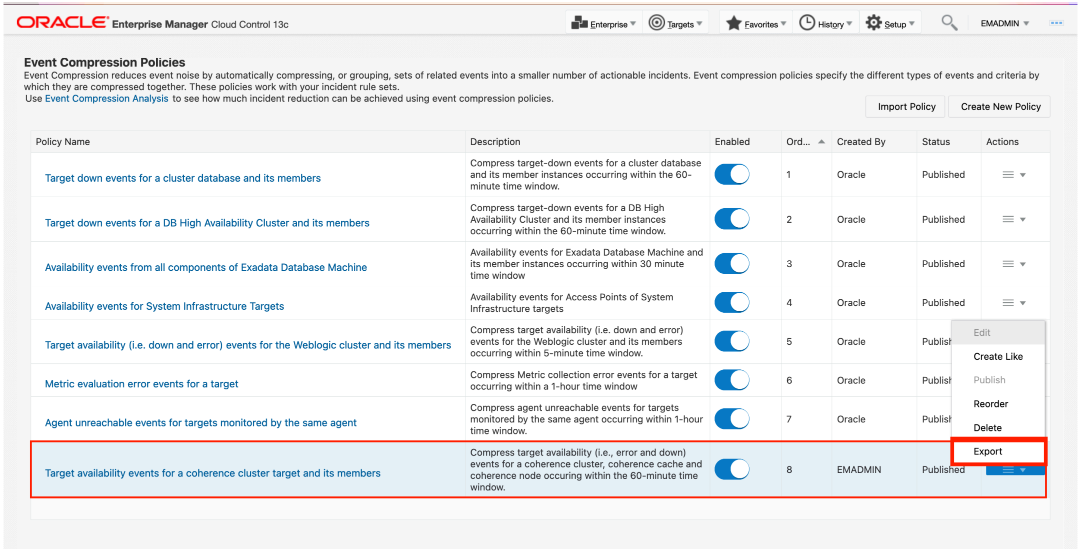

Event Compression Policies

Event Compression reduces event noise by automatically compressing, or grouping, sets of related events into a smaller number of actionable incidents. Event compression policies specify the different types of events and criteria by which they are compressed together. These policies work with your incident rule sets.
Use Event Compression Analysis to see how much incident reduction can be achieved using event compression policies.

Policy Name	Description	Enabled	Ord...	Created By	Status	Actions
Target down events for a cluster database and its members	Compress target-down events for a cluster database and its member instances occurring within the 60-minute time window.		1	Oracle	Published	
Target down events for a DB High Availability Cluster and its members	Compress target-down events for a DB High Availability Cluster and its member instances occurring within the 60-minute time window.		2	Oracle	Published	
Availability events from all components of Exadata Database Machine	Availability events for Exadata Database Machine and its member instances occurring within 30 minute time window		3	Oracle	Published	
Availability events for System Infrastructure Targets	Availability events for Access Points of System Infrastructure targets		4	Oracle	Published	
Target availability (i.e. down and error) events for the Weblogic cluster and its members	Compress target availability (i.e. down and error) events for the Weblogic cluster and its members occurring within 5-minute time window.		5	Oracle	Publish	
Metric evaluation error events for a target	Compress Metric collection error events for a target occurring within a 1-hour time window		6	Oracle	Publish	
Agent unreachable events for targets monitored by the same agent	Compress agent unreachable events for targets monitored by the same agent occurring within 1-hour time window.		7	Oracle	Publish	
Target availability events for a coherence cluster target and its members	Compress target availability (i.e., error and down) events for a coherence cluster, coherence cache and coherence node occuring within the 60-minute time window.		8	EMADMIN	Published	

2. Enter the a file name, and click **Ok**.

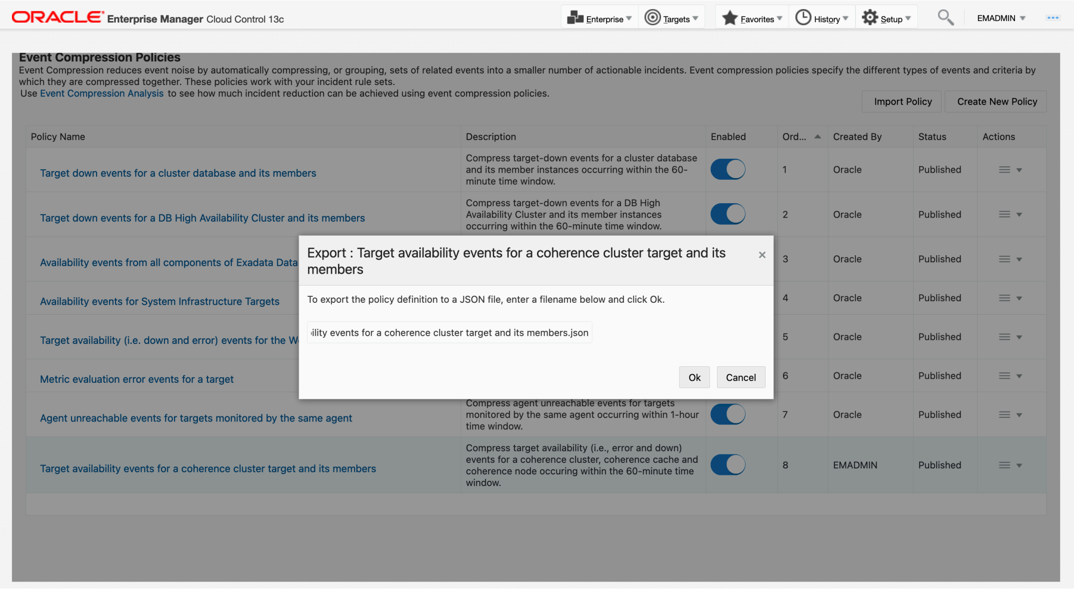

Event Compression Policies

Event Compression reduces event noise by automatically compressing, or grouping, sets of related events into a smaller number of actionable incidents. Event compression policies specify the different types of events and criteria by which they are compressed together. These policies work with your incident rule sets.
Use Event Compression Analysis to see how much incident reduction can be achieved using event compression policies.

Policy Name	Description	Enabled	Ord... ▲	Created By	Status	Actions
Target down events for a cluster database and its members	Compress target-down events for a cluster database and its member instances occurring within the 60-minute time window.	●	1	Oracle	Published	≡ ▾
Target down events for a DB High Availability Cluster and its members	Compress target-down events for a DB High Availability Cluster and its member instances occurring within the 60-minute time window.	●	2	Oracle	Published	≡ ▾
Availability events from all components of Exadata Data			3	Oracle	Published	≡ ▾
Availability events for System Infrastructure Targets			4	Oracle	Published	≡ ▾
Target availability (i.e. down and error) events for the W			5	Oracle	Published	≡ ▾
Metric evaluation error events for a target			6	Oracle	Published	≡ ▾
Agent unreachable events for targets monitored by the same agent	Compress agent unreachable events for targets monitored by the same agent occurring within 1-hour time window.	●	7	Oracle	Published	≡ ▾
Target availability events for a coherence cluster target and its members	Compress target availability (i.e., error and down) events for a coherence cluster, coherence cache and coherence node occuring within the 60-minute time window.	●	8	EMADMIN	Published	≡ ▾

Export : Target availability events for a coherence cluster target and its members ×

To export the policy definition to a JSON file, enter a filename below and click Ok.

ility events for a coherence cluster target and its members.json

3. After the policy is succesfully exported a confirmation pop-up will show.

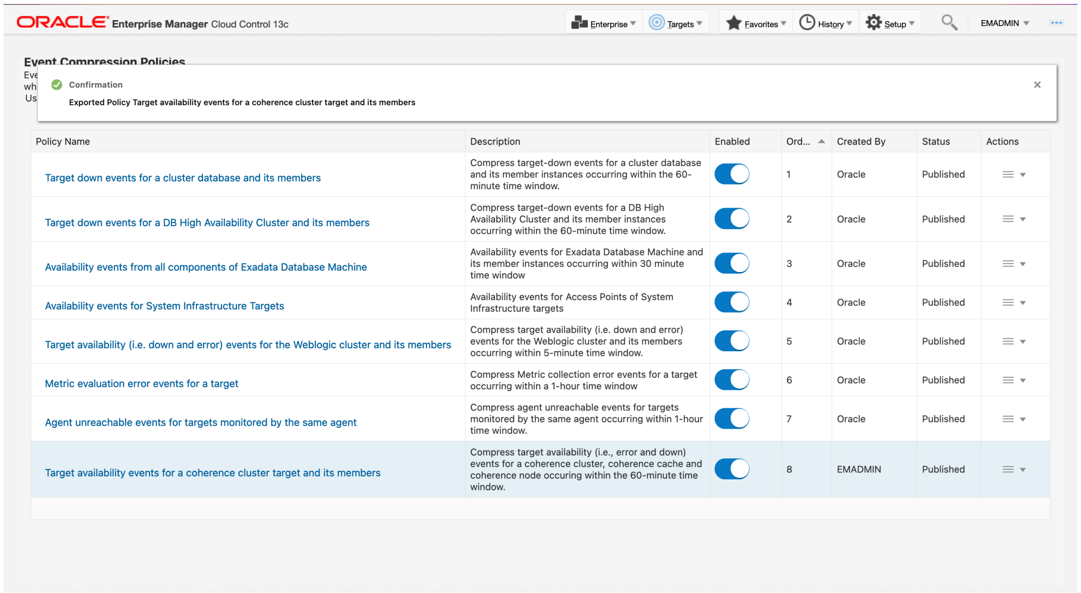

Event Compression Policies

Eve
whi
Us

Confirmation ×

Exported Policy Target availability events for a coherence cluster target and its members

Policy Name	Description	Enabled	Ord... ▲	Created By	Status	Actions
Target down events for a cluster database and its members	Compress target-down events for a cluster database and its member instances occurring within the 60-minute time window.	●	1	Oracle	Published	≡ ▾
Target down events for a DB High Availability Cluster and its members	Compress target-down events for a DB High Availability Cluster and its member instances occurring within the 60-minute time window.	●	2	Oracle	Published	≡ ▾
Availability events from all components of Exadata Database Machine	Availability events for Exadata Database Machine and its member instances occurring within 30 minute time window	●	3	Oracle	Published	≡ ▾
Availability events for System Infrastructure Targets	Availability events for Access Points of System Infrastructure targets	●	4	Oracle	Published	≡ ▾
Target availability (i.e. down and error) events for the Weblogic cluster and its members	Compress target availability (i.e. down and error) events for the Weblogic cluster and its members occurring within 5-minute time window.	●	5	Oracle	Published	≡ ▾
Metric evaluation error events for a target	Compress Metric collection error events for a target occurring within a 1-hour time window	●	6	Oracle	Published	≡ ▾
Agent unreachable events for targets monitored by the same agent	Compress agent unreachable events for targets monitored by the same agent occurring within 1-hour time window.	●	7	Oracle	Published	≡ ▾
Target availability events for a coherence cluster target and its members	Compress target availability (i.e., error and down) events for a coherence cluster, coherence cache and coherence node occuring within the 60-minute time window.	●	8	EMADMIN	Published	≡ ▾

Rule-based Event Compression

Rule-based Event Compression is an older method of automatically compressing multiple events into a single incident.

> **Note:**
>
> Oracle recommends using the newer Event Compression Policies method as it is easier to configure. You only need to configure it once and it will apply to all rules, as opposed to individual rules.

Some situations where it is beneficial to deal with multiple events as a single incident are:

- You want automatic consolidation of all *Tablespace Used (%)* alerts across all tablespaces for a specific database into a single incident.
- You want automatic consolidation of all Metric Collection Errors for a target into a single incident.
- You want automatic consolidation of all SOA composite Target Down events within a WebLogic Domain into a single incident.

For convenience, Enterprise Manager provides out-of-box rules that automatically compress related events into single incidents. These rules address some of the most common conditions where event grouping could be helpful.

- Target down for RAC database instances.
- Metric collection errors for a target.
- Configuration standard violations for a rule on a target.

Creating an Incident Compression Rule Set

The following example shows you how to create an incident rule set that generates a single incident (compressed), and notifies an administrator of via email that a compressed incident has been generated.

1. Click **Setup** > **Incidents** > **Incident Rules**.

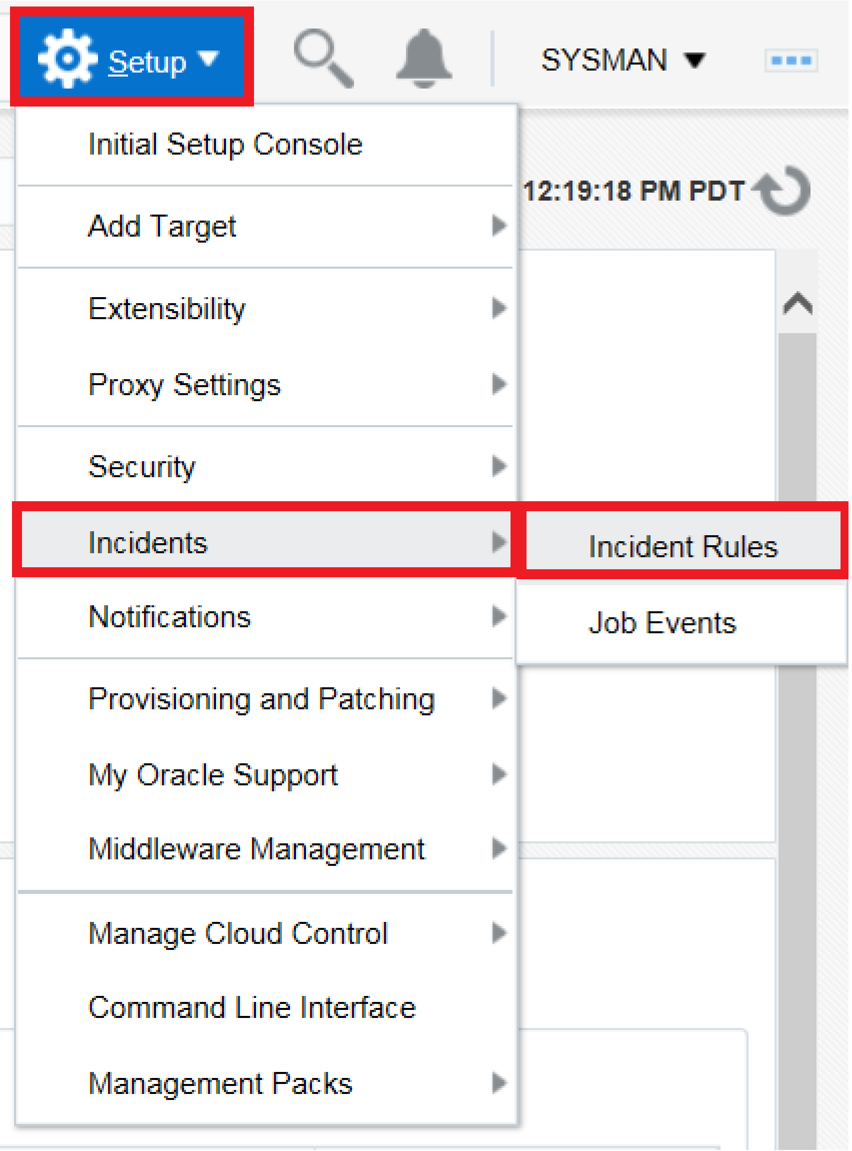

2. Click **Create Rule Set**.

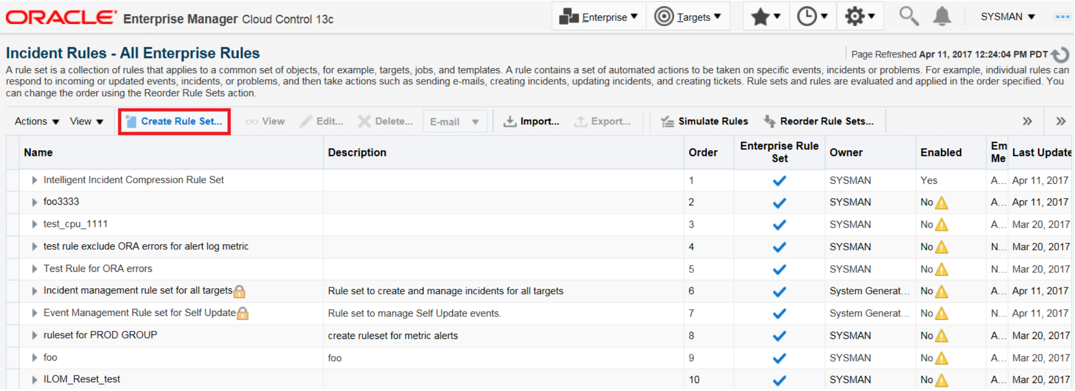

Incident Rules - All Enterprise Rules

Page Refreshed **Apr 11, 2017 12:24:04 PM PDT**

A rule set is a collection of rules that applies to a common set of objects, for example, targets, jobs, and templates. A rule contains a set of automated actions to be taken on specific events, incidents or problems. For example, individual rules can respond to incoming or updated events, incidents, or problems, and then take actions such as sending e-mails, creating incidents, updating incidents, and creating tickets. Rule sets and rules are evaluated and applied in the order specified. You can change the order using the Reorder Rule Sets action.

Name	Description	Order	Enterprise Rule Set	Owner	Enabled	Em Me	Last Update
Intelligent Incident Compression Rule Set		1	✓	SYSMAN	Yes	A...	Apr 11, 2017
foo3333		2	✓	SYSMAN	No ⚠	A...	Apr 11, 2017
test_cpu_1111		3	✓	SYSMAN	No ⚠	A...	Mar 20, 2017
test rule exclude ORA errors for alert log metric		4	✓	SYSMAN	No ⚠	N...	Mar 20, 2017
Test Rule for ORA errors		5	✓	SYSMAN	No ⚠	N...	Mar 20, 2017
Incident management rule set for all targets	Rule set to create and manage incidents for all targets	6	✓	System Generat...	No ⚠	A...	Apr 11, 2017
Event Management Rule set for Self Update	Rule set to manage Self Update events.	7	✓	System Generat...	No ⚠	N...	Apr 11, 2017
ruleset for PROD GROUP	create ruleset for metric alerts	8	✓	SYSMAN	No ⚠	A...	Mar 20, 2017
foo	foo	9	✓	SYSMAN	No ⚠	A...	Mar 20, 2017
ILOM_Reset_test		10	✓	SYSMAN	No ⚠	A...	Mar 20, 2017

3. Enter the name of the Rule Set and select the target the rule applies to. In this example, a Database Instance target is selected.

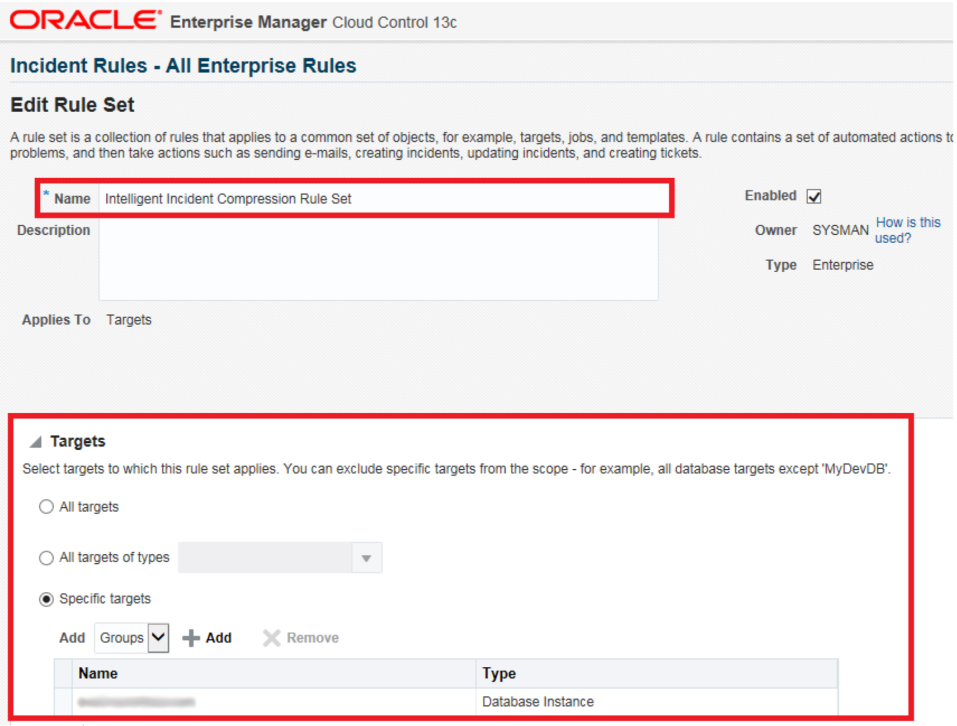

ORACLE Enterprise Manager Cloud Control 13c

Incident Rules - All Enterprise Rules

Edit Rule Set

A rule set is a collection of rules that applies to a common set of objects, for example, targets, jobs, and templates. A rule contains a set of automated actions t problems, and then take actions such as sending e-mails, creating incidents, updating incidents, and creating tickets.

* Name	Intelligent Incident Compression Rule Set
Description	
Applies To	Targets
Enabled	☑
Owner	SYSMAN
Type	Enterprise

Targets

Select targets to which this rule set applies. You can exclude specific targets from the scope - for example, all database targets except 'MyDevDB'.

○ All targets

○ All targets of types

● Specific targets

Name	Type
[redacted]	Database Instance

4. Scroll down to the **Rules** section and click **Create**.

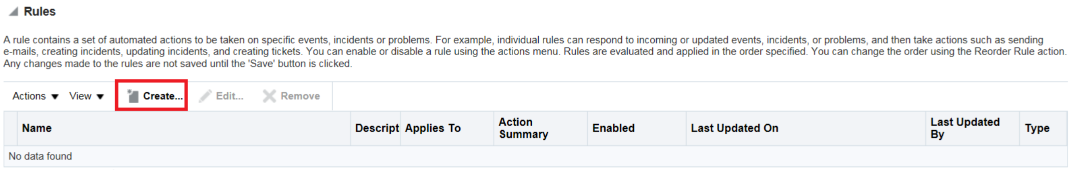

Rules

A rule contains a set of automated actions to be taken on specific events, incidents or problems. For example, individual rules can respond to incoming or updated events, incidents, or problems, and then take actions such as sending e-mails, creating incidents, updating incidents, and creating tickets. You can enable or disable a rule using the actions menu. Rules are evaluated and applied in the order specified. You can change the order using the Reorder Rule action. Any changes made to the rules are not saved until the 'Save' button is clicked.

Name	Descript	Applies To	Action Summary	Enabled	Last Updated On	Last Updated By	Type
No data found							

5. In the **Select Type of Rule to Create** popup window, leave the default (*Incoming events and updates to events*) selected and click **Continue**.

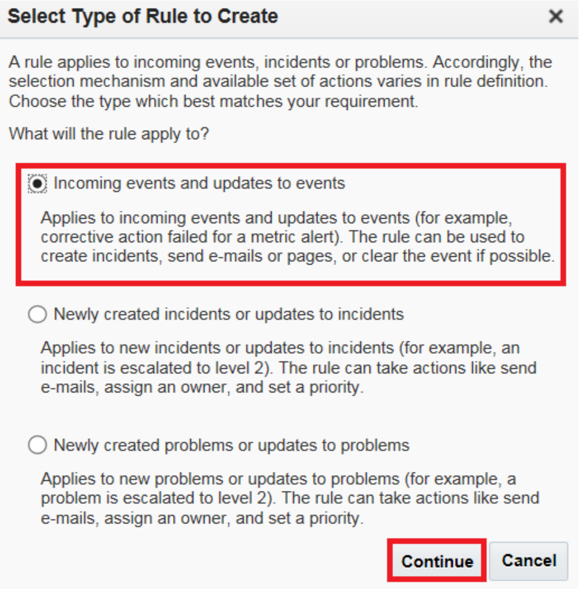

Select Type of Rule to Create ✕

A rule applies to incoming events, incidents or problems. Accordingly, the selection mechanism and available set of actions varies in rule definition. Choose the type which best matches your requirement.

What will the rule apply to?

● Incoming events and updates to events

Applies to incoming events and updates to events (for example, corrective action failed for a metric alert). The rule can be used to create incidents, send e-mails or pages, or clear the event if possible.

○ Newly created incidents or updates to incidents

Applies to new incidents or updates to incidents (for example, an incident is escalated to level 2). The rule can take actions like send e-mails, assign an owner, and set a priority.

○ Newly created problems or updates to problems

Applies to new problems or updates to problems (for example, a problem is escalated to level 2). The rule can take actions like send e-mails, assign an owner, and set a priority.

6. Select **Target Availability** from the Type drop down menu and select **Specific events of type Target Availability**. Click **Add**.

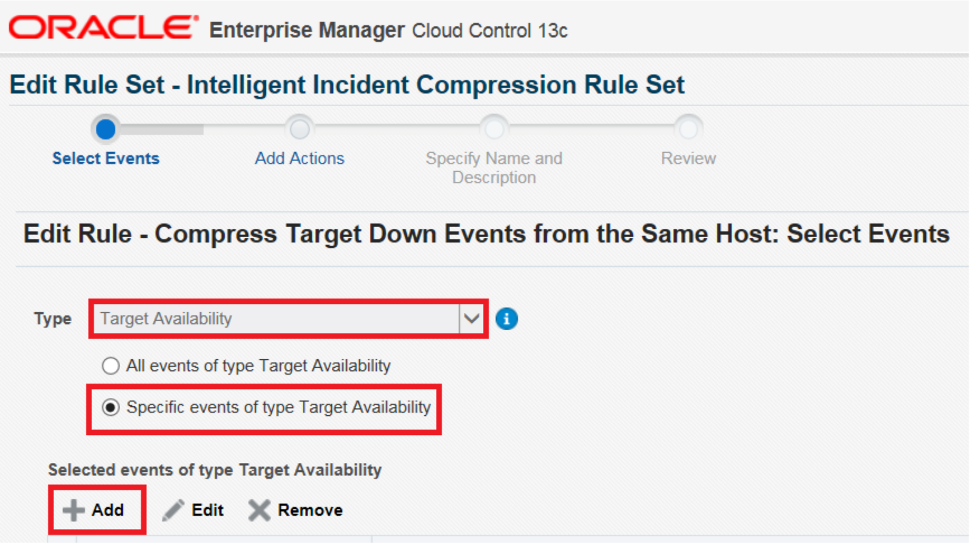

ORACLE Enterprise Manager Cloud Control 13c

Edit Rule Set - Intelligent Incident Compression Rule Set

Select Events
Add Actions
Specify Name and Description
Review

Edit Rule - Compress Target Down Events from the Same Host: Select Events

Type: Target Availability

○ All events of type Target Availability

● Specific events of type Target Availability

Selected events of type Target Availability

7. Check the “Down” checkbox and click **OK**.

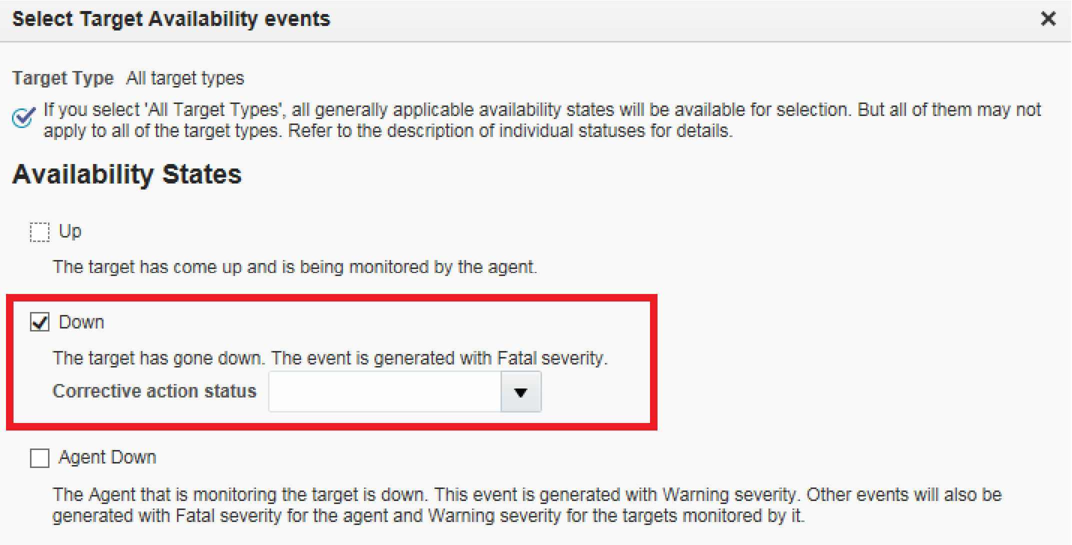

Select Target Availability events ✕

Target Type All target types

If you select 'All Target Types', all generally applicable availability states will be available for selection. But all of them may not apply to all of the target types. Refer to the description of individual statuses for details.

Availability States

□ Up

The target has come up and is being monitored by the agent.

☑ Down

The target has gone down. The event is generated with Fatal severity.

Corrective action status

□ Agent Down

The Agent that is monitoring the target is down. This event is generated with Warning severity. Other events will also be generated with Fatal severity for the agent and Warning severity for the targets monitored by it.

8. Click **Next**.

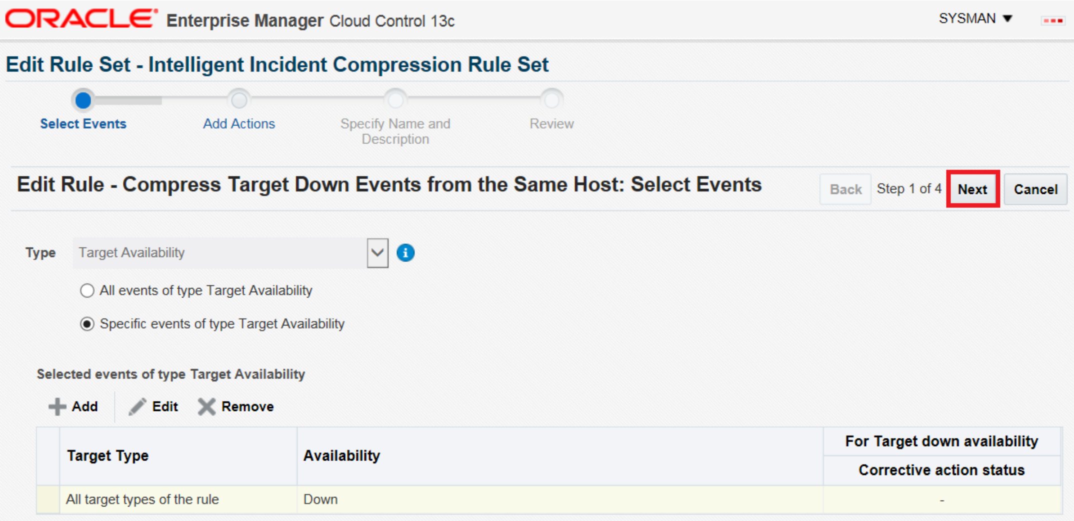

Edit Rule Set - Intelligent Incident Compression Rule Set

Edit Rule - Compress Target Down Events from the Same Host: Select Events

Type Target Availability

○ All events of type Target Availability

● Specific events of type Target Availability

Selected events of type Target Availability

Target Type	Availability	For Target down availability
		Corrective action status
All target types of the rule	Down	-

9. Click **Add** to add Conditional Actions.

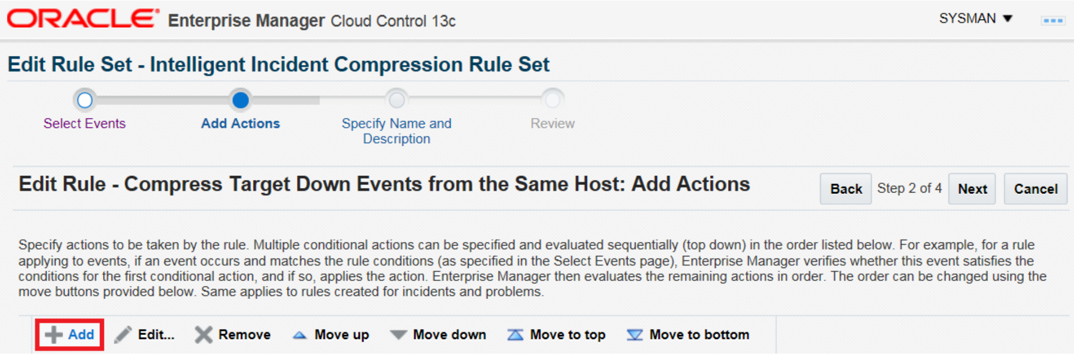

ORACLE Enterprise Manager Cloud Control 13c
SYSMAN ▾

Edit Rule Set - Intelligent Incident Compression Rule Set

Select Events
Add Actions
Specify Name and Description
Review

Edit Rule - Compress Target Down Events from the Same Host: Add Actions

Back
Step 2 of 4
Next
Cancel

Specify actions to be taken by the rule. Multiple conditional actions can be specified and evaluated sequentially (top down) in the order listed below. For example, for a rule applying to events, if an event occurs and matches the rule conditions (as specified in the Select Events page), Enterprise Manager verifies whether this event satisfies the conditions for the first conditional action, and if so, applies the action. Enterprise Manager then evaluates the remaining actions in order. The order can be changed using the move buttons provided below. Same applies to rules created for incidents and problems.

10. Select the conditions for compressing the events and click **Continue**. In this example we are compressing events from targets on the same host.

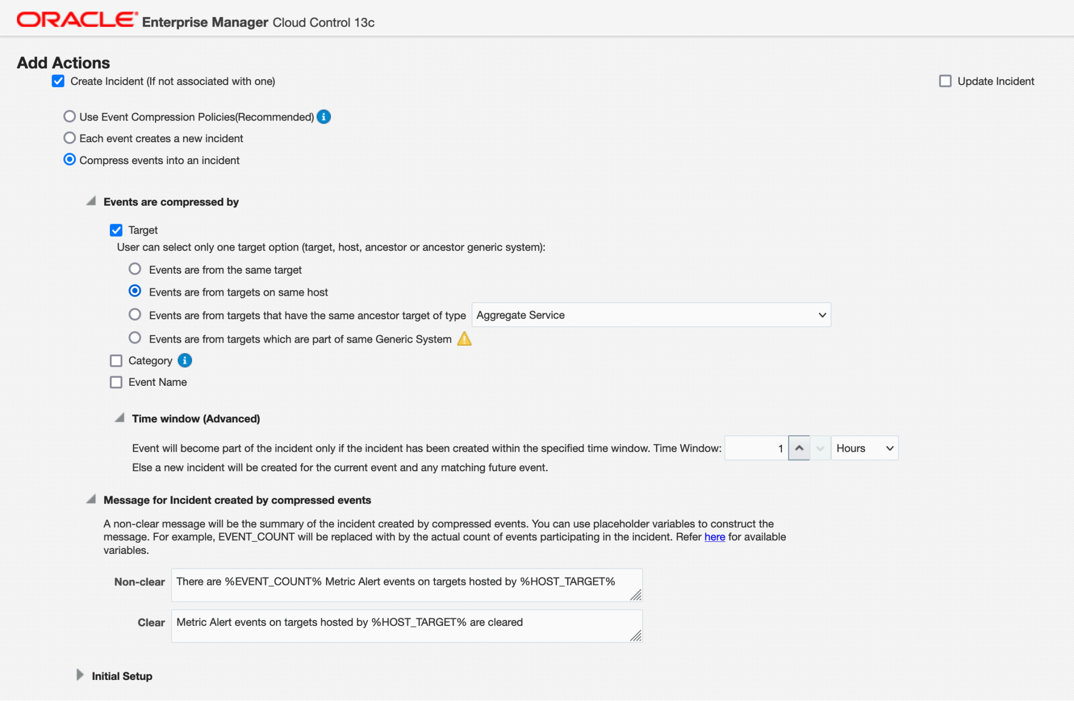

Add Actions

☑ Create Incident (If not associated with one)

☐ Update Incident

○ Use Event Compression Policies(Recommended)

○ Each event creates a new incident

● Compress events into an incident

Events are compressed by

☑ Target

User can select only one target option (target, host, ancestor or ancestor generic system):

○ Events are from the same target

● Events are from targets on same host

○ Events are from targets that have the same ancestor target of type: Aggregate Service

○ Events are from targets which are part of same Generic System

☐ Category

☐ Event Name

Time window (Advanced)

Event will become part of the incident only if the incident has been created within the specified time window. Time Window: 1 Hours

Else a new incident will be created for the current event and any matching future event.

Message for Incident created by compressed events

A non-clear message will be the summary of the incident created by compressed events. You can use placeholder variables to construct the message. For example, EVENT_COUNT will be replaced with by the actual count of events participating in the incident. Refer here for available variables.

Non-clear: There are %EVENT_COUNT% Metric Alert events on targets hosted by %HOST_TARGET%

Clear: Metric Alert events on targets hosted by %HOST_TARGET% are cleared

Initial Setup

11. Click **OK** in the popup window.

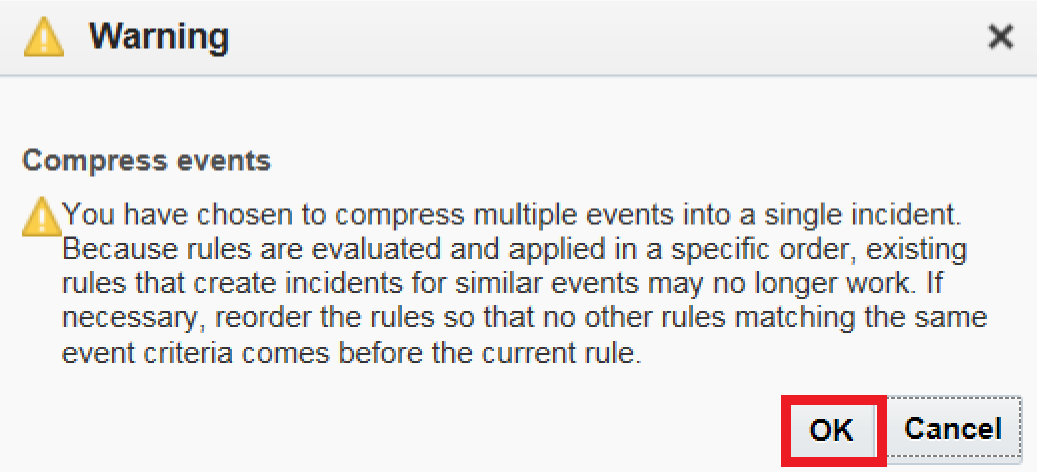

Warning

Compress events

You have chosen to compress multiple events into a single incident. Because rules are evaluated and applied in a specific order, existing rules that create incidents for similar events may no longer work. If necessary, reorder the rules so that no other rules matching the same event criteria comes before the current rule.

12. Click **Next**.

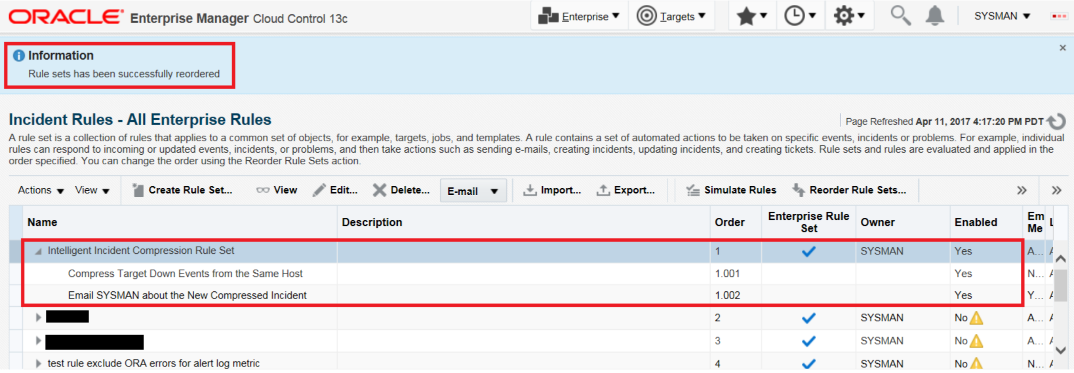

Information
Rule sets has been successfully reordered

Incident Rules - All Enterprise Rules

Page Refreshed **Apr 11, 2017 4:17:20 PM PDT**

A rule set is a collection of rules that applies to a common set of objects, for example, targets, jobs, and templates. A rule contains a set of automated actions to be taken on specific events, incidents or problems. For example, individual rules can respond to incoming or updated events, incidents, or problems, and then take actions such as sending e-mails, creating incidents, updating incidents, and creating tickets. Rule sets and rules are evaluated and applied in the order specified. You can change the order using the Reorder Rule Sets action.

Name	Description	Order	Enterprise Rule Set	Owner	Enabled	Em Me
Intelligent Incident Compression Rule Set		1	✓	SYSMAN	Yes	A...
Compress Target Down Events from the Same Host		1.001			Yes	N...
Email SYSMAN about the New Compressed Incident		1.002			Yes	Y...
[redacted]		2	✓	SYSMAN	No ⚠	A...
[redacted]		3	✓	SYSMAN	No ⚠	A...
test rule exclude ORA errors for alert log metric		4	✓	SYSMAN	No ⚠	N...

13. Provide a name for the rule and click **Next**.

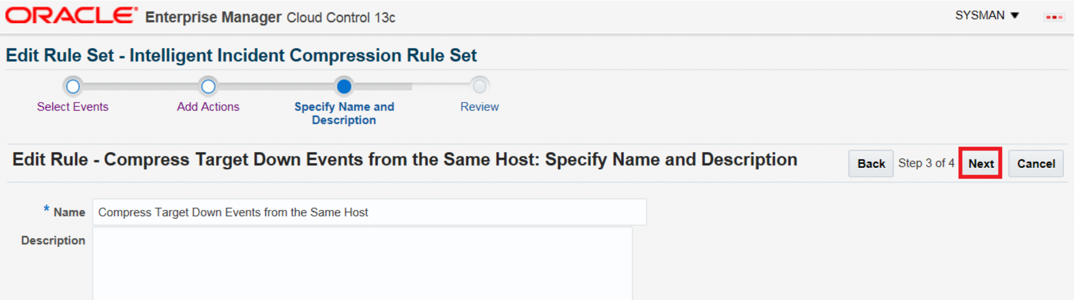
ORACLE Enterprise Manager Cloud Control 13c
SYSMAN

Edit Rule Set - Intelligent Incident Compression Rule Set

Select Events
Add Actions
Specify Name and Description
Review

Edit Rule - Compress Target Down Events from the Same Host: Specify Name and Description

Back
Step 3 of 4
Next
Cancel

* Name: Compress Target Down Events from the Same Host

Description:

14. Click **Continue** to add the rule.

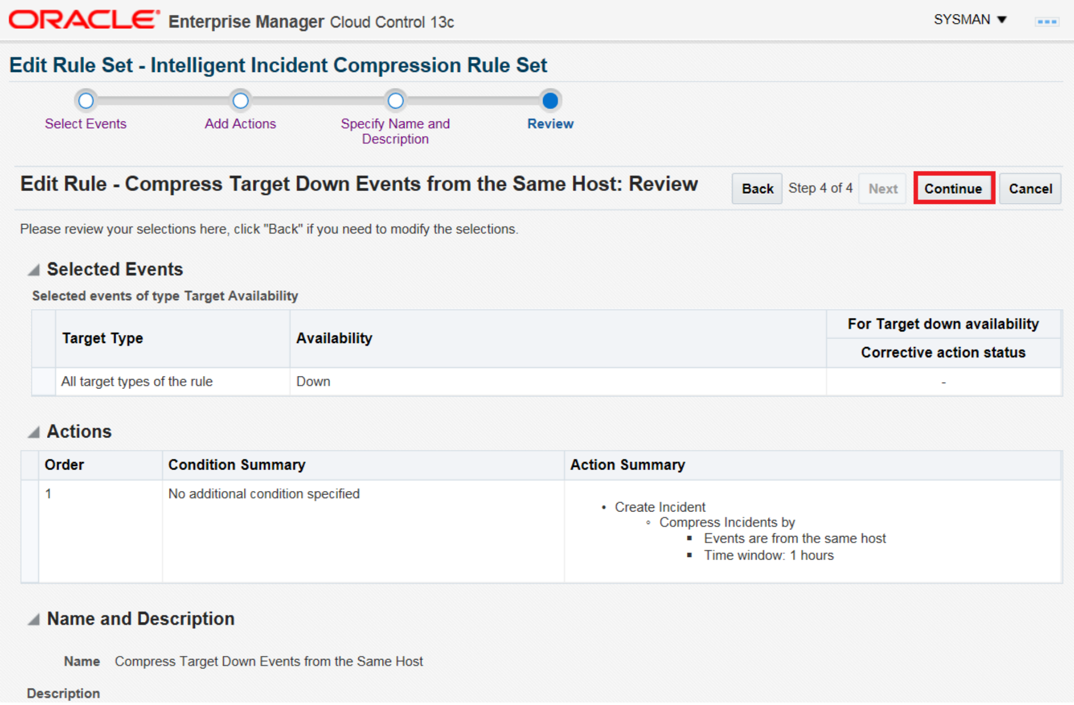
ORACLE Enterprise Manager Cloud Control 13c
SYSMAN

Edit Rule Set - Intelligent Incident Compression Rule Set

Select Events
Add Actions
Specify Name and Description
Review

Edit Rule - Compress Target Down Events from the Same Host: Review

Back
Step 4 of 4
Next
Continue
Cancel

Please review your selections here, click "Back" if you need to modify the selections.

Selected Events

Selected events of type Target Availability

Target Type	Availability	For Target down availability
		Corrective action status
All target types of the rule	Down	-

Actions

Order	Condition Summary	Action Summary
1	No additional condition specified	• Create Incident ◦ Compress Incidents by ▪ Events are from the same host ▪ Time window: 1 hours

Name and Description

Name: Compress Target Down Events from the Same Host

Description:

Next, you will create a new rule to notify an administrator for the compressed incident.

15. Click **Create** to create a new rule to configure email notification.

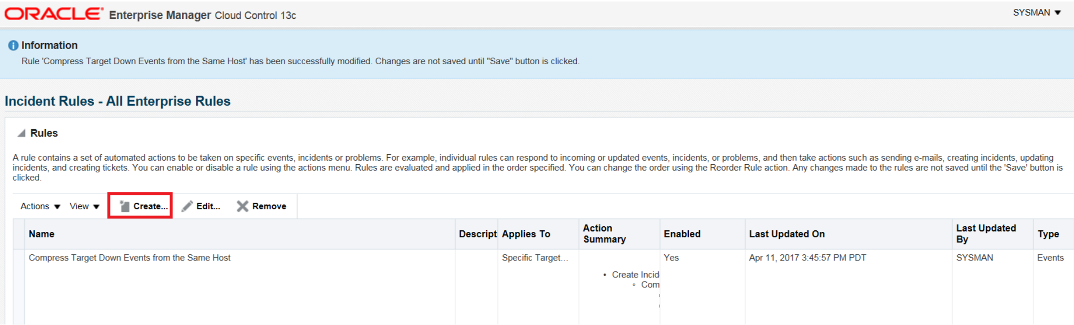

ORACLE Enterprise Manager Cloud Control 13c

Information

Rule 'Compress Target Down Events from the Same Host' has been successfully modified. Changes are not saved until "Save" button is clicked.

Incident Rules - All Enterprise Rules

Rules

A rule contains a set of automated actions to be taken on specific events, incidents or problems. For example, individual rules can respond to incoming or updated events, incidents, or problems, and then take actions such as sending e-mails, creating incidents, updating incidents, and creating tickets. You can enable or disable a rule using the actions menu. Rules are evaluated and applied in the order specified. You can change the order using the Reorder Rule action. Any changes made to the rules are not saved until the 'Save' button is clicked.

Name	Descript	Applies To	Action Summary	Enabled	Last Updated On	Last Updated By	Type
Compress Target Down Events from the Same Host		Specific Target...	• Create Incid ◦ Com	Yes	Apr 11, 2017 3:45:57 PM PDT	SYSMAN	Events

16. Select **Newly created incidents or updates to incidents** and click **Continue**.

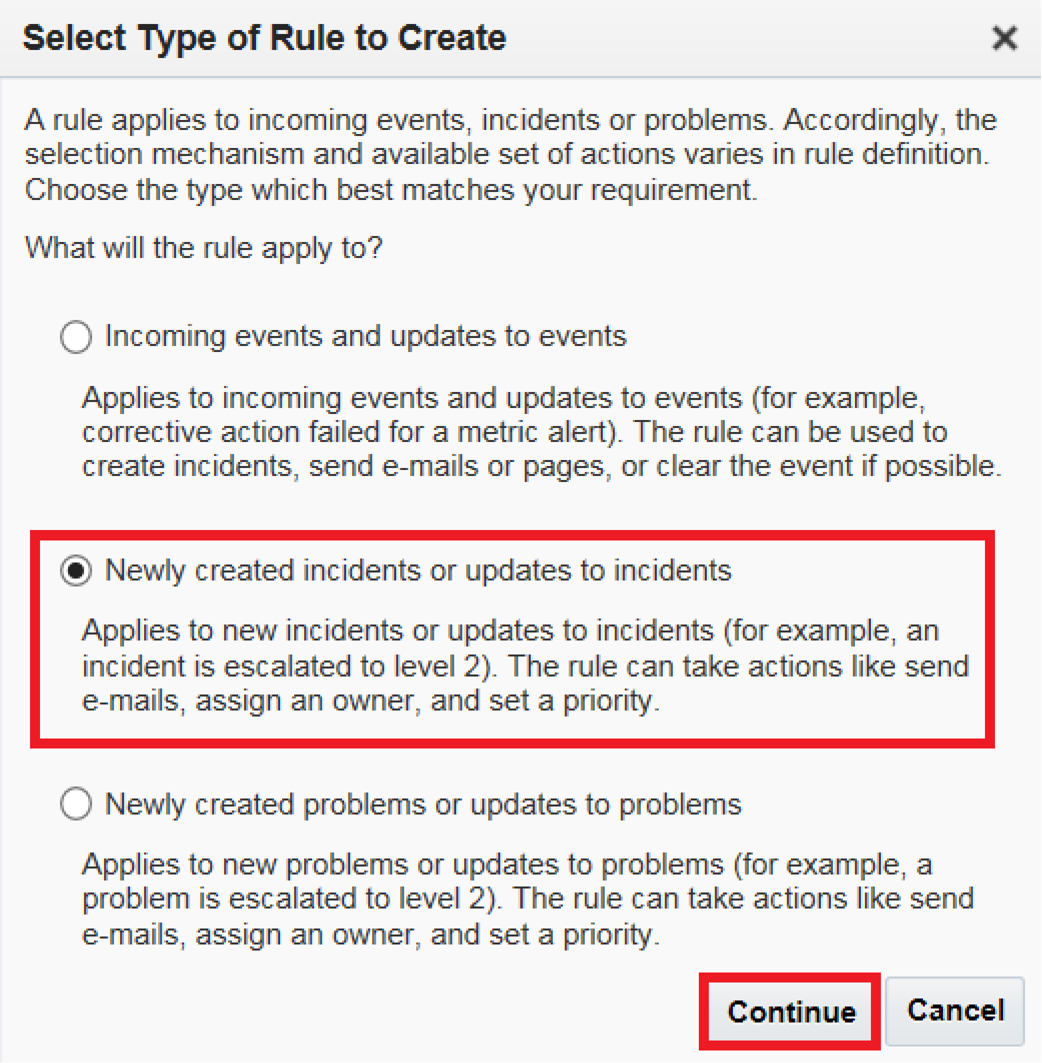

Select Type of Rule to Create

A rule applies to incoming events, incidents or problems. Accordingly, the selection mechanism and available set of actions varies in rule definition. Choose the type which best matches your requirement.

What will the rule apply to?

○ Incoming events and updates to events

Applies to incoming events and updates to events (for example, corrective action failed for a metric alert). The rule can be used to create incidents, send e-mails or pages, or clear the event if possible.

● Newly created incidents or updates to incidents

Applies to new incidents or updates to incidents (for example, an incident is escalated to level 2). The rule can take actions like send e-mails, assign an owner, and set a priority.

○ Newly created problems or updates to problems

Applies to new problems or updates to problems (for example, a problem is escalated to level 2). The rule can take actions like send e-mails, assign an owner, and set a priority.

17. Select **Specific incidents**.
    a. Select the checkbox for **Rules that created the incidents** and select the rule that you just created.
    b. Select the checkbox for **Status** and select **New**.
    c. Click **Next**.

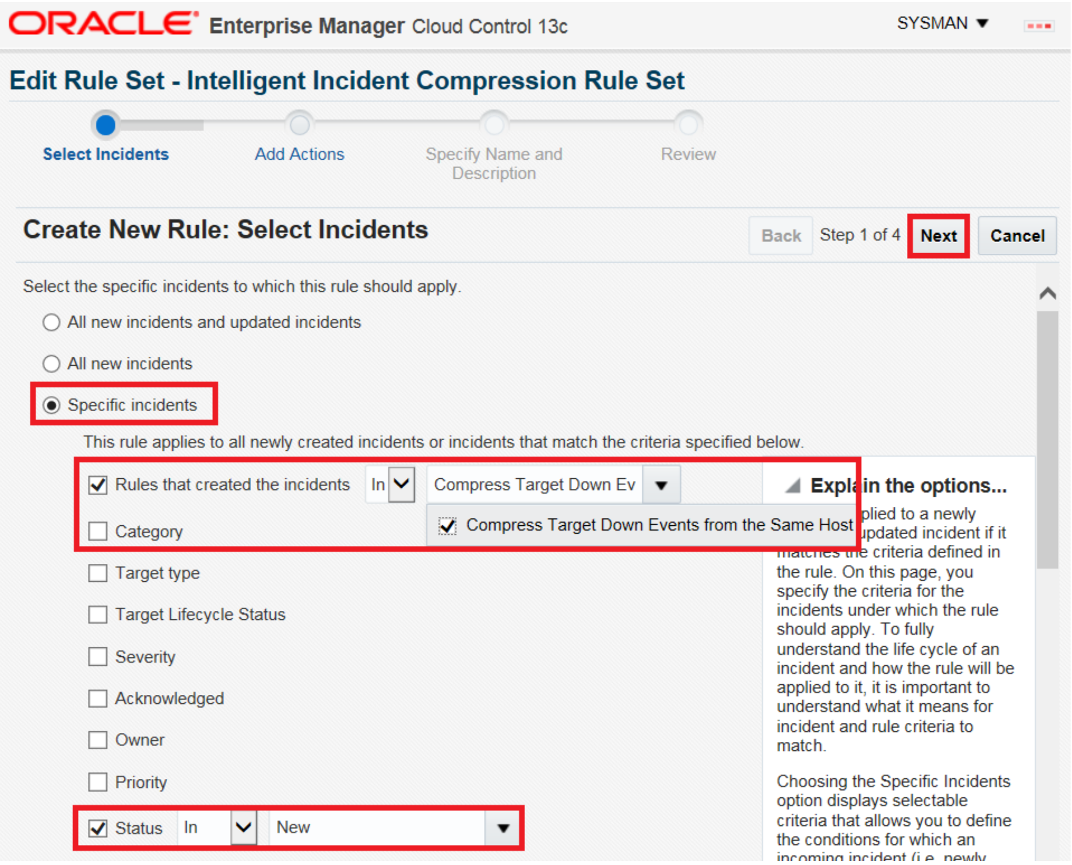

ORACLE Enterprise Manager Cloud Control 13c
SYSMAN

Edit Rule Set - Intelligent Incident Compression Rule Set

Select Incidents
Add Actions
Specify Name and Description
Review

Create New Rule: Select Incidents

Back
Step 1 of 4
Next
Cancel

Select the specific incidents to which this rule should apply.

○ All new incidents and updated incidents

○ All new incidents

● Specific incidents

This rule applies to all newly created incidents or incidents that match the criteria specified below.

☑ Rules that created the incidents In | Compress Target Down Ev
☑ Compress Target Down Events from the Same Host

□ Category

□ Target type

□ Target Lifecycle Status

□ Severity

□ Acknowledged

□ Owner

□ Priority

☑ Status In | New

Explain the options...

...plied to a newly ...updated incident if it matches the criteria defined in the rule. On this page, you specify the criteria for the incidents under which the rule should apply. To fully understand the life cycle of an incident and how the rule will be applied to it, it is important to understand what it means for incident and rule criteria to match.

Choosing the Specific Incidents option displays selectable criteria that allows you to define the conditions for which an incoming incident (i.e. newly

18. Click **Add** to add a Conditional Action.

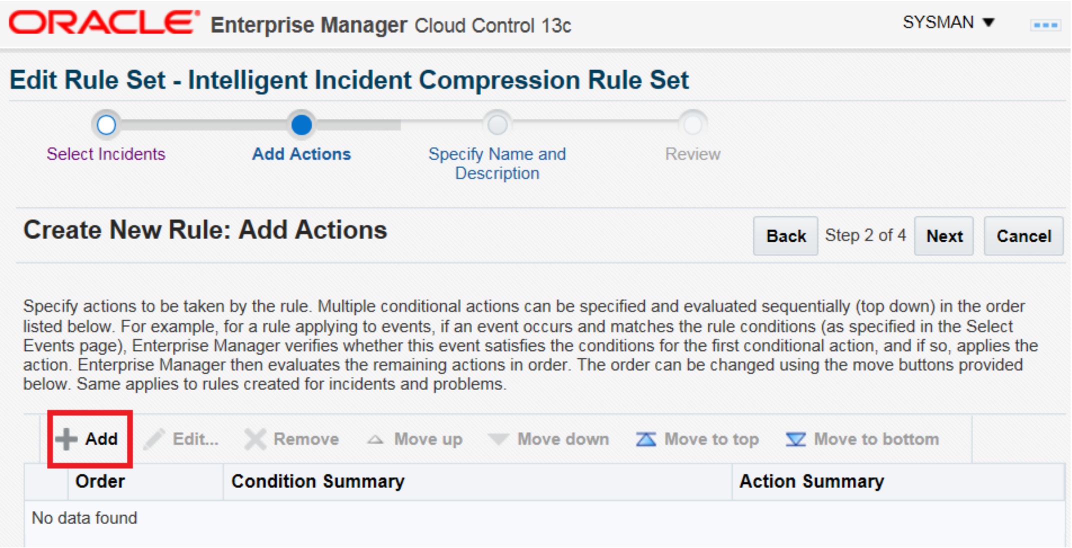

ORACLE Enterprise Manager Cloud Control 13c
SYSMAN

Edit Rule Set - Intelligent Incident Compression Rule Set

Select Incidents
Add Actions
Specify Name and Description
Review

Create New Rule: Add Actions

Back
Step 2 of 4
Next
Cancel

Specify actions to be taken by the rule. Multiple conditional actions can be specified and evaluated sequentially (top down) in the order listed below. For example, for a rule applying to events, if an event occurs and matches the rule conditions (as specified in the Select Events page), Enterprise Manager verifies whether this event satisfies the conditions for the first conditional action, and if so, applies the action. Enterprise Manager then evaluates the remaining actions in order. The order can be changed using the move buttons provided below. Same applies to rules created for incidents and problems.

Order	Condition Summary	Action Summary
No data found		

19. Populate Conditional Actions as shown in the screen shot below. Change the notified user according to your requirement.

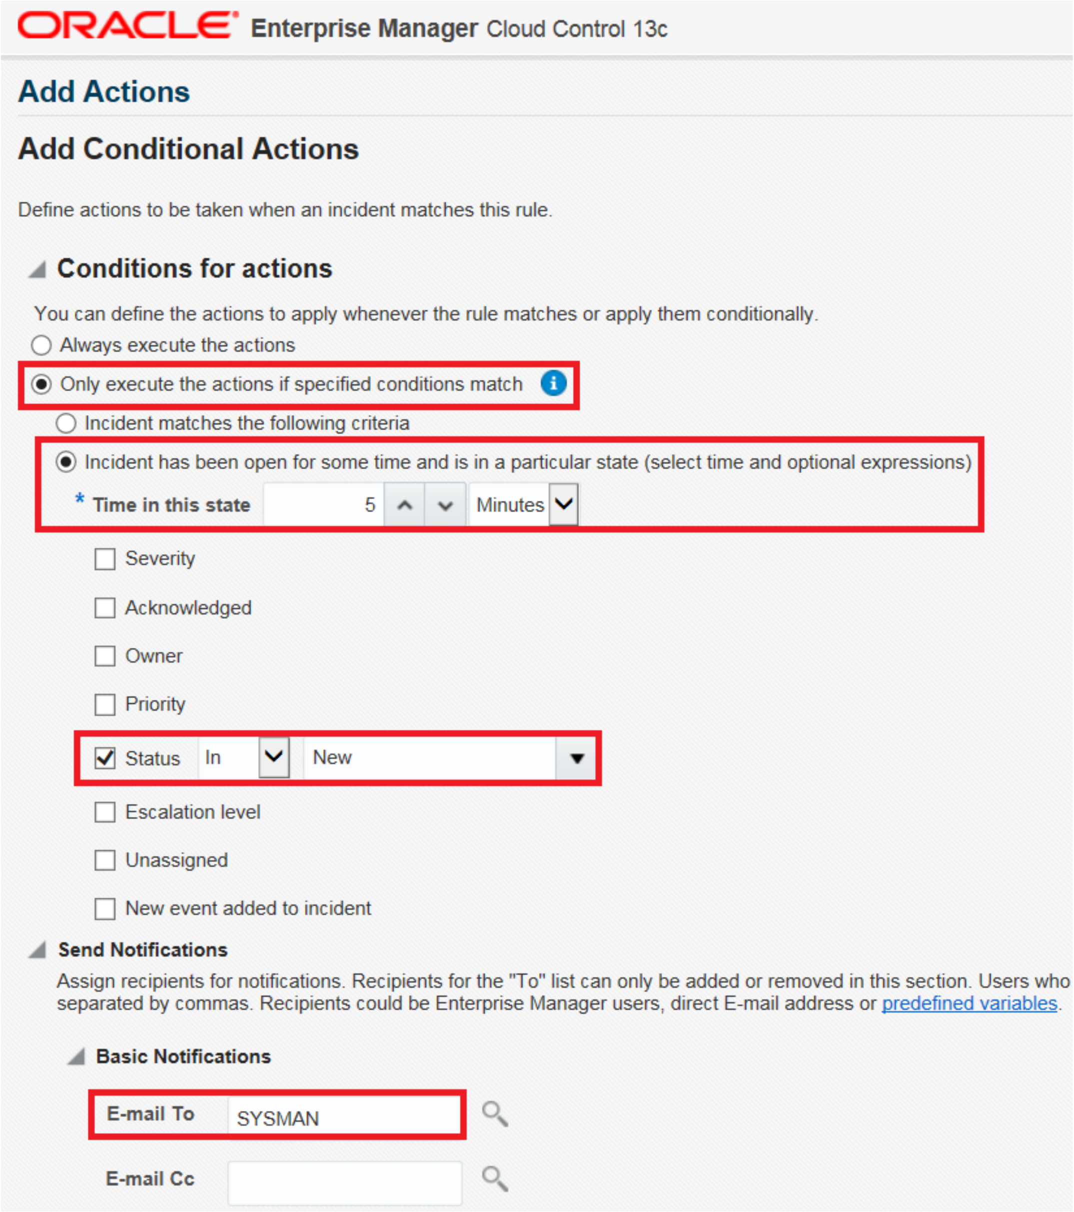

Add Actions

Add Conditional Actions

Define actions to be taken when an incident matches this rule.

Conditions for actions

You can define the actions to apply whenever the rule matches or apply them conditionally.

○ Always execute the actions

● Only execute the actions if specified conditions match

○ Incident matches the following criteria

● Incident has been open for some time and is in a particular state (select time and optional expressions)

* Time in this state: 5 Minutes

□ Severity

□ Acknowledged

□ Owner

□ Priority

☑ Status In New

□ Escalation level

□ Unassigned

□ New event added to incident

Send Notifications

Assign recipients for notifications. Recipients for the "To" list can only be added or removed in this section. Users who
separated by commas. Recipients could be Enterprise Manager users, direct E-mail address or predefined variables.

Basic Notifications

E-mail To: SYSMAN

E-mail Cc:

20. Click **Next**.

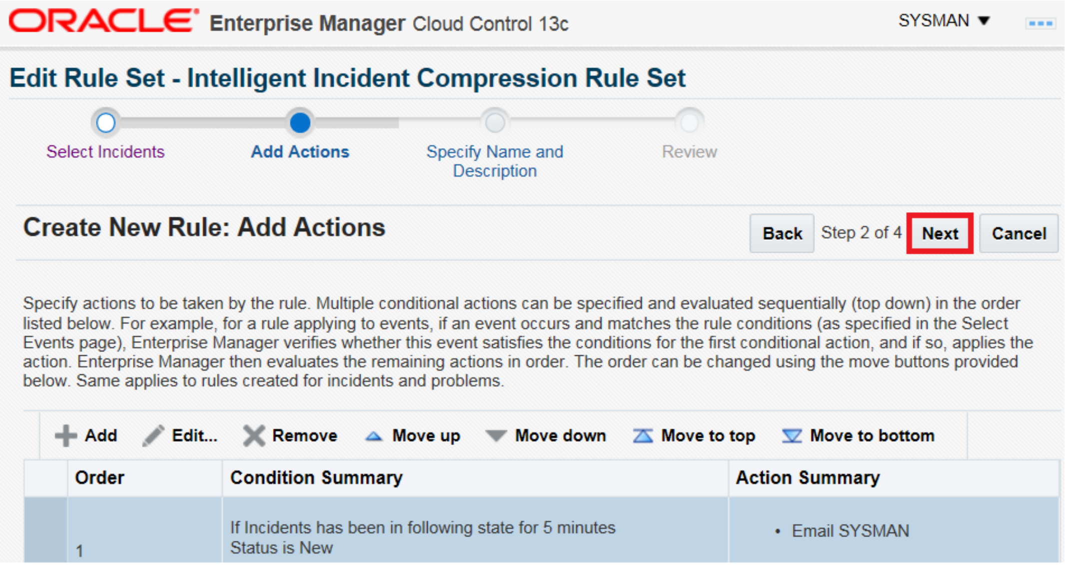

ORACLE Enterprise Manager Cloud Control 13c
SYSMAN ▼

Edit Rule Set - Intelligent Incident Compression Rule Set

Select Incidents
Add Actions
Specify Name and Description
Review

Create New Rule: Add Actions

Back
Step 2 of 4
Next
Cancel

Specify actions to be taken by the rule. Multiple conditional actions can be specified and evaluated sequentially (top down) in the order listed below. For example, for a rule applying to events, if an event occurs and matches the rule conditions (as specified in the Select Events page), Enterprise Manager verifies whether this event satisfies the conditions for the first conditional action, and if so, applies the action. Enterprise Manager then evaluates the remaining actions in order. The order can be changed using the move buttons provided below. Same applies to rules created for incidents and problems.

Order	Condition Summary	Action Summary
1	If Incidents has been in following state for 5 minutes Status is New	• Email SYSMAN

21. Enter a name for the rule and click **Next**.

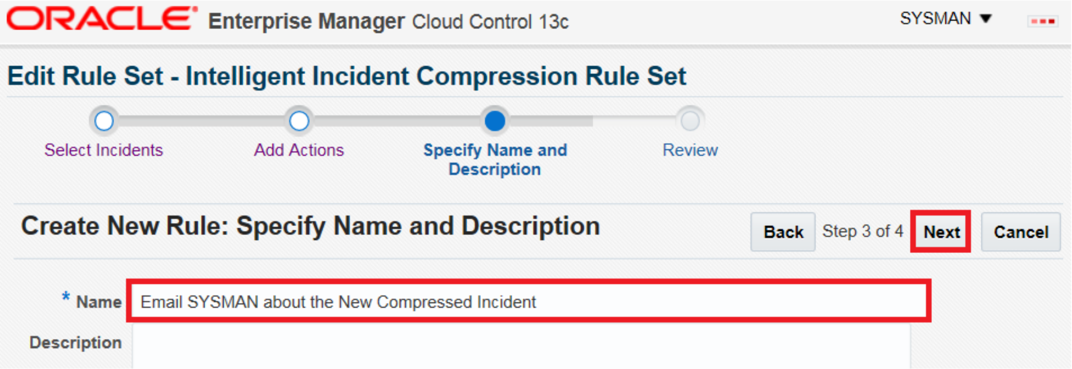

ORACLE Enterprise Manager Cloud Control 13c
SYSMAN ▼

Edit Rule Set - Intelligent Incident Compression Rule Set

Select Incidents
Add Actions
Specify Name and Description
Review

Create New Rule: Specify Name and Description

Back
Step 3 of 4
Next
Cancel

* Name: Email SYSMAN about the New Compressed Incident

Description:

22. Click **Continue**.

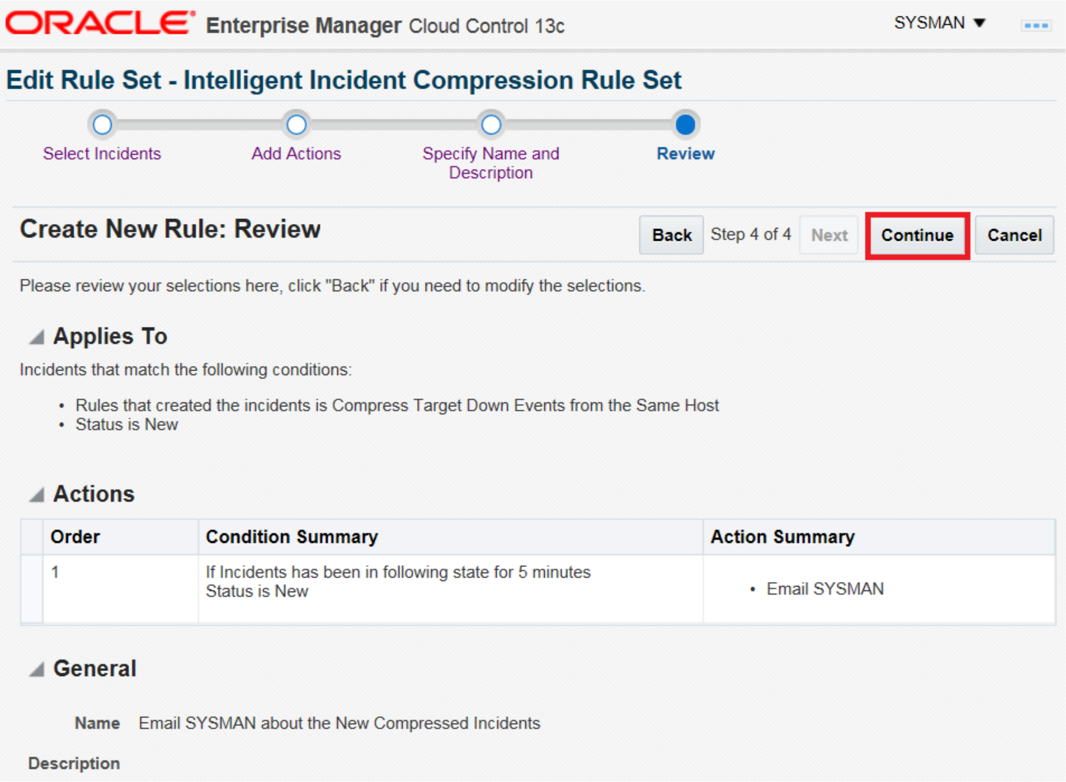

ORACLE Enterprise Manager Cloud Control 13c SYSMAN ▼

Edit Rule Set - Intelligent Incident Compression Rule Set

Create New Rule: Review

Please review your selections here, click "Back" if you need to modify the selections.

◢ Applies To

Incidents that match the following conditions:

- Rules that created the incidents is Compress Target Down Events from the Same Host
- Status is New

◢ Actions

Order	Condition Summary	Action Summary
1	If Incidents has been in following state for 5 minutes Status is New	• Email SYSMAN

◢ General

Name Email SYSMAN about the New Compressed Incidents

Description

23. Click **Save** to save the rule set.

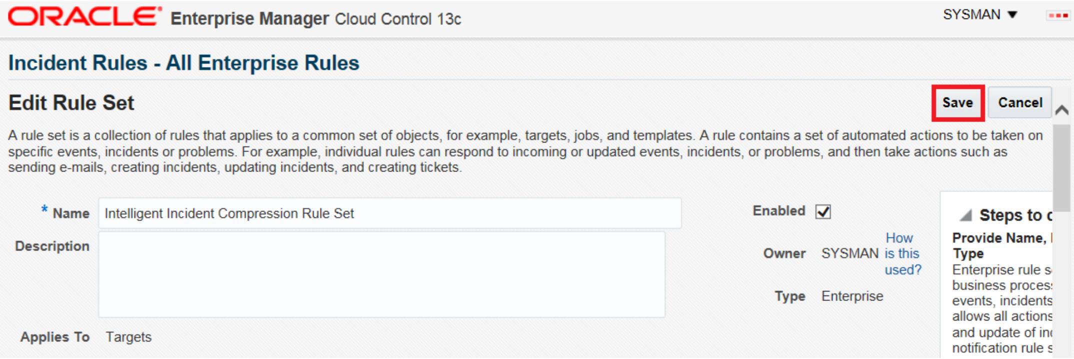

ORACLE Enterprise Manager Cloud Control 13c SYSMAN ▼

Incident Rules - All Enterprise Rules

Edit Rule Set

A rule set is a collection of rules that applies to a common set of objects, for example, targets, jobs, and templates. A rule contains a set of automated actions to be taken on specific events, incidents or problems. For example, individual rules can respond to incoming or updated events, incidents, or problems, and then take actions such as sending e-mails, creating incidents, updating incidents, and creating tickets.

* Name	Intelligent Incident Compression Rule Set
Description	
Applies To	Targets
Enabled	☑
Owner	SYSMAN How is this used?
Type	Enterprise

◢ Steps to c
Provide Name, I
Type
Enterprise rule s
business proces
events, incidents
allows all actions
and update of in
notification rule s

24. Highlight the rule set and click **Reorder Rule Sets**.

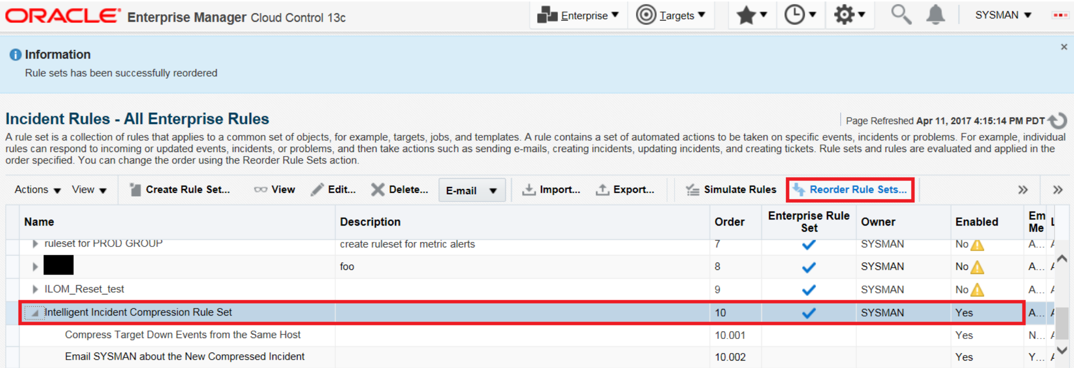

ORACLE Enterprise Manager Cloud Control 13c

Information
Rule sets has been successfully reordered

Incident Rules - All Enterprise Rules

Page Refreshed **Apr 11, 2017 4:15:14 PM PDT**

A rule set is a collection of rules that applies to a common set of objects, for example, targets, jobs, and templates. A rule contains a set of automated actions to be taken on specific events, incidents or problems. For example, individual rules can respond to incoming or updated events, incidents, or problems, and then take actions such as sending e-mails, creating incidents, updating incidents, and creating tickets. Rule sets and rules are evaluated and applied in the order specified. You can change the order using the Reorder Rule Sets action.

Name	Description	Order	Enterprise Rule Set	Owner	Enabled	Em Me
▸ ruleset for PROD GROUP	create ruleset for metric alerts	7	✓	SYSMAN	No ⚠	A...
▸ [redacted]	foo	8	✓	SYSMAN	No ⚠	A...
▸ ILOM_Reset_test		9	✓	SYSMAN	No ⚠	A...
◢ Intelligent Incident Compression Rule Set		10	✓	SYSMAN	Yes	A...
Compress Target Down Events from the Same Host		10.001			Yes	N...
Email SYSMAN about the New Compressed Incident		10.002			Yes	Y...

25. Move the rule set to the top and click **OK**.

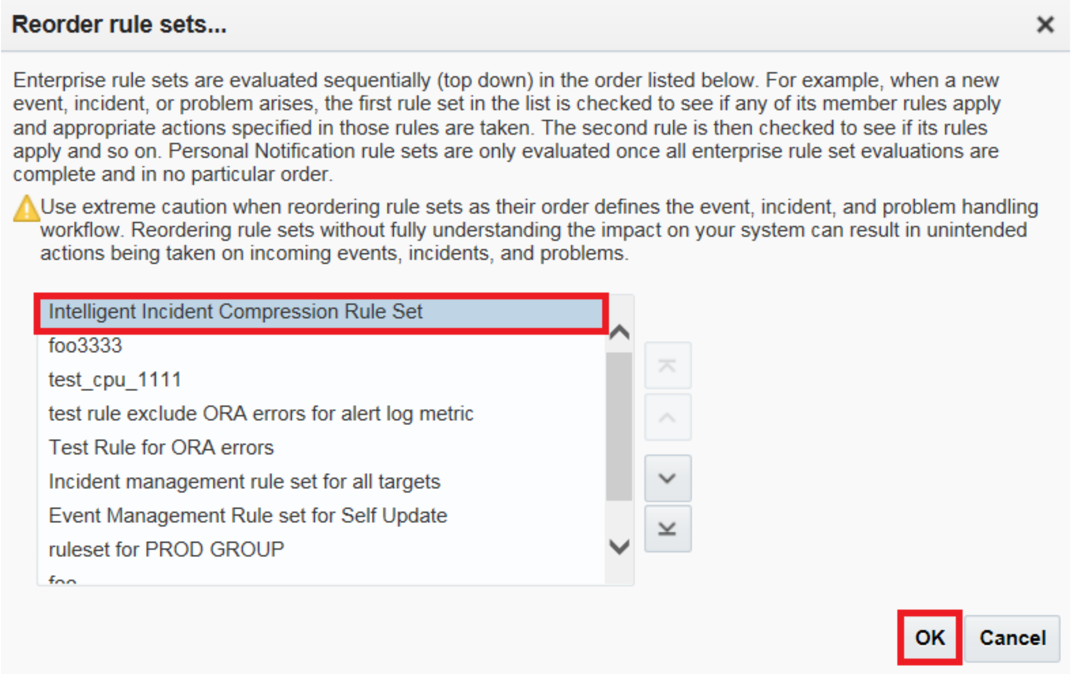

Reorder rule sets...

Enterprise rule sets are evaluated sequentially (top down) in the order listed below. For example, when a new event, incident, or problem arises, the first rule set in the list is checked to see if any of its member rules apply and appropriate actions specified in those rules are taken. The second rule is then checked to see if its rules apply and so on. Personal Notification rule sets are only evaluated once all enterprise rule set evaluations are complete and in no particular order.

⚠ Use extreme caution when reordering rule sets as their order defines the event, incident, and problem handling workflow. Reordering rule sets without fully understanding the impact on your system can result in unintended actions being taken on incoming events, incidents, and problems.

Intelligent Incident Compression Rule Set
foo3333
test_cpu_1111
test rule exclude ORA errors for alert log metric
Test Rule for ORA errors
Incident management rule set for all targets
Event Management Rule set for Self Update
ruleset for PROD GROUP
foo

Once the new rule set is moved to the top of the list, it will be evaluated first by the rules engine.

Event Prioritization

When working in a large enterprise, it is conceivable that when systems are under heavy load, a large number of incidents and events may be generated. All of these need to be processed in a timely and efficient manner in accordance with your business priorities. An effective prioritization scheme is needed to determine which events/incidents should be resolved first.

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. Lifecycle Status is a target property that specifies a target's operational status. You can set/view a target's Lifecycle Status from the UI (from a target's **Target Setup** menu, select **Properties**). You can set target Lifecycle Status properties across multiple targets simultaneously by using the Enterprise Manager Command Line Interface (EM CLI) set_target_property_value verb.

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
- Stage
- Test
- Development (lowest priority)

Incident/Event Type

- Availability events (highest priority)
- Non-informational events.
- Informational events

Root Cause Analysis (RCA) and Target Down Events

Root Cause Analysis (RCA) tries to identify the root causes of issues that cause operational events. Beginning with Enterprise Manager Could Control 12.1.0.3, Incident Manager automatically performs RCA over *target down events*, thus actively identifying whether the target down event is the cause or symptom of other target down events.The term target down event specifically pertains to Target Availability events that are raised when the targets are detected to be down.

How RCA Works

RCA is an ongoing process that identifies whether a target down event is root cause or symptom. It uses the Causal Analysis Update attribute of the event to store the results of its analysis, i.e. identifying whether or not the target down event is root cause or symptom. Whenever a new target availability event comes in, RCA is automatically performed on the incoming event and existing target down events that are related to it. Based on the analysis, it updates the Causal Analysis Update attribute value if the incoming event is a target down event. It also updates the Causal Analysis Update attribute for the related target down events if there is a change.

Two types of target relationships are used for identifying the related targets: *dependency* and *containment*.

When one target depends on another target for its availability, dependency relationship exists between them. For example, J2EE application target depends on the WebLogic Server target over which it is deployed.

The causal analysis update attribute is used only for target down events (such as a Target Availability event for target down) and can have be assigned any one of the following values by the RCA process:

- *Symptom* -- The target down event has been caused by another target down event.
- *Cause* - The target down event has caused another target down event and it is not the symptom of any other target down event.
- *Root Cause* - The target down event has caused another target down event and it is not the symptom of any other target down event.
- *N/A* - Root cause analysis is not applicable to this event. Root cause analysis applies to target down events only.
- *Not a cause and not a symptom* - The target down event is not a root cause and not a symptom of other target down events. This is shown in Incident Manager as a dash (-).

The following rules describe the RCA process:

- **Rule 1**: Down event on a non-container target (a target that does not have members) is marked as the cause if a dependent target is down and it is not symptom of other target down events.

  Examples:

  - You have J2EE applications deployed on a standalone WebLogic Server. If both J2EE application and WebLogic Server targets are down, the WebLogic Server down event is the cause for the J2EE applications deployed on it.
  - You have a J2EE application deployed on couple of WebLogic Servers, which are part of a WebLogic Cluster. If one WebLogic Server is down along with its J2EE application, then the WebLogic Server down event is the cause of the J2EE application target down. This assumes the WebLogic Cluster is not down.
- **Rule 2**: Down event on a non-container target (a target that does not have members) is marked as a symptom if a target it depends on is down or if the target containing it is down.

  Examples:

  - You have a J2EE application deployed on a standalone WebLogic Server. If both J2EE application and WebLogic Server targets are down, J2EE application down event is the symptom of WebLogic Server being down.
  - You have a couple of WebLogic Servers which are part of a WebLogic Cluster. Each WebLogic Server has a J2EE application deployed on it. If the WebLogic Cluster is down, this means both WebLogic Servers are down. Consequently, the J2EE applications that are deployed on these servers are also down. The WebLogic Server down events would be marked as the causes of the WebLogic Cluster being down. See Rule 3 for details.
  - You have a couple of RAC database instance targets that are part of a cluster database target. If the cluster database is down, then all RAC instances are also down. The RAC instance down events would be marked as the causes of cluster database being down. See Rule 3 for details..
- **Rule 3**: Down event on a container target is marked as symptom down if all member targets are down and any target containing it is not down.

  Examples:

  - You have a couple of WebLogic Servers, which are part of a standalone WebLogic Cluster. A WebLogic Cluster down event would be marked as symptom, if both the WebLogic Servers are down.
  - You have a couple of RAC database instance targets that are part of a cluster database target. The cluster database target down event would be marked as a symptom, if both database instances are down.
- **Rule 4**: Down event on a container target is marked as symptom if the target containing it is down.

  Example:

  You have a couple of WebLogic Clusters that are part of a WebLogic Domain target. If the WebLogic domain is down, this means the WebLogic Clusters are also down. The WebLogic Cluster target down events would be the cause of WebLogic Domain being down. The WebLogic Domain down event would be marked as symptom.

Leveraging RCA Results in Incident Rule Sets

As described above, RCA is an ongoing process which results in marking target down events as *cause*, *symptom* or *neither* as new target down events come in and are processed. So a target down event may be marked as a cause or symptom as it comes in or after some time when RCA has analyzed additional event information.

Most datacenters automatically create incidents for target down events since these are important events that need to be resolved right away. This is recommended best practice and also implemented by the out-of-the-box rule sets. However, in terms of notifying response teams or creating trouble tickets, it is not desirable to do so for symptom incidents. Some datacenters may also choose to not create incidents for symptom events.

So the RCA results can be leveraged to do the following:

1. Notify or create tickets only for non-symptom events:

   This can be achieved in 2 ways:

   - Create two separate event rules , one event rule to create incidents for all relevant events, but take no further action (no notification or ticket creation) and another one to create incidents for non-symptom events only and also send notifications and create tickets. See "Creating Incidents On Non-symptom Events" for instructions.
   - Create an event rule that creates incidents for all target down events. Create another rule to update the incident priority, send notifications and create tickets only for incidents stemming from non-symptom events. Once the incident priority is set to say "Urgent", customer can also create additional incident rules to take additional actions on the Urgent priority incidents. See "Creating a Rule to Update Incident Priority for Non-symptom Events".

2. Only create incidents after a suitable wait for events that are not initially marked as neither a cause nor a symptom:

   As mentioned previously, RCA is an iterative process whereby incoming target down events are continually being evaluated, resulting in updates to causal analysis state of existing events. Over a period of time (minutes), a target down event that was initially marked as a root cause may or may not remain a root cause depending on other incoming target down events. The original target down event may later be classified as a symptom.

   To avoid prematurely creating an incident and opening a ticket for an event which may later turn out to be a symptom event, you can set up your rules as follows:

   - In addition to the rules already defined in the previous step, create an additional event rule to act upon RCA updates to events and when the RCA update indicates that the event is marked as a symptom, lower the priority of the incident to "Low". This will also send an update to the ticket automatically. This is recommended. See "Introducing a Time Delay" for instructions.

     OR

   - To allow time for target down events to be reported, analyzed, and then acted upon (such as creating an incident or updating an incident), you can add a delay in the rule actions. This is useful when customer have some tolerance to take action after some minimum delay (typically 5 minutes).

3. Only create incidents for non-symptom events.

   Some datacenters may choose not to create any incidents for symptom events. This can be achieved by changing the rules to only create incidents for events marked as cause or neither a cause nor symptom. See "Creating Incidents On Non-symptom Events" for instructions.

   Please note that, even in this approach, it is possible that an event that was originally marked as cause or neither a cause nor symptom, may be marked as a symptom when more information is received. Customers can use an approach similar to that of the second option in step 2 to build some delay in creating the incidents. Even with this, it is still feasible but a bit unlikely, that newer information shows up after the pre-set delay and ends up marking the event as symptom. So it is recommended to use the approach of setting incident priority and using that as a way to manage workflow.

Leveraging RCA Results in Incident Manager

You can use the RCA results to focus on the non-symptom incidents in Incident Manager. This involves using the Causal Analysis Update incident attribute when creating custom views.

1. From the **Enterprise** menu, select **Monitoring**, and then **Incidents**. The Indicent Manager page displays.
2. From the Views region, click **Create**. The Search page displays.
3. Click **Add Fields...** and then choose **Causal analysis update**. The Causal analysis update displays as additional search criteria.

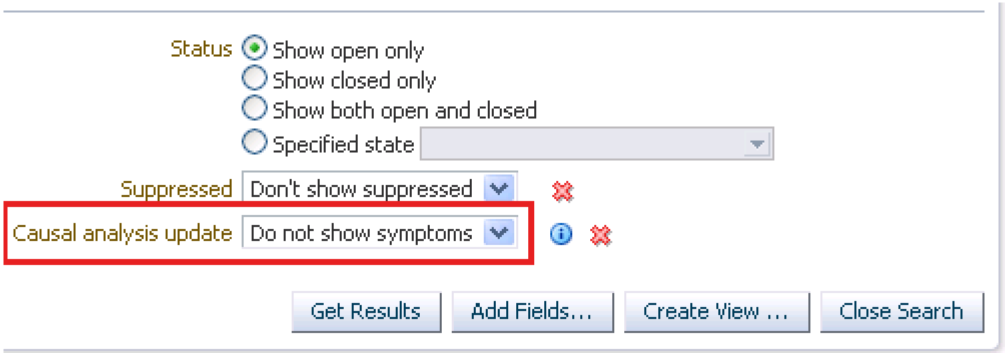

Status ● Show open only
○ Show closed only
○ Show both open and closed
○ Specified state

Suppressed: Don't show suppressed

Causal analysis update: Do not show symptoms

4. Choose 'Do Not Show Symptoms' from the list of available criteria. This will automatically exclude incidents that have been marked as 'symptom'. Incidents that are not marked as symptom or root cause will be included as long as it matches any other criteria you may have specified.

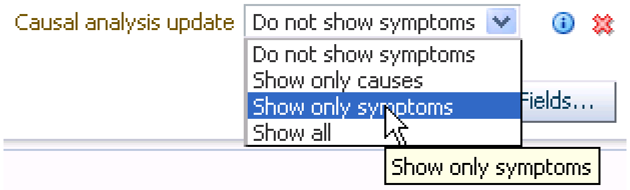

Causal analysis update: Do not show symptoms

5. Click **Create View**, enter a **View Name** when prompted, and then click **OK**.

Showing RCA Results in an Incident Detail

An incident that is a root cause or symptom will be identified prominently as part of the details of the incident in Incident Manager. In addition, in case the incident is a symptom, a **Causes** section will be added to identify the root cause(s) of the incident. In case the incident has, in turn, caused other target down incidents, an **Impacted Targets** section will also be added to show the targets that have been affected, that is. other targets that are down as a result of the original target down.

Leveraging RCA Results in the System Dashboard

In the System Dashboard, you can use the RCA results to exclude symptom incidents from the Incidents table so administrators can focus their attention on incidents that are root cause or have not been caused by other target down events.

To exclude Symptom Incidents:

1. In the System Dashboard, click on the **View** option that is accessible from the upper left hand corner of the **Incidents and Problems** table.
2. Choose the option to 'Exclude symptoms'. Alternatively, you can also choose the option 'Cause only' show only shows target down incidents that have been identified as cause of other target down incidents. Regardless of the option chosen, incidents that have not been marked as symptom or root cause will continue to be displayed.

Creating a Rule to Update Incident Priority for Non-symptom Events

1. Create an event rule to select only non-symptom events.

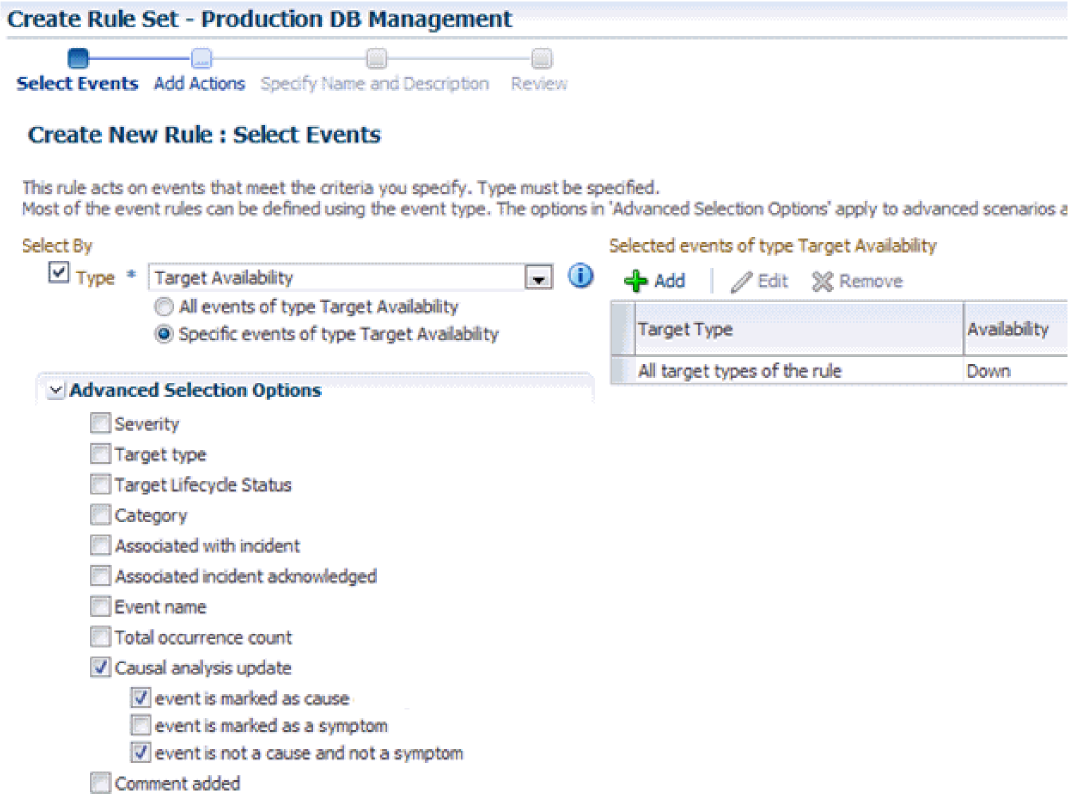

Create Rule Set - Production DB Management

Create New Rule : Select Events

This rule acts on events that meet the criteria you specify. Type must be specified.
Most of the event rules can be defined using the event type. The options in 'Advanced Selection Options' apply to advanced scenarios a

Select By

☑ Type * Target Availability

○ All events of type Target Availability
● Specific events of type Target Availability

Selected events of type Target Availability

Target Type	Availability
All target types of the rule	Down

Advanced Selection Options

- □ Severity
- □ Target type
- □ Target Lifecycle Status
- □ Category
- □ Associated with incident
- □ Associated incident acknowledged
- □ Event name
- □ Total occurrence count
- ☑ Causal analysis update
  - ☑ event is marked as cause
  - □ event is marked as a symptom
  - ☑ event is not a cause and not a symptom
- □ Comment added

2. When adding an action, select the priority to be set for incidents associated with the non-symptom events selected above.

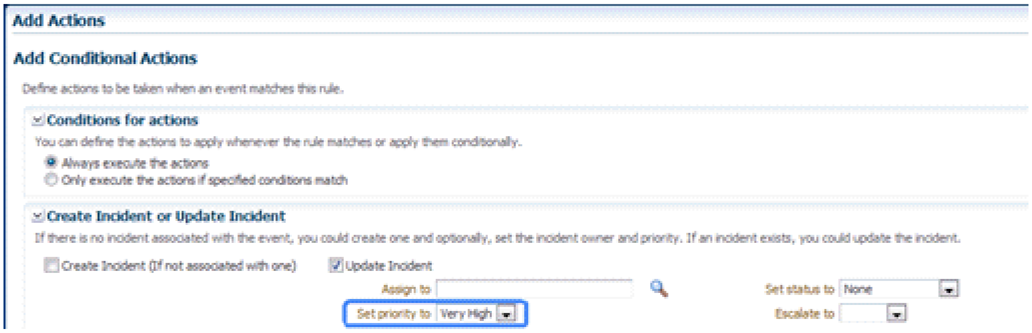

Add Actions

Add Conditional Actions

Define actions to be taken when an event matches this rule.

Conditions for actions

You can define the actions to apply whenever the rule matches or apply them conditionally.

● Always execute the actions
○ Only execute the actions if specified conditions match

Create Incident or Update Incident

If there is no incident associated with the event, you could create one and optionally, set the incident owner and priority. If an incident exists, you could update the incident.

□ Create Incident (If not associated with one) ☑ Update Incident

Assign to

Set priority to Very High

Set status to None

Escalate to

Creating Incidents On Non-symptom Events

You can leverage Incident Manager's Root Cause Analysis (RCA) capability by creating rule sets that generate incidents for non-symptom, target down events. For monitoring situations where a high number of symptom target down events are generated, but only a few non-symptom target down events, you can create/modify a rule set that generates incidents and send notifications only for non-symptom events.

To create a rule set that generates incidents for this monitoring condition, you need to create two event rules (one for each of the RCA filters):

- *Event Rule 1*: Generate incidents for **all relevant events**, but take no further action (no notification or ticket creation). The event is marked as a cause.
- *Event Rule 2*: Generate incidents for **non-symptom events only** and also send notifications and create tickets. The event is not a cause and not a symptom.

To create the event rules to handle non-symptom target down events, navigate to the Incident Rules - All Enterprise Rules page (Setup—>Incidents—>Incident Rules). From here, you can create a new rule set (click **Create Rule Set…**) or edit an existing rule set (click **Edit…**).

To create a rule that generates incidents for all relevant events:

1. From the Rules region of the Create Rule Set/Edit Rule Set page, click **Create ...** The Select Type of Rule to Create dialog appears.
2. Select **Incoming events and updates to events**.

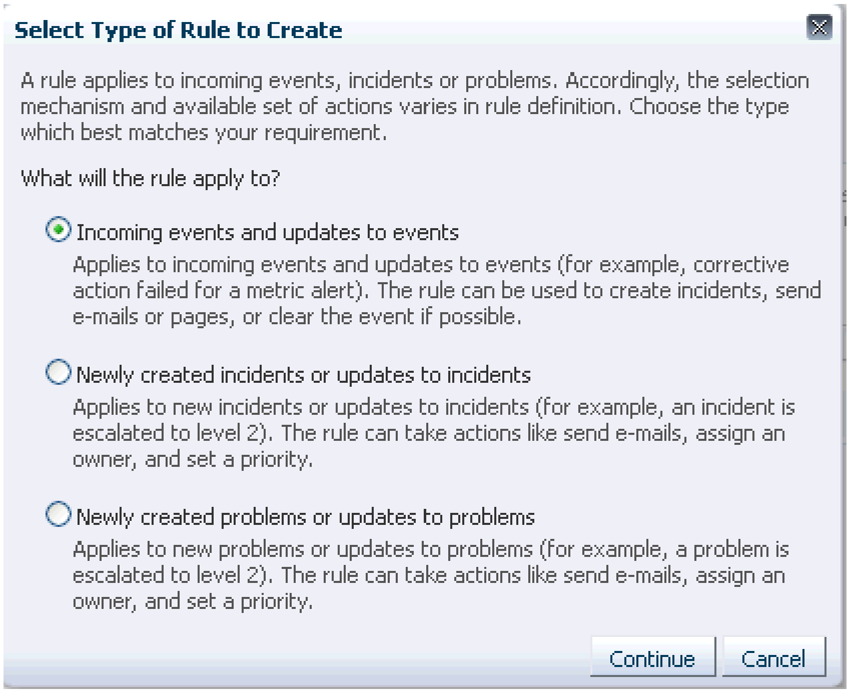

Select Type of Rule to Create

A rule applies to incoming events, incidents or problems. Accordingly, the selection mechanism and available set of actions varies in rule definition. Choose the type which best matches your requirement.

What will the rule apply to?

● Incoming events and updates to events
Applies to incoming events and updates to events (for example, corrective action failed for a metric alert). The rule can be used to create incidents, send e-mails or pages, or clear the event if possible.

○ Newly created incidents or updates to incidents
Applies to new incidents or updates to incidents (for example, an incident is escalated to level 2). The rule can take actions like send e-mails, assign an owner, and set a priority.

○ Newly created problems or updates to problems
Applies to new problems or updates to problems (for example, a problem is escalated to level 2). The rule can take actions like send e-mails, assign an owner, and set a priority.

3. Click **Continue**. The Create New Rule: Select Events dialog displays. Select **Target Availability**.

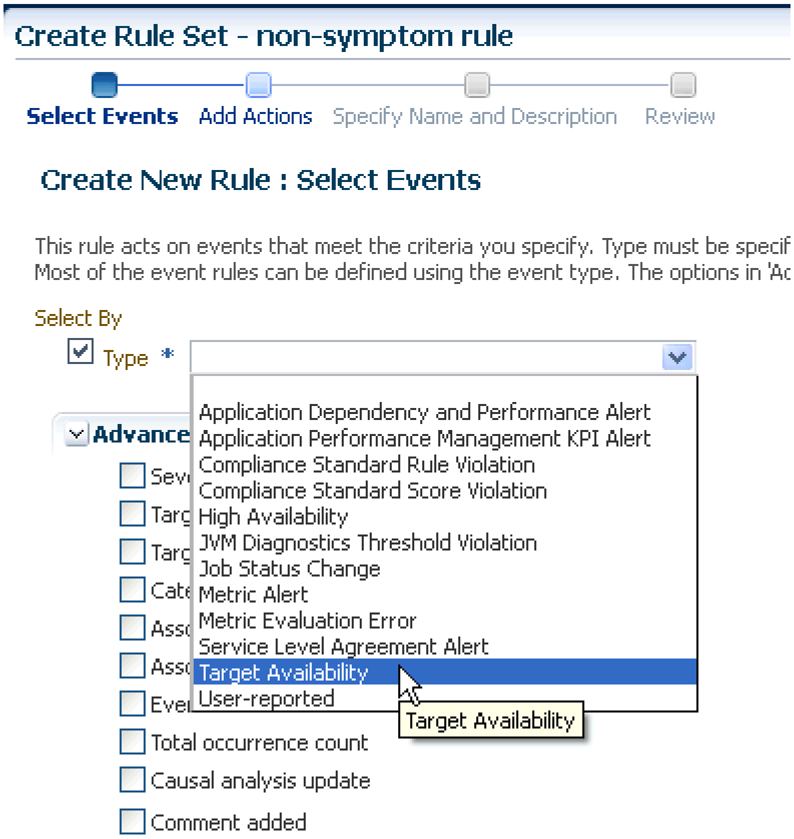

Create Rule Set - non-symptom rule

Create New Rule : Select Events

This rule acts on events that meet the criteria you specify. Type must be specif
Most of the event rules can be defined using the event type. The options in 'Ac

Select By

☑ Type *

- Application Dependency and Performance Alert
- Application Performance Management KPI Alert
- Compliance Standard Rule Violation
- Compliance Standard Score Violation
- High Availability
- JVM Diagnostics Threshold Violation
- Job Status Change
- Metric Alert
- Metric Evaluation Error
- Service Level Agreement Alert
- Target Availability
- User-reported

□ Total occurrence count

□ Causal analysis update

□ Comment added

4. In the Advanced Selection Options region, choose **Causal analysis update**. Three causal event options display:
   - *Event is marked as cause*: A target down is considered a cause if other targets depending on it are down.
   - *Event is marked as a symptom*: A target down is considered a symptom if a target it depends on is also down.
   - *Event is not a cause and not a symptom*: A target down is neither a cause or symptom.

> **Note:**
>
> Note: By selecting an option, you filter out extraneous target down events and focus on those target availability events that pertain to targets with interdependencies.

5. Select **event is marked as a cause** and click **Next**.
6. On the Create New Rule : Add Actions page, click **Add**. The Add Conditional Actions page displays.
7. In the Create Incident or Update Incident region, choose Create Incident (if not associated with one) and click **Continue**.
8. Complete the remaining Create Rule Set wizard pages to return to the Create Rule Set/ Edit Rule Set page.

   Next, you need to create a rule that generates incidents for **non-symptom events only** and also send notifications.

9. Repeat steps 1-4.
10. Select **event is not a cause and not a symptom** and click **Next**.
11. On the Create New Rule: Add Actions page, click **Add**. The Add Conditional Actions page displays.
12. In the Create Incident or Update Incident region, choose **Create Incident (if not associated with one)**.
13. In the Send Notifications region, complete the requisite notification details and click **Continue**. The Edit Rule Set page displays with the newly defined action listed in the table.
14. Complete the remaining Create Rule Set wizard pages to return to the Create Rule Set/ Edit Rule Set page. At this point, the two RCA event rules will have been added to the rule set.
15. Click **Save** to save the changes to the rule set.

Introducing a Time Delay

As mentioned previously, Incident Manager RCA is an iterative process whereby incoming target down events are continually being evaluated, resulting in updates to causal analysis states. Over a period of time (minutes), a root cause may or may not remain a root cause depending on incoming target down events. The original target down event may later be classified as a symptom. To allow time for target down events to be reported, analyzed, and then acted upon (such as creating an incident), you can define an event evaluation time delay when creating a rule set.

In the previous example, where incidents are created for non-symptom events, without a time delay in the rule, there could potentially be an incident created for a non-symptom event that eventually becomes a symptom.

To add a time delay to the rule:

1. From the *Create Rule Set* wizard *Add Actions* page, click **Add** or **Edit** (modify an existing rule). The *Add Conditional Actions* page displays.
2. In the *Conditions for Actions* region, choose **Only execute the actions if specified conditions match**. A list of conditions displays.
3. Choose **Event has been open for specified duration**.
4. Specify the desired time delay.
5. Click **Continue** and complete the remaining steps in the wizard.

Moving from Enterprise Manager 10/11g to 12c and Greater

Beginning with Enterprise Manager 12c, incident management 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 emails based on specific monitoring conditions, you will continue to receive those emails without interruption. To take advantage of 12c features, however, you may need to perform additional migration tasks.

> **Note:**
>
> Alerts that were generated pre-12c will still be available. For example, critical metric alerts will be available as critical incidents.

Rules

When you migrate to Enterprise Manger 12c, all of your existing notification rules are automatically converted to 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.

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

- Appendix A, " Overview of Notification in Enterprise Manager Cloud Control" section "Migrating Notification Rules to Rule Sets" in the *Enterprise Manager Cloud Control Upgrade Guide*.
- Chapter 29 "Updating Rules" in the *Enterprise Manager Cloud Control Upgrade Guide*.

Privilege Requirements

The *Create Enterprise Rule Set* resource privilege is now required in order to edit/create enterprise rule sets and rules contained 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 *Create Enterprise Rule Set* resource privilege if they wish to create new rules. Enterprise Manager Super Administrators, by default, can edit and create rule sets.

Monitoring: Common Tasks

The following sections provide "how-to" examples illustrating common tasks for incident/ monitoring setup and usage.

- Setting Up a Mail Server for Notifications
- Sending Email for Metric Alerts
- Sending SNMP Traps for Metric Alerts
- Sending Events to an Event Connector
- Sending Email to Different Email Addresses for Different Periods of the Day

Sending Email for Metric Alerts

Task

Configure Enterprise Manager to send email to administrators when a metric alert threshold is reached. In this example, you want to send an email notification when a metric alert is raised when CPU Utilization reaches Critical severity.

User Roles

- IT Operator/Manager
- Enterprise Manager Administrator

Prerequisites

- Set up an Email Gateway that allows Enterprise Manager to send email to administrators.

  For more information, see Setting Up a Mail Server for Notifications.

- Metric thresholds have been set for CPU Utilization.
- User's Enterprise Manager account has been granted the appropriate privileges to manage incidents from his managed system.

  For information, see Setting Up Administrators and Privileges.
- User's Enterprise Manager account has notification preferences (email and schedule). This is required not just for the administrator who is creating/editing a rule, but also for any user who is being notified as a result of the rule action.

  For more information, see Setting Up a Notification Schedule.

How to do it:

1. From the **Setup** menu, select **Incidents,** then select **Incident Rules.**
2. Click **Create Rule Set.**
3. Enter a name and description for the rule set.
4. In the Targets tab, select **All targets that the rule set owner can view**.

> **Note:**
>
> *Having the rule set apply to specific targets/group.*
>
> Although we have chosen to have the rule set apply to all targets in this example, alternatively, you can have a rule set apply only to specific targets or groups.
>
> To do this:
>
> a. From the Targets tab, select **Specific targets**.
>
> b. From the Add drop-down menu, choose **Groups** or **Targets**
>
> c. Click **Add**. The Target selector dialog displays.
>
> d. Either search for a target/group name or select one from the table.
>
> e. Click **Select** once you have chosen the targets/groups of interest. The dialog closes and the targets appear in the Specific Targets list.

5. In the Rules tab, click **Create.** The Select Type of Rule to Create dialog appears.

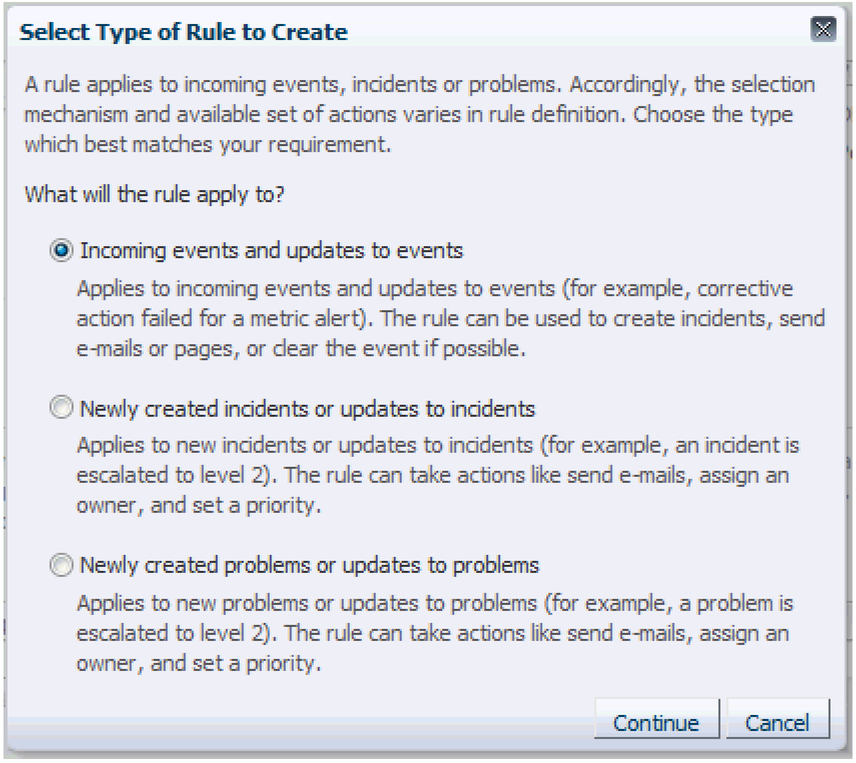

Select Type of Rule to Create

A rule applies to incoming events, incidents or problems. Accordingly, the selection mechanism and available set of actions varies in rule definition. Choose the type which best matches your requirement.

What will the rule apply to?

● Incoming events and updates to events

Applies to incoming events and updates to events (for example, corrective action failed for a metric alert). The rule can be used to create incidents, send e-mails or pages, or clear the event if possible.

○ Newly created incidents or updates to incidents

Applies to new incidents or updates to incidents (for example, an incident is escalated to level 2). The rule can take actions like send e-mails, assign an owner, and set a priority.

○ Newly created problems or updates to problems

Applies to new problems or updates to problems (for example, a problem is escalated to level 2). The rule can take actions like send e-mails, assign an owner, and set a priority.

6. Select **Incoming events and updates to events,** and click **Continue.**
7. On the Select Events page, set the criteria for events based on which the rule should act. In this case, choose **Metric Alert** from the drop down list.

   Click **Next.**
8. Select the **Specific events of type Metric Alert** option. A metric selection area displays.

   In this example, we only want to send notifications for CPU % Utilization greater reaches the defined Critical threshold.

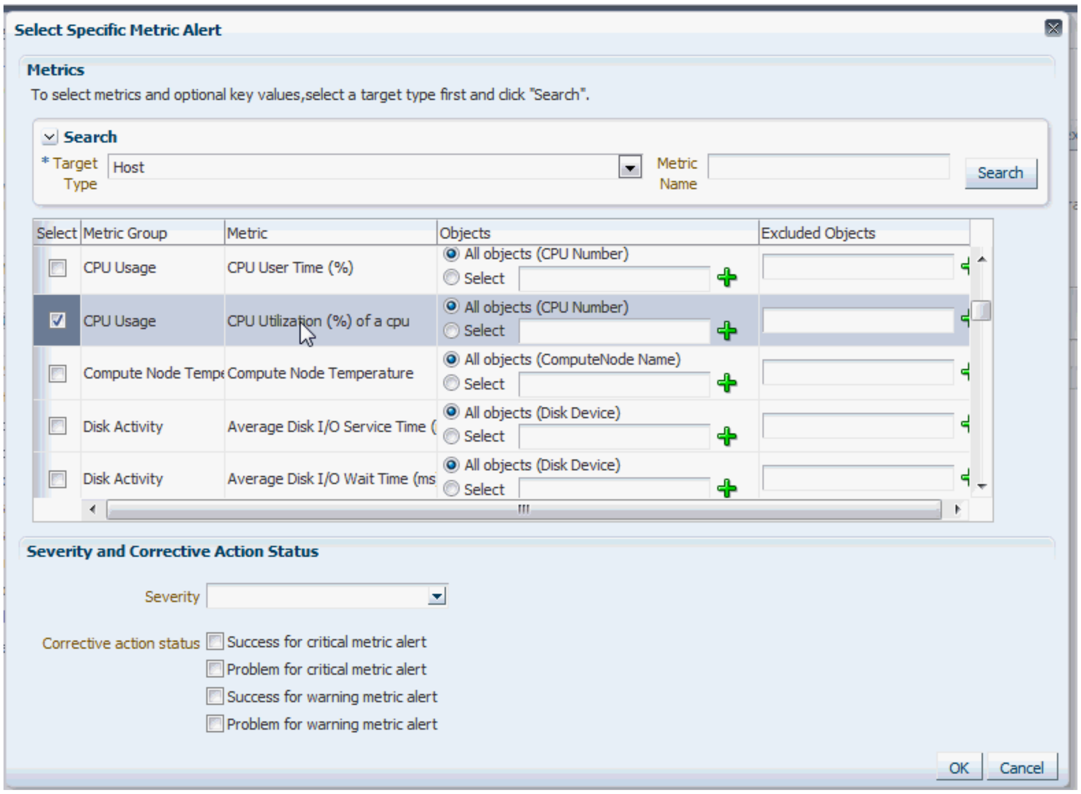

Select Specific Metric Alert

Metrics

To select metrics and optional key values,select a target type first and click "Search".

Search

* Target Type: Host   Metric Name: Search

Select	Metric Group	Metric	Objects	Excluded Objects
□	CPU Usage	CPU User Time (%)	● All objects (CPU Number) ○ Select	
☑	CPU Usage	CPU Utilization (%) of a cpu	● All objects (CPU Number) ○ Select	
□	Compute Node Tempe	Compute Node Temperature	● All objects (ComputeNode Name) ○ Select	
□	Disk Activity	Average Disk I/O Service Time (	● All objects (Disk Device) ○ Select	
□	Disk Activity	Average Disk I/O Wait Time (ms	● All objects (Disk Device) ○ Select	

Severity and Corrective Action Status

Severity:

Corrective action status
□ Success for critical metric alert
□ Problem for critical metric alert
□ Success for warning metric alert
□ Problem for warning metric alert

9. Choose Severity **Critical** from the drop down menu.

   Click **OK**.

10. Click **Next**.

11. On the Add Actions page, click **Add** and add actions to be taken by the rule. In the Notifications section, enter the email addresses where the notifications must be send. Click **Next.**

    Multiple conditional actions can be specified and evaluated sequentially (top down) in the order you add them.

> **Note:**
>
> *Sending email notifications to mailing list.*
>
> In addition to specifying email addresses, you may also specify defined Enterprise Manager administrators. Mailing distribution lists can also be specified to notify entire categories of users. Using mailing lists allows you to change who gets notified without having to update individual rule sets.

12. On the Specify Name and Description page, enter a name and description for the rule. Click **Next.**

13. On the review page, review the details, and click **Continue.**

14. On the Create Rule Set page, click **Save.**

What you have accomplished:

At this point, you have created a new rule set that will send an administrator email a notification whenever the CPU Utilization reaches the Critical metric threshold. To subscribe to this rule set, see Subscribing to Receive Email from a Rule for further instructions.

What's Next?

- How Do I Set Up Email Notifications for Other Administrators
- Add/Update/Delete Email Addresses and Define a Notification Schedule
- Responding and Working on a Simple Incident

Sending SNMP Traps for Metric Alerts

Task

You want to configure Enterprise Manager to send event information (for example, a metric alert) via SNMP trap to an HP Openview console. This is done in two phases

1. Create a notification method to send the SNMP Trap.
2. Create an incident rule to send an SNMP trap when a metric alert is raised.

User Roles

- Enterprise Manager Administrator

Prerequisites

- User must have Super Administrator privileges.

  For more information, see Setting Up a Mail Server for Notifications.

How to do it:

Create a notification method based on an SNMP Trap.

For instructions, see Sending SNMP Traps to Third Party Systems

Create an incident rule to send an SNMP trap when a metric alert is raised.

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.

   The Incident Rules - All Enterprise Rules page displays.
2. On the Incident Rules - All Enterprise Rules page, click **Create Rule Set...** The Create Rule Set page displays.
3. Enter the rule set **Name**, a brief **Description**, and select the type of source object the rule **Applies to** (Targets).
4. Click on the **Rules** tab and then click **Create...**
5. On the Select Type of Rule to Create dialog, select **Incoming events and updates to events** and then click **Continue**.
6. On the Select Events page, set the criteria for events based on which the rule should act. In this case, choose **Metric Alert** from the drop down list.

   Click **Next.**
7. Select the **Specific events of type Metric Alert** option. A metric selection area displays:

   In this example, we only want to send notifications for CPU % Utilization greater reaches the defined Critical threshold.

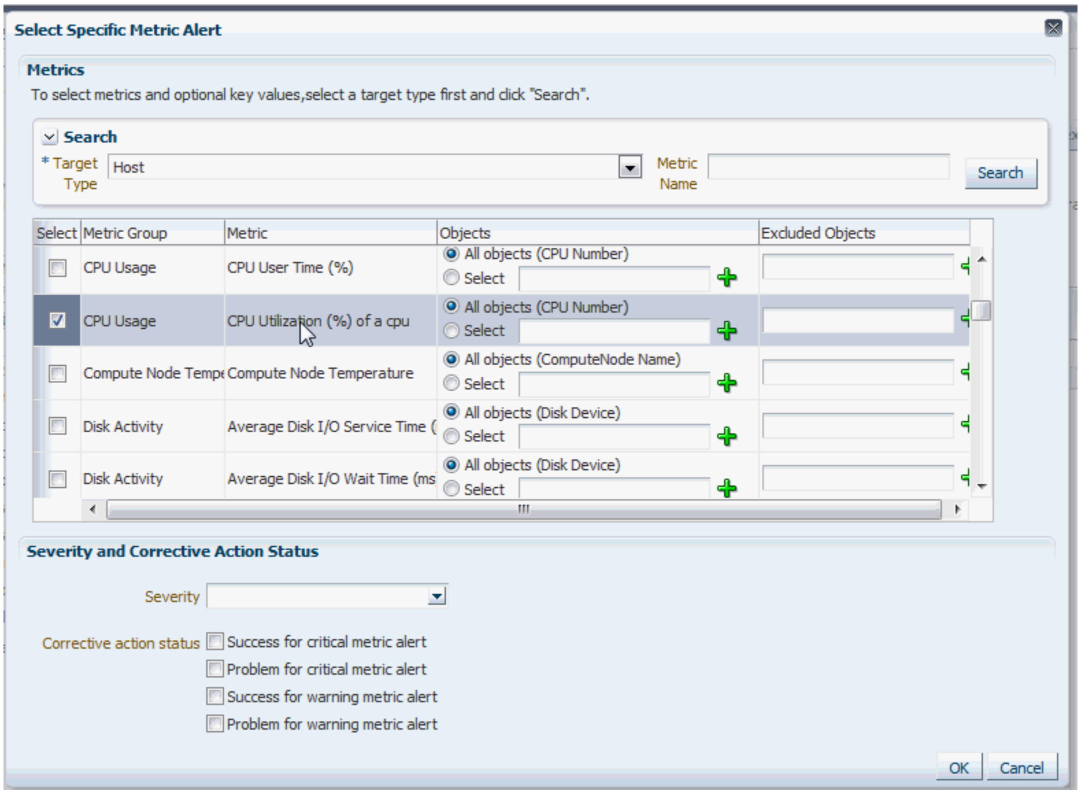

Select Specific Metric Alert

Metrics

To select metrics and optional key values,select a target type first and click "Search".

Search

* Target Type: Host

Metric Name:

Select	Metric Group	Metric	Objects	Excluded Objects
□	CPU Usage	CPU User Time (%)	● All objects (CPU Number) ○ Select	
☑	CPU Usage	CPU Utilization (%) of a cpu	● All objects (CPU Number) ○ Select	
□	Compute Node Tempe	Compute Node Temperature	● All objects (ComputeNode Name) ○ Select	
□	Disk Activity	Average Disk I/O Service Time (	● All objects (Disk Device) ○ Select	
□	Disk Activity	Average Disk I/O Wait Time (ms	● All objects (Disk Device) ○ Select	

Severity and Corrective Action Status

Severity:

Corrective action status
□ Success for critical metric alert
□ Problem for critical metric alert
□ Success for warning metric alert
□ Problem for warning metric alert

8. Choose Severity **Critical** from the drop down menu.

   Click **OK**.

9. Click **Next**.

10. On the Create New Rule : Add Actions page, click **Add**. The Add Conditional Actions page displays.

11. In the Notifications section, under Advanced Notifications, select an existing SNMP trap notification method.

    For information on creating SNMP trap notification methods, see Sending SNMP Traps to Third Party Systems.

12. Click **Continue** to return to the Create New Rule : Add Actions page.

13. Click **Next** to go to the Create New Rule : Specify Name and Description page.

14. Specify a rule name and a concise description and then click **Next**.

15. Review the rule definition and then click **Continue** add the rule to the rule set. A message displays indicating the rule has been added to the rule set but has not yet been saved. Click **OK** to close the message.

16. Click **Save** to save the rule set. A confirmation is displayed. Click **OK** to close the message.

What you have accomplished:

At this point, you have created an incident rule set that instructs Enterprise Manager to send an SNMP trap to a third-party system whenever a metric alert is raised (%CPU Utilization).

What's next?

- Subscribing to Receive Email from a Rule
- Searching for Incidents

Sending Events to an Event Connector

Task

You want to send event information from Enterprise Manager to IBM Tivoli Netcool/OMNIbus using a connector. To do so, you must create an incident rule that invokes the IBM Tivoli Netcool/OMNIbus Connector connector.

User Roles

- System Administrator
- IT Operator

Prerequisites

- User must have the Create Enterprise Rule Set resource privilege and at least View privileges on the targets where events are to be forward to Netcool/OMNIbus.

  For more information, see Setting Up a Mail Server for Notifications.

- The IBM Tivoli Netcool/OMNIbus connector must be installed and configured.

  For more information, see the Oracle® Enterprise Manager IBM Tivoli Netcool/OMNIbus Connector Installation and Configuration Guide.

How to do it:

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.

   The Incident Rules - All Enterprise Rules page displays.

2. Click **Create Rule Set.**
3. Enter a name and description for the rule set.
4. In the Targets tab, select **All targets that the rule set owner can view**.

   > **Note:**
   >
   > Although we have chosen to have the rule set apply to all targets in this example, you can alternatively have a rule set apply only to specific targets or groups.
   >
   > To do this:
   >
   > a. From the Targets tab, select **Specific targets**.
   >
   > b. From the Add drop-down menu, choose **Groups** or **Targets**
   >
   > c. Click **Add**. The Target selector dialog displays.
   >
   > d. Either search for a target/group name or select one from the table.
   >
   > e. Click **Select** once you have chosen the targets/groups of interest. The dialog closes and the targets appear in the Specific Targets list.

5. In the Rules tab, click **Create.** The Select Type of Rule to Create dialog appears.

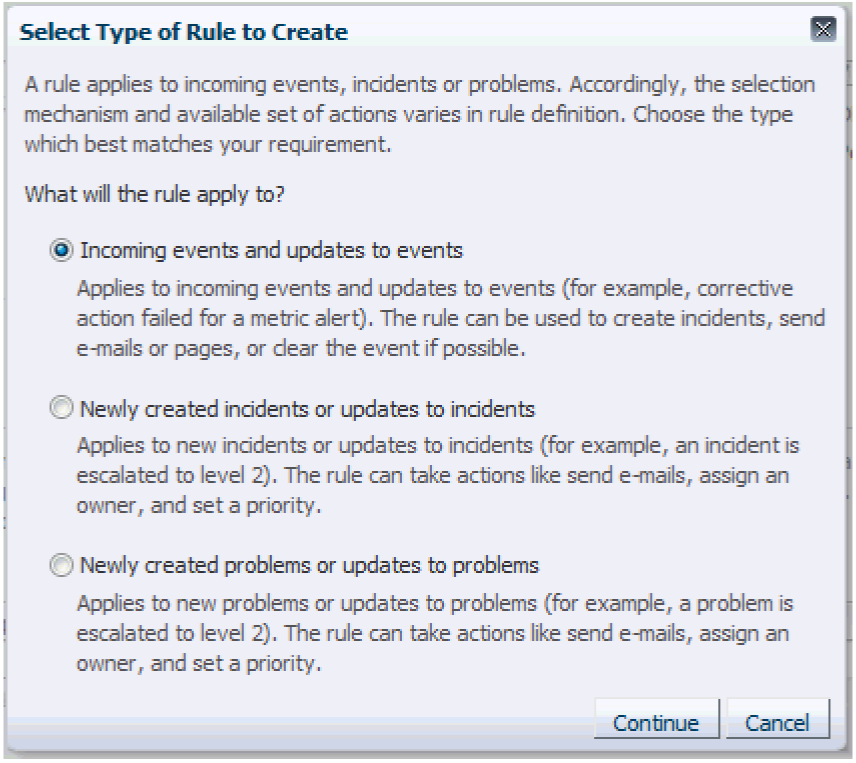

Select Type of Rule to Create

A rule applies to incoming events, incidents or problems. Accordingly, the selection mechanism and available set of actions varies in rule definition. Choose the type which best matches your requirement.

What will the rule apply to?

● Incoming events and updates to events

Applies to incoming events and updates to events (for example, corrective action failed for a metric alert). The rule can be used to create incidents, send e-mails or pages, or clear the event if possible.

○ Newly created incidents or updates to incidents

Applies to new incidents or updates to incidents (for example, an incident is escalated to level 2). The rule can take actions like send e-mails, assign an owner, and set a priority.

○ Newly created problems or updates to problems

Applies to new problems or updates to problems (for example, a problem is escalated to level 2). The rule can take actions like send e-mails, assign an owner, and set a priority.

6. Select **Incoming events and updates to events,** and click **Continue.**
7. On the Select Events page, set the criteria for events based on which the rule should act. In this case, choose **Metric Alert** from the drop down list.

   Click **Next.**
8. Select the **Specific events of type Metric Alert** option. A metric selection area displays:

   In this example, we only want to send notifications for CPU % Utilization greater reaches the defined Critical threshold.

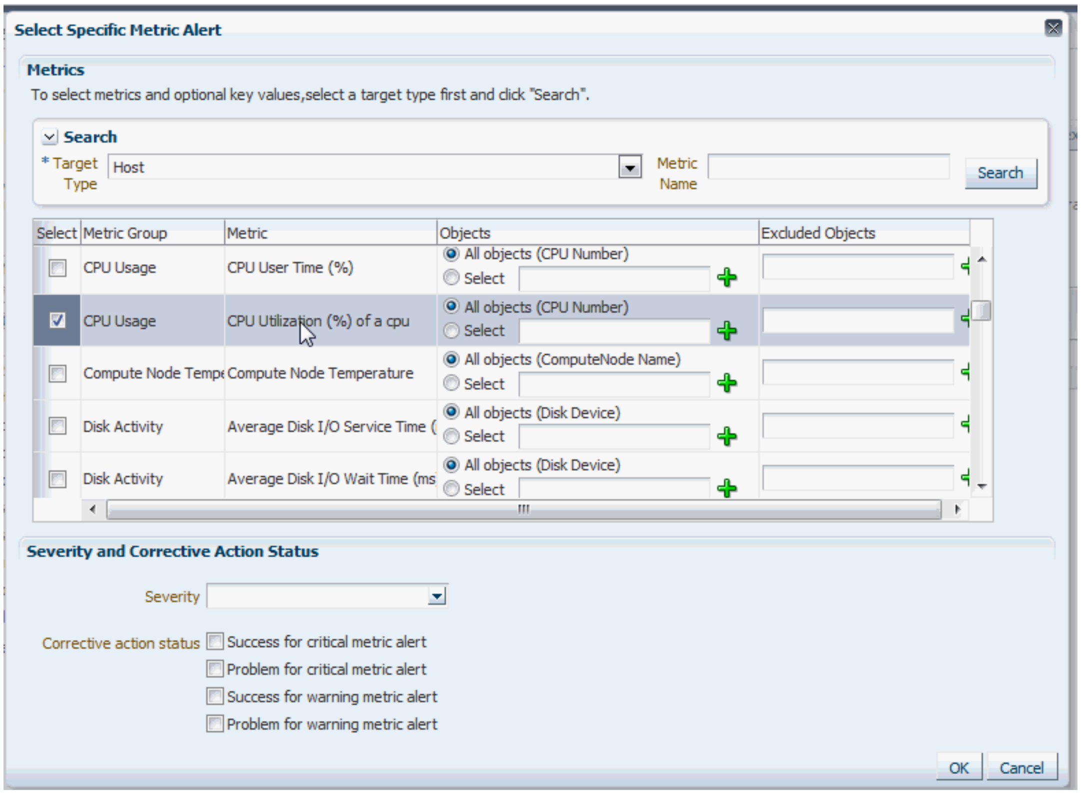

Select Specific Metric Alert

Metrics

To select metrics and optional key values,select a target type first and click "Search".

Search

* Target Type: Host

Metric Name:

Select	Metric Group	Metric	Objects	Excluded Objects
□	CPU Usage	CPU User Time (%)	● All objects (CPU Number) ○ Select	
☑	CPU Usage	CPU Utilization (%) of a cpu	● All objects (CPU Number) ○ Select	
□	Compute Node Tempe	Compute Node Temperature	● All objects (ComputeNode Name) ○ Select	
□	Disk Activity	Average Disk I/O Service Time (	● All objects (Disk Device) ○ Select	
□	Disk Activity	Average Disk I/O Wait Time (ms	● All objects (Disk Device) ○ Select	

Severity and Corrective Action Status

Severity:

Corrective action status:
□ Success for critical metric alert
□ Problem for critical metric alert
□ Success for warning metric alert
□ Problem for warning metric alert

9. Choose Severity **Critical** from the drop down menu.

   Click **OK**.

10. Click **Next**. The Add Actions page displays.

11. Click **Add**. The Add Conditional Actions page displays.

12. Select one or more connector instances listed in the Forward to Event Connectors section and, click **>** button to add the connector to the Selected Connectors list and then click **Continue**.The Add Actions page appears again and lists the new action.

13. Click **Next**. The Specify Name and Description page displays.

14. Enter a name and description for the rule, then click **Next**. The Review page displays.

15. Click **Continue** if everything appears correct.

    An information pop-up appears that states, "Rule has been successfully added to the current rule set. Newly added rules are not saved until the Save button is clicked."

    You can click **Back** and make corrections to the rule if necessary.

What you have accomplished:

At this point, you have created a rule that invokes the IBM Tivoli Netcool/OMNIbus Connector connector when a metric alert is raised.

What's next?

Subscribing to Receive Email from a Rule

Sending Email to Different Email Addresses for Different Periods of the Day

Task

Your worldwide IT department operates 24/7. Support responsibility rotates to different data centers across the globe depending on the time of day. When Enterprise Manager sends an email notification, you want it sent to the administrator currently on duty (normal work day), which in this situation changes depending on the time of day.

There are four adminstrators to handle Enterprise Manager notification:

- ADMIN_ASIA
- ADMIN_EU
- ADMIN_UK
- ADMIN_US

You want the notifications to be sent to specific administrators during their normal work hours.

User Roles

- System Administrator
- IT Operator

Prerequisites

- Email addresses have been defined for all administrators you want to send email nofifications.

  For more information, see "Defining Email Addresses".
- You must have Super Administrator privileges.
- All administrators who are to receive email notifications have been defined.

How to do it:

1. From the **Setup** menu, select **Notifications**, then select **My Notification Schedule**.

   The Notification Schedule page displays.
2. Specify the administrator who's notification schedule you wish to edit and click **Change**. The selected administrator's notification schedule displays. You can click the search icon (magnifying glass) for a list of available administrators.
3. Click **Edit Schedule Definition**. The Edit Schedule Definition: Time Period page displays. The Edit Existing Schedule option is chosen by default. If necessary, modify the rotation schedule.
4. Click **Continue**. The Edit Schedule Definition: Email Addresses page displays.
5. Follow the instructions on the Edit Schedule Definition: Email Addresses page to adjust the administrator's notification schedule as required.
6. Click **Finish** once the notification schedule changes for the selected administrator are have been made. You are returned to the Notification Schedule page.
7. Repeat this process (steps two through six) for each administrator until all four administrators' notification schedules are in sync with their normal workdays.

What you have accomplished:

You have created a notification schedule where administrators in different time zones across the globe are only sent alert notifications during their assigned work hours.

What's next?

" Subscribing to Receive Email from a Rule "

6 Using 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 Using 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. In Enterprise Manager 13.5 RU15, the Enterprise Manager notification system also supports sending incident or events to third-party applications through Webhooks. 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 Email
- Sending Notifications Using OS Commands and Scripts
- Sending Notifications Using PL/SQL Procedures
- Sending SNMP Traps to Third Party Systems
- Sending Notifications Using Webhooks and Slack
- Management Information Base (MIB)
- Passing Corrective Action Status Change Information
- Passing Job Execution Status Information
- Passing User-Defined Target Properties to Notification Methods
- Troubleshooting Notifications
- EMOMS Properties
- Passing Event, Incident, Problem Information to an OS Command or Script
- Passing Information to a PL/SQL Procedure

Setting Up Notifications

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

Setting Up a Mail Server for Notifications

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

You specify the Outgoing Mail (SMTP) server on the Notification Methods page. To display the Notification Methods page, from the **Setup** menu, select **Notifications**, then select **Mail Servers**.

> **Note:**
>
> You must have Super Administrator privileges in order to configure the Enterprise Manager notifications system. This includes:
>
> - Setting up the SMTP server
> - Defining notification methods
> - Customizing notification email formats

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 email address you want to use to send your email 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 email address if any problem is encountered during the sending of an email notification. Example 6-1 shows sample notification method entries.

> **Note:**
>
> The email address you specify on this page is not the email address to which the notification is sent. You will have to specify the email address (where notifications will be sent) from the Password and Email page. From the **Setup** menu, choose **MyPreferences** and then **Enterprise Manager Password & Email**.
>
> As standard practice, each user should have their own email address.

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

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

Setting the Enterprise Manager Console URL when Using an SLB

If you have a multi-OMS environment with a Server Load Balancer (SLB) configured for the OMS instances, you should update the console URL to ensure that any emails from Enterprise Manager direct you to the Enterprise Manager console through the SLB URL and not the specific OMS URL from which the email may have originated.

To change the console URL:

1. From the **Setup** menu, select **Manage the Manager**, and then **Health Overview**. The Management Services and Repository page displays.
2. On the Management Services and Repository page, in the Overview section, click **Add/ Edit** against the *Console URL* label.

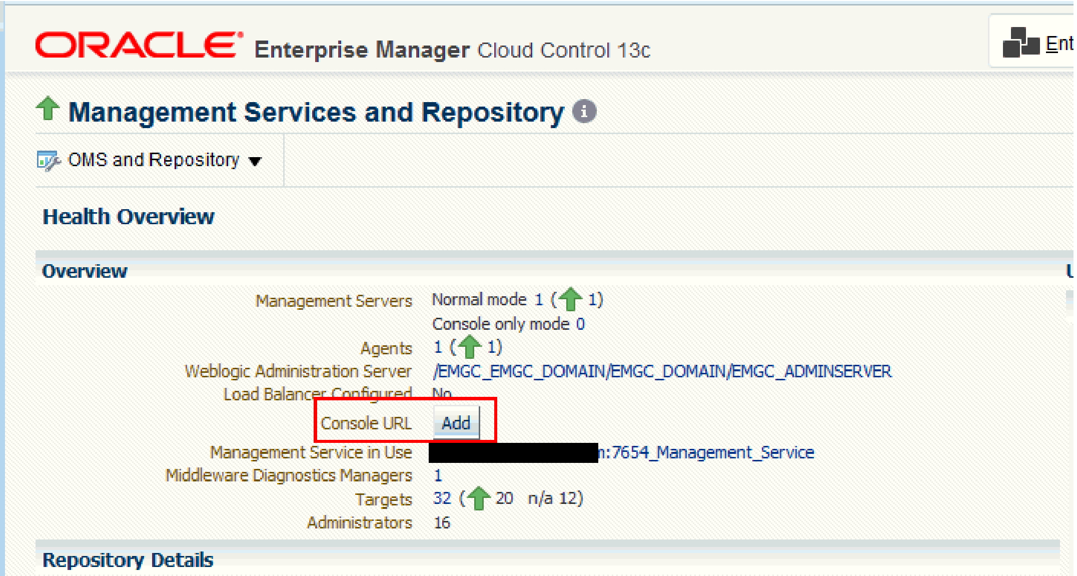
ORACLE Enterprise Manager Cloud Control 13c
Ent
Management Services and Repository
OMS and Repository
Health Overview
Overview
Repository Details

Management Servers: Normal mode 1 (1)
Console only mode 0
Agents: 1 (1)
Weblogic Administration Server: /EMGC_EMGC_DOMAIN/EMGC_DOMAIN/EMGC_ADMINSERVER
Load Balancer Configured: No
Console URL: Add
Management Service in Use: [redacted]:7654_Management_Service
Middleware Diagnostics Managers: 1
Targets: 32 (20 n/a 12)
Administrators: 16

The *Console URL* page displays.

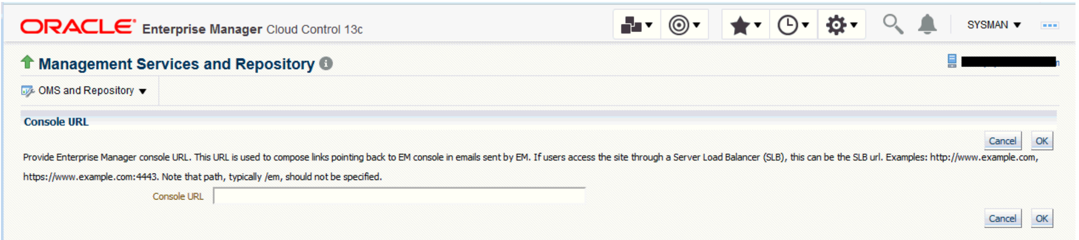
ORACLE Enterprise Manager Cloud Control 13c
SYSMAN
Management Services and Repository
OMS and Repository
Console URL
Cancel
OK

[redacted]

Provide Enterprise Manager console URL. This URL is used to compose links pointing back to EM console in emails sent by EM. If users access the site through a Server Load Balancer (SLB), this can be the SLB url. Examples: http://www.example.com, https://www.example.com:4443. Note that path, typically /em, should not be specified.

Console URL:

3. Modify the Console URL to the SLB URL.

   Examples:

   http://www.example.com

   https://www.example.com:4443.

   Note that path, typically */em*, should not be specified.

4. Click **OK**.

Example 6-1 Mail Server Settings

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

Setting Up Email for Yourself

If you want to receive notifications by email, you will need to specify your email address(s) in the Password & Email page (from the **Setup** menu, select **MyPreferences**, then select **Enterprise Manager Password & Email**). In addition to defining notification email addresses, you associate the notification message format (long, short, pager) to be used for your email address.

Setting up email involves three steps:

Step 1: Define an email addresses.

Step 2: Set up a Notification Schedule.

Step 3: Subscribe to incident rules in order to receive emails.

Defining Email Addresses

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

To add an email address:

1. From *username* drop-down menu, select **Enterprise Manager Password & Email**.
2. Click **Add Another Row** to create a new email entry field in the **Email Addresses** table.
3. Specify the email associated with your Enterprise Manager account. All email notifications you receive from Enterprise Manager will be sent to the email addresses you specify.

   For example, `user1@myco.com`

   Select the *Email Type* (message format) for your email address. *Email (Long)* sends a HTML formatted email that contains detailed information. Example 6-2 shows a typical notification that uses the long format.

   Email (Short) and *Pager(Short)* (Example 6-3) send a concise, text email that is limited to a configurable number of characters, thereby allowing the email 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 email address.

Example 6-2   Long Email Notification for Metric Alerts

```
Target type=Host Target name=machine6140830.example.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=FilesystemsMetric=Filesystem Space Available (%)Metric value=54.39Key Value=/Key
Column 1=Mount PointRule Name=NotifRuleSet1,Event rule1 Rule Owner=SYSMAN
```

Example 6-3   Short Email 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 Email(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 email to page. Beginning with Enterprise Manager 12c, the notification system allows you to tag email addresses explicitly as 'page' or 'email'. Explicit system differentiation between these two notification methods allows you to take advantage of the multiple action capability of incident rules. For example, the email versus page distinction is required in order to send you an email 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 email.

> **Note:**
>
> To receive a Page, an administrator should be added to the Page Notification option in the Incident Rule.

Setting Up a Notification Schedule

Once you have defined your email notification addresses, you will need to define a notification schedule. For example, if your email addresses are user1@myco.com, user2@myco.com, user3@myco.com, you can choose to use one or more of these email addresses for each time period in your notification schedule. Only email addresses that have been specified with your user preferences (**Enterprise Manager Password and Email** page) can be used in the notification schedule.

> **Note:**
>
> When you enter email 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 email 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 email 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, select **Notifications**, then select **My Notification Schedule**.
2. Follow the directions on the Notification Schedule page to specify when you want to receive emails.

Subscribe to Receive Email 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 email) 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, email 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 email, 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, select **Incidents**, then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, click on the rule set containing incident escalation rule in question and click **Edit...**. Rules are created in the context of a 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. In the Rules section of the Edit Rule Set page, highlight the escalation rule and click **Edit...**.
4. Navigate to the Add Actions page.
5. Select the action that escalates the incident and click **Edit...**

6. In the Notifications section, add the DBA to the **Email cc** list.
7. Click **Continue** and then navigate back to the **Edit Rule Set** page and click **Save**.

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 email notifications for the conditions defined in the rule using the subscribe procedure shown in the previous section.

The out-of-box incident rule set for all targets does not generate emails for *warning* alerts by default.

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, select **Incidents**, then select **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, then click **Create**.
5. Choose the incoming incident, event or problem to which you want the rule to apply. See "Setting Up Rule Sets" 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.

Setting Up Email for Other Administrators

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

Step 1: Ensure Each Administrator Account has an Associated Email Address

Each administrator to which you want to send email notifications must have a valid email address.

1. From the **Setup** menu, select **Security** and then **Administrators**.
2. For each administrator, define an email address. This sets up a 24x7 notification schedule for this user that uses all the email addresses specified. By default, this adds the *Email ID* with type set to *Email Long*. It is not possible to specify the *Email Type* option here.

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 email 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 email address for the first time, you should review and edit the notification schedule as needed.

1. From the **Setup** menu, select **Notifications**, then select **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: Email Addresses page appears.
5. Follow the directions on the Edit Schedule Definition: Email 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, select **Incidents**, then select **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 **Email To** or **Email Cc** field in the **Basic Notification** region.
6. Click **Continue**, click **Next**, click **Next**, click **Continue**, and finally click **Save**.

Email Customization

Enterprise Manager allows Super Administrators to customize global email 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 emails by selecting from a wide variety of information content.

To customize an email:

1. From the **Setup** menu, select **Notifications,** then select **Customize Email Formats**.
2. Choose the **Type** and **Format**.
3. Click **Customize**. The Customize Email Template page displays.

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

Email Customization Reference

The following reference summarizes the semantics and component syntax of the pseudo-language used to define emails. The pseudo-language provides you with a simple, yet flexible way to customize email 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 email. 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 email.
- HTML is not supported.

Reserved Words and Operators

The following table lists all reserved words and operators used when modifying email scripts.

Table 6-1 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 email content. The text will be displayed in the email 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 email.

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 email 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 6-2 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 email. If more than one attribute is applicable, only the left-most attribute is used.

Table 6-2 (Cont.) Control Structures

<table>
<tr><th>Control Structure</th><th>Description</th></tr>
<tr><td>IF</td><td>

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]
```

</td></tr>
</table>

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 email script, it will convert the “<" and “>" characters of HTML tags into encoded format (&lt; and &gt;). This ensures that the HTML tag is not treated as HTML by the destination system.

Examples

Email customization template scripts support three main operators.

- Comparison operators: EQ/NEQ/GT/LT/GE/LELogic operators: AND/ORPipeline operator: |

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 (email, OS command, PL/SQL procedure, and SNMP trap).

Configuring Repeat Notifications Globally

To enable repeat notifications for a notification method (globally), select the **Send Repeat Notifications** option on the Notification Methods page . In addition to setting the maximum number of repeat notifications, you can also set the time interval at which the notifications are sent.

> **Note:**
>
> 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.

Configuring Repeat Notifications Via Incident Rules

Setting repeat notifications globally at the notification method level may not 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.

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

Non-Email Repeat Notifications

For non-email repeat notifications (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.

Extending Notification Beyond Email

Notification Methods are the mechanisms by which notifications are sent. Enterprise Manager Super Administrators can set up email notifications by configuring the 'email' 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 email notifications through incident rules
- Create other custom notification methods using OS and PL/SQL scripts and SNMP traps.
- Set global repeat notifications.

Sending Notifications Using OS Commands and Scripts

Notification system can invoke a custom script when an incident rule matches the OS Command advanced notification action. A 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 *oracle.sysman.core.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.

> **Note:**
>
> Running an OS command such as "sudo" for receiving notifications will fail because the command does not have read permission of the OMS account. The OMS account must have read permission over the OS command in order to send notifications.
>
> To overcome the permissions problem, embed the command in a wrapper script that is readable by the OMS administrator account. Once the command is contained within the wrapper script, you then specify this script in place of the OS command.

Registering a Custom Script

In order to use a custom script, you must first register the script with the notification system. This is performed in four steps:

1. Define your OS command or script.
2. Deploy the script on each Management Service host.
3. Register your OS Command or Script as a new Notification Method.
4. Assign the notification method to an incident rule.

Example 6-4 Changing the oracle.sysman.core.notification.os_cmd_timeout emoms Property

```
emctl set property -name oracle.sysman.core.notification.os_cmd_timeout value 30
```

Example 6-5 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
```

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"
- " Passing Corrective Action Execution Status to an OS Command or Script"
- " Passing Job Execution Status to an OS Command or Script"

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.

> **Note:**
>
> Both scripts and OS Commands should be specified using absolute paths. 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, eventually progressing to 30 minutes. From here, the procedure is retried every 30 minutes until the notification is a 24 hours old. The notification will be then be purged.

Example 6-4 shows the parameter in emoms.properties that controls how long the OS Command can execute without being killed by the Management Service. This prevents OS Commands from running for an inordinate length of time and blocks the delivery of other notifications. By default the command is allowed to run for 30 seconds before it is killed. The *oracle.sysman.core.notification.os_cmd_timeout* emoms property can be configured to change the default timeout value.

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, select **Notifications**, then select **Notification Methods**. From this page, you can define a new notification based on the 'OS Command' type. See "Sending Notifications Using OS Commands and Scripts".

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 6-5 shows information required for the notification method.

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

Step 4: Assign the notification method to an incident 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** and then **Incident Rules**. The Incident Rules page provides access to all available rule sets.

For detailed reference information on passing event, incident, and problem information to an OS Command or script, see "Passing Event, Incident, Problem Information to an OS Command or Script".

Script Examples

The sample OS script shown in Example 6-6 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 6-6 Sample OS Command Script

```
#!/bin/ksh

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

Example 6-7 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 6-7 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 6-8 shows a script that sends an alert to an HP OpenView console from Enterprise Manager. When a metric alert is triggered, the Enterprise Manager displays the alert. The HP OpenView script is then called, invoking opcmsg and forwarding the information to the HP OpenView management server.

Example 6-8   HP OpenView Script

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

Migrating pre-12c OS Command Scripts

This section describes how to map pre-12c OS Command notification shell environment variables to 13c OS Command shell environment variables.

> **Note:**
>
> Pre-12*c* notification rules only map to event level rules. Mapping to incident level rules is not permitted.

> **Note:**
>
> Policy Violations are no longer supported beginning with the Enterprise Manager 12*c* 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 6-3   Pre-12c/13c metric_alert Environment Variable Mapping

Pre-12c Environment Variable	Corresponding 13c 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

Table 6-3 (Cont.) Pre-12c/13c metric_alert Environment Variable Mapping

Pre-12c Environment Variable	Corresponding 13c Environment Variables
REPEAT_COUNT	REPEAT_COUNT
RULE_NAME	RULESET_NAME
RULE_OWNER	RULE_OWNER
SEVERITY	SEVERITY
TARGET_NAME	TARGET_NAME
TARGET_TYPE	TARGET_TYPE
TIMESTAMP	EVENT_REPORTED_TIME
METRIC_VALUE	VALUE
VIOLATION_CONTEXT	EVENT_CONTEXT_ATTRS
VIOLATION_GUID	SEVERITY_GUID
POLICY_RULE	No mapping, Obsolete as of Enterprise Manager 12c release.

To obtain KEY_VALUE_NAME and KEY_VALUE, 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_COLUMN_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 6-4 pre-12c/12c target_availability Environment Variable Mappings

Pre-12c Environment Variable	Corresponding 13c 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

Table 6-4 (Cont.) pre-12c/12c target_availability Environment Variable Mappings

Pre-12c Environment Variable	Corresponding 13c Environment Variables
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 'job_status_change'.

Table 6-5 pre-12c/13c job_status_change Environment Variable Mappings

Pre-12c Environment Variable	Corresponding 13c 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 6-6 pre-12c/13c Corrective Action Environment Variable Mappings

Pre-12c Environment Variable	Corresponding 13c Environment Variables
NUM_TARGETS	1
TARGET_NAME1	TARGET_NAME
TARGET_TYPE1	TARGET_TYPE
JOB_NAME	CA_JOB_NAME
JOB_OWNER	CA_JOB_OWNER
JOB_STATUS	CA_JOB_STATUS
JOB_TYPE	CA_JOB_TYPE

Notification Type Mapping

Table 6-7 pre-12c/13c notif_type Mappings

notif_type (13c)	notif_type (Pre-12c)
NOTIF_NORMAL	1
NOTIF_REPEAT	4
NOTIF_DURATION	9
NOTIF_RETRY	2

Sending Notifications Using PL/SQL Procedures

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

> **Note:**
>
> When upgrading from pre-12c to 13c versions of Enterprise Manager, existing pre-12c PL/SQL advanced notification methods will continue to work without modification. You should, however, update the procedures to use new signatures.
>
> New PL/SQL advanced notification methods created with Enterprise Manager 13c must use the new signatures documented in the following sections.

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

Defining a PL/SQL-based Notification Method

Creating a PL/SQL-based notification method consists of four steps:

1. Define the PL/SQL procedure.
2. Create the PL/SQL procedure on the Management Repository.
3. Register your PL/SQL procedure as a new notification method.
4. Assign the notification method to an incident rule.

Example 6-9   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
```

Example 6-10   PL/SQL Script

```
-- Assume log_table is created by following DDL
-- CREATE TABLE log_table (message VARCHAR2(4000)) ;
-- Define PL/SQL notification method for Events
CREATE OR REPLACE PROCEDURE log_table_notif_proc(s IN GC$NOTIF_EVENT_MSG)
IS
  l_categories gc$category_string_array;
  l_category_codes gc$category_string_array;
  l_attrs gc$notif_event_attr_array;
  l_ca_obj gc$notif_corrective_action_job;
BEGIN
  INSERT INTO log_table VALUES ('notification_type: ' || s.msg_info.notification_type);
  INSERT INTO log_table VALUES ('repeat_count: ' || s.msg_info.repeat_count);
  INSERT INTO log_table VALUES ('ruleset_name: ' || s.msg_info.ruleset_name);
  INSERT INTO log_table VALUES ('rule_name: ' || s.msg_info.rule_name);
  INSERT INTO log_table VALUES ('rule_owner: ' || s.msg_info.rule_owner);
  INSERT INTO log_table VALUES ('message: ' || s.msg_info.message);
  INSERT INTO log_table VALUES ('message_url: ' || s.msg_info.message_url);
  INSERT INTO log_table VALUES ('event_instance_guid: ' ||
s.event_payload.event_instance_guid);
  INSERT INTO log_table VALUES ('event_type: ' || s.event_payload.event_type);
  INSERT INTO log_table VALUES ('event_name: ' || s.event_payload.event_name);
  INSERT INTO log_table VALUES ('event_msg: ' || s.event_payload.event_msg);
  INSERT INTO log_table VALUES ('source_obj_type: ' ||
s.event_payload.source.source_type);
  INSERT INTO log_table VALUES ('source_obj_name: ' ||
s.event_payload.source.source_name);
  INSERT INTO log_table VALUES ('source_obj_url: ' || s.event_payload.source.source_url);
  INSERT INTO log_table VALUES ('target_name: ' || s.event_payload.target.target_name);
  INSERT INTO log_table VALUES ('target_url: ' || s.event_payload.target.target_url);
  INSERT INTO log_table VALUES ('severity: ' || s.event_payload.severity);  INSERT INTO
log_table VALUES ('severity_code: ' || s.event_payload.severity_code);
  INSERT INTO log_table VALUES ('event_reported_date: ' ||
to_char(s.event_payload.reported_date, 'D MON DD HH24:MI:SS'));

  l_categories := s.event_payload.categories;
  IF l_categories IS NOT NULL
  THEN
    FOR c IN 1..l_categories.COUNT
    LOOP
      INSERT INTO log_table VALUES ('category ' || c || ' - ' || l_categories(c));
    END LOOP;
  END IF;
  l_category_codes := s.event_payload.category_codes;
  IF l_categories IS NOT NULL
  THEN
    FOR c IN 1..l_category_codes.COUNT
    LOOP
      INSERT INTO log_table VALUES ('category_code ' || c || ' - ' ||
l_category_codes(c));
    END LOOP;
  END IF;

  l_attrs := s.event_payload.event_attrs;
  IF l_attrs IS NOT NULL
  THEN
    FOR c IN 1..l_attrs.COUNT
    LOOP
        INSERT INTO log_table VALUES ('EV.ATTR name=' || l_attrs(c).name || ' value=' ||
l_attrs(c).value || ' nls_value=' || l_attrs(c).nls_value);    END LOOP;
  END IF;

COMMIT ;
END ;
/
```

Example 6-11   PL/SQL Script to Log Events to a Table

```
CREATE TABLE event_log (
  notification_type      VARCHAR2(32),
  repeat_count           NUMBER,
  ruleset_name           VARCHAR2(256),
  rule_owner             VARCHAR2(256),
  rule_name              VARCHAR2(256),
  message                VARCHAR2(4000),
  message_url            VARCHAR2(4000),
  event_instance_guid    RAW(16),
  event_type             VARCHAR2(20),
  event_name             VARCHAR2(512),
  event_msg              VARCHAR2(4000),
  categories             VARCHAR2(4000),
  source_obj_type        VARCHAR2(120),
  source_obj_name        VARCHAR2(256),
  source_obj_url         VARCHAR2(4000),
  severity               VARCHAR2(128),
  severity_code          VARCHAR2(32),
  target_name            VARCHAR2(256),
  target_type            VARCHAR2(128),
  target_url             VARCHAR2(4000),
  host_name              VARCHAR2(256),
  timezone               VARCHAR2(64),
  occured                DATE,
  ca_guid                RAW(16),
  ca_name                VARCHAR2(128),
  ca_owner               VARCHAR2(256),
  ca_type                VARCHAR2(256),
  ca_status              VARCHAR2(64),
  ca_status_code         NUMBER,
  ca_job_step_output     VARCHAR2(4000),
  ca_execution_guid      RAW(16),
  ca_stage_change_guid   RAW(16)
)
;
```

```
CREATE OR REPLACE PROCEDURE log_event(s IN GC$NOTIF_EVENT_MSG)
IS
   l_categories gc$category_string_array;
   l_ca_obj gc$notif_corrective_action_job;
   l_categories_new VARCHAR2(1000);
BEGIN
    -- save event categories
   l_categories := s.event_payload.categories;
        IF l_categories IS NOT NULL
   THEN
     FOR c IN 1..l_categories.COUNT
     LOOP
       l_categories_new := (l_categories_new|| c || ' - ' || l_categories(c)||',');
     END LOOP;
   END IF;

   -- save event message
   IF s.msg_info.notification_type = 'NOTIF_CA' AND s.event_payload.corrective_action IS
NOT NULL
   THEN
     l_ca_obj := s.event_payload.corrective_action;
      INSERT INTO event_log (notification_type, repeat_count, ruleset_name, rule_name,
rule_owner, message, message_url, event_instance_guid, event_type, event_name,
event_msg, categories, source_obj_type, source_obj_name, source_obj_url, severity,
severity_code, target_name, target_type, target_url, host_name, timezone, occured,
ca_guid, ca_name, ca_owner, ca_type, ca_status, ca_status_code, ca_job_step_output,
ca_execution_guid, ca_stage_change_guid)
      VALUES (s.msg_info.notification_type, s.msg_info.repeat_count,
s.msg_info.ruleset_name, s.msg_info.rule_name,s.msg_info.rule_owner, s.msg_info.message,
s.msg_info.message_url, s.event_payload.event_instance_guid, s.event_payload.event_type,
s.event_payload.event_name, s.event_payload.event_msg, l_categories_new,
s.event_payload.source.source_type, s.event_payload.source.source_name,
s.event_payload.source.source_url, s.event_payload.severity,
s.event_payload.severity_code, s.event_payload.target.target_name,
s.event_payload.target.target_type, s.event_payload.target.target_url,
s.event_payload.target.host_name, s.event_payload.target.target_timezone,
s.event_payload.occurrence_date, l_ca_obj.JOB_GUID, l_ca_obj.JOB_NAME,
l_ca_obj.JOB_OWNER, l_ca_obj.JOB_TYPE, l_ca_obj.JOB_STATUS, l_ca_obj.JOB_STATUS_CODE,
l_ca_obj.JOB_STEP_OUTPUT, l_ca_obj.JOB_EXECUTION_GUID,
l_ca_obj.JOB_STATE_CHANGE_GUID);   ELSE
      INSERT INTO event_log (notification_type, repeat_count, ruleset_name, rule_name,
rule_owner, message, message_url, event_instance_guid, event_type, event_name,
event_msg, categories, source_obj_type, source_obj_name, source_obj_url, severity,
severity_code, target_name, target_type, target_url, host_name, timezone, occured,
ca_guid, ca_name, ca_owner, ca_type, ca_status, ca_status_code, ca_job_step_output,
ca_execution_guid, ca_stage_change_guid)
      VALUES (s.msg_info.notification_type, s.msg_info.repeat_count,
s.msg_info.ruleset_name, s.msg_info.rule_name, s.msg_info.rule_owner,
s.msg_info.message, s.msg_info.message_url, s.event_payload.event_instance_guid,
s.event_payload.event_type, s.event_payload.event_name, s.event_payload.event_msg,
l_categories_new, s.event_payload.source.source_type,
s.event_payload.source.source_name, s.event_payload.source.source_url,
s.event_payload.severity, s.event_payload.severity_code,
s.event_payload.target.target_name, s.event_payload.target.target_type,
s.event_payload.target.target_url, s.event_payload.target.host_name,
s.event_payload.target.target_timezone, s.event_payload.occurrence_date,
null,null,null,null,null,null,null,null,null);
   END IF;
   COMMIT;
END log_event;
/
```

Example 6-12 PL/SQL Script to Log Incidents to a Table

```
CREATE TABLE incident_log (
  notification_type       VARCHAR2(32),
  repeat_count            NUMBER,
  ruleset_name            VARCHAR2(256),
  rule_owner              VARCHAR2(256),
  rule_name               VARCHAR2(256),
  message                 VARCHAR2(4000),
  message_url             VARCHAR2(4000),
  incident_id             VARCHAR2(128),
  ticket_url              VARCHAR2(4000),
  assoc_event_cnt         NUMBER,
  severity                VARCHAR2(128),
  severity_code           VARCHAR2(32),
  priority                VARCHAR2(128),
  priority_code           VARCHAR2(32),
  status                  VARCHAR2(32),
  categories              VARCHAR2(1000),
  target_name             VARCHAR2(256),
  target_type             VARCHAR2(128),
  host_name               VARCHAR2(256),
  timezone                VARCHAR2(64),
  occured                 DATE
)
;
   CREATE OR REPLACE PROCEDURE log_incident(s IN GC$NOTIF_INCIDENT_MSG)
IS
   l_src_info_array GC$NOTIF_SOURCE_INFO_ARRAY;
   l_src_info  GC$NOTIF_SOURCE_INFO;
   l_categories gc$category_string_array;
   l_target_obj GC$NOTIF_TARGET;
   l_target_name VARCHAR2(256);
   l_target_type VARCHAR2(256);
   l_target_timezone VARCHAR2(256);
   l_hostname VARCHAR2(256);
   l_categories_new VARCHAR2(1000);
BEGIN
     -- Save Incident categories
  IF l_categories IS NOT NULL
   THEN
     FOR c IN 1..l_categories.COUNT
     LOOP
       l_categories_new := (l_categories_new|| c || ' - ' || l_categories(c)||',');
     END LOOP;
   END IF;

   -- GET target info
   l_src_info_array := s.incident_payload.incident_attrs.source_info_arr;
   IF l_src_info_array IS NOT NULL
   THEN
     FOR I IN 1..l_src_info_array.COUNT
     LOOP
       IF l_src_info_array(I).TARGET IS NOT NULL
       THEN
         l_target_name := l_src_info_array(I).TARGET.TARGET_NAME;
         l_target_type := l_src_info_array(I).TARGET.TARGET_TYPE;
         l_target_timezone := l_src_info_array(I).TARGET.TARGET_TIMEZONE;
         l_hostname := l_src_info_array(I).TARGET.HOST_NAME;
      END IF;
     END LOOP;
   END IF;
   -- save Incident notification message    INSERT INTO incident_log(notification_type,
repeat_count, ruleset_name, rule_owner, rule_name, message, message_url, incident_id,
ticket_url, assoc_event_cnt, severity, severity_code, priority, priority_code, status,
categories, target_name, target_type, host_name, timezone, occured)
   VALUES (s.msg_info.notification_type, s.msg_info.repeat_count,
s.msg_info.ruleset_name, s.msg_info.rule_owner, s.msg_info.rule_name,
s.msg_info.message, s.msg_info.message_url, s.incident_payload.incident_attrs.id,
s.incident_payload.ticket_url, s.incident_payload.assoc_event_count,
s.incident_payload.incident_attrs.severity,
s.incident_payload.incident_attrs.severity_code,
s.incident_payload.incident_attrs.priority,
s.incident_payload.incident_attrs.priority_code,
s.incident_payload.incident_attrs.STATUS, l_categories_new, l_target_name,
l_target_type, l_hostname,l_target_timezone,
s.incident_payload.incident_attrs.creation_date);
   COMMIT;
END log_incident;
/
```

Example 6-13   PL/SQL Script to Log Problems to a Table

```
CREATE TABLE problem_log (
  notification_type      VARCHAR2(32),
  repeat_count           NUMBER,
  ruleset_name           VARCHAR2(256),
  rule_owner             VARCHAR2(256),
  rule_name              VARCHAR2(256),
  message                VARCHAR2(4000),
  message_url            VARCHAR2(4000),
  problem_key            VARCHAR2(850),
  assoc_incident_cnt     NUMBER,
  problem_id             NUMBER,
  owner                  VARCHAR2(256),
  severity               VARCHAR2(128),
  severity_code          VARCHAR2(32),
  priority               VARCHAR2(128),
  priority_code          VARCHAR2(32),
  status                 VARCHAR2(32),
  categories             VARCHAR2(1000),
  target_name            VARCHAR2(256),
  target_type            VARCHAR2(128),
  host_name              VARCHAR2(256),   timezone               VARCHAR2(64),
  occured                DATE
)
;
     CREATE OR REPLACE PROCEDURE log_problem(s IN GC$NOTIF_PROBLEM_MSG)
IS
   l_src_info_array GC$NOTIF_SOURCE_INFO_ARRAY;
   l_src_info  GC$NOTIF_SOURCE_INFO;
   l_categories gc$category_string_array;
   l_target_obj GC$NOTIF_TARGET;
   l_target_name VARCHAR2(256);
   l_target_type VARCHAR2(256);
   l_target_timezone VARCHAR2(256);
   l_hostname VARCHAR2(256);
   l_categories_new VARCHAR2(1000);
BEGIN
     -- Save Problem categories
  l_categories := s.problem_payload.problem_attrs.categories;
  IF l_categories IS NOT NULL
```

```
   THEN
     FOR c IN 1..l_categories.COUNT
     LOOP
       l_categories_new := (l_categories_new|| c || ' - ' || l_categories(c)||',');
     END LOOP;
   END IF;

   -- GET target info
   l_src_info_array := s.problem_payload.problem_attrs.source_info_arr;
   IF l_src_info_array IS NOT NULL
   THEN
     FOR I IN 1..l_src_info_array.COUNT
     LOOP
       IF l_src_info_array(I).TARGET IS NOT NULL
       THEN
         l_target_name := l_src_info_array(I).TARGET.TARGET_NAME;
         l_target_type := l_src_info_array(I).TARGET.TARGET_TYPE;
         l_target_timezone := l_src_info_array(I).TARGET.TARGET_TIMEZONE;
         l_hostname := l_src_info_array(I).TARGET.HOST_NAME;
      END IF;
     END LOOP;
   END IF;

  -- save Problem notification message
   INSERT INTO problem_log(notification_type, repeat_count, ruleset_name, rule_owner,
rule_name, message, message_url, problem_key, assoc_incident_cnt, problem_id, owner,
severity, severity_code, priority, priority_code, status, categories, target_name,
target_type, host_name, timezone, occured)
   VALUES (s.msg_info.notification_type, s.msg_info.repeat_count,
s.msg_info.ruleset_name, s.msg_info.rule_owner, s.msg_info.rule_name,
s.msg_info.message, s.msg_info.message_url, s.problem_payload.problem_key,
s.problem_payload.ASSOC_INCIDENT_COUNT, s.problem_payload.problem_attrs.id,
s.problem_payload.problem_attrs.owner,
s.problem_payload.problem_attrs.severity,
s.problem_payload.problem_attrs.severity_code, s.problem_payload.problem_attrs.PRIORITY,
s.problem_payload.problem_attrs.PRIORITY_CODE, s.problem_payload.problem_attrs.status,
l_categories_new, l_target_name, l_target_type, l_hostname,l_target_timezone,
s.problem_payload.problem_attrs.CREATION_DATE);
   COMMIT;
END log_problem;
/
```

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"

"Passing Corrective Action Status Change Information"

" Passing Job Execution Status Information"

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** and then **Notification Methods** to access the Notification Methods page. From this page, you can define a new notification based on 'PL/SQL Procedure'. See Sending Notifications Using PL/SQL Procedures.

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 6-9.

Figure 6-1 illustrates how to add a PL/SQL-based notification method from the Enterprise Manager UI.

Figure 6-1 Adding a PL/SQL Procedure

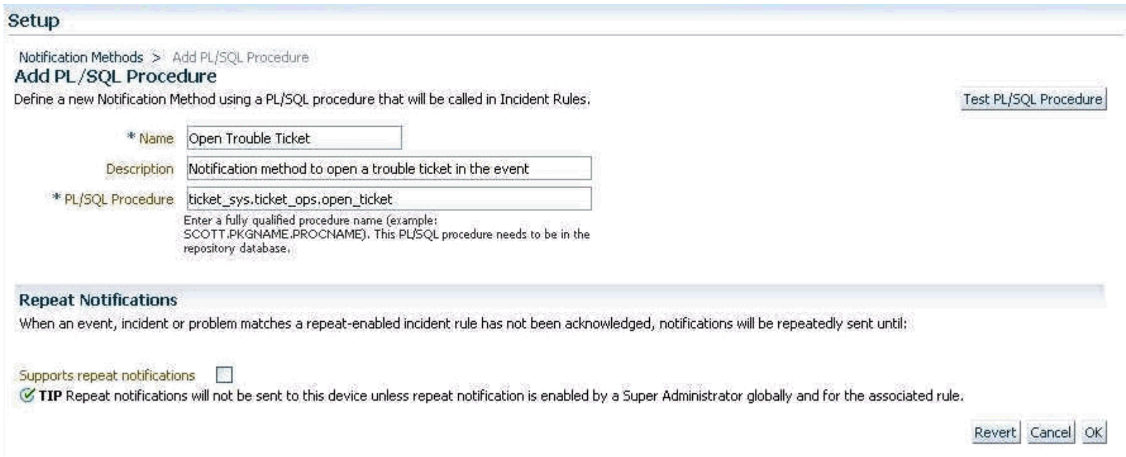

Add PL/SQL Procedure

Define a new Notification Method using a PL/SQL procedure that will be called in Incident Rules.

* Name: Open Trouble Ticket

Description: Notification method to open a trouble ticket in the event

* PL/SQL Procedure: ticket_sys.ticket_ops.open_ticket

Enter a fully qualified procedure name (example: SCOTT.PKGNAME.PROCNAME). This PL/SQL procedure needs to be in the repository database.

Repeat Notifications

When an event, incident or problem matches a repeat-enabled incident rule has not been acknowledged, notifications will be repeatedly sent until:

Supports repeat notifications □

TIP Repeat notifications will not be sent to this device unless repeat notification is enabled by a Super Administrator globally and for the associated rule.

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, select **Incidents** and then select **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.

See "Passing Information to a PL/SQL Procedure" for more information about how incident, event, and problem information is passed to the PLSQL procedure.

Migrating Pre-12c PL/SQL Advanced Notification Methods

Pre-12*c* notifications map to event notifications in Enterprise Manager 12*c*. The event types metric_alert, target_availability and job_status_alert correspond to the pre-12c notification functionality.

> **Note:**
>
> Policy Violations are no longer available beginning with Enterprise Manager 12*c*.

This section describes the mapping between Enterprise Manager 13c 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 13c PL/SQL notification payload.Please note that Policy Violations are no longer supported in the 13c 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 6-8 Metric Alert Mapping

MGMT_NOTIFY_SEVERITY	13c 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$notif_event_payload.severity_code, see Table 6-9.
MESSAGE	gc$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$notif_event_attr.value where its name=' metric_guid' in gc$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

The following example illustrates how to obtain similar pre-12c KEY_VALUE and KEY_VALUE_NAME from an Enterprise Manager 13c notification payload.

Example 6-14 Extracting KEY_VALUE and KEY_VALUE_NAME

```
-- Get the pre-12c KEY_VALUE and KEY_VALUE_NAME from an Enterprise Manager 13c
-- notification payload
-- parameters
--   IN Parameters:
--      event_msg : The event notification payload
--   OUT Parameters
--      key_value_name_out : the KEY_VALUE_NAME backward compitable to pre-12c
--                           notification payload
--      key_value_out      : the KEY_VALUE backward compitable to pre-12c
--                           notification payload
--
CREATE OR REPLACE PROCEDURE get_pre_12c_key_value(
          event_msg IN GC$NOTIF_EVENT_MSG,
          key_value_name_out OUT VARCHAR2,
          key_value_out OUT VARCHAR2)
IS

  l_key_columns MGMT_SHORT_STRING_ARRAY := MGMT_SHORT_STRING_ARRAY();
  l_key_column_values MGMT_MEDIUM_STRING_ARRAY := MGMT_MEDIUM_STRING_ARRAY();
  l_key_value VARCHAR2(1790) := NULL;
  l_num_keys NUMBER := 0;
  l_attrs gc$notif_event_attr_array;
  l_key_value_name VARCHAR2(512);
BEGIN
  l_attrs := event_msg.event_payload.event_attrs;
  key_value_name_out := NULL;
  key_value_out := NULL;

  IF l_attrs IS NOT NULL AND
     l_attrs.COUNT > 0
  THEN
    l_key_columns.extend(7);
    l_key_column_values.extend(7);
    FOR c IN 1..l_attrs.COUNT
    LOOP
      CASE l_attrs(c).name
        WHEN 'num_keys' THEN
          BEGIN
            l_num_keys := to_number(l_attrs(c).value);
          EXCEPTION
          WHEN OTHERS THEN
            -- should never happen, but guard against it l_num_keys := 0;
          END;
        WHEN 'key_value' THEN
          l_key_value := substr(l_attrs(c).nls_value,1,1290);
        WHEN 'key_column_1' THEN
          l_key_columns(1) := substr(l_attrs(c).nls_value,1,64);
        WHEN 'key_column_2' THEN
          l_key_columns(2) := substr(l_attrs(c).nls_value,1,64);
        WHEN 'key_column_3' THEN
          l_key_columns(3) := substr(l_attrs(c).nls_value,1,64);
        WHEN 'key_column_4' THEN
          l_key_columns(4) := substr(l_attrs(c).nls_value,1,64);
        WHEN 'key_column_5' THEN
          l_key_columns(5) := substr(l_attrs(c).nls_value,1,64);
        WHEN 'key_column_6' THEN
          l_key_columns(6) := substr(l_attrs(c).nls_value,1,64);
        WHEN 'key_column_7' THEN
          l_key_columns(7) := substr(l_attrs(c).nls_value,1,64);
        WHEN 'key_column_1_value' THEN
          l_key_column_values(1) := substr(l_attrs(c).nls_value,1,256);
        WHEN 'key_column_2_value' THEN
          l_key_column_values(2) := substr(l_attrs(c).nls_value,1,256);
        WHEN 'key_column_3_value' THEN
          l_key_column_values(3) := substr(l_attrs(c).nls_value,1,256);
        WHEN 'key_column_4_value' THEN
          l_key_column_values(4) := substr(l_attrs(c).nls_value,1,256);
        WHEN 'key_column_5_value' THEN
```

```
          l_key_column_values(5) := substr(l_attrs(c).nls_value,1,256);
        WHEN 'key_column_6_value' THEN
          l_key_column_values(6) := substr(l_attrs(c).nls_value,1,256);
        WHEN 'key_column_7_value' THEN
          l_key_column_values(7) := substr(l_attrs(c).nls_value,1,256);
        ELSE
          NULL;
      END CASE;
    END LOOP;

    -- get key_value and key_value_name when l_num_keys > 0

    IF l_num_keys > 0
    THEN
      -- get key value name
      IF l_key_columns IS NULL OR l_key_columns.COUNT = 0
      THEN
        key_value_name_out := NULL;
      ELSE
        l_key_value_name := NULL;
        FOR i in l_key_columns.FIRST..l_num_keys
        LOOP
          IF i > 1
          THEN
            l_key_value_name := l_key_value_name || ';';
          END IF;
          l_key_value_name := l_key_value_name || l_key_columns(i);
        END LOOP;
        key_value_name_out := l_key_value_name;
      END IF;
            -- get key_value
      IF l_num_keys = 1
      THEN
        key_value_out := l_key_value;
      ELSE
        l_key_value := NULL;
        IF l_key_column_values IS NULL OR l_key_column_values.COUNT = 0
        THEN
          key_value_out := NULL;
        ELSE
          FOR i in l_key_column_values.FIRST..l_num_keys
          LOOP
            IF i > 1
            THEN
              l_key_value := l_key_value || ';';
            END IF;
            l_key_value := l_key_value || l_key_column_values(i);
                        END LOOP;
          --  max length for key value in pre-12c = 1290
          key_value_out := substr(l_key_value,1,1290);
        END IF;
      END IF;
    END IF; -- l_num_keys > 0
  END IF;  -- l_attrs IS NOT NULL
END get_pre_12c_key_value;
/
```

When the event type is metric_alert:

Use the following severity code mapping from 13c to pre-12c when the event type is *metric_alert.*

Table 6-9 Severity Code Mapping

13c 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
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 6-10 Target Availability Mapping

MGMT_NOTIFY_SEVERITY	13c 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 6-11 Job Status Change Mapping

MGMT_NOTIFY_JOB	13c 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 6-8

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 6-12 Corrective Action Mapping

MGMT_NOTIFY_CORRECTIVE_ACTION	13c Notification Payload
JOB_NAME	gc$ notif_corrective_action_job.job_name
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.

Sending SNMP Traps to Third Party Systems

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

> **Note:**
>
> In order for a third-party system to interpret traps sent by the OMS, the omstrap.v1 file must first be loaded into the third-party SNMP console. For more information about this file and its location, see "MIB Definition".
>
> The Enterprise Manager 13c version of the MIB file incorporates the 10g and 11g MIB content, thus ensuring backward compatibility with earlier Enterprise Manager releases.

Enterprise Manager supports both SNMP Version 1 and Version 3 traps. The traps are described by the MIB definition shown in Enterprise Manager MIB Definition. See "Management Information Base (MIB)" for an explanation of how the MIB works. If you are using Enterprise Manager 13c, see Interpreting Variables of the Enterprise Manager MIB and Enterprise Manager MIB Definition. If you are upgrading from a pre-12c version of Enterprise Manager, see SNMP Trap Mappings for specific version mappings.

For Enterprise Manager 12c, SNMP traps are delivered for event notifications. Starting in Enterprise Manager 13.5 RU16, SNMP trap notifications are supported for both events and incidents. SNMP trap notifications are not supported for problems.

> **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.

SNMP Version 1 Versus SNMP Version 3

SNMP Version 3 shares the same basic architecture of Version 1, but adds numerous enhancements to SNMP administration and security. The primary enhancement relevant to Enterprise Manager involves additional security levels that provide both authentication and privacy as well as authorization and access control.

User-based Security Model (USM)

USM defines the security-related procedures followed by an SNMP engine when processing SNMP messages. Enterprise Manager SNMP V3 support takes advantage of this added SNMP message-level security enhancement to provide a secure messaging environment.

USM protects against two primary security threats:

- Modification of information: The modification threat is the danger that some unauthorized entity may alter in-transit SNMP messages generated on behalf of an authorized principal in such a way as to effect unauthorized management operations, including falsifying the value of an object.
- Masquerade: The masquerade threat is the danger that management operations not authorized for some user may be attempted by assuming the identity of another user that has the appropriate authorizations.

For both SNMP versions, the basic methodology for setting up Enterprise Manager advanced notifications using SNMP traps remains the same:

1. **Define the notification method based on an SNMP trap**.
2. **Assign the notification method to an incident rule**.

Working with SNMP V3 Trap Notification Methods

The procedure for defining an SNMP V3 trap notification method differs slightly from that of V1. Beginning with Enterprise Manager Release 12.1.0.4, a separate interface consolidates key information and configuration functionality pertaining to SNMP V3 trap notification methods. The SNMP V3 Trap interface helps guide you through the process of creating SNMP notification methods, enabling the OMS to send SNMP traps, and defining user security settings for SNMP trap notifications.

Configuring the OMS to Send SNMP Trap Notifications

Before creating an SNMP Trap notification method, you must enable at least one OMS is your environment to handle SNMP Trap notifications. For SNMP V3, the OMS serves as an SNMP Agent which sends traps to the SNMP Manager that is monitoring all SNMP Agents deployed in the network.

1. From the **Setup** menu, select **Notifications** and then **SNMP V3 Traps**. The Getting Started page displays. This page documents the high-level workflow for configuring Enterprise Manager to send traps to third-party SNMP Managers.

2. Click the **Configuration** tab. The Configuration page displays.
3. In the OMS Configuration region, select the OMS you wish to enable.
4. Check the following for each OMS and make changes, if necessary:
   - OMS requires a port for SNMPv3 traps. Check if the default port can be used by OMS.
   - OMS requires a unique Engine ID for sending traps. By default, it is being generated from the host name and port.
5. Click **Enable**.

Creating/Editing an SNMP V3 Trap Notification Method

Once an OMS has been enabled to send SNMP traps notifications, the next step is to create a notification method than can be used by an incident rule.

1. From the **Setup** menu, select **Notifications** and then **SNMP V3 Traps**. The Getting Started page displays.

> **Note:**
>
> If want to edit an existing Notification Method, select the desired method from the Notification Methods region and click **Edit**.

2. Click the **Configuration** tab. The Configuration page displays.

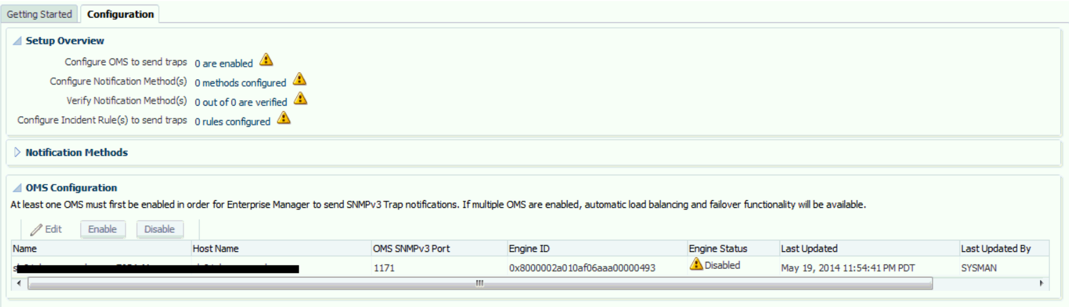

Configure OMS to send traps: 0 are enabled
Configure Notification Method(s): 0 methods configured
Verify Notification Method(s): 0 out of 0 are verified
Configure Incident Rule(s) to send traps: 0 rules configured

At least one OMS must first be enabled in order for Enterprise Manager to send SNMPv3 Trap notifications. If multiple OMS are enabled, automatic load balancing and failover functionality will be available.

Name	Host Name	OMS SNMPv3 Port	Engine ID	Engine Status	Last Updated	Last Updated By
[redacted]	[redacted]	1171	0x8000002a010af06aaa00000493	⚠ Disabled	May 19, 2014 11:54:41 PM PDT	SYSMAN

3. From the Notification Methods region, click **Create**. The SNMPv3 Traps: Create Notification Method page displays.
4. Enter the requisite Notification Method definition parameters. Note: You can enable Repeat Notifications at this point.
5. If you choose to create a new User Security Model entry, from the User Security Model region, ensure the **Create New** option is chosen.
   a. Specify a **Username** that uniquely identifies the credential. SNMP V3 allows multiple usernames to be set in an SNMP Agent as well as SNMP Manager applications.
   b. Select a **Security Level** from the drop-down menu. Available parameters become available depending on the security level. There are three levels from which to choose:

      AuthPriv (Authentication + Privacy:) The sender's identity must be confirmed by the receiver (authentication). SNMP V3 messages are encrypted by the sender and must be decrypted by the receiver (privacy).

      AuthNoPriv (Authentication only): The receiver must authenticate the sender's identity before accepting the message.

NoAuthNoPriv (no security): Neither sender identity confirmation nor message encryption is used.

c. For AuthPriv and AuthNoPriv security levels, choose a the desired **Authentication Protocol**. Two authentication protocols are available:

   Secure Hash Algorithm (SHA)

   Message Digest algorithm (MD5)

   The authentication protocols are used to build the message digest when the message is authenticated.

   Privacy Protocol (used for the AuthPriv security level) is used to encrypt/decrypt messages. USM uses the Data Encryption Standard (DES). The **Privacy Password** is used in conjunction with the Privacy Protocol. the privacy password on both the SNMP Agent and SNMP Manager must match in order for encryption/decryption to succeed.

If you have already have predefined User Security Model entries, choose the **Use Existing** option and select one of the USM entries from the drop-down menu. USM entries are listed by username.

> **Note:**
>
> Ensure that the USM credentials are identical in OMS and the external trap receiver. If they do not match, Enterprise Manager will still send the SNMP trap, but the trap will not be received. If the USM credentials are invalid, Enterprise Manager will still send the SNMP trap, however, the trap will not be received as the incorrect credentials will result in an authentication error at the SNMP receiver. This type of authentication error will not be apparent from the Enterprise Manager console.

6. Once you have entered the requisite Notification Method and USM parameters, click **Save**. The newly created notification method appears in the Notification Method region of the Configuration page.

> **Note:**
>
> Once you have defined the SNMP V3 Trap notification method, you must add it to a rule. See Creating a Rule to Send SNMP Traps for Events to Third Party Systems or Creating a Rule to Send SNMP Traps for Incidents to Third Party Systems for instructions.

Editing a User Security Model Entry

You can add USM entries at any time.

1. From the **Setup** menu, select **Notifications**, and then **SNMP V3 Traps**.
2. Click on the **Configurations** tab.
3. From the User Security Model Entries region, click **Create**. The User Security Model Entries dialog displays.

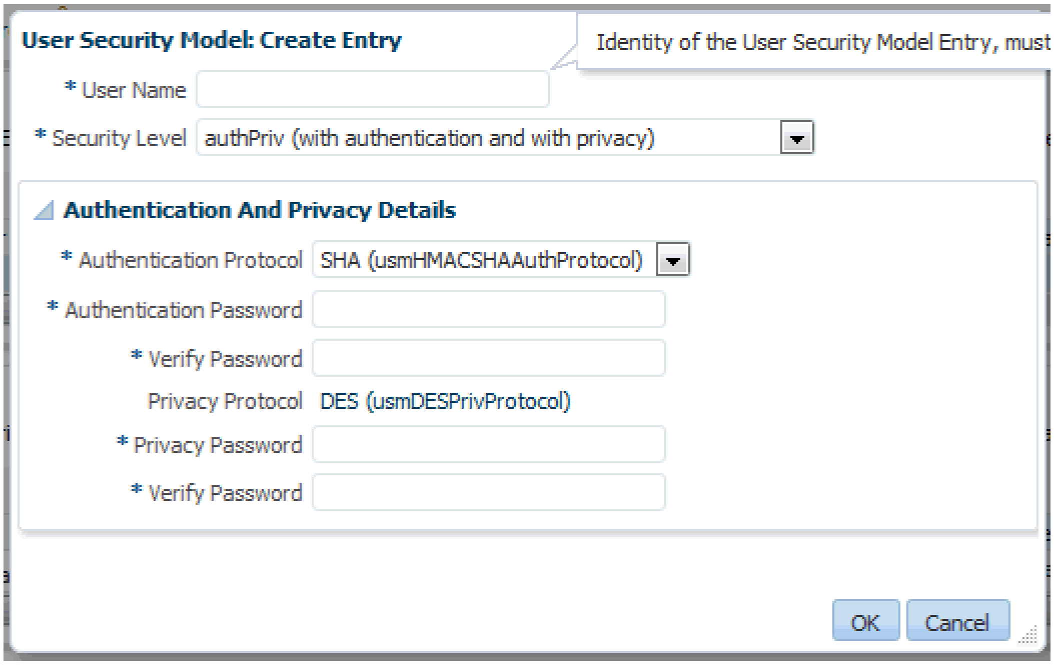

User Security Model: Create Entry

Identity of the User Security Model Entry, must

* User Name

* Security Level: authPriv (with authentication and with privacy)

Authentication And Privacy Details

* Authentication Protocol: SHA (usmHMACSHAAuthProtocol)

* Authentication Password

* Verify Password

Privacy Protocol: DES (usmDESPrivProtocol)

* Privacy Password

* Verify Password

4. Specify a **Username** that uniquely identifies the credential. SNMP V3 allows multiple usernames to be set in an SNMP Agent as well as SNMP Manager applications.

5. Select a **Security Level** from the drop-down menu. Available parameters become available depending on the security level. There are three levels from which to choose:

   AuthPriv (Authentication + Privacy:) The sender's identity must be confirmed by the receiver (authentication). SNMP V3 messages are encrypted by the sender and must be decrypted by the receiver (privacy).

   AuthNoPriv (Authentication only): The receiver must authenticate the sender's identity before accepting the message.

   NoAuthNoPriv (no security): Neither sender identity confirmation nor message encryption is used.

6. For AuthPriv and AuthNoPriv security levels, choose a the desired **Authentication Protocol**. Two authentication protocols are available:

   Secure Hash Algorithm (SHA)

   Message Digest algorithm (MD5)

   The authentication protocols are used to build the message digest when the message is authenticated.

   Privacy Protocol (used for the AuthPriv security level) is used to encrypt/decrypt messages. USM uses the Data Encryption Standard (DES). The **Privacy Password** is used in conjunction with the Privacy Protocol. the privacy password on both the SNMP Agent and SNMP Manager must match in order for encryption/decryption to succeed.

7. Click **OK**.

   The new USM username will appear in the User Security Model Entries table. When creating new SNMP V3 Trap notification methods, the USM username will appear as a selectable option from the **Existing Entries** drop-down menu.

   After editing the USM, you should verify the change via the notification methods that use it.

Viewing Available SNMP V3 Trap Notification Methods

To view available SNMP V3 Trap notification methods:

1. From the **Setup** menu, select **Notifications**, and then **SNMP V3 Traps**.
2. Click on the **Configurations** tab.
3. The Notification Methods region displays existing SNMP V3 Trap notification methods.

Deleting an SNMP V3 Trap Notification Method

To delete available SNMP V3 Trap notification methods:

1. From the **Setup** menu, select **Notifications**, and then **SNMP V3 Traps**.
2. Click on the **Configurations** tab.
3. From the Notification Methods region, select an existing SNMP V3 Trap notification method.
4. Click **Delete**.

Creating an SNMP V1 Trap

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

Log in to Enterprise Manager as a Super Administrator. From the **Setup** menu, select **Notifications** and then select **Scripts and SNMPv1 Traps**.

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. As shown in, the SNMP host will receive your SNMP traps.

Figure 6-2 SNMP Trap Required Information

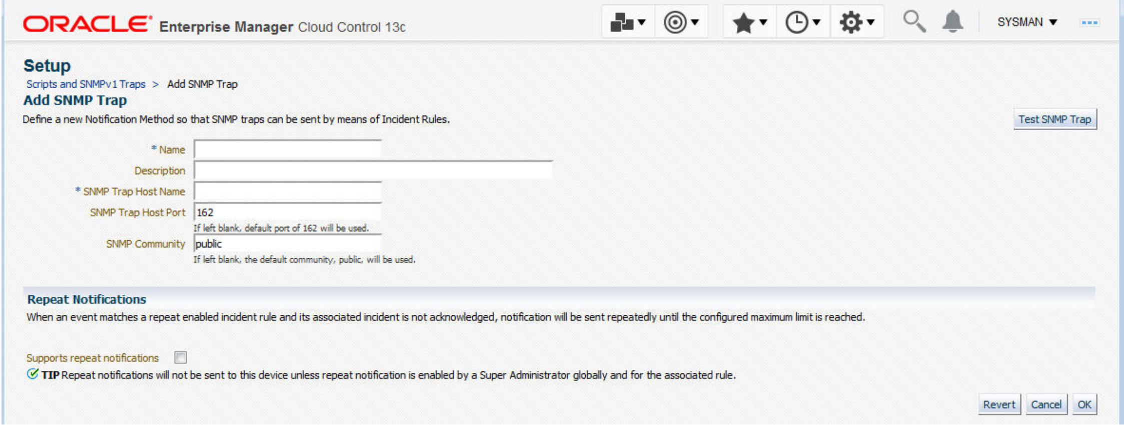

Setup

Scripts and SNMPv1 Traps > Add SNMP Trap

Add SNMP Trap

Define a new Notification Method so that SNMP traps can be sent by means of Incident Rules.

Test SNMP Trap

* Name	
Description	
* SNMP Trap Host Name	
SNMP Trap Host Port	162 If left blank, default port of 162 will be used.
SNMP Community	public If left blank, the default community, public, will be used.

Repeat Notifications

When an event matches a repeat enabled incident rule and its associated incident is not acknowledged, notification will be sent repeatedly until the configured maximum limit is reached.

Supports repeat notifications □

✓ TIP Repeat notifications will not be sent to this device unless repeat notification is enabled by a Super Administrator globally and for the associated rule.

Revert Cancel OK

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

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

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. For instructions on setting up events or incidents rules using SNMP traps, see Creating a Rule to Send SNMP Traps for Events to Third Party Systems or Creating a Rule to Send SNMP Traps for Incidents to Third Party Systems respectively.

Example SNMP Trap Implementation

In this scenario, you want to identify the unique issues from the SNMP traps that are sent. Keep in mind that all events that are related to the same issue are part of the same event sequence. Each event sequence has a unique identification number.

An event sequence is a sequence of related events that represent the life of a specific issue from the time it is detected and an event is raised to the time it is fixed and a corresponding *clear* event is generated. For example, a warning metric alert event is raised when the CPU utilization of a host crosses 80%. This starts the event sequence representing the issue *CPU Utilization of the host is beyond normal level*. Another critical event is raised for the same issue when the CPU utilization goes above 90% and the event is added to the same event sequence. After a period of time, the CPU utilization returns to a normal level and a *clear* event is raised. At this point, the issue is resolved and the event sequence is closed.

The SNMP trap sent for this scenario is shown in Example 6-15. Each piece of information is sent as a variable embedded in the SNMP Trap.

This following example illustrates how OIDs are used during the lifecycle of an event. Here, for one event (while the event is open), the event sequence OID remains the same even though the event severity changes.

The OID for the event sequence is:

```
1.3.6.1.4.1.111.15.3.1.1.42.1: C77AE9E578F00773E040F00A6D242F90
```

The OID for the event severity code is:

```
1.3.6.1.4.1.111.15.3.1.1.6.1: CRITICAL
```

When the event clears, these OIDs show the same event sequence with a different severity code:

The OID for the event sequence is:

```
1.3.6.1.4.1.111.15.3.1.1.42.1: C77AE9E578F00773E040F00A6D242F90
```

The OID for the event severity code is:

```
1.3.6.1.4.1.111.15.3.1.1.6.1: CLEAR
```

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

Example 6-15 SNMP Trap

```
**************V1 TRAP***[1]*****************
Community : public
Enterprise :1.3.6.1.4.1.111.15.2
Generic :6
Specific :3
TimeStamp :67809
Agent adress :10.240.36.109
```

```
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: CPU Utilization is 92.658%, crossed warning (80) or
critical (90) threshold.
1.3.6.1.4.1.111.15.3.1.1.4.1: https://sampleserver.oracle.com:5416/em/redirect?
pageType=sdk-core-event-console-detailEvent&issueID=C77AE9E578F00773E040F00A6D242F90
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
1.3.6.1.4.1.111.15.3.1.1.8.1:
1.3.6.1.4.1.111.15.3.1.1.9.1:
1.3.6.1.4.1.111.15.3.1.1.10.1: Aug 17, 2012 3:26:36 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:cpuUtil
1.3.6.1.4.1.111.15.3.1.1.15.1: 281
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: sampleserver.oracle.com
1.3.6.1.4.1.111.15.3.1.1.22.1: https://sampleserver.oracle.com:5416/em/redirect?
pageType=TARGET_HOMEPAGE&targetName=sampleserver.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: sampleserver.oracle.com
1.3.6.1.4.1.111.15.3.1.1.25.1: SYSMAN
1.3.6.1.4.1.111.15.3.1.1.26.1:
1.3.6.1.4.1.111.15.3.1.1.27.1: 5.8.0.0.0
1.3.6.1.4.1.111.15.3.1.1.28.1: Operating System=Linux, Platform=x86_64,
1.3.6.1.4.1.111.15.3.1.1.29.1:
1.3.6.1.4.1.111.15.3.1.1.30.1:
1.3.6.1.4.1.111.15.3.1.1.31.1:
1.3.6.1.4.1.111.15.3.1.1.32.1:
1.3.6.1.4.1.111.15.3.1.1.33.1:
1.3.6.1.4.1.111.15.3.1.1.34.1:
1.3.6.1.4.1.111.15.3.1.1.35.1:
1.3.6.1.4.1.111.15.3.1.1.36.1:
1.3.6.1.4.1.111.15.3.1.1.37.1:
1.3.6.1.4.1.111.15.3.1.1.38.1:
1.3.6.1.4.1.111.15.3.1.1.39.1: SnmpNotifRuleset
1.3.6.1.4.1.111.15.3.1.1.40.1: SnmpNotifRuleset,SnmpNotifEvent
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: C77AE9E578F00773E040F00A6D242F90
1.3.6.1.4.1.111.15.3.1.1.43.1:
1.3.6.1.4.1.111.15.3.1.1.44.1:
1.3.6.1.4.1.111.15.3.1.1.45.1:
1.3.6.1.4.1.111.15.3.1.1.46.1: CPU Utilization is 92.658%, crossed warning (80) or
critical (90) threshold., Incident created by rule (Name = Incident management Ruleset
for all targets, Incident creation Rule for metric alerts.; Owner = <SYSTEM>).
1.3.6.1.4.1.111.15.3.1.1.61.1: Metric GUID=0C71A1AFAC2D7199013837DA35522C08
1.3.6.1.4.1.111.15.3.1.1.62.1: Severity GUID=C77AE9E578EC0773E040F00A6D242F90
1.3.6.1.4.1.111.15.3.1.1.63.1: Cycle GUID=C77AE9E578EC0773E040F00A6D242F90
1.3.6.1.4.1.111.15.3.1.1.64.1: Collection 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=CPU 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: Metric value=92.658
1.3.6.1.4.1.111.15.3.1.1.69.1: Key Value=
1.3.6.1.4.1.111.15.3.1.1.70.1:
1.3.6.1.4.1.111.15.3.1.1.71.1:
1.3.6.1.4.1.111.15.3.1.1.72.1:
```

```
1.3.6.1.4.1.111.15.3.1.1.73.1:
1.3.6.1.4.1.111.15.3.1.1.74.1:
1.3.6.1.4.1.111.15.3.1.1.75.1:
1.3.6.1.4.1.111.15.3.1.1.76.1:
1.3.6.1.4.1.111.15.3.1.1.77.1:
1.3.6.1.4.1.111.15.3.1.1.78.1:
1.3.6.1.4.1.111.15.3.1.1.79.1:
1.3.6.1.4.1.111.15.3.1.1.80.1:
1.3.6.1.4.1.111.15.3.1.1.81.1:
1.3.6.1.4.1.111.15.3.1.1.82.1:
1.3.6.1.4.1.111.15.3.1.1.83.1:
1.3.6.1.4.1.111.15.3.1.1.84.1: Number of keys=0
1.3.6.1.4.1.111.15.3.1.1.85.1:
**************END V1 TRAP******************
```

SNMP Traps: Moving from Previous Enterprise Manager Releases to 12c and Greater

> **Note:**
>
> When you upgrade from a pre-Enterprise Manager 12c release to 12c and greater, SNMP advanced notification methods defined using previous versions of Enterprise Manager (pre-12c) will continue to function without modification.

For Enterprise Manager 11g and earlier, there were two types of SNMP traps:

- Alerts
- Job Status

Beginning with Enterprise Manager 12c there is now a single, comprehensive SNMP trap type that covers all available event types such as metric alerts, target availability, compliance standard violations, or job status changes. For more information about pre-12*c* to 12*c* SNMP trap mappings, see SNMP Trap Mappings . Traps will conform to the older Enterprise Manager MIB definition. Hence, pre-Enterprise Manager 12c traps will continue to be sent. See SNMP Trap Mappings for more information.

Also, for Enterprise Manager 12c, size of SNMP trap has increased in order to accommodate all event types and provide more comprehensive information. By default, the maximum SNMP packet size is 5120 bytes. If the third party system has a limit in the size of SNMP trap it can receive, you can change the default size of SNMP trap that Enterprise Manager sends. To change the default packet size, set this *emoms* `oracle.sysman.core.notification.snmp_packet_length` parameter, and then bounce the OMS.

> **Note:**
>
> When limiting the SNMP trap packet size, Oracle recommends not setting the oracle.sysman.core.notification.snmp_packet_length parameter any lower than 3072 bytes (3K).

The Enterprise Manager 12c MIB includes all pre-Enterprise Manager 12c MIB definitions. Hence, if you have an Enterprise Manager 12c MIB in your third party system, you can receive SNMP traps from both pre-Enterprise Manager 12c as well as Enterprise Manager 12c sites. For detailed information on version mapping, see SNMP Trap Mappings .

Sending Notifications Using Webhooks and Slack

Starting with Enterprise Manager 13c Release 5 RU 15, using the EM Generic Webhook, you can send event and incident notifications to any application that accepts webhook requests and can parse payload specific content. The webhook payload format depends on the type of the event or incident. Users will have to provide the structure for the payload to map EM to the external software.

To access the **Webhooks and Slack** page navigate to the **Setup** drop-down menu, hover over **Notifications** and locate the **Webhooks and Slack** option.

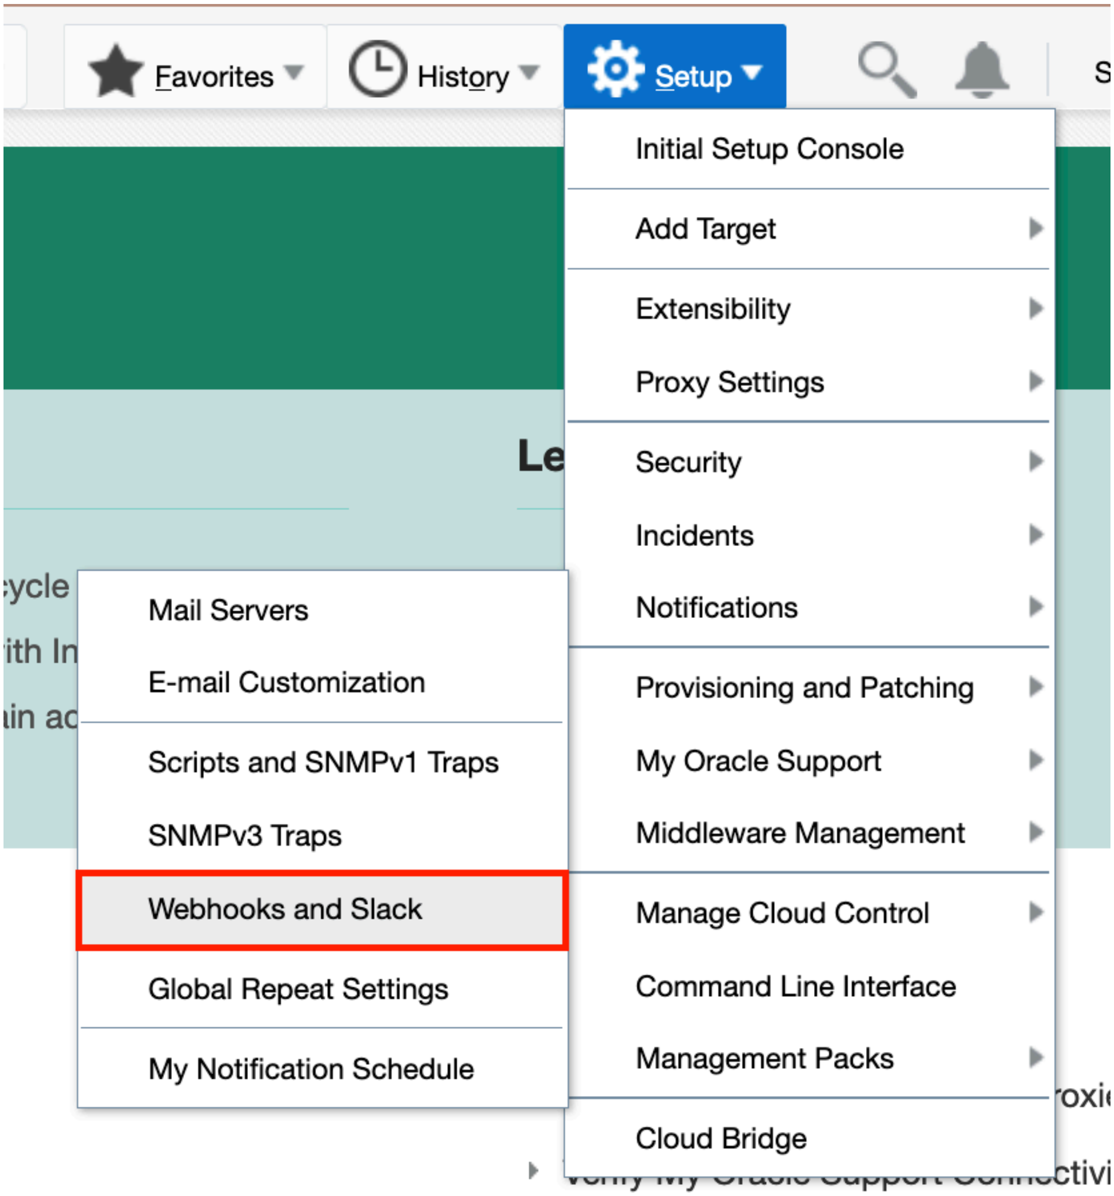

To use Webhooks or Slack as a notification method, you must enable it through Incident Rules. When on the **Add Actions** page for an incident or an event rule, locate the **Advanced Notifications** section to enable your preferred notification method (e.g., Slack).

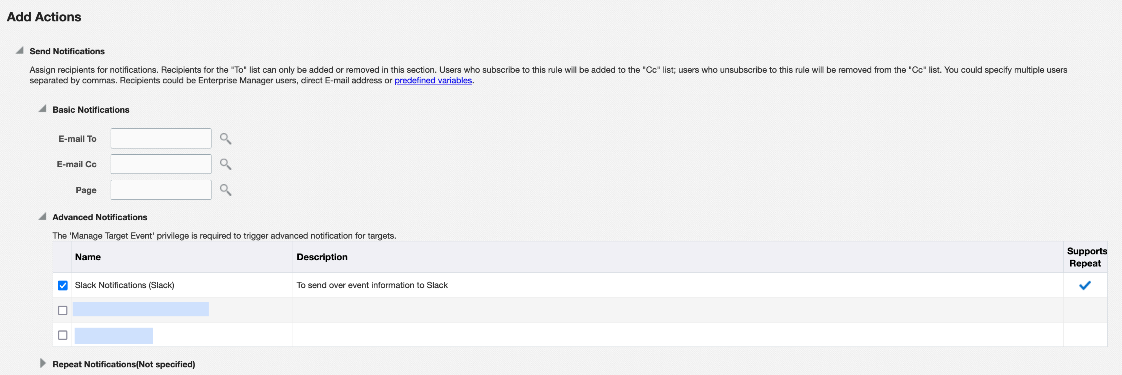

Add Actions

Send Notifications

Assign recipients for notifications. Recipients for the "To" list can only be added or removed in this section. Users who subscribe to this rule will be added to the "Cc" list; users who unsubscribe to this rule will be removed from the "Cc" list. You could specify multiple users separated by commas. Recipients could be Enterprise Manager users, direct E-mail address or predefined variables.

Basic Notifications

E-mail To

E-mail Cc

Page

Advanced Notifications

The 'Manage Target Event' privilege is required to trigger advanced notification for targets.

	Name	Description	Supports Repeat
☑	Slack Notifications (Slack)	To send over event information to Slack	✓
☐			
☐			

Repeat Notifications(Not specified)

Webhooks

Using a Webhook, you can configure EM to send events or incidents data directly into an application of your choice.

Add a Webhook Notification Method:

1. Get the incoming Webhook URL for your application.
2. Click **Add Notification Method** in the **Webhooks and Slack** page.
3. Select from the drop-down list to what you want the notification method to apply (**Incidents** or **Events**), and click **Webhook**:
   - **Incidents**

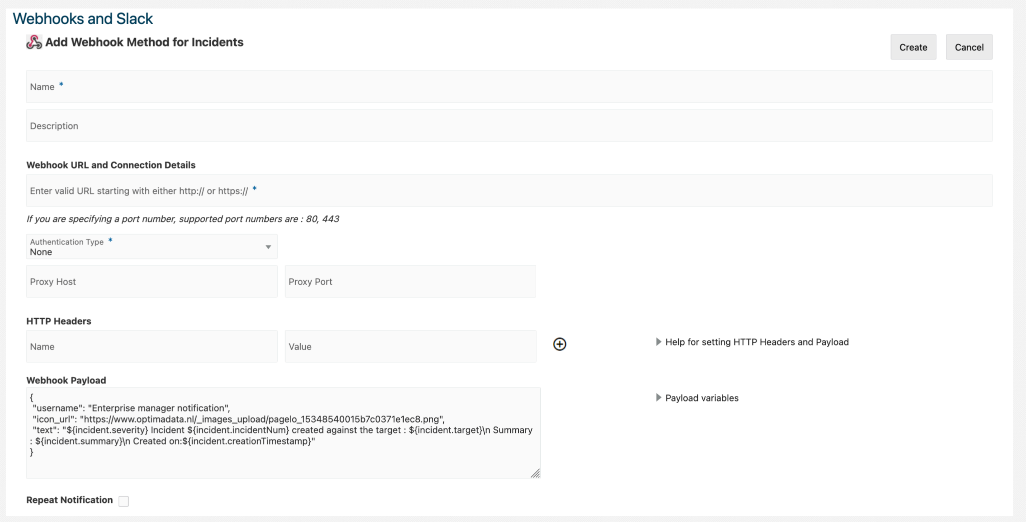

Webhooks and Slack

Add Webhook Method for Incidents

Name *

Description

Webhook URL and Connection Details

Enter valid URL starting with either http:// or https:// *

If you are specifying a port number, supported port numbers are : 80, 443

Authentication Type *
None

Proxy Host

Proxy Port

HTTP Headers

Name

Value

Webhook Payload

```
{
 "username": "Enterprise manager notification",
 "icon_url": "https://www.optimadata.nl/_images_upload/pagelo_15348540015b7c0371e1ec8.png",
 "text": "${incident.severity} Incident ${incident.incidentNum} created against the target : ${incident.target}\n Summary
: ${incident.summary}\n Created on:${incident.creationTimestamp}"
}
```

Repeat Notification ☐

   - **Events**

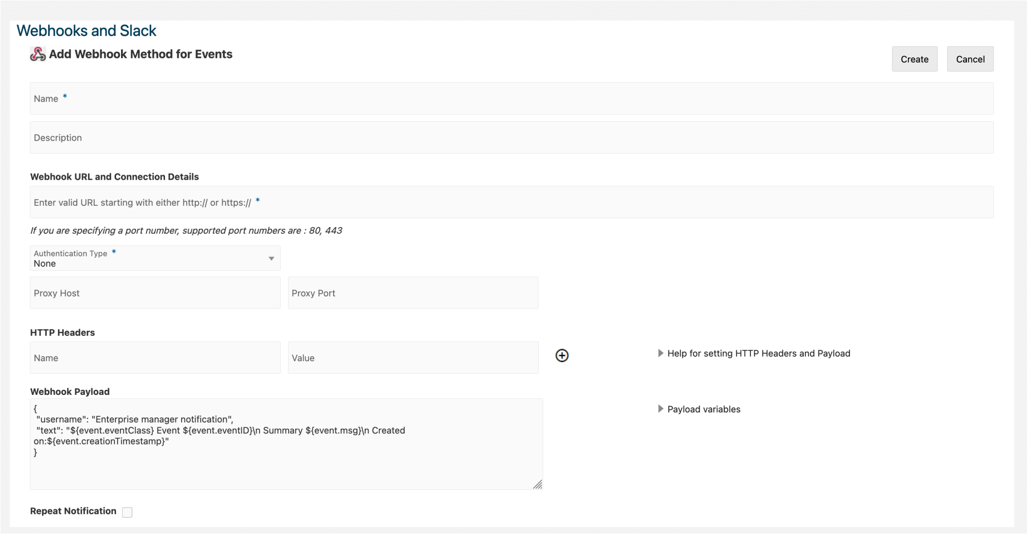
Webhooks and Slack
Add Webhook Method for Events
Create
Cancel
Name *
Description
Webhook URL and Connection Details
Enter valid URL starting with either http:// or https:// *
Proxy Host
Proxy Port
HTTP Headers
Name
Value
Help for setting HTTP Headers and Payload
Webhook Payload
Payload variables
Repeat Notification □

If you are specifying a port number, supported port numbers are : 80, 443

Authentication Type * None

```
{
 "username": "Enterprise manager notification",
 "text": "${event.eventClass} Event ${event.eventID}\n Summary ${event.msg}\n Created
on:${event.creationTimestamp}"
}
```

4. Enter a name and description for your notification method.
5. Enter the application’s incoming Webhook URL, and optionally other connection details.
6. Optionally, modify the **Webhook Payload**. You can see the available payload variables by opening the **Payload variable** list on the right.
7. Click **Create**.

Test a Webhooks notification method

1. From the **Webhooks and Slack Page**, select the icon under the **Actions** menu for the notification method you want to test, and select **Test Notification Method**.
2. Check the name of the notification method is correct in the **Test Notification Method** dialog box and select **Test**.
3. Check the status of the test is **Succeeded**.

Assign a Webhooks notification method to an incident rule

1. From the **Setup** menu, navigate to **Incidents**, and select **Incidents Rules**.
2. **Create a rule set** or select an existing rule set, and select **Edit**.
3. In the **Rules** section, select a rule and select **Edit**.
4. Review and edit as needed the incidents to which the rules apply, and select **Next**.
5. Select an action and select **Edit**.
6. From the **Send Notification** section, navigate to **Advance Notifications**, and select the Slack notification method to be assigned. Select **Continue**.
7. Select **Next**.
8. Provide a **Name** and a **Description** for the rule set, and select **Next**.
9. Review the rule set, edit as needed, and select **Continue**.
10. Select **Save**.

Slack

Using a Slack Webhook, you can configure EM to send events or incidents data directly into a Slack channel.

How to set up Slack Notifications - YouTube Video

Add a Slack notification method

1. Get the Slack channel’s incoming Webhook URL.

> **Note:**
>
> You must be the workspace owner to get the incoming webhook URL.

2. Click **Add Notification Method** in the **Webhooks and Slack** page.
3. Select from the drop-down list to what you want the notification method to apply (**Incidents** or **Events**), and click **Slack**:
   - **Incidents**

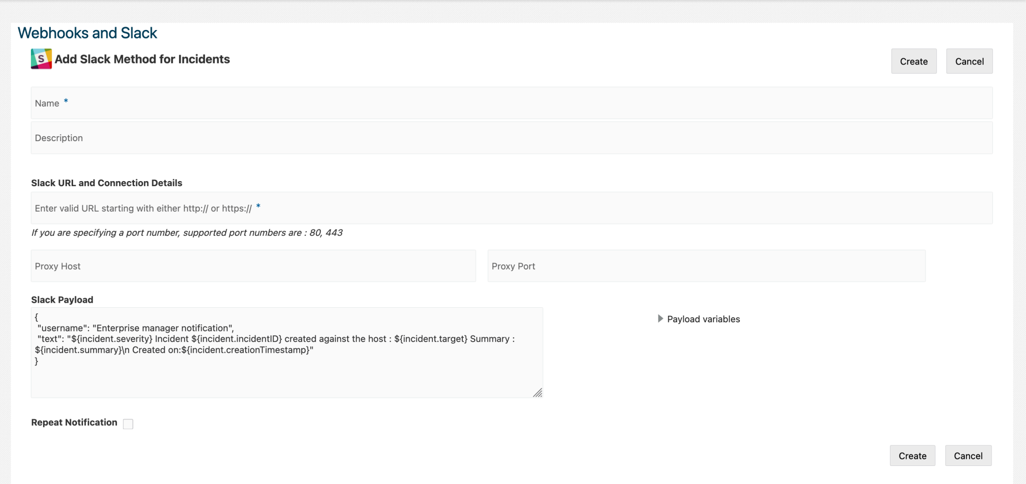

If you are specifying a port number, supported port numbers are : 80, 443

```
{
 "username": "Enterprise manager notification",
 "text": "${incident.severity} Incident ${incident.incidentID} created against the host : ${incident.target} Summary :
${incident.summary}\n Created on:${incident.creationTimestamp}"
}
```

   - **Events**

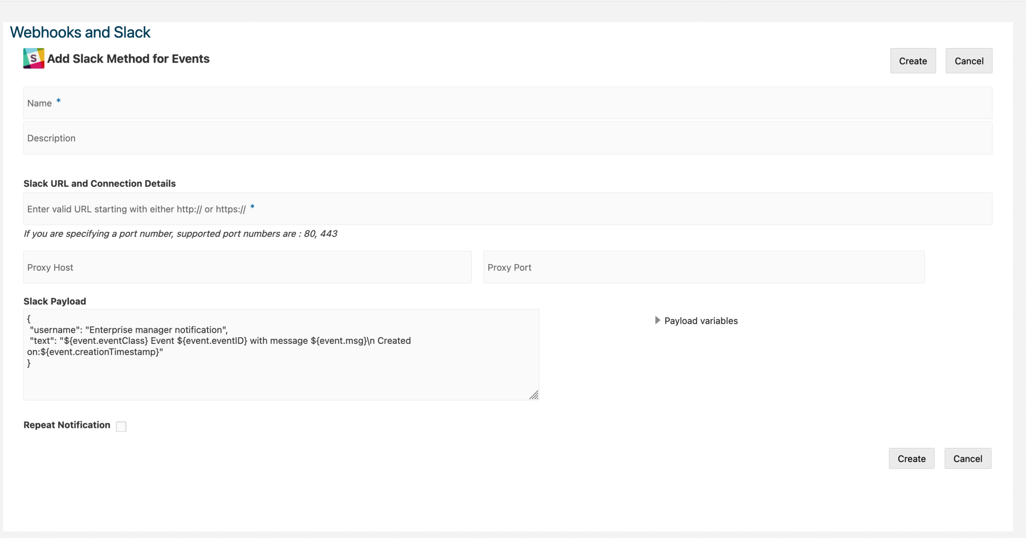

Webhooks and Slack

Add Slack Method for Events

Name *

Description

Slack URL and Connection Details

Enter valid URL starting with either http:// or https:// *

If you are specifying a port number, supported port numbers are : 80, 443

Proxy Host

Proxy Port

Slack Payload

```
{
 "username": "Enterprise manager notification",
 "text": "${event.eventClass} Event ${event.eventID} with message ${event.msg}\n Created
on:${event.creationTimestamp}"
}
```

Repeat Notification □

4. Enter a name and description for your notification method.
5. Enter the Slack channel's incoming Webhook URL, and optionally other connection details.

> **Note:**
>
> Make sure the Slack channel's incoming Webhook URL is in the `https://hooks.slack.com/services/...` format.

6. Optionally, modify the **Slack Payload**. You can see the available payload variables by opening the **Payload variable** list on the right.
7. Click **Create**.

Test a Slack notification method

1. From the **Webhooks and Slack Page**, select the icon under the **Actions** menu for the notification method you want to test, and select **Test Notification Method**.
2. Check the name of the notification method is correct in the **Test Notification Method** dialog box and select **Test**.
3. Check the status of the test is **Succeeded**, and navigate to the Slack channel used to see the sample test incident notification.

Assign a Slack notification method to an incident rule

1. From the **Setup** menu, navigate to **Incidents**, and select **Incidents Rules**.
2. **Create a rule set** or select an existing rule set, and select **Edit**.
3. In the **Rules** section, select a rule and select **Edit**.
4. Review and edit as needed the incidents to which the rules apply, and select **Next**.
5. Select an action and select **Edit**.
6. From the **Send Notification** section, navigate to **Advance Notifications**, and select the Slack notification method to be assigned. Select **Continue**.
7. Select **Next**.

8. Provide a **Name** and a **Description** for the rule set, and select **Next**.
9. Review the rule set, edit as needed, and select **Continue**.
10. Select **Save**.

Management Information Base (MIB)

Enterprise Manager 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.

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.

MIB Definition

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

OMS_HOME/network/doc/omstrap.v1

> **Note:**
>
> The omstrap.v1 file is compatible with both SNMP V1 and SNMP V3.

The file *omstrap.v1* is the OMS MIB.

For more information, see Interpreting Variables of the Enterprise Manager MIB.

A hardcopy version of omstrap.v1 can be found in Enterprise Manager MIB Definition.

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

For Enterprise Manager 12c, SNMP traps are delivered for event notifications. Starting in Enterprise Manager 13.5 RU16, SNMP trap notifications are supported for both events and incidents. SNMP trap notifications are not supported for 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.

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 1, is not included in these MIB variable descriptions. Since Oracle has implemented all MIB variables as CURRENT, this value does not vary.

Variable Name

Syntax
Maps to the SYNTAX element of SNMP MIB definition, Version 1.

Max-Access
Maps to the MAX-ACCESS element of SNMP MIB definition, Version 1.

Status
Maps to the STATUS element of SNMP MIB definition, Version 1.

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 1.

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.

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.

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 6-13 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

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 6-45 for details.

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

Table 6-14 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

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 Sending Notifications Using OS Commands and Scripts and Sending Notifications Using PL/SQL Procedures.

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 6-15 lists all corrective action status change attributes that can be passed:

Table 6-15 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 6-16 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

Table 6-16 (Cont.) Job Status Codes

Name	Datatype	Value
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 6-16 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;
/
```

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 6-17 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.

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 email format is used.

Figure 6-3   Host Target Properties

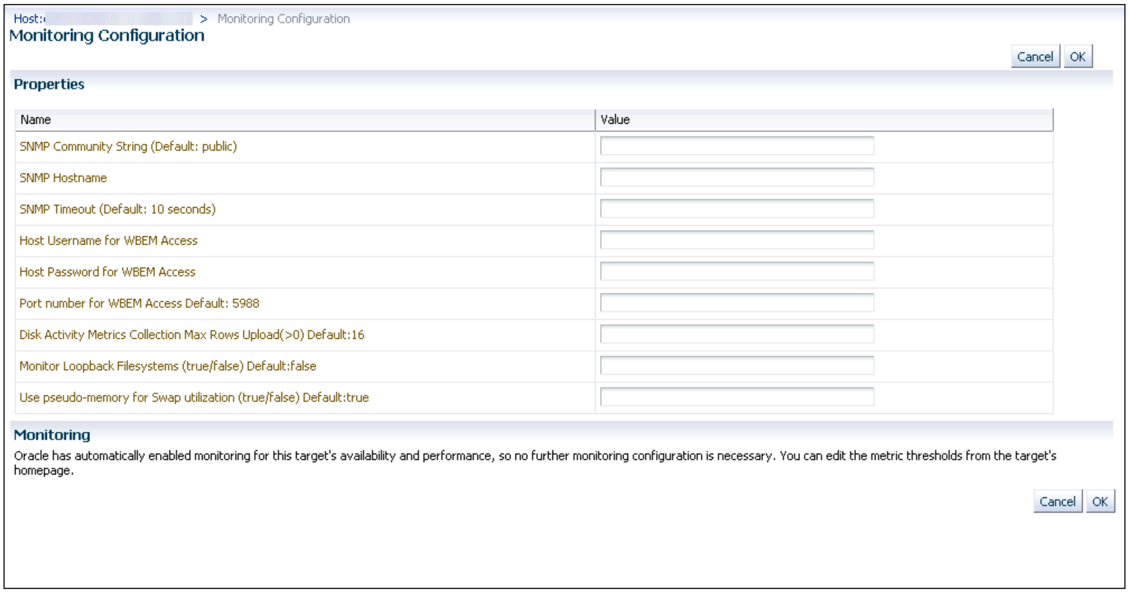

Host:[redacted] > Monitoring Configuration

Monitoring Configuration

Properties

Name	Value
SNMP Community String (Default: public)	
SNMP Hostname	
SNMP Timeout (Default: 10 seconds)	
Host Username for WBEM Access	
Host Password for WBEM Access	
Port number for WBEM Access Default: 5988	
Disk Activity Metrics Collection Max Rows Upload(>0) Default:16	
Monitor Loopback Filesystems (true/false) Default:false	
Use pseudo-memory for Swap utilization (true/false) Default:true	

Monitoring

Oracle has automatically enabled monitoring for this target's availability and performance, so no further monitoring configuration is necessary. You can edit the metric thresholds from the target's homepage.

Notification Reference

This section contains the following reference material:

- EMOMS Properties
- Passing Event, Incident, Problem Information to an OS Command or Script
- Passing Information to a PL/SQL Procedure
- Passing Information to Webhooks and Slack
- Troubleshooting Notifications

EMOMS Properties

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

Table 6-18   emoms Properties for Notifications

Property Name	Default Value	Description
oracle.sysman.core.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.
oracle.sysman.core.notification.cmds_per_minute	100	OS Command delivery limits per minute. The Notification system uses this value to throttle number of OS Command delivery per minutes.

Table 6-18 (Cont.) emoms Properties for Notifications

Property Name	Default Value	Description
oracle.sysman.core.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.
oracle.sysman.core.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.
oracle.sysman.core.notification.traps_per_minute	250	SNMP delivery limits per minute. The Notification system uses this value to control the number of SNMP trap per minutes.
oracle.sysman.core.notification.locale.plsql	OMS Locale	This property specifies the Locale delivered by advanced PL/SQL notification. The customer can define this property to overwrite the default Locale where the OMS is installed. Valid Locales: • en (English) • de (German) • es (Spanish) • fr (French) • it (Italian) • ja (Japanese) • ko (Korean) • pt_br (Portuguese, Brazilian) • zh_cn (Chinese, simplified) • zh_tw (Chinese, traditional)
oracle.sysman.core.notification.locale.email	OMS Locale	This property specifies the Locale delivered by Email. Customer can define this property to overwrite the default Locale where the OMS is installed. Valid Locales: • en (English) • de (German) • es (Spanish) • fr (French) • it (Italian) • ja (Japanese) • ko (Korean) • pt_br (Portuguese, Brazilian) • zh_cn (Chinese, simplified) • zh_tw (Chinese, traditional)

Table 6-18 (Cont.) emoms Properties for Notifications

Property Name	Default Value	Description
oracle.sysman.core.notification.locale.oscmd	OMS Locale	This property specifies the Locale delivered by OS Command. Customer can define this property to overwrite the default Locale where the OMS is installed. Valid Locales: • en (English) • de (German) • es (Spanish) • fr (French) • it (Italian) • ja (Japanese) • ko (Korean) • pt_br (Portuguese, Brazilian) • zh_cn (Chinese, simplified) • zh_tw (Chinese, traditional)
oracle.sysman.core.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 the OMS is installed. Valid Locales: • en (English) • de (German) • es (Spanish) • fr (French) • it (Italian) • ja (Japanese) • ko (Korean) • pt_br (Portuguese, Brazilian) • zh_cn (Chinese, simplified) • zh_tw (Chinese, traditional)
oracle.sysman.core.notification.oscmd.max_env_var_length	512	The maximum length of OS Common environment variable value.
oracle.sysman.core.notification.snmp.max_oid_length	2560	The maximum length of SNMP OID value.
oracle.sysman.core.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. Setting the value higher will use more system resources, but will deliver more notifications.
oracle.sysman.core.notification.max_delivery_threads	24	The maximum number of active threads in the thread pool when the system is in the high activities. This value should greater than em.notification.min_delivery_threads. Setting the value higher will use more system resources and deliver more notifications.

Table 6-18 (Cont.) emoms Properties for Notifications

Property Name	Default Value	Description
oracle.sysman.core.notification.short_format_length	>=1 (155)	The size limit of the total number of characters in short email format. The customer should modify this property value to fit their email or pager limit content size. The email subject is restricted to a maximum of 80 characters for short email format notifications.
oracle.sysman.core.notification.snmp_packet_length	<=1 (5120)	The maximum size of SNMP Protocol Data unit.
oracle.sysman.core.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).
oracle.sysman.core.notification.emails_per_connection	>=1 (20)	The maximum number of emails delivered to same email gateway before switching to the next available email gateway (assumes customers have configured multiple email gateways). This property is used for email gateway load balance.
oracle.sysman.core.notification.short_format	both, subject, body (both)	Use short format on both subject and body, subject only, or body only..
oracle.sysman.core.notification.send_prior_status_after_agent_unreachable_clears	True	By default , a notification is sent indicating a target's status whenever the monitoring Agent comes out of *unreachable* status, even if the target's status has not changed. Use this emoms property to enable (True)/disable (False) the duplicate target status notification. To disable duplicate target status notifications, set this property to *False*: 1. emctl set property oracle.sysman.core.notification.send_prior_status_after_agent_unreachable_clears -value false 2. Restart the OMS. To enable duplicate target status notifications, set the property to *True*. 1. emctl set property oracle.sysman.core.notification.send_prior_status_after_agent_unreachable_clears -value true 2. Restart the OMS.

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 `emctl get/set/delete/list` property command.

Note:

The following commands require an OMS restart in order for the changes to take place.

Get Property Command

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

Set Property Command

```
emctl set  property -name oracle.sysman.core.notification.short_format_length -value 155
```

Emoms Properties Entries for a Short Email Format

```
emctl set  property -name oracle.sysman.core.notification.short_format_length -value 155
emctl set property -name oracle.sysman.core.notification.short_format -value both
```

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 6-19 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 6-20 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

Table 6-21 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 6-21 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 6-22 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

Table 6-22 (Cont.) Event Notification-Specific Environment Variables

Environment Variable	Description
SEQUENCE_ID	The unique event sequence identifier. An event sequence may consist of one or more events. All events in this sequence have the same event sequence ID.
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
RCA_DETAILS	If RCA is associated with this events.
CURRENT_OCCURRENCE_COUNT	Total number of occurrences of the event in the current collection period. This attribute only applies to de-duplicated events.
CURRENT_FIRST_OCCUR_DATE	Time stamp when the event first occurred in the current collection period. This attribute only applies to de-duplicated events.
CURRENT_LAST_OCCUR_DATE_DESC	Time stamp when the e vent last occurred in the current collection period. This attribute only applies to de-duplicated events.

Table 6-23 lists the environment variables for the incident associated with an event. They are populated when the event is associated with an incident.

Table 6-23 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.

Table 6-24 lists the common environment variables related to the Source Object. They are populated when $SOURCE_OBJ_TYPE is not TARGET.

Table 6-24 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.

Table 6-24 (Cont.) Source Object-Related Environment Variables

Environment Variable	Description
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.

Table 6-25 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 6-25 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.
TARGET_LIFECYCLE_STATUS	Life Cycle Status of the target. Possible values: Production, Mission Critical, Stage, Test, and Development. It is null if not defined.
TARGET_VERSION	Target Version of the target

Environment Variables Specific to Event Types

Events are classified into multiple types. For example, the mertc_alert event type is used for modeling metric alerts. You can use SQL queries to list the event types in your deployment as well as their event-specific payload. The following SQL example can be used to list all internal event type names that are registered in Enterprise Manager.

```
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;
```

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

Table 6-26 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.
SEVERITY_GUID	The GUID of the severity record associated with this metric alert.
CYCLE_GUID	A unique identifier for a metric alert cycle, which starts from the time the metric alert is initially generated until the time it is clear.
VALUE	Value of the metric when the event triggered.

Table 6-27 Environment variables specific to Target Availability Event Type

Environment Variable	Description
AVAIL_SEVERITY	The transition severity that resulted in the status of the target to change to the current availability status. Possible Values for AVAIL_SEVERITY • 15 (Target Up) • 25 (Target Down) • 115 (Agent Unreachable, Cleared) • 125 (Agent Unreachable) • 215 (Blackout Ended) • 225 (Blackout Started) • 315 (Collection Error Cleared) • 325 (Collection Error) • 425 (No Beacons Available) • 515 (Status Unknown)
AVAIL_SUB_STATE	The substatus of a target for the current status.
CYCLE_GUID	A unique identifier for a metric alert cycle, which starts from the time the metric alert is initially generated until the time it is clear.
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.

Table 6-27 (Cont.) Environment variables specific to Target Availability Event Type

Environment Variable	Description
TARGET_STATUS_NLS	The translated current availability status of the target.

Table 6-28 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. For possible values, see Table 6-16.
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 Enterprise Manager console:

1. From the **Setup** menu, select **Notifications**, then select **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 6-29 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, High, Medium, Low, None
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.
ADR_INCIDENT_ID	Automatic Diagnostic Reposito ry (ADR) Incident ID: A unique numeric identifier for the ADR Incident. An ADR I ncident is an occurrence of a Problem.
ADR_IMPACT	Impact of the Automatic Diagnostic Repository (ADR) Incident.
ADR_ECID	Execution Context ID (ECID) associated with the associated Automatic Diagnostic Repository (ADR) incident. An ECID i s a globally unique identifier used to tag and track a single call through the Oracle software stack. It is used to correlate problems that could occur across multiple tiers of the stack.

Table 6-29 (Cont.) Incident-Specific Environment Variables

Environment Variable	Description
ASSOC_PROBLEM_KEY	Problem key associated with the Automatic Diagnostic Repository (ADR) incident. Problems are critical error s in an Oracle product. The Problem key is a text string that describes the prob lem. It includes an error code and in some cases, other error-specific values.

Table 6-30 Environment Variables Specific to Compressed Incident Notifications

Variable Name	Description
MESSAGE_URL	Clicking on this link will navigate to compressed incident details in Enterprise Manager.
ASSOC_EVENT_COUNT	Total compressed events in the incident
INC_COMPR_EVT_1_MSG INC_COMPR_EVT_2_MSG INC_COMPR_EVT_3_MSG INC_COMPR_EVT_4_MSG INC_COMPR_EVT_5_MSG	Event Message of the five latest events that are part of the compressed incident. Latest event message is INC_COMPR_EVT_5_MSG.
INC_COMPR_EVT_1_TYPE INC_COMPR_EVT_2_TYPE INC_COMPR_EVT_3_TYPE INC_COMPR_EVT_4_TYPE INC_COMPR_EVT_5_TYPE	Event Type of the five latest events that are part of the compressed incident. Latest event type is INC_COMPR_EVT_5_TYPE.
INC_COMPR_EVT_1_SEV INC_COMPR_EVT_2_SEV INC_COMPR_EVT_3_SEV INC_COMPR_EVT_4_SEV INC_COMPR_EVT_5_SEV	Event Severity of the five latest events that are part of the compressed incident. SEVERITY: WARNING, CRITICAL, FATAL, CLEAR Latest event severity is INC_COMPR_EVT_5_SEV.
INC_COMPR_EVT_1_REP_TIME INC_COMPR_EVT_2_REP_TIME INC_COMPR_EVT_3_REP_TIME INC_COMPR_EVT_4_REP_TIME INC_COMPR_EVT_5_REP_TIME	Event Creation Date of the five latest events that are part of the compressed incident. Latest event creation date is INC_COMPR_EVT_5_REP_TIME.
INC_COMPR_EVT_1_TGT_NAME INC_COMPR_EVT_2_TGT_NAME INC_COMPR_EVT_3_TGT_NAME INC_COMPR_EVT_4_TGT_NAME INC_COMPR_EVT_5_TGT_NAME	The Target Name in the five latest events that are part of the compressed incident. Latest target name is INC_COMPR_EVT_5_TGT_NAME.
INC_COMPR_EVT_1_TGT_TYPE INC_COMPR_EVT_2_TGT_TYPE INC_COMPR_EVT_3_TGT_TYPE INC_COMPR_EVT_4_TGT_TYPE INC_COMPR_EVT_5_TGT_TYPE	The Target Type in the five latest events that are part of the compressed incident. Latest target type is INC_COMPR_EVT_5_TGT_TYPE.

Table 6-31 lists the associated problem's environment variables, when the incident is associated with a problem.

Table 6-31 Associated Problem Environment Variables for Incidents

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 6-32 Problem-Specific Environment Variables

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

Table 6-32   (Cont.) Problem-Specific Environment Variables

Environment Variable	Description
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.
BUG_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.

Environment Variables Related to Event Sources

The number of event sources is set by the EVENT_SOURCE_COUNT environment variable. Using the EVENT_SOURCE_COUNT information, a 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 6-33   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.

Table 6-34 lists the environment variables for a target of *x*th Event Source.

Table 6-34 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.

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 three 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 three 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 6-35 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 6-36 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 6-37 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.
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 6-38 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 6-39 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.
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*)

Contains the incident specific attributes, associated problem and ticket information.

Table 6-40 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.

gc$notif_problem_payload (*payload specific to problems*)

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

Table 6-41 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*)

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

Table 6-42 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.

Table 6-42 (Cont.) Payload Common to Incidents and Problems

Attribute	Datatype	Additional Information
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.
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 incident and problem payloads*)

Represents the associated incident summary in the event payload, or associated problem summary in the incident payload, respectively.

Table 6-43 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.
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. The notification system delivers up to 10 categories.

gc$notif_event_context_array

gc$notif_event_context_array provides information about the additional diagnostic data that was captured at event detection time. Note that notification system delivers up to 200 elements from the captured event context. Each element of this array is of the type *gc$notif_event_context*.

gc$notif_event_context: This object represents the detail of event context data which is additional contextual information captured by the source system at the time of event generation that may have diagnostic value. The context for an event should consist of a set of keys and values along with data type (Number or String only).

Table 6-44 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

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 6-45 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 Enterprise Manager. For more information on status codes, see Table 6-14.
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

Provides access to the multiple sources to which an incident or a problem could be related. 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 6-46 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

Used for referencing source objects other than a job target.

Table 6-47 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 6-48 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 6-48 (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

Array of *gc$notif_udtp* type with a maximum size of 20.

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

gc$notif_udtp

Used for referencing User-defined target properties. UDTP should consist of a set of property key names and property values.

Table 6-49 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

Notification Payload Elements Specific to Event Types

gc$notif_event_attr_array

Array of *gc$notif_event_attr* is used for referencing event-specific attributes. The array has a maximum size of 25. Each element of the array is of type *gc$notif_event_attr* (used for referencing event type-specific attributes).

Table 6-50 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.

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

```
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
order by event_type, env_var_name;
```

You should convert the attribute name to upper case before using the name for comparison.

There is an attribute variable payload specific to each event type that can be accessed from a *gc$notif_event_attr_array* database type. The following tables list notification attributes for the most critical event types. You should convert the attribute name to uppercase before using the name for comparison.

Table 6-51 Environment variables specific to Metric Alert Event Type

Environment Variable	Description
COLL_NAME	The 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_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_DESCRIPTION	Brief description of the metric.
METRIC_GROUP	The name of the metric.
NUM_KEYS	The number of key metric columns in the metric.
SEVERITY_GUID	The GUID of the severity record associated with this metric alert.
VALUE	Value of the metric when the event triggered.

Table 6-52 Environment variables specific to Target Availability Event Type

Environment Variable	Description
AVAIL_SEVERITY	The transition severity (0-6) that resulted in the status of the target to change to the current availability status. Possible Values for AVAIL_SEVERITY • 0 (Target Down) • 1 (Target Up) • 2 (Target Status Error) • 3 (Agent Down) • 4 (Target Unreachable) • 5 (Target Blackout) • 6 (Target Status Unknown)
AVAIL_SUB_STATE	The substatus of a target for the current status.
CYCLE_GUID	A unique identifier for a metric alert cycle, which starts from the time the metric alert is initially generated until the time it is clear.
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.

Table 6-53 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.

Table 6-53 (Cont.) Environment variables specific to Job Status Change event type

Environment Variable	Description
EXEC_STATUS_CODE	Execution status code of job execution. For possible values, see Table 6-16.
STATE_CHANGE_GUID	Unique ID of last status change

Example 6-17 PL/SQL Script: Event Type Payload Elements

```
-- log_table table is created by following DDL to demostrate how to access
-- event notification payload GC$NOTIF_EVENT_MSG.

CREATE TABLE log_table (message VARCHAR2(4000)) ;

-- Define PL/SQL notification method for Events
CREATE OR REPLACE PROCEDURE log_table_notif_proc(s IN GC$NOTIF_EVENT_MSG)
IS
  l_categories gc$category_string_array;
  l_category_codes gc$category_string_array;
  l_attrs gc$notif_event_attr_array;
  l_ca_obj gc$notif_corrective_action_job;
BEGIN
  INSERT INTO log_table VALUES ('notification_type: ' || s.msg_info.notification_type);
  INSERT INTO log_table VALUES ('repeat_count: ' || s.msg_info.repeat_count);
  INSERT INTO log_table VALUES ('ruleset_name: ' || s.msg_info.ruleset_name);
  INSERT INTO log_table VALUES ('rule_name: ' || s.msg_info.rule_name);
  INSERT INTO log_table VALUES ('rule_owner: ' || s.msg_info.rule_owner);
  INSERT INTO log_table VALUES ('message: ' || s.msg_info.message);
  INSERT INTO log_table VALUES ('message_url: ' || s.msg_info.message_url);
  INSERT INTO log_table VALUES ('event_instance_guid: ' ||
s.event_payload.event_instance_guid);
  INSERT INTO log_table VALUES ('event_type: ' || s.event_payload.event_type);
  INSERT INTO log_table VALUES ('event_name: ' || s.event_payload.event_name);
  INSERT INTO log_table VALUES ('event_msg: ' || s.event_payload.event_msg);
  INSERT INTO log_table VALUES ('source_obj_type: ' ||
s.event_payload.source.source_type);
  INSERT INTO log_table VALUES ('source_obj_name: ' ||
s.event_payload.source.source_name);
  INSERT INTO log_table VALUES ('source_obj_url: ' || s.event_payload.source.source_url);
  INSERT INTO log_table VALUES ('target_name: ' || s.event_payload.target.target_name);
  INSERT INTO log_table VALUES ('target_url: ' || s.event_payload.target.target_url);
  INSERT INTO log_table VALUES ('severity: ' || s.event_payload.severity);
  INSERT INTO log_table VALUES ('severity_code: ' || s.event_payload.severity_code);
  INSERT INTO log_table VALUES ('event_reported_date: ' ||
to_char(s.event_payload.reported_date, 'D MON DD HH24:MI:SS'));

  IF s.event_payload.target.TARGET_LIFECYCLE_STATUS IS NOT NULL
  THEN
    INSERT INTO log_table VALUES ('target lifecycle_status: ' ||
s.event_payload.target.TARGET_LIFECYCLE_STATUS);
  END IF;

  -- Following block illustrates the list of category codes to which the event
  -- belongs.

  l_category_codes := s.event_payload.category_codes;
  IF l_categories IS NOT NULL
  THEN
    FOR c IN 1..l_category_codes.COUNT
    LOOP
      INSERT INTO log_table VALUES ('category_code ' || c || ' - ' ||
l_category_codes(c));
    END LOOP;
  END IF;

  --
  -- Each event type has a specific set of attributes modeled. Examples of
  -- event types include metric_alert, target_availability, job_status_change.
  -- Following block illustrates how to access the attributes for job_status change
  -- event type
  --
  IF s.event_payload.event_type = 'job_staus_chage'
  THEN
    l_attrs := s.event_payload.event_attrs;
    IF l_attrs IS NOT NULL
    THEN
      FOR c IN 1..l_attrs.COUNT
      LOOP
        INSERT INTO log_table VALUES ('EV.ATTR name=' || l_attrs(c).name || ' value=' ||
l_attrs(c).value || ' nls_value=' || l_attrs(c).nls_value);
      END LOOP;
    END IF;
  END IF;

  -- Following block illustrates how to access corrective action job's attributes  IF
s.msg_info.notification_type = GC$NOTIFICATION.NOTIF_CA AND
s.event_payload.corrective_action IS NOT NULL
  THEN
    l_ca_obj := s.event_payload.corrective_action;
    INSERT INTO log_table VALUES ('CA JOB_GUID: ' || l_ca_obj.JOB_GUID);
    INSERT INTO log_table VALUES ('CA JOB_NAME: ' || l_ca_obj.JOB_NAME);
    INSERT INTO log_table VALUES ('CA JOB_OWNER: ' || l_ca_obj.JOB_OWNER);
    INSERT INTO log_table VALUES ('CA JOB_TYPE: ' || l_ca_obj.JOB_TYPE);
    INSERT INTO log_table VALUES ('CA JOB_STATUS: ' || l_ca_obj.JOB_STATUS);
    INSERT INTO log_table VALUES ('CA JOB_STATUS_CODE: ' || l_ca_obj.JOB_STATUS_CODE);
    INSERT INTO log_table VALUES ('CA JOB_STEP_OUTPUT: ' || l_ca_obj.JOB_STEP_OUTPUT);
    INSERT INTO log_table VALUES ('CA JOB_EXECUTION_GUID: ' ||
l_ca_obj.JOB_EXECUTION_GUID);
    INSERT INTO log_table VALUES ('CA JOB_STATE_CHANGE_GUID: ' ||
l_ca_obj.JOB_STATE_CHANGE_GUID);
    INSERT INTO log_table VALUES ('CA OCCURRED_DATE: ' || l_ca_obj.OCCURRED_DATE);  END
IF;

COMMIT ;
END ;
/
```

Passing Information to Webhooks and Slack

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.

Token Variables Specific to Event Notifications

Incident Payload

Token Name	Description
`${incident.incidentID}`	ID of the incident
	Display ID of the incident
`${incident.severity}`	Severity code of the incident
	Name of the target
`${incident.targetType}`	Target type
`${incident.creationTimestamp}`	Incident creation timestamp
`${incident.summary}`	Summary message for the incident
`${incident.resolutionStatus}`	Resolution state for the incident
`${incident.owner}`	Owner of the incident

Event Payload

Token Name	Description
`${event.eventID}`	ID for the event
`${event.eventClass}`	Event class
`${event.msg}`	Message for the event
`${event.severity}`	Severity code for the event
`${event.target}`	Name of the target
`${event.creationTimestamp}`	Evet creation timestamp
`${event.targetType}`	Target type

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.

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.

Email Notifications

- Make sure an email gateway is set up under the Notification Methods page of Setup. The Sender's email address should be valid. Clicking the Test button will send an email to the Sender's email address. Make sure this email is received. Note that the Test button ignores any Notification Schedule.

- Make sure an email address is set up. Clicking the Test button will send an email to specified address and you should make sure this email is received. Note that the Test button ignores any Notification Schedule.
- Make sure an email schedule is defined. No emails 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 email and notification methods are assigned to the rule.

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, select **Management Services and Repository** to view these errors.
- Check for any delivery errors. You can view them from Incident Manager. From the **Enterprise** menu, select **Monitoring**, then select **Incident Manager**. The details will give the reason why the notification was not delivered.

Notification System Trace Messages

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

Tracing is configured by setting the *log4j.category.oracle.sysman.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 -name log4j.category.oracle.sysman.em.notification -value
DEBUG -module logging
```

Note: The system will prompt you for the SYSMAN password.

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@myco.com

2006-11-08 03:18:47,006 [DeliveryThread-EMAIL1] INFO  em.notification run.227 -
Notification handled for SYSMAN/admin@myco.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@myco.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@myco.com
2011-08-17 14:02:20,921 [DeliveryThread-OSCMD14] DEBUG notification.pbs logp.251 -
Notification handled for SYSMAN, OSCMD, 8
```

Email Errors

The SMTP gateway is not set up correctly:

```
Failed to send email to my.admin@myco.com: For email 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.oracle.com: Sending failed;
nested exception is:
javax.mail.MessagingException: Unknown SMTP     host: badhost.example.com;
```

Invalid email address:

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

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

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 (machineb10.oracle.com_Management_Service) (SYSMAN, myrule )
```

No execute permission on executable:

```
Error calling /bin/myscript: java.io.IOException: /bin/myscript: cannot execute
(machineb10.oracle.com_Management_Service) (SYSMAN, myrule )
```

Timeout because OS Command ran too long:

```
Timeout occurred running /bin/myscript (machineb10.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.

SNMP Trap Errors

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

The OMS will not report an error if the SNMP trap cannot reach the third party SNMP console as this is sent via UDP. If the SNMP trap encounters problems when trying to reach the third party SNMP console, possible SNMP trap problems include: invalid host name, port, community for a machine running an SNMP Console or a network issue such as a firewall problem.

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

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 Enterprise Manager console.

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

System Broadcasts

Enterprise Manager allows you to broadcast important instantly viewable system messages to Enterprise Manager consoles throughout your managed environment. These messages can be directed to specific users or all Enterprise Manager users. This feature can be useful when notifying users that Enterprise Manager is about to go down, when some part of your managed infrastructure has been updated, or when there is a system emergency.

> **Note:**
>
> Only Super Administrators can send system broadcasts.

Setting System Broadcast Preferences

Before you can send a system broadcast, you must first set your broadcast preferences.

1. Log in to Enterprise Manager as a Super Administrator.
2. From the <USERNAME> menu, select **Preference** and then **System Broadcast**. The System Broadcast User Preferences UI displays.

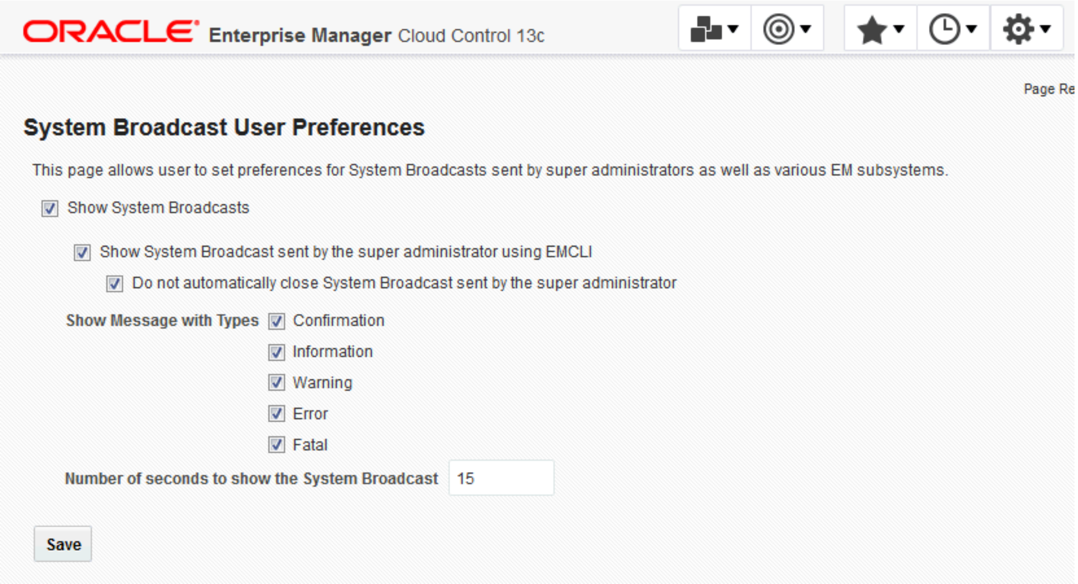

System Broadcast User Preferences

This page allows user to set preferences for System Broadcasts sent by super administrators as well as various EM subsystems.

☑ Show System Broadcasts

☑ Show System Broadcast sent by the super administrator using EMCLI

☑ Do not automatically close System Broadcast sent by the super administrator

Show Message with Types ☑ Confirmation

☑ Information

☑ Warning

☑ Error

☑ Fatal

Number of seconds to show the System Broadcast 15

Save

3. Check the desired broadcast message preferences then click **Save**.

   Whenever you send a system broadcast message, these are the preferences that will be used.

   > **Note:**
   >
   > The *Number of seconds to show the System Broadcast* setting will only work when the *Do not automatically close System Broadcast sent by the super administrator* option is disabled.

Creating a System Broadcast

Once your preferences are set, you use the EM CLI verb *send_system_broadcast* to send a system broadcast message.

```
emcli send_system_broadcast
        -toOption="ALL|SPECIFIC"
        [-to="comma separated user names"]
        [-messageType="INFO|CONF|WARN|ERROR|FATAL" (default is INFO)]
        -message="message details"
```

Options

- *toOption*

  Enter the value ALL to send the broadcast message to all users logged into the Enterprise Manager Console. Or enter SPECIFIC to send System Broadcast to users specified by *-to*.

- *to*

  Comma-separated list of users who are to receive the broadcast message. This option can only be used if the *-toOption* is set to *SPECIFIC*.

- *messageType*

  Type of System Broadcast, it can be one of following types

  - INFO *(Information)*

  - CONF *(Confirmation)*
  - WARN *(Warning)*
  - ERROR
  - FATAL
- *message*

  Message to be sent in the System Broadcast. The message has a maximum of 200 characters.

Example:

In this example, you want to broadcast an informational message indicating that you will be bringing down Enterprise Manager within an hour in order to perform an emergency patching operation.

```
emcli send_system_broadcast -messageType="INFO" -toOption="ALL" -message="Enterprise
Manager will be taken down in an hour for an emergency patch"
```

7

Using Dynamic Runbooks

Dynamic Runbooks allow you to capture procedural knowledge and expertise in diagnosing and resolving issues from your subject matter experts and store them in Enterprise Manager for ready access and execution by other Enterprise Manager users. Dynamic Runbooks are supported for Incidents and Metrics.

This chapter covers the following topics:

- About Dynamic Runbooks
- Working with Dynamic Runbooks
- Creating Dynamic Runbooks
- Using Dynamic Runbooks
- Exporting and Importing Dynamic Runbooks
- Oracle Provided Runbooks
- Advanced Operations and Configuration

Feature Updates

Feature/Update	Release
Metric Steps Support for Related Targets	Enterprise Manager 13*c* Release 5 Update 21 (13.5.0.21)
Launch Runbooks from the All Metrics Page	Enterprise Manager 13*c* Release 5 Update 21 (13.5.0.21)
Support Links to EM Pages in Step Instructions and Note Steps	Enterprise Manager 13*c* Release 5 Update 21 (13.5.0.21)
Save Step Output into Runbook Variables	Enterprise Manager 13*c* Release 5 Update 15 (13.5.0.18)
Comments in the Runbook Session	Enterprise Manager 13*c* Release 5 Update 15 (13.5.0.17)
Runbook Variables as bind parameters on SQL	Enterprise Manager 13*c* Release 5 Update 15 (13.5.0.17)
Warning and Critical Thresholds in Metric Step chart	Enterprise Manager 13*c* Release 5 Update 15 (13.5.0.17)
Oracle Provided Runbooks and OS Command Step	Enterprise Manager 13*c* Release 5 Update 15 (13.5.0.15)
Import/Export and Note Step External Links	Enterprise Manager 13*c* Release 5 Update 11 (13.5.0.11)
Relevant Runbooks	Enterprise Manager 13*c* Release 5 Update 10 (13.5.0.10)
Create Like	Enterprise Manager 13*c* Release 5 Update 8 (13.5.0.8)
Dynamic Runbooks Initial Release	Enterprise Manager 13*c* Release 5 Update 7 (13.5.0.7)

About Dynamic Runbooks

Background and Concepts

Dynamic Runbooks are documented procedures that IT Operations staff follow to resolve an issue. The steps in the Dynamic Runbook are often based on the operational experience of subject matter experts who are responsible for providing a consistent way to respond to and resolve issues in a timely manner.

Dynamic Runbooks typically consist of a set of ordered instructions (steps) that include looking at metric data, correlating data across targets, and executing SQL against the Enterprise Manager repository or target databases to resolve the issue. Many, if not all such tasks, can be performed entirely within Enterprise Manager. With Dynamic Runbooks, users can execute these steps inside Enterprise Manager.

Terminology

- *Runbook*
  Set of ordered instructions, data, and SQL designed to provide a standard way to diagnose and respond to an incident.
- *Runbook Session*
  Execution of a Runbook by a user
- *Runbook Variables*
  Context values applied to a Runbook. These could be: metric name, target name, target type, etc. These can be populated from an incident or metric in which the Runbook session is executed, or they could be an output of a previous Runbook step, i.e. the value for a Runbook variable can be manually filled in by the session user based on the output from one of the previous steps. Instructions for how this should be manually done from a previous step can be covered via the description in the step.

Prerequisites

You need to set the following OMS properties in order for the Repository SQL and Target SQL Runbook steps to work:

- `oracle.sysman.db.restfulapi.executesql.repository.query.enable`
- `oracle.sysman.db.restfulapi.executesql.target.query.enable`
- `oracle.sysman.db.restfulapi.executesql.target.update.enable`
- `oracle.sysman.db.restfulapi.executesql.throttle.max.req.per.user.interval.sec`
- `oracle.sysman.db.restfulapi.executesql.throttle.max.req.per.user`
- `oracle.sysman.db.restfulapi.executesql.throttle.max.concurrent.request`

Examples:

```
emctl set property -name
oracle.sysman.db.restfulapi.executesql.repository.query.enable -value true -
sysman_pwd "<your password>"
emctl set property -name
oracle.sysman.db.restfulapi.executesql.target.query.enable -value true -
sysman_pwd "<your password>"
emctl set property -name
oracle.sysman.db.restfulapi.executesql.target.update.enable -value true -
```

```
sysman_pwd  "<your password>"
emctl set property -name
oracle.sysman.db.restfulapi.executesql.throttle.max.req.per.user.interval.sec
-value 60 -sysman_pwd "<your password>"
emctl set property -name
oracle.sysman.db.restfulapi.executesql.throttle.max.req.per.user -value 30 -
sysman_pwd "<your password>"
emctl set property -name
oracle.sysman.db.restfulapi.executesql.target.update.enable -value true -
sysman_pwd "<your password>"
emctl set property -name
oracle.sysman.db.restfulapi.executesql.throttle.max.concurrent.request -value
20 -sysman_pwd "<your password>"
```

Run any SQL on Database and *Connect Target* privilege is required on the database in order to run a *Target SQL* type step.

Working with Dynamic Runbooks

At a high level, working with Dynamic Runbooks involves the following workflows:

Create a New Dynamic Runbook

You can create a Runbook in one of three ways:

You can go to Incident Manager, choose the incident for which you want to create a Runbook, and then click on the **Create Runbook** link.

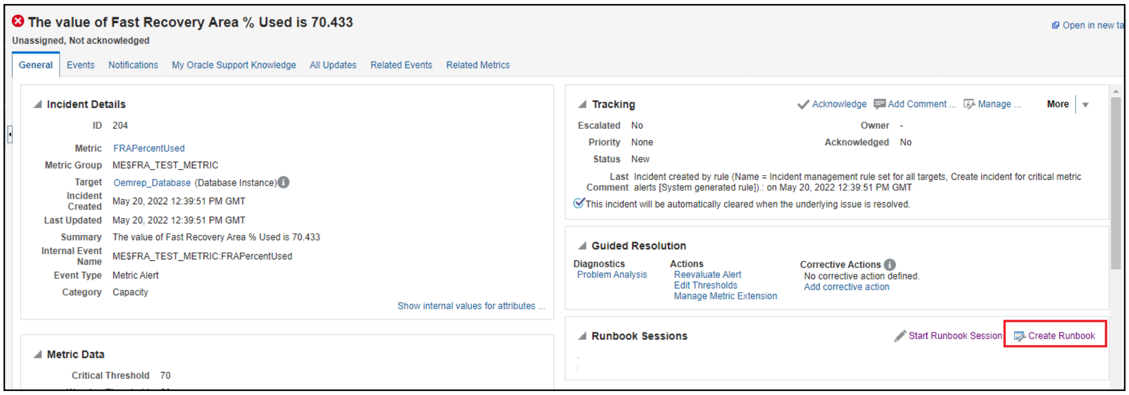

The value of Fast Recovery Area % Used is 70.433

Unassigned, Not acknowledged

Incident Details

ID 204

Metric FRAPercentUsed

Metric Group ME$FRA_TEST_METRIC

Target Oemrep_Database (Database Instance)

Incident Created May 20, 2022 12:39:51 PM GMT

Last Updated May 20, 2022 12:39:51 PM GMT

Summary The value of Fast Recovery Area % Used is 70.433

Internal Event Name ME$FRA_TEST_METRIC:FRAPercentUsed

Event Type Metric Alert

Category Capacity

Metric Data

Critical Threshold 70

Tracking

Escalated No

Priority None

Status New

Owner -

Acknowledged No

Last Comment Incident created by rule (Name = Incident management rule set for all targets, Create incident for critical metric alerts [System generated rule]).: on May 20, 2022 12:39:51 PM GMT

This incident will be automatically cleared when the underlying issue is resolved.

Guided Resolution

Diagnostics
Problem Analysis

Actions
Reevaluate Alert
Edit Thresholds
Manage Metric Extension

Corrective Actions
No corrective action defined.
Add corrective action

Runbook Sessions

Go to the All Metrics page of a target homepage and select the metric for which you would like to create a runbook. In the Runbook Sessions section, click on Create Runbook link.

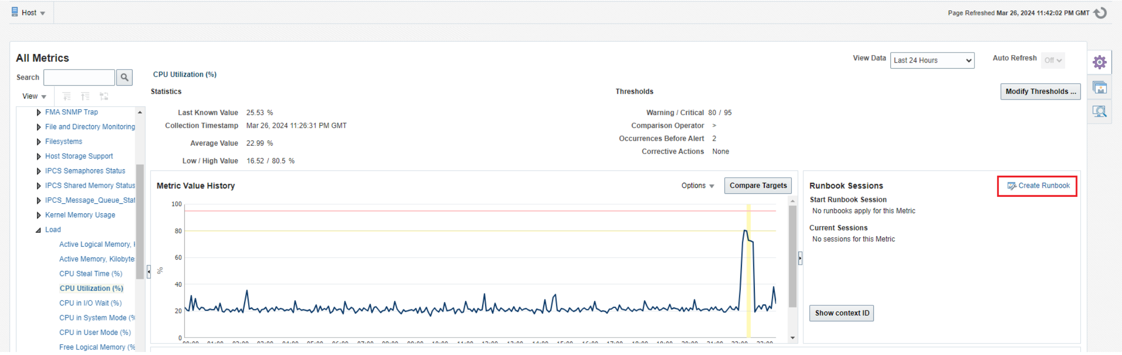

Page Refreshed Mar 26, 2024 11:42:02 PM GMT

View Data: Last 24 Hours
Auto Refresh: Off

CPU Utilization (%)

Statistics

Last Known Value 25.53 %
Collection Timestamp Mar 26, 2024 11:26:31 PM GMT
Average Value 22.99 %
Low / High Value 16.52 / 80.5 %

Thresholds

Warning / Critical 80 / 95
Comparison Operator >
Occurrences Before Alert 2
Corrective Actions None

Runbook Sessions

Start Runbook Session
No runbooks apply for this Metric

Current Sessions
No sessions for this Metric

Alternatively, you create a Runbook by creating a Runbook Draft and then specify the context and the context ID for which you would like to create a Runbook. The context ID can either be an Incident context or Metric context. For Incidents, get the Incident ID from Incident Manager. For Metrics, get the metric context from the Runbook Sessions section associated with the metric in the target's All Metrics page.

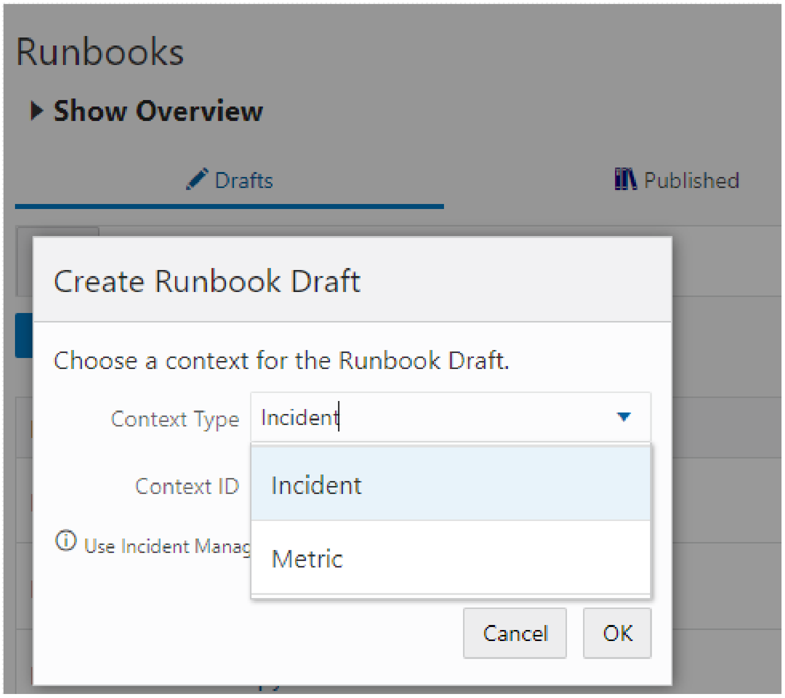

Create Runbook Draft

Choose a context for the Runbook Draft.

Context Type: Incident

Context ID

Use Incident Manag

In all these ways, choosing the incident or metric will provide the context that will be used to develop and test the Runbook steps. Once the Runbook is complete, you can publish the Runbook.

Create a New Version of a Dynamic Runbook

You create a new version of a published Dynamic Runbook by creating a Runbook draft from the Runbook, editing the steps and re-publishing the Runbook. When you publish the new version of the Runbook, the new published version replaces the prior version.

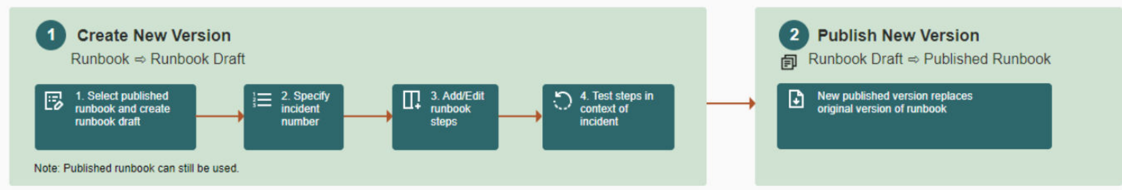

1 Create New Version

Runbook ⇨ Runbook Draft

1. Select published runbook and create runbook draft

2. Specify incident number

3. Add/Edit runbook steps

4. Test steps in context of incident

Note: Published runbook can still be used.

2 Publish New Version

Runbook Draft ⇨ Published Runbook

New published version replaces original version of runbook

Use a Published Dynamic Runbook

By executing a Dynamic Runbook from the context of an incident or metric, a Runbook session is created.

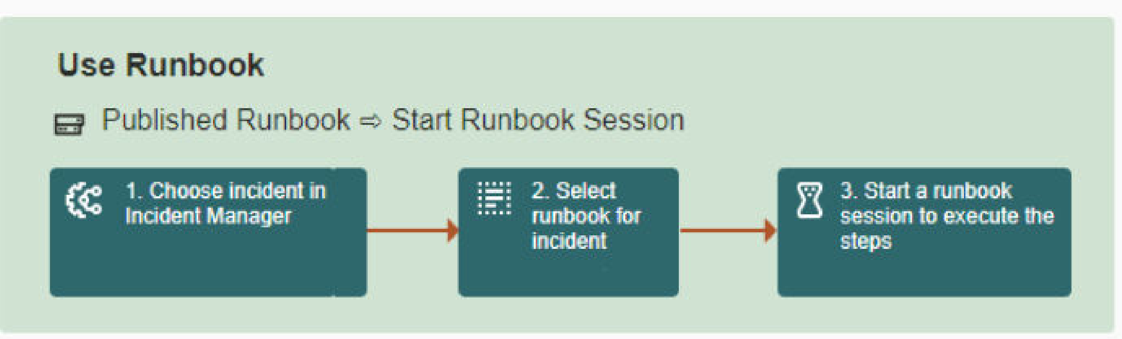

Use Runbook

Published Runbook ⇨ Start Runbook Session

1. Choose incident in Incident Manager

2. Select runbook for incident

3. Start a runbook session to execute the steps

Types of Dynamic Runbook Users

There are different types of Dynamic Runbook users that play specific roles in the creation, running, and administrative maintenance of Runbooks.

- *Runbook Author*
  This user can:
  - Create and publish Runbooks
  - Show/hide published Runbooks

  Create Runbooks privilege is required.
- *Runbook User*
  This user can:
  - Execute published Runbooks (start Runbook session)
  - If a user can see an incident, they can execute the Runbook session

  All Enterprise Manager Users are Runbook users.
- *Runbook Administrator*
  This user can:
  - Delete and show/hide on any Runbook
  - Show/hide published Runbooks
  - List and delete Runbook sessions marked done that have been created by other users
  - Extend expiry of any runbook session marked done
  - View Runbook session data without access to incidents/jobs

  Manage Runbooks privilege is required.

> **Note:**
>
> The Enterprise Manager Super Administrator has complete access to all facets of the Runbook lifecycle (creation, administration, editing, publication)

Accessing Dynamic Runbooks

From the Enterprise Manager console, you can access Runbooks from the following locations:

Runbooks List Page

From the **Enterprise** menu, select **Monitoring** and then **Runbooks**. The Runbooks page displays all Published Runbooks. If you have the *Create Runbooks* privilege, you will also see available Runbook drafts. You can initiate Runbook draft creation from this page. However, the recommended way to create an new Runbook is from Incident Manager, in context of the incident for which you want to create a Runbook.

Runbook Sessions Page

From the **Enterprise** menu, select **Monitoring** and then **Runbook Sessions**. Here, you'll see a list of Runbook sessions you own (both active and done). To start a new Runbook session, go to Incident Manager, locate the incident for which you would like to start a Runbook session, then click on the option to **Start Runbook Session**.

Target's All Metrics Page

From any target homepage menu, select Monitoring then All Metrics. When you choose a metric from the list of metrics, a Runbook Sessions section will be available next to the metric chart. From here you can start a Runbook session from any available runbooks for that metric or create a new Runbook in context of that metric.

Incident Manager

From the **Enterprise** menu, select **Monitoring** and then **Incident Manager**. When you click on a particular incident, you'll see any Runbook sessions created for that incident. You'll also have the option to view the top (up to three) Runbooks that may be relevant to the incident (see Relevant Runbooks for more information). From here, you'll be able to start a Runbook session or create a new Runbook within the context of the incident.

All Metrics

In the target instance homepage, from the target menu, select Monitoring and then All Metrics. When you click on a particular metric in a metric group, you will see a Runbooks section similar to the one you would see in Incident Manager with all of the same options for starting a Runbook session or creating a new Runbook draft.

Creating Dynamic Runbooks

Any Runbook Author can create a Dynamic Runbook

How to create a Dynamic Runbook - YouTube Video

1. To create a Dynamic Runbook for an incident, from the Enterprise menu, select Monitoring and then Incident Manager. View incident details by clicking its row and then go to the **Runbook Sessions** region. If there is a Runbook session that is currently active, it will be listed.

To create a Dynamic Runbook for a metric, locate the metric in the **All Metrics** page of the appropriate target. You should see a Runbook Sessions region next to the metric's performance charts.

To create a Dynamic Runbook without an existing incident or metric, from the **Runbooks** list page, select **Create Draft** and choose **Universal Context**. Universal context does not require any specific context ID, as it is not dependent on any external artifact. Even before an incident is raised, some metric value reaches critical threshold or incident/metric does not apply at all (for jobs/agents/other areas), universal context Runbooks can help triage and resolve underlying issues.

2. In the Runbook Sessions region, select **Create Runbook**. The *New Runbook* authoring UI displays in the context of the incident or metric, thus allowing you to create Runbook steps.
3. By default, the authoring UI opens with the *Overview & Prerequisites* as the first step. Here, you enter a description of what the new Runbook will accomplish and any prerequisites that are required to complete all Runbook steps. To aid readability, the Overview & Prerequisites step supports simple formatting via markdown notation. Click on the information icon (“I”) in the text entry area to view a list of markdown elements supported by the text editor. The markdown formatting will be applied when you save the step.
4. Click **Save Step**.
5. From the **Add a Step** drop-down menu, select the desired step type you want to define. There are five step types:
   - *Note*: Use the Note step to provide instructions to the user. You can use markdown formatting with this step and also include Runbook variables. Refer to the Runbooks UI for details on the Runbook variables
   - *Metric Data*: Show a graphical view of a metric over time
   - *Repository SQL*: Allows you to run a SQL statement against documented Enterprise Manager repository schema views (MGMT$)

     > **Note:**
     >
     > Accessing other elements from the Enterprise Manager repository schema is not supported.

   - *Target SQL*: Allows you to execute SQL against any database target in Enterprise Manager
   - *OS Commands and Scripts*: Allows you to execute a single operation, as you would do on the command line or to specify an interpreter to be used to execute a script which you enter as a part of the step definition.

   For more information about step types and how to use them, see Step Types.
6. When you select a step type, the step authoring UI displays. Enter a step title and description of what the step will do, instructions, and any other information regarding interpretation of results.
7. By default, the step will be prepopulated with settings from the context of the incident or metric. Depending on the step type, you can select other step parameters such as Metric Group Name by clicking on the current selection to display the pop-up selector and choose a new value from the drop-down list.

> **Note:**
>
> At the right, you can choose from a list of variables pulled from the incident context. If none of the variables meet your requirement, you can create your own variable by clicking **Add a Runbook Variable**.

When defining a step, you can assign data from the output of the step to populate values of existing author-defined Runbook variables.

Starting with Enterprise Manager 13c Release 5 Update 21 (13.5.0.21) details, such as target name and internal target type, for a different target instance, can be set directly or through the use of Runbook Variables.

8. Links to other steps (Optional)

   Using markdown language, links to other steps can be added using `(GOTO:`*`step number`*`)`.

   Here is an example which adds a link to another step based on the step input using an if-else clause:

   ```
   If the value is less than 100, (GOTO:5)
   Else, move ahead to the next step.
   ```

9. **Save to Variable** (Optional)

   When defining a step, you can dynamically assign data from the output of the step to populate values of author-defined Runbook variables, following the steps below:

   > **Note:**
   >
   > A Runbook Variable can be assigned a value from only one cell of one single step's output.

   a. Click **Save to Variable** to open the tab.

   b. From the output of the step, click the cell containing the value you want to save to a variable. Once you do so, you will see a **Save to Runbook Variable** dialog appear.

   c. In the **Save to Runbook Variable** dialog box, click to open the dropdown list providing the available variables, and select the one you want to contain the data, then click **Save**. If you are saving a value to a Runbook Variable from another step, your author-defined Runbook Variable will not show up in the dropdown list. If the list is empty or you want to create a new variable, click **Create New Runbook Variable** at the top of the list, and use the pop-up window to create a new variable.

   The **Saved Variables** list, on the right side of the **Save to Variables** tab, shows the variables that take data from the step output.

10. **Links to EM Pages in Step Instructions** (Optional)

    Use the MarkDown language to specify a link to an EM page. To construct the link, navigate to an EM page and then copy the part of the URL which begins with " / " after the host and port name. For example, if the url is `http://my.emsite.com:1234/em/faces/si-host-home?type=oracle_database&target=Oemrep_Database` then you could use it as a link as [Target Home Page](/em/faces/si-host-home?)

Other examples:

```
[Runbook Page](/em/faces/runbookPage)
```

```
[Target Page](/em/faces/core-uifwk-targetSearch-home)
```

You can use also use variables in your URLs. For example, if your target `Oemrep_Database` is `type=oracle_database`, they can be put into variables `target_type` and `target_name` respectively such that the link would be `[Target Home Page](/em/faces/si-host-home?type=$target_type&target=$target_name)`

11. Click **Run** to execute the step and view any rendered visualizations such as charts or tables defined by the current parameter settings. You can *run* multiple tests to ensure that your settings provide the desired results. Once the step is run, you will see that any values selected to be saved to Runbook Variables will now be assigned to those Runbook Variables.
12. Click **Save Step**. The step is now ready to use and will be added to the list of steps.
13. Once you've finished adding steps, you've created a Runbook draft as it's still editable. It will appear in the Runbooks list page. (From the **Enterprise** menu, select **Monitoring** and then **Runbooks**). In draft mode, only the author can view/edit the Runbook.

Step Types

You can create the following types of steps:

- Note Step
- Metric
- Repository SQL
- Target SQL
- OS Command and Scripts

> **Note:**
>
> The following step limitations apply to all step types:
>
> - 1MB of data (configurable via the `oracle.sysman.core.runbooks.maxStepDataSizeKB` is the OMS property)
> - 20 steps per Runbook (configurable via the `oracle.sysman.core.runbooks.maxSteps` OMS property)

- **Note Step**
  Any type of text meant to provide information or instruction pertaining to the Runbook. The Note can be plain text, but can accommodate simple formatting via markdown language. A markdown reference table showing supported formatting is provided when creating a Note step.

  Input:

  1. Any type of text
  2. Basic formatting of text (markdown)

   a. bulleted lists (1 level)
   b. numbered lists (1 level)
   c. bold text
   d. italic text
   e. bold, italic text
   f. Header (h3, h4, h5)
   g. horizontal rule
   h. blockquote

3. Runbook Variables
4. External Links
   You can specify three types of external links within a Note step using the following markdown notation:

   Links to a specific URL: `[Link Text](http//www.anothercompany.com)`

   Links to My Oracle Support notes (only the MOS ID is required): `[MOS Link Text](MOS:MOS ID)`. For example `[Enterprise Manager configuration issues](MOS:1234567.1)`

   Links to EM pages: `[Target Page](/em/faces/core-uifwk-targetSearch-home)`

   > **Note:**
   >
   > External linking is available starting with Enterprise Manager 13c Release 5 Update 11 (13.5.0.11).

Output:

Formatted text with Runbook Variables substituted out for values.

Example:

Remember to have a named credential for the target database **$ora_target_name** on which the FRA incident has occurred. The Note step will show:Remember to have a named credential for the target database **orcl_database** on which the FRA incident has occurred.

- **Metric**
  This step shows the time series chart for the specified metric. Defaults to the metric from passed in context, but it can be changed. When the Runbook is launched in context of an incident/event, the chart also shows when the event occurred.

  With Oracle Enterprise Manager 13c Release 5 Update 17 (13.5.0.17), the chart shows lines representing the current warning and critical thresholds for the metric. This allows you to evaluate the metric trend against warning and critical thresholds.

  Starting with Enterprise Manager 13c Release 5 Update 21 (13.5.0.21) details, such as target name and internal target type, for a different target instance, are available to be typed directly in the target selector input box or set through the use of Runbook Variables.

  > **Note:**
  >
  > Configuration metrics are not supported.

Input

1. Target name/specification
2. Metric Group
3. Metric
4. Metric Key Values
5. Time range notation:
   You can also specify a metric time span by providing start and end dates. You can make those times relative to *now* or *evt_time*.
   - *now* corresponds to current time.
   - *evt_time* corresponds to the date and time when the event occurred.

   Well known Runbook variable named ***ora_event_time*** from the incoming context corresponds to the ***evt_time***. Use the keyword ***now*** or ***evt_time*** followed by **+** or **-** and finally an integer and **m** for minutes, **h** for hours, or **d** for days.

   Examples:

   `now` : current time when the step is run

   `now - 6h` : 6 hours prior to the current time

   `evt_time + 3d` : 3 days after the time the event occurred

Output:

1. Time series line chart display
2. Legend

> **Note:**
>
> If the Runbook was launched in context of an incident/event, the chart will show when the incident occurred so long as the specified time range includes when the incident happened.

- **Repository SQL**
  You provide a SQL statement to be run against MGMT$ views in the repository schema. You are limited to running the SQL against the repository view tables (MGMT$ or GC$). You can use BIND parameters. There is a limit on step size data: The maximum limit for a SQL query is 20 columns and 1000 rows, or 1 MB of data (whichever limit is reached first).

  Using Bind Parameters

  You can bind values from Runbook Variables into your SQL statement using a colon followed by the Runbook Variable reference. For instance:

  ```
  select target_name from MGMT$TARGETS where target_type
  = :$ora_target_type_name
  ```

  This will create a bind parameter in the SQL statement and use whatever the value of `$ora_target_type_name` is as the value for that bind parameter when the SQL is executed.

  Starting with Enterprise Manager 13c Release 5 Update 21 (13.5.0.21) you can use a repository SQL step to find the details of another target instance (target name, and internal target type). These values can be assigned to Runbook variables, that can be used in the metric step edit screen to drive the values of the step.

SQL Query Resultset formatting

1. If the query column list contains a clob, wrap the CLOB column_name in TO_CHAR() function.
   Example:

   ```
   SELECT em_varchar_col, to_char(clob_column) from MGMT$CLOB_TABLE
   ```

2. If the query column list contains a timestamp, wrap the date column name in a TO_CHAR() function and provide the proper date format mask.
   Example:

   ```
   SELECT em_varchar_col, to_char(em_timestamp_col,'DD-MON-YYYY
   HH24:MI:SS') from MGMT$DATE_TABLE
   ```

3. If a user defined variable is a string that represents a date and the variable is used in the query WHERE clause, wrap the variable name in a TO_TIMESTAMP() or TO_DATE() function and provide the corresponding date format mask.
   Example:

   ```
   SELECT em_varchar_col,to_char(em_timestamp_col,'DD-MON-YYYY
   HH24:MI:SS') FROM MGMT$DATE_TABLE where em_timestamp_col
   < to_date(:$1_user_defined_date_variable,'DD-MON-YYYY')
   ```

- **Target SQL**
  Ability to run predefined SQL (single statement only) on any database target. In addition to executing select statements, you can run Data Manipulation Language (DML) and Data Definition Language (DDL). The named credential is required whether it is a select statement or DDL/DML.

  Pre-requisites:

  In order to run the Target SQL step on a database target, you will need to have these privileges on your database targets:

  1. Run any sql on Database
  2. Connect Target
  3. Access to the named credential referenced in the step

  Input

  1. Target Name
  2. Target Type
  3. SQL
  4. Named Credential

  Using Runbook Variables

  You can use Runbook Variables in your Target SQL queries, DML or DDL statements.

  1. SQL queries (SELECT statements):

     Using Bind Parameters you can bind values from Runbook Variables into your SQL statement using a colon followed by the Runbook Variable reference.

Example:

```
select target_name from MGMT$TARGETS where target_type
=:$ora_target_type_name
```

This will create a bind parameter in the SQL statement and use whatever the value of $ora_target_type_name is as the value for that bind parameter when the SQL is executed.

2. DML or DDL statements

   You can use Runbook Variables in your DML or DDL statements by using a colon followed by the Runbook Variable, such as `:$ora_target_name` or `:$fra_size`.

   Example:

   ```
   alter system set db_recovery_file_dest_size=:$fra_size  scope=both
   ```

   It will substitute the value of the variable before executing the DML or DDL statement.

Using ISO 8601 formatted Runbook variables in SQL Query Where Clause

If a Runbook variable's value is in *ISO 8601 date* format (such as 2022-03-01T19:21:00.000Z) and you need to use it in a *where* clause on a database table/view with date column:

1. Use *to_utc_timestamp_tz* function to convert the value into a database *timestamp with time zone* value.
2. Value from #1 can be converted to another time zone (such as target time zone) as needed.
3. *cast* the output from #1 or #2 into *date* type, and use it as part of the *where* clause to join with tables having *date* column.
4. Following examples assume the *utc_time1* Runbook variable is set to a value in *ISO 8601 date format* (such as '2022-03-01T19:21:00.000Z').

   Examples:

   Example #1: Shows how to cast the variable value to date without taking timezone conversion into account.

   ```
   select to_char(min(collection_timestamp),'YYYY-MM-DD HH24:MI:SS') as
   min_time, to_char(max(collection_timestamp), 'YYYY-MM-DD HH24:MI:SS')
   as max_time, count(1) as total_count
   from mgmt$metric_current
   where collection_timestamp >=
   cast(TO_UTC_TIMESTAMP_TZ(:$ora_event_time_utc) as date)
   ```

   Example #2: Shows how to convert the variable value to target timezone first, then cast it to date;

   ```
   select to_char(min(collection_timestamp),'YYYY-MM-DD HH24:MI:SS') as
   min_time, to_char(max(collection_timestamp), 'YYYY-MM-DD HH24:MI:SS')
   as max_time, count(1) as total_count
   from mgmt$metric_current
   ```

```
where collection_timestamp >=
cast((TO_UTC_TIMESTAMP_TZ(:$ora_event_time_utc) at time zone
TIMEZONE_REGION) as date)
```

Output

1. Tabular display of return data

- **OS Commands and Scripts**
  Allows you to execute a single operation, as you would do on the command line on the host of the specified target. or to execute a script which you enter as a part of the step definition using the specified interpreter.

  Pre-requisites:

  In order to run the OS Command and Scripts step, you will need to have these privileges:

  1. View privilege on the target, target's host and the agent that monitors the target's host
  2. Access to the named credential referenced in the step

  Input

  1. Target Name
  2. Target Type
  3. Named Credential
  4. Command Execution Type:

     Single Operation

     a. Command

     > **Note:**
     >
     > The interpreter will be defaulted for whatever the OS is on the host.

     Script

     > **Note:**
     >
     > For Windows platform, **Script** as a Command Execution Type is currently not supported.

     a. Script
     b. Interpreter

     > **Note:**
     >
     > If left blank, the interpreter is defaulted according to the OS on the target host where the script is executed.

  Output:

  Returns the output from running the command or script on the target host.

Creating Runbook Variables

Runbook Variables are populated with values from the incident details when a session or draft is created. These values can be used by steps for display in the instruction text or as a part of the step definition which gets data or performs some task. This allows a Runbook session to provide information and perform tasks based on context specific to that session. And for a draft, it allows an author to define a Runbook using the Runbook Variables rather than hardcoded values.

For values which are not included as context when the draft is created, the author can create their own Runbook Variables. This Runbook Variable can be populated in these ways:

- The runbook Variable can have a static value which remains the same across the sessions that will eventually be created on that Runbook
- The Runbook Variable can require the session user to provide a value each time a new session is created for the Runbook. For instance, for target DB SQL steps, a named credential is required. Creating a Runbook Variable for this named credential and requiring the value to be set by the session user ensures that the SQL as a part of that step will be executed with a named credential the session user has access to for that session's target DB.
- The Runbook Variable can also by dynamically populated with data from the output of a step. This feature makes the use of Runbook Variables more powerful and is available starting with Oracle Enterprise Manager 13c Release 5 Update 18. As an example, you may have a step that determines the filesystem used by the database. The author can save the output of that step into a Runbook Variable, and then use that Runbook Variable in the next step that determines the available free space of the filesystem. Note that a Runbook Variable can save data from a specific value cell of a step's output, hence if you need to save multiple data values, you'll need to save these in multiple Runbook Variables.

There are some restrictions around the naming of Runbook Variables an author can create:

1. It cannot start with `ora_`.
2. It can contain only letters, numbers, and underscores.
3. It cannot be longer than 64 characters.
4. Additionally, for credentials, `_cred` should be included in the name (this allows us to filter for the specific case of named credentials used in the target DB step)

When creating a Runbook Variable, the author will provide:

1. A name, which is how the Runbook Variable will be referenced in the step (for instance, `user_target_db_cred`)
2. A display name, which is a short, user readable name (for instance, Target Database Named Credential)
3. A description (optionally), which will be displayed in the session UI when the session user is asked to fill in a value (for instance, **Provide a named credential for this target database**)
4. Value to be used in the draft, which allows you to create the Runbook Draft with a value the author can use to execute the step
5. Options for **Value used in a Runbook Session**:
   There are options for how the value will be populated in a Runbook Session. The table below has a summary of these options and usage guidelines.

Value used in a Runbook Session	When to use this option	How the runbook variable value is populated	Can runbook session user change it?	Usage Notes
Specified by session user or by running a step	Use this option when you want the runbook session user to specify the value of the runbook variable. Typically, this is done by running a runbook step to populate the runbook variable. For advanced users, the session user can manually enter the runbook value if needed.	In a runbook session, the user runs a step to populate the runbook variable or he types in a value for the variable.	Yes - by typing in a value or by running a step	The initial value of the variable is empty (null).
Value provided by author or by running a step	Use this option if you don't want the session user the ability to manually specify the runbook variable value. Instead, you want the session user to run a step to populate the value or you want to define the value as part of the runbook definition.	The runbook author defines the value as part of runbook definition or he has the user run a step to populate it.	Yes - only by running a step, if the runbook author creates a step that populates the runbook variable	If the author expects the value of the variable to be populated by running a step, the initial value of the variable is null.

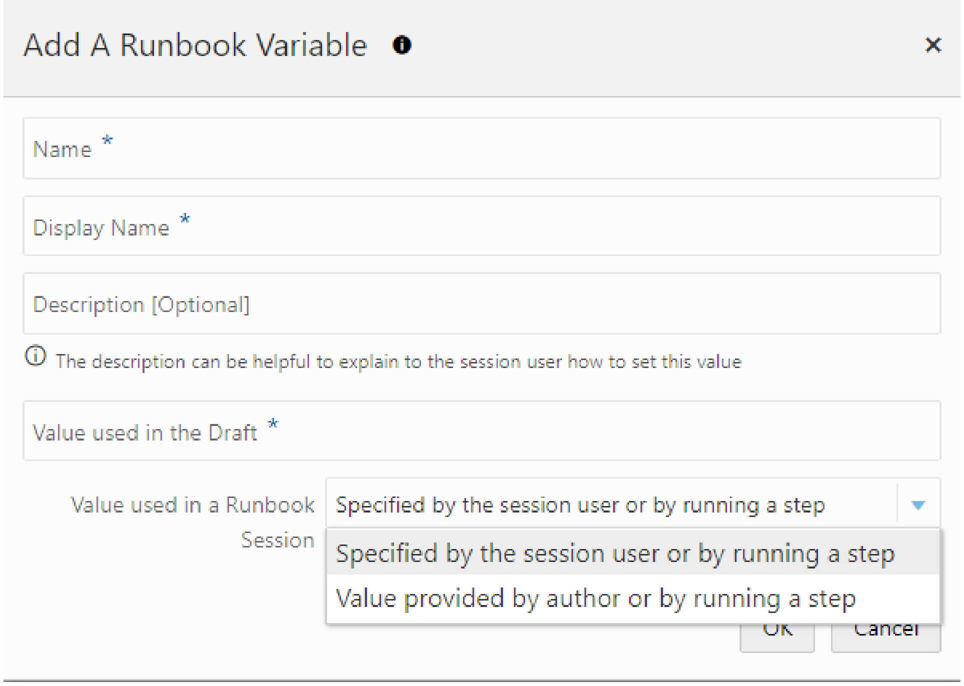

Add A Runbook Variable

Name *

Display Name *

Description [Optional]

The description can be helpful to explain to the session user how to set this value

Value used in the Draft *

Value used in a Runbook Session: Specified by the session user or by running a step

Specified by the session user or by running a step
Value provided by author or by running a step
OK
Cancel

Publishing Runbooks

Once you are finished defining your Runbook draft, you're ready to publish it. Once published it will become available to all Runbook users. Only the author of the Runbook draft can publish it.

Navigate to the Runbooks list page (**Enterprise**->**Monitoring**->**Runbooks**) and go to the **Drafts** tab. From the Actions menu for your Runbook, select **Publish**. Select **Keep Draft** if the Runbook is still in development or the draft is still relevant, and click **OK**. Click on the **Published** tab to view your Runbook in the published Runbook list.

Creating New Versions of Runbooks

- Create a draft from a published Runbook. Navigate to the Runbooks page. From the **Actions** menu for the desired Runbook, select **Create New Version**.
- While the draft is being edited, the published version of the Runbook remains available for users to continue to create sessions.
- The Runbook author will be required to provide a context ID to do so.
- When the new draft is published, it replaces the previous published Runbook after user confirmation.
- This new published Runbook will be used to create all future Runbook sessions. It has no impact on prior Runbook sessions.

Creating New Runbooks from Existing Runbooks

To save time and effort, you can create a new Dynamic Runbook from an existing one by using the *Create Like* action. *Create Like* can be used on either published or draft Runbooks.

Using Published Dynamic Runbooks

> **Note:**
>
> You do not have to be the owner of the published Runbook to use the Create Like action.

1. From the Enterprise menu, select **Monitoring** and then **Runbooks**.
2. Click on the **Published** tab to view the table of published Runbooks.
3. From the **Actions** drop-down menu for the chosen Runbook, select **Create Like**. The Create Like dialog is displayed.
4. Enter a context ID to use to create the new draft Runbook.
5. Click **OK**. The Runbook editor screen is displayed with a copy of the published runbook in draft mode. As the owner of this draft, you can edit this runbook in the context of the specified context ID.

Using Draft Dynamic Runbooks

Since you own the draft version of the Runbook, a more complete copy can be made. You do not need to provide an context ID since the selected draft Runbook will already have that information.

1. From the Enterprise menu, select **Monitoring** and then **Runbooks**.

2. Click on the **Drafts** tab to view the table of draft Runbooks that you own.
3. From the **Actions** drop-down menu for the chosen Runbook, select **Create Like**. The Create Like confirmation dialog is displayed.
4. Click **OK**. The Runbook editor page displays the new draft Runbook.

Other Operations

After you've published the Runbook, you can perform other operations on it from the **Actions** menu.

You can:

- Hide/Unhide the Runbook to control availability of published Runbooks for public use.
- Delete the published Runbook.

> **Note:**
>
> These operations can only be performed by the Runbook author or Runbook administrator.

Adding External Links

You can add external links to information that may assist a future Runbook user with additional diagnostic or contextual information.

Creating a Rule to designate a Dynamic Runbook for an Incident

After creating and publishing a Runbook, you may want to designate specific Runbooks to be used for an incident. The designation of the specific Runbook for an incident can be done as part of Event Rule actions:

1. Start creating an event rule.
2. In the **Add Actions** page, select **Create Incident**.

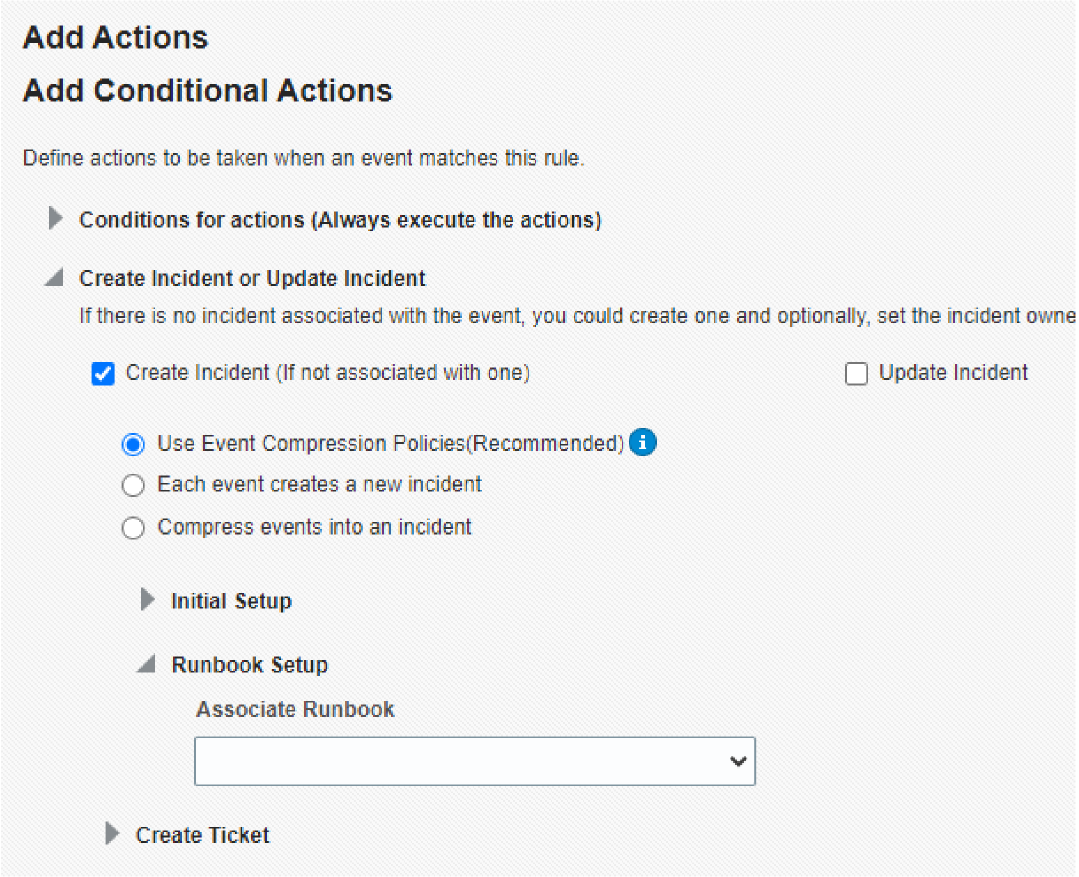

Add Actions

Add Conditional Actions

Define actions to be taken when an event matches this rule.

▶ Conditions for actions (Always execute the actions)

◢ Create Incident or Update Incident

If there is no incident associated with the event, you could create one and optionally, set the incident owne

☑ Create Incident (If not associated with one) □ Update Incident

● Use Event Compression Policies(Recommended)

○ Each event creates a new incident

○ Compress events into an incident

▶ Initial Setup

◢ Runbook Setup

Associate Runbook

▶ Create Ticket

3. Click the **search bar** in the **Associate Runbook** section.
4. Search and select the Runbook you want to designate for the incident, and click **Select**.

When the incident rule is applied to create an incident, the Runbook associated with it in the step above will be a part of the incident and appear in the section below under **Start Runbook Session**. It will be the only entry with text **Recommended By Incident Rule**. All other relevant Runbooks are still accessible through the **All Relevant Runbooks...** link.

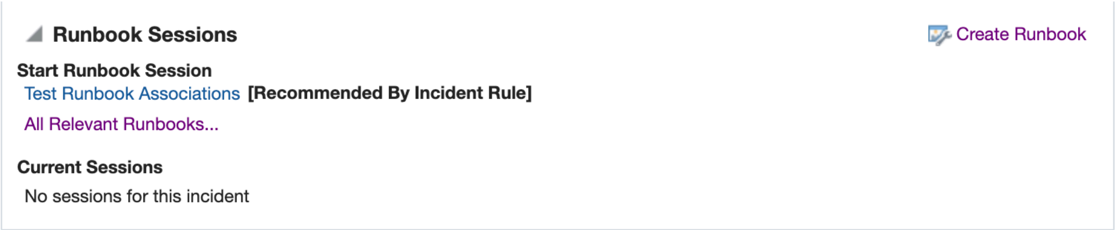

◢ Runbook Sessions

Create Runbook

Start Runbook Session

Test Runbook Associations [Recommended By Incident Rule]

All Relevant Runbooks...

Current Sessions

No sessions for this incident

Runbooks associated through the incident rule will always appear at the top of the list and have the text **[Recommended by Incident Rule]** next to the Runbook name. The relevance score will be automatically 100% (full bar is shown) and the text will show **Recommended**.

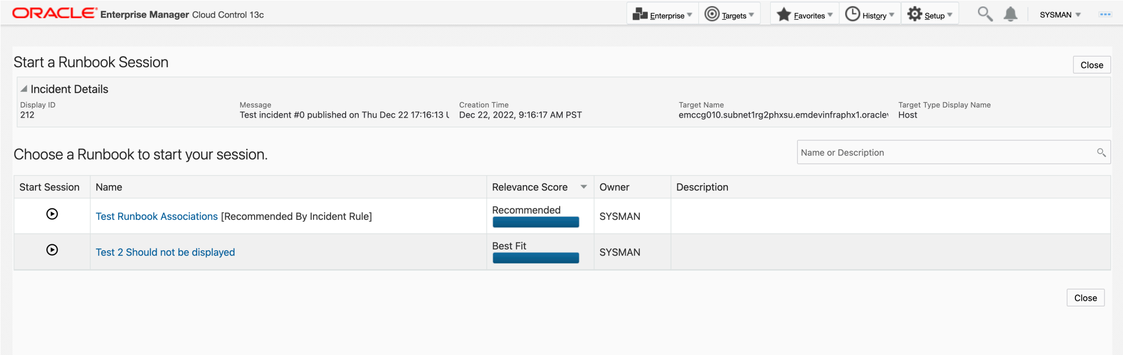

Start a Runbook Session

Incident Details

Display ID	Message	Creation Time	Target Name	Target Type Display Name
212	Test incident #0 published on Thu Dec 22 17:16:13 l	Dec 22, 2022, 9:16:17 AM PST	emccg010.subnet1rg2phxsu.emdevinfraphx1.oracle	Host

Choose a Runbook to start your session.

Start Session	Name	Relevance Score	Owner	Description
⊙	Test Runbook Associations [Recommended By Incident Rule]	Recommended	SYSMAN	
⊙	Test 2 Should not be displayed	Best Fit	SYSMAN	

Using Dynamic Runbooks

As a Runbook user, you can create a Runbook session from Incident Manager or from the All Metrics page.

How to start a Dynamic Runbook Session - YouTube Video

1. Navigate to the Runbook Sessions region, and click **Start Runbook Session**.
   The *Start a Runbook Session* page appears and displays a table where you can select a Runbook with which to start a session.

   > **Note:**
   >
   > From the Incident Manager Runbook Sessions region, you can also view a list of Runbooks that may be relevant to the specific incident (see Relevant Runbooks for more information). To access one of the Relevant Runbooks, click **All Relevant Runbooks…** to display the *Start a Runbook Session* page listing the Relevant Runbooks.

2. Click **Start Session**. The Runbook session page opens for that incident or metric. Incident or metric context details are listed at the top of the page.
3. The Overview & Prerequisites step is intended to show what the Runbook is used for and any prerequisites required of the user before executing subsequent steps in the Runbook. Ensure all prerequisites are met before beginning the Runbook session. Once you've completed this step, you can click on the checkbox to indicate the step has been completed.
4. For each step, click the play icon to execute the operation and review live diagnostic information produced by the step. If you see a gear icon next to a Runbook step, that means user input is required before the step can be executed. An example of user input is a *named credential* for the target database. Once the input is specified, the gear icon will change to the play icon, which will indicate the step can be executed.

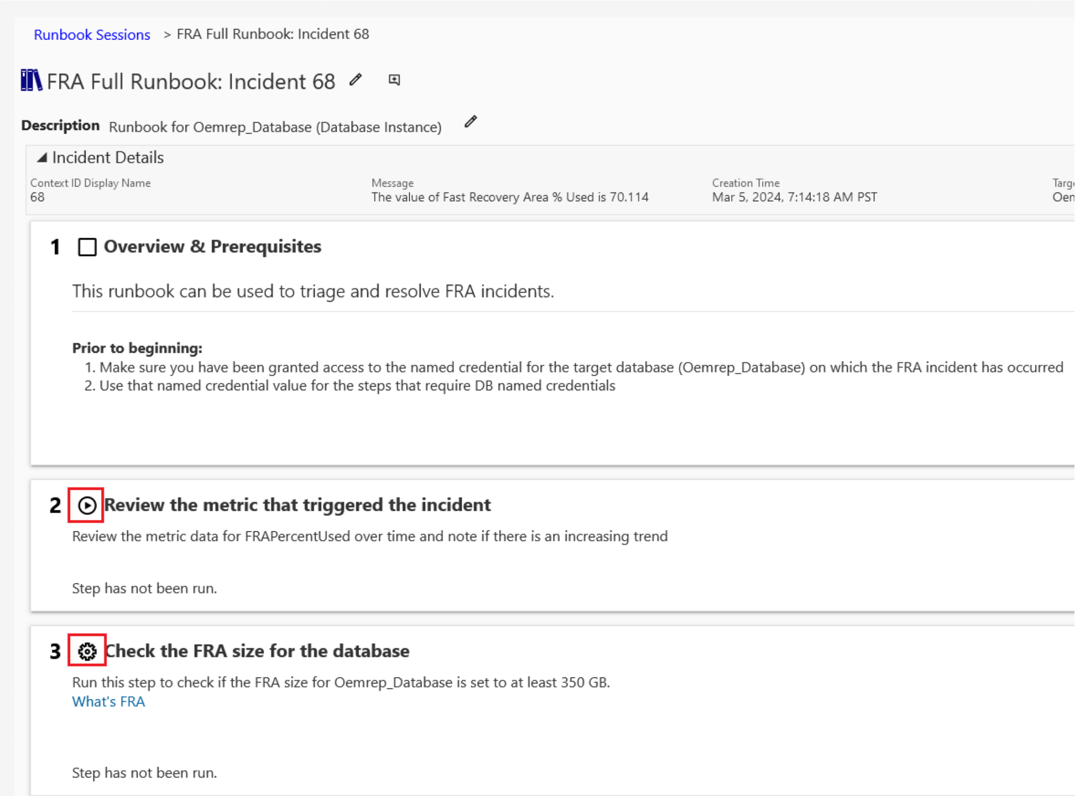

Runbook Sessions > FRA Full Runbook: Incident 68

FRA Full Runbook: Incident 68

Description Runbook for Oemrep_Database (Database Instance)

Incident Details

Context ID Display Name	Message	Creation Time	Targ
68	The value of Fast Recovery Area % Used is 70.114	Mar 5, 2024, 7:14:18 AM PST	Oen

1 □ **Overview & Prerequisites**

This runbook can be used to triage and resolve FRA incidents.

Prior to beginning:

1. Make sure you have been granted access to the named credential for the target database (Oemrep_Database) on which the FRA incident has occurred
2. Use that named credential value for the steps that require DB named credentials

2 **Review the metric that triggered the incident**

Review the metric data for FRAPercentUsed over time and note if there is an increasing trend

Step has not been run.

3 **Check the FRA size for the database**

Run this step to check if the FRA size for Oemrep_Database is set to at least 350 GB.
What's FRA

Step has not been run.

At the right is a list of Runbook Variables that pertain to the incident context.

Starting with Enterprise Manager 13c Release 5 Update 21 (13.5.0.21) metric steps for other target instances will be available.

> **Note:**
>
> The final step should be a Success Criteria step to allow the session user to verify that the Runbook steps performed by the user resolved the issue.

5. If you don’t complete the Runbook session, you can leave it open and come back to it. Only you, as the session owner will be able to return to continue running steps and complete the Runbook session. Note that there is a 14 day purge policy, which means you have 14 days to complete the Runbook session before it is purged from the system.

> **Note:**
>
> The default purge policy can be increased by setting the `oracle.sysman.core.runbooks.runbookSessionMinLifeDays` OMS property.

6. Once you have finished with the Runbook session, Click **Mark as Done**.
   - This places the Runbook session into read-only mode. No more edits or running steps allowed.
   - Once marked done, the session can be saved for a longer period of time by extending the expiration date in the Runbooks Sessions list page.

Extend the Expiry Date

- Sessions are kept for a default amount of time and automatically purged
- The time can be extended if the session is marked as done.
- The default time a session is saved is 14 days, default extended time is 45 days from when session is started.

> **Note:**
>
> The default purge policy can be increased by setting the `oracle.sysman.core.runbooks.runbookSessionMaxLifeDays` OMS property.

Exporting and Importing Dynamic Runbooks

Beginning with Oracle Enterprise Manager 13c Release 5 Update 11 (13.5.0.11), you can export and import published Runbook definitions, thus significantly enhancing Runbook reuse and sparing administrators from having to develop Runbooks from scratch.

> **Note:**
>
> You can only export/import published Runbooks. Runbook drafts and Runbook sessions cannot be used as they have contextual properties that are specific to the Enterprise Manager system in which they are authored. They will not be portable across Enterprise Manager installations.

Usage Guidelines

Exporting/importing published Runbooks across Enterprise Manager versions:

- Export from an older Enterprise Manager version and import to a newer Enterprise Manager version will always work.
- Export from a newer Enterprise Manager version and import to an older Enterprise Manager version will only work if step types are compatible (supported in both versions).

The following table lists the currently available step types and the corresponding minimum compatible Enterprise Manager release.

Step Type	Applies to EM release
Note	Enterprise Manager 13c Release 5 Update 7 and later
Target SQL	Enterprise Manager 13c Release 5 Update 7 and later
Repository SQL	Enterprise Manager 13c Release 5 Update 7 and later
Metric	Enterprise Manager 13c Release 5 Update 7 and later

Republish Runbooks authored before Enterprise Manager 13c Release 5 Update 11

For published Runbooks created using Enterprise Manager 13c Release 5 Update 7 thru 10, you must republish the Runbooks you plan on exporting/importing. This ensures that Runbook properties are up-to-date and compatible across Enterprise Manager deployments.

> **Note:**
>
> To republish a Runbook, simply create a new draft from an existing published Runbook, optionally edit the draft, and then publish Runbook.

Required User Privileges

In order to export or import a published Runbook, you must have at least one of the following privileges:

- CREATE_RUNBOOK—Allows you to create and publish Runbooks.
- MANAGE_RUNBOOK—Allows you to perform administrative operations, such as deleting/ showing/hiding Runbooks, or viewing Runbook session data.

For more information, see Types of Dynamic Runbook Users.

Export a Published Runbook

1. Navigate to the Published page. From the Enterprise menu, select **Monitoring** and then **Runbooks**.
2. Select **Export** from the *Actions* menu for the Runbook you want to export. An Export dialog displays with the default *Export Filename*. By default, the name of the currently published Runbook is used, however, you can enter a new filename.
3. Click **OK**. A Runbook definition JSON file is saved to your local file system.

Import a Published Runbook

1. Navigate to the Published page. From the Enterprise menu, select **Monitoring** and then **Runbooks**.
2. Click **Import**. The *Import a Published Runbook* dialog displays
3. Either drag and drop a Runbook definition file directly into the dialog, or click the Drag and Drop region to open a file browser to select a file. The imported definition file will immediately appear in the *Published* Runbook table.
4. Click **Close**.

> **Note:**
>
> The user importing a Runbook definition becomes the new owner of that Runbook.

Oracle Provided Runbooks

Beginning with Oracle Enterprise Manager 13c Release 5 Update 15 (13.5.0.15), you can use these Runbooks to help triage and resolve issues in your Enterprise Manager environment.

Dynamic Runbook	Description
Loader issues causing Agent backoff requests	This Runbook helps you diagnose and resolve incidents that the EM administrator receives when many agents have been asked to back off their uploading of data to the Oracle Management Server (OMS).
DB Tablespace Used Metric Time Out Error	This Runbook helps diagnose metric collection errors for the Oracle database Tablespace Space Used (%) metric.
Job Suspended Agent Not Ready	This Runbook helps troubleshoot job executions that are in status Suspended Agent Not Ready.
Triage Notification Backlog	Notification sub system is part of EM, and it is essential for notifications delivered by EM. This runbook can be used for triaging/fixing its issues.

To help find the best Oracle provided Runbooks for your needs, here are some groups of questions you might have, together with the Runbook that responds to them, respectively:

- **Loader issues causing Agent backoff requests**
  - Why are my metric charts not showing the latest data?
  - Why am I not seeing alerts coming from my agents?
  - How do I fix this incident: Health Check: Loader is in FAILURE
  - Why are my agents backing off?
  - How do I fix my large loader backlog?
- **DB Tablespace Used Metric Time Out Error**
  - How do I fix problemTbsp:ORA-00604: error occurred at recursive SQL level 1 ORA-01013: user requested cancel of current operation
  - How do I fix this issue: Timeout during collection: task timeout: 3600000 MILLISECONDS
  - How do I fix metric collection errors with the tablespace metric?
- **Job Suspended Agent Not Ready**
  - Why is my job in status "Suspended Agent Not Ready" state?
  - Is the agent ready to run my jobs?
- **Triage Notification Backlog**
  - Why is my notification system not working?
  - Is it moving at all, or is it just being slow?)3.
  - Which method is slow/stuck (pl/sql, email or os command)?
  - What is the current backlog and when did it start growing?

Advanced Operations and Configuration

The following topics cover areas that may be of interest to Runbook administrators and authors.

Dropping a User and the Impact on Dynamic Runbooks

Deleting an Enterprise Manager user will affect Dynamic Runbooks in the following ways:

1. Published Runbooks
   1. When a user is deleted, published Runbooks owned by the deleted user can be reassigned to any other Enterprise Manager User
      1. The reassigned user does not need *Create Runbooks* privilege. It can be granted later
   2. When the deleted user owns published Runbooks and no reassigned user is specified, then Runbooks are reassigned to the repository owner
2. Draft Runbooks
   1. Draft (unpublished) Runbooks owned by the deleted user are always deleted
   2. Any data saved as a part of the Draft Runbook will be deleted as well
3. Runbook Sessions
   1. Runbook sessions are not reassigned to any user, even if a reassign user is specified
   2. Open Runbook sessions owned by the deleted user are automatically marked as DONE
   3. DONE Runbook sessions owned by the deleted user will be dropped by the automated purge policy when they expire, as are any other sessions.

Configurable Properties

Runbook OMS Properties

The following OMS properties are used to configure Dynamic Runbook. See EMCTL Commands for OMS for information on setting OMS properties.

- **Max Step Data Size Limit**
  Name: `oracle.sysman.core.Runbooks.maxStepDataSizeKB`

  Description: The maximum data size allowed for a Runbook step in kilobytes.

  Default: 1024 (1 MB)

  Min: 100 (100KB)

  Max: 10240 (10MB)
- **Max Steps Allowed**
  Name: `oracle.sysman.core.Runbooks.maxSteps`

  Description: The maximum number of steps allowed for a Runbook.

  Default: 20

  Min: 10

  Max: 50
- **Minimum days of life for session**

Name: `oracle.sysman.core.Runbooks.RunbookSessionMinLifeDays`

Description: Initial number of days a session will be kept before being purged.

Default: 14

Min: 14

Max: 30

- **Maximum days of life for session**
  Name: `oracle.sysman.core.Runbooks.RunbookSessionMaxLifeDays`

  Description: Maximum number of days allowed for a Runbook session before it is purged.

  Default: 45

  Min: 45

  Max: 90

- **Max Output Size Returned By OS Command Step**
  Name: `oracle.sysman.core.runbooks.oscmd.maxRunStepOutputSize.KB`

  Description: The maximum data size that os command run step returns in kilobytes.

  Default: 1024 (1 MB)

  Min: 10 (100KB)

  Max: 10240 (10MB)

- **Max Time To Run OS Command Step**
  Name: `oracle.sysman.core.runbooks.oscmd.maxRunStepTime.seconds`

  Description: The maximum time allowed for os command run step REST API in seconds.

  Default: 120 (2 minutes)

  Min: 15 (15 seconds)

  Max: 600 (10 minutes)

Repository SQL Throttling

To minimize impact on Enterprise Manager performance, you can enable SQL throttling. See Repository Session (SQL) Throttling for more information.

Relevant Runbooks

A Runbook session can be started directly from the Incident Details page for a relevant Runbook, or you can click on **All Relevant Runbooks...** and a list of Runbooks ranked by *Relevance Score* is available from which to choose. Any Runbook that applies for the incident will appear in the *Start a Runbook Session* page table, paginated if necessary.

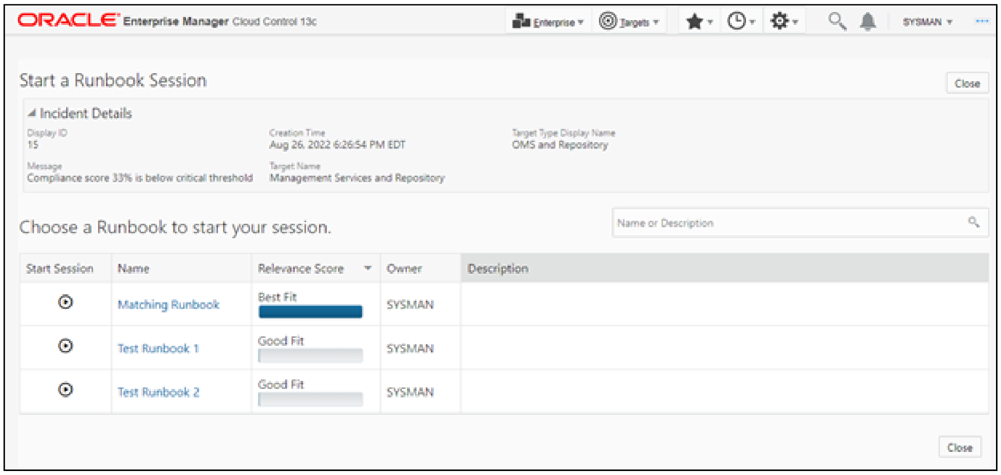

Start a Runbook Session

Close

Incident Details

Display ID
15

Creation Time
Aug 26, 2022 6:26:54 PM EDT

Target Type Display Name
OMS and Repository

Message
Compliance score 33% is below critical threshold

Target Name
Management Services and Repository

Choose a Runbook to start your session.

Name or Description

Start Session	Name	Relevance Score	Owner	Description
⊙	Matching Runbook	Best Fit	SYSMAN	
⊙	Test Runbook 1	Good Fit	SYSMAN	
⊙	Test Runbook 2	Good Fit	SYSMAN	

Close

From this page, you can start a Runbook session that best matches the resolution requirements of the incident using the Relevance Score as a guide. As shown above, Relevance Scores are conveniently sorted in ascending order.

About Relevance Scores

A Relevance Score is a measure of how well a Runbook matches the current incident. The greater level of matching, the higher the Relevance Score. A Runbook draft is created in context of a specific incident and carries forward some of those values when it is published. The Runbook carries parts of this incident's context even when it is published and available to use for any incident. That initial context tells the published Runbook which properties it provides are required and which are optional.

When a Relevance Score is created, the context from the incident you are currently viewing is compared to the published Runbook. The following is used to generate the score:

- All *required* properties must be present. This is required to create a Runbook Session for the current incident.
- Optional properties are evaluated. The Relevance Score is increased if an optional property exists. The score is further increased if the property value from the published Runbook matches the current incident context.
- Runbook variables used in steps are evaluated. Similar to optional properties, the Relevance Score is increased if a variable used in a step appears in the information provided by the incident. The score is further increased if the incident and variable values match.

> **Note:**
>
> By default, Runbooks published prior to Enterprise Manager 13c Release 5 Update 10 (13.5.0.10) will have lower Relevance Scores. To bring them up to parity with newer Runbooks, you need to republish these older Runbooks. Republishing will cause the system to recompute the Relevance Score.

Follow these steps to quickly republish your older Runbooks.

As the Runbook owner:

1. Locate the published version of the Runbook and create a new version. It will create a Runbook draft. There is no need to change anything. The Runbook draft will be saved automatically.
2. Next, immediately publish this new draft. A *Publish Runbook Draft* dialog will appear, asking you if you want to replace the existing published Runbook with this new version. Click **OK** to accept this option. Your Runbook will then be published.

8

Using Blackouts

This chapter covers the following topics:

- Blackouts and Notification Blackouts
- Working with Blackouts/Notification Blackouts
- Controlling Blackouts Using the Command Line Utility
- About Blackouts Best Effort

Blackouts and Notification Blackouts

Blackouts and Notification Blackouts help you maintain monitoring accuracy during target maintenance windows by providing you with the ability to suspend various Enterprise Manager monitoring functions for the duration of the maintenance period. For example, when bringing down targets for upgrade or patching, you may not want that downtime included as part of the collected metric data or have it affect a Service Level Agreement (SLA).

Blackout/Notification Blackout functionality is available from both the Enterprise Manager console as well as via the Enterprise Manager command-line interface (EMCLI).

About Blackouts

Blackouts allow you to suspend monitoring on one or more targets in order to perform maintenance operations. Blackouts, also known as patching blackouts, ensure that the target is not changed during the period of the blackout so that a maintenance operation on the actual target will not be affected. During this period, the Agent does not perform metric data collection on the target and no notifications will be raised for the target. Blackouts will allow Enterprise Manager jobs to run on the target during the blackout period by default. Optionally, job runs can be prevented during the blackout period.

A blackout can be defined for individual target(s), a group of multiple targets that reside on different hosts, or for all targets on a host. The blackout can be scheduled to run immediately or in the future, or stop after a specific duration. Blackouts can be created on an as-needed basis, or scheduled to run at regular intervals. If, during the maintenance period, the administrator discovers that he needs more (or less) time to complete his maintenance tasks, he can easily extend (or stop) the blackout that is currently in effect. Blackout functionality is available from both the Enterprise Manager console as well as via the Enterprise Manager command-line interface (EMCLI). EMCLI is often useful for administrators who would like to incorporate the blacking out of a target within their maintenance scripts.

Why use blackouts?

Blackouts allow you to collect accurate monitoring data. For example, you can stop data collections during periods where a managed target is undergoing routine maintenance, such as a database backup or hardware upgrade. If you continue monitoring during these periods, the collected data will show trends and other monitoring information that are not the result of normal day-to-day operations. To get a more accurate, long-term picture of a target's performance, you can use blackouts to exclude these special-case situations from data analysis.

Blackout Access

Enterprise Manager administrators that have at least Blackout Target privileges on all Selected Targets in a blackout will be able to create, edit, stop, or delete the blackout.

In case an administrator has at least Blackout Target privileges on all Selected Targets (targets directly added to the blackout), but does not have Blackout Target privileges on some or all of the Dependent Targets, then that administrator will be able to edit, stop, or delete the blackout. For more information on Blackout access, see "About Blackouts Best Effort".

About Notification Blackouts

Notification Blackouts are solely for suppressing the notifications on targets during the Notification Blackout duration. The Agent continues to monitor the target under Notification Blackout and the OMS will show the actual target status along with an indication that the target is currently under Notification Blackout. Events will be generated as usual during a Notification Blackout. Only the event notifications are suppressed.

The period of time under which the target is in Notification Blackout is not used to calculate the target's Service Level Agreement (SLA).

To place a target under Notification Blackout, you need to have at least Blackout Target privilege on the target.

There are two types of Notification Blackouts:

- **Maintenance Notification Blackout**: The target is under a planned maintenance and administrators do not want to receive any notifications during this period. Since the target is brought down deliberately for maintenance purposes, the Notification Blackout duration should not be considered while calculating the availability percentage and SLA. In this scenario, an administrator should create a maintenance Notification Blackout.
- **Notification-only Notification Blackout**: The target is experiencing an unexpected down time such as a server crash. While the administrator is fixing the server, they do not want to receive alerts as they are already aware of the issue and are currently working to resolve it. However, the availability percentage computation should consider the actual target status of the Notification Blackout duration and the SLA should be computed accordingly. In this scenario, the administrator should create a Notification-only Notification Blackout.

By default, when a Notification Blackout is created, it is a maintenance Notification Blackout (the *Under Maintenance* option will be selected by default and the administrator will need to select the *Non-maintenance* option in order to create a regular *Notification-only* Notification Blackout.

A Notification Blackout can be defined for individual target(s), a group of multiple targets that reside on different hosts, or for all targets on a host. The Notification Blackout can be scheduled to run immediately or in the future, or stop after a specific duration. Notification Blackouts can be created on an as-needed basis, or scheduled to run at regular intervals. If, during the maintenance period, the administrator discovers that he needs more (or less) time to complete his maintenance tasks, he can easily extend (or stop) the Notification Blackout that is currently in effect. Notification Blackout functionality is available from both the Enterprise Manager console as well as via the Enterprise Manager command-line interface (EMCLI). EMCLI is often useful for administrators who would like to incorporate the blacking out of a target within their maintenance scripts.

Notification Blackout Access

Enterprise Manager administrators that have at least Blackout Target privileges on all Selected Targets in a Notification Blackout will be able to create, edit, stop, or delete the Notification Blackout.

In case an administrator has at least Blackout Target privileges on all Selected Targets (targets directly added to the Notification Blackout), but does not have Blackout Target privileges on some or all of the *Dependent Targets*, then that administrator will be able to edit, stop, or delete the Notification Blackout.

Working with Blackouts/Notification Blackouts

Blackouts allow you to collect accurate monitoring data. For example, you can stop data collections during periods where a managed target is undergoing routine maintenance, such as a database backup or hardware upgrade. If you continue monitoring during these periods, the collected data will show trends and other monitoring information that are not the result of normal day-to-day operations. To get a more accurate, long-term picture of a target's performance, you can use blackouts to exclude these special-case situations from data analysis.

Performing maintenance operations such as create, edit, or stop a blackout, as well as retrieve various information and statistics on your defined blackouts, can be done using Blackout REST API. For more information, see Blackout APIs.

Creating Blackouts/Notification Blackouts

Blackouts/Notification Blackouts allow you to suspend monitoring on one or more managed targets.

To create a Blackout/Notification Blackout:

1. From the **Enterprise** menu, select Monitoring and then **Blackouts**.
2. From the table, click **Create**. Blackout and Notification Blackout selection dialog displays.
3. Choose either Blackout or Notification Blackout and click **Create**. The Create Blackout/ Notification Blackout page displays.

   When creating a Notification Blackout, you will also be able to specify the type of Notification Blackout via the *Maintenance Window* options.

   - Under maintenance. Target downtime is excluded from Availability(%) calculations.
   - Non-maintenance. Any target downtime will impact Availability(%) calculations.
4. Enter the requisite parameters for the new blackout/Notification Blackout and then click Submit.

Editing Blackouts/Notification Blackouts

Blackouts allow you to suspend monitoring on one or more managed targets.

To edit a Blackout:

1. From the **Enterprise** menu, select **Monitoring** and then **Blackouts**.
2. If necessary, use the **Search** and display options to show the blackouts you want to change in the blackouts table.
3. Select the desired Blackout/Notification Blackout. Details are displayed. click **Edit**. The Edit Blackout/Notification Blackout page displays.

4. Make the desired changes and click **Submit**.

Note: Enterprise Manager also allows you to edit blackouts after they have already started.

Viewing Blackouts/Notification Blackouts

To view information and current status of a blackout:

1. From the **Enterprise** menu, select **Monitoring** and then **Blackouts**.
2. If necessary, you can use the **Search** and display options to show the blackouts you want to view in the blackouts table.
3. Select the desired Blackout/Notification Blackout. Details are displayed.

Viewing Blackouts from Target Home Pages

For most target types, you can view a Blackout/Notification Blackout information from the target home page for any target currently under Blackout/Notification Blackout. The Blackout/ Notification Blackout Summary region provides pertinent Blackout/Notification Blackout status information for that target.

Viewing Blackout/Notification Blackouts from Groups and Systems Target Administration Pages

For Groups and Systems, you can view Blackout/Notification Blackout information about the number of active/scheduled Blackouts/Notification Blackouts on a group/system and its member targets.

Purging Blackouts/Notification Blackouts that have Ended

When managing a large number of targets, the number of completed Blackouts/Notification Blackouts, or those Blackouts/Notification Blackouts that have been ended by an administrator can become quite large. Removing these ended Blackouts/Notification Blackouts facilitates better search and display for current Blackouts/Notification Blackouts.

To purge ended Blackouts/Notification Blackouts from Enterprise Manager:

1. From the **Enterprise** menu, select **Monitoring** and then **Blackouts**.
2. Use the search criteria to filter for the desired targets.
3. From the **Blackout Timeframe** drop-down menu, select **History**.
4. In the table, select the ended Blackouts/Notification Blackouts you want to remove and click **Delete**. The Delete Blackout/Brownout confirmation page appears.
5. Click **Delete** to complete the purge process.

Retroactive Blackouts and Outages

If a target is brought down for maintenance and a blackout is not scheduled for that target, the maintenance period would be reflected as target downtime, thus a negative impact on the target’s availability history. This would be a problem for say a target’s service-level agreement (SLA) where the collected metrics that define the level of expected service would be inaccurate.

Enterprise Manager lets you remedy this situation by allowing you to define blacklouts and outages retroactively.

Retroactive Blackouts

As mentioned previously, retroactive blackouts can be used to specify past maintenance periods where the administrator has forgotten to set a blackout. The retroactive period will be used to adjust the target's availability (%) period by excluding these periods from availability (%) calculations, thus increasing a target's availability (%). Retroactive blackouts can be created either from the console or using EM CLI.

The following sequence of events illustrates the typical scenario for a retroactive blackout.

- The target is brought down for maintenance.
- The target availability % goes down from 100% => 80%.
- The target is brought back up after the maintenance work has been completed.
- The administrator realizes no blackout was created for the maintenance period.
- The administrator creates a retroactive blackout for the maintenance period.
- The target availability goes back up from 80% ==> 100%.

Creating a Retroactive Blackout from the Enterprise Manager Console

1. From the **Enterprise** menu, select **Monitoring** and then **Blackouts**. The Blackouts page displays.
2. Click **Settings**. The Settings page displays.

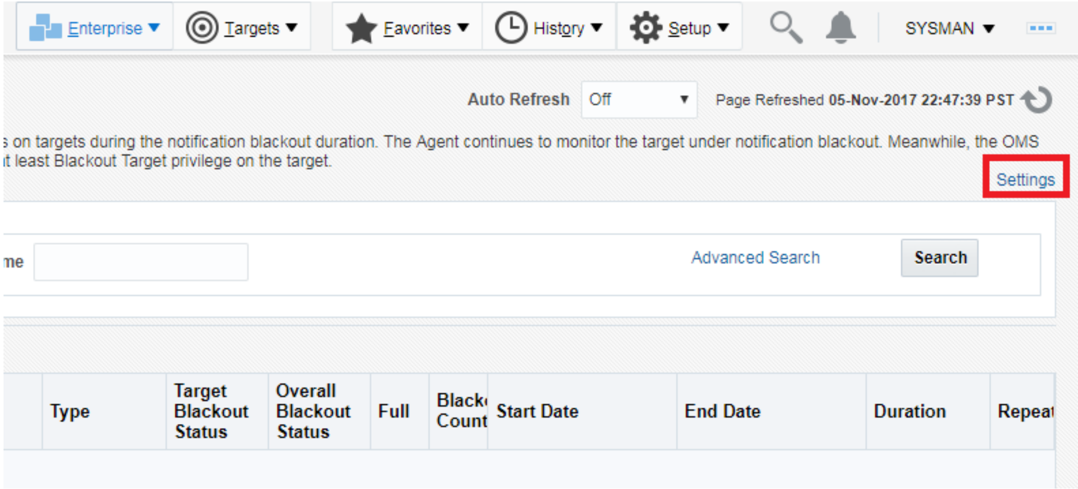

Auto Refresh Off

Page Refreshed 05-Nov-2017 22:47:39 PST

s on targets during the notification blackout duration. The Agent continues to monitor the target under notification blackout. Meanwhile, the OMS
t least Blackout Target privilege on the target.

Type	Target Blackout Status	Overall Blackout Status	Full	Black Count	Start Date	End Date	Duration	Repeat

3. At the bottom of the Settings page, check the **Enable Retroactive Blackout Feature** option box.

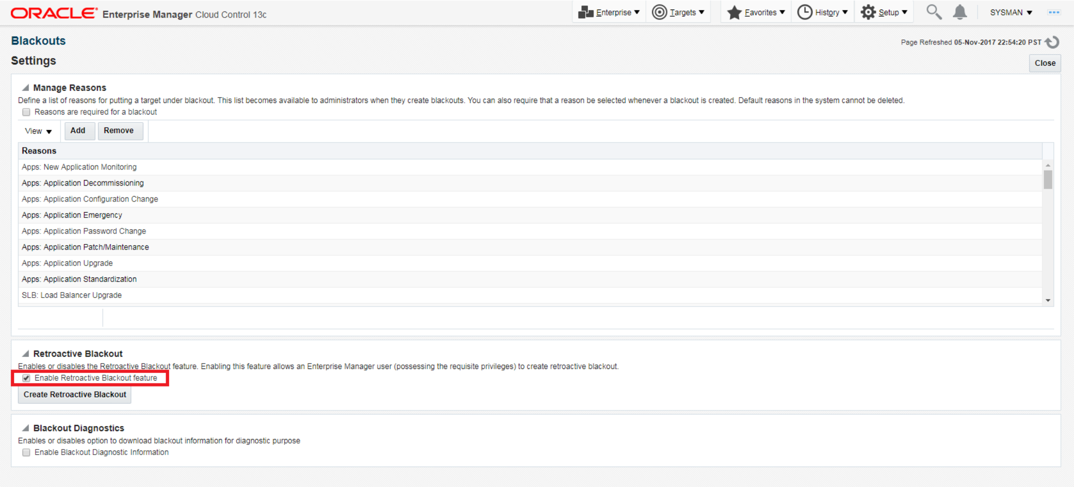

Blackouts

Settings

Page Refreshed 05-Nov-2017 22:54:20 PST

Close

Manage Reasons

Define a list of reasons for putting a target under blackout. This list becomes available to administrators when they create blackouts. You can also require that a reason be selected whenever a blackout is created. Default reasons in the system cannot be deleted.

☐ Reasons are required for a blackout

View ▾ | Add | Remove

Reasons
Apps: New Application Monitoring
Apps: Application Decommissioning
Apps: Application Configuration Change
Apps: Application Emergency
Apps: Application Password Change
Apps: Application Patch/Maintenance
Apps: Application Upgrade
Apps: Application Standardization
SLB: Load Balancer Upgrade

Retroactive Blackout

Enables or disables the Retroactive Blackout feature. Enabling this feature allows an Enterprise Manager user (possessing the requisite privileges) to create retroactive blackout.

☑ Enable Retroactive Blackout feature

Create Retroactive Blackout

Blackout Diagnostics

Enables or disables option to download blackout information for diagnostic purpose

☐ Enable Blackout Diagnostic Information

4. Click **Create Retroactive Blackout**. The Create Retroactive Blackout page displays.

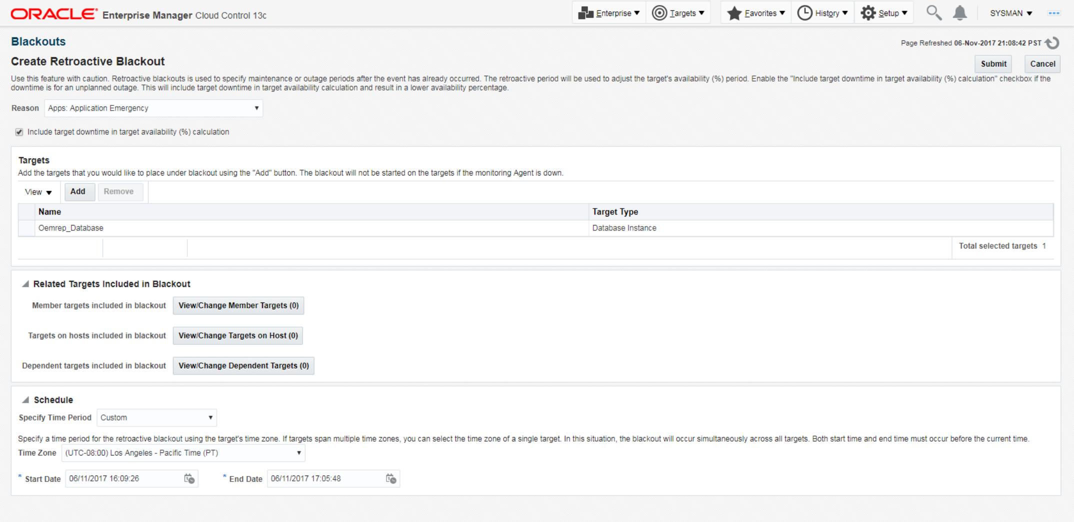

Blackouts

Create Retroactive Blackout

Page Refreshed 06-Nov-2017 21:08:42 PST

Submit | Cancel

Use this feature with caution. Retroactive blackouts is used to specify maintenance or outage periods after the event has already occurred. The retroactive period will be used to adjust the target's availability (%) period. Enable the "Include target downtime in target availability (%) calculation" checkbox if the downtime is for an unplanned outage. This will include target downtime in target availability calculation and result in a lower availability percentage.

Reason: Apps: Application Emergency

☑ Include target downtime in target availability (%) calculation

Targets

Add the targets that you would like to place under blackout using the "Add" button. The blackout will not be started on the targets if the monitoring Agent is down.

View ▾ | Add | Remove

Name	Target Type
Oemrep_Database	Database Instance

Total selected targets 1

Related Targets Included in Blackout

Member targets included in blackout: View/Change Member Targets (0)

Targets on hosts included in blackout: View/Change Targets on Host (0)

Dependent targets included in blackout: View/Change Dependent Targets (0)

Schedule

Specify Time Period: Custom

Specify a time period for the retroactive blackout using the target's time zone. If targets span multiple time zones, you can select the time zone of a single target. In this situation, the blackout will occur simultaneously across all targets. Both start time and end time must occur before the current time.

Time Zone: (UTC-08:00) Los Angeles - Pacific Time (PT)

* Start Date: 06/11/2017 16:09:26

* End Date: 06/11/2017 17:05:48

5. Enter the requisite information and then click **Submit**.

Retroactive Outages

If a monitored target goes down (outage) and Enterprise Manager does not detect it, the target availability percentage will be inaccurate. In this situation, the availability percentage will be too high. To remedy this inaccuracy, Enterprise Manager lets you specify this outage retroactively. A retroactive outage is essentially a retroactive blackout that specifies target downtime should be included as part of the availability calculation.

Retroactive outage can be created either from the console or using EM CLI.

The following sequence of events illustrates the typical scenario for a retroactive outage.

- The target goes down for an unknown reason.
- Enterprise Manager does not detect the target's down state. Target availability remains at 100%.
- The administrator brings the target back up.

- The administrator creates a retroactive blackout with the *Include target downtime in target availability (%) calculation* option selected for the unplanned outage period.
- Target availability goes down from 100% ==> 84%.

Creating a Retroactive Outage from the Enterprise Manager Console

1. From the **Enterprise** menu, select **Monitoring** and then **Blackouts**. The Bl

   ackouts page displays.
2. Click **Settings**. The Settings page displays.
3. Click **Create Retroactive Blackout**. The Create Retroactive Blackout page displays.

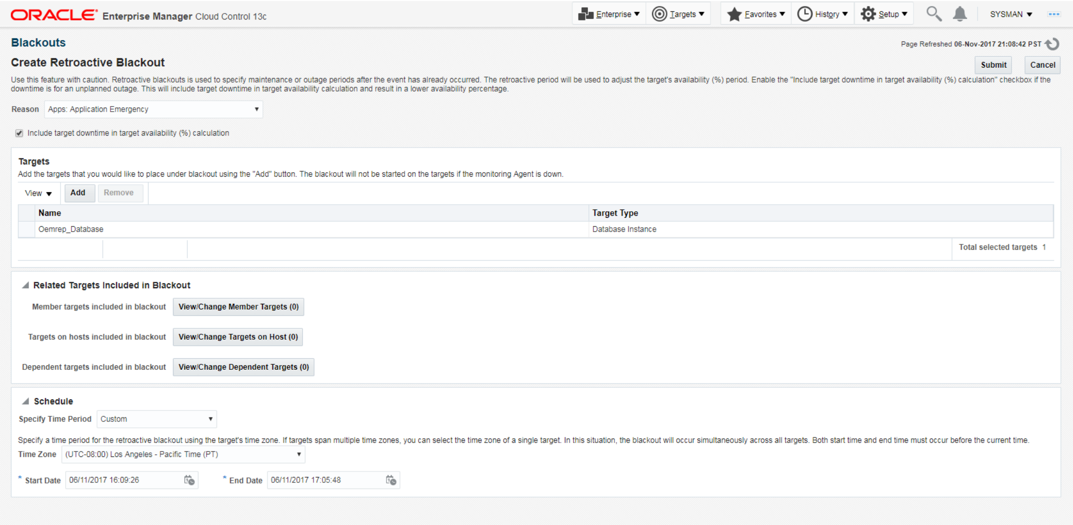
ORACLE Enterprise Manager Cloud Control 13c
Enterprise ▼
Targets ▼
Favorites ▼
History ▼
Setup ▼
SYSMAN ▼
Submit
Cancel
View ▼
Add
Remove
View/Change Member Targets (0)
View/Change Targets on Host (0)
View/Change Dependent Targets (0)

Blackouts

Page Refreshed 06-Nov-2017 21:08:42 PST

Create Retroactive Blackout

Use this feature with caution. Retroactive blackouts is used to specify maintenance or outage periods after the event has already occurred. The retroactive period will be used to adjust the target's availability (%) period. Enable the "Include target downtime in target availability (%) calculation" checkbox if the downtime is for an unplanned outage. This will include target downtime in target availability calculation and result in a lower availability percentage.

Reason: Apps: Application Emergency

☑ Include target downtime in target availability (%) calculation

Targets

Add the targets that you would like to place under blackout using the "Add" button. The blackout will not be started on the targets if the monitoring Agent is down.

Name	Target Type
Oemrep_Database	Database Instance

Total selected targets 1

Related Targets Included in Blackout

Member targets included in blackout

Targets on hosts included in blackout

Dependent targets included in blackout

Schedule

Specify Time Period: Custom

Specify a time period for the retroactive blackout using the target's time zone. If targets span multiple time zones, you can select the time zone of a single target. In this situation, the blackout will occur simultaneously across all targets. Both start time and end time must occur before the current time.

Time Zone: (UTC-08:00) Los Angeles - Pacific Time (PT)

* Start Date: 06/11/2017 16:09:26 * End Date: 06/11/2017 17:05:48

4. Ensure that the **Include target downtime in target availability (%) calculation** option is checked.

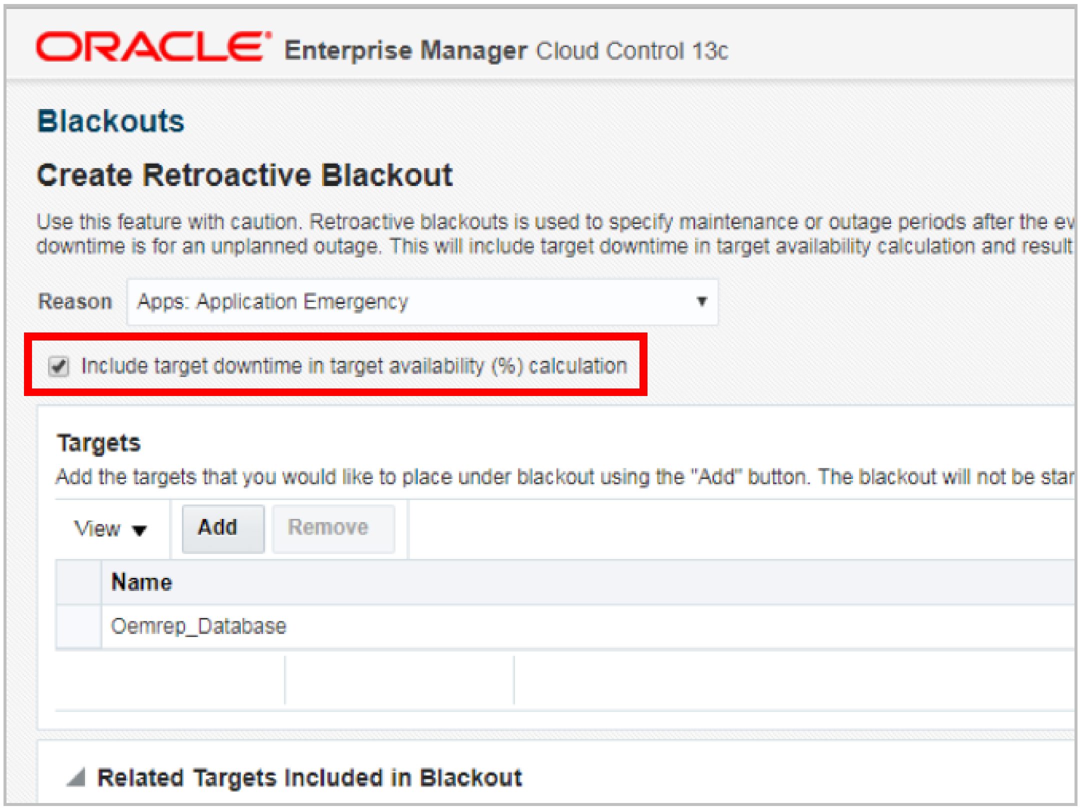

ORACLE Enterprise Manager Cloud Control 13c

Blackouts

Create Retroactive Blackout

Use this feature with caution. Retroactive blackouts is used to specify maintenance or outage periods after the ev
downtime is for an unplanned outage. This will include target downtime in target availability calculation and result

Reason: Apps: Application Emergency

☑ Include target downtime in target availability (%) calculation

Targets

Add the targets that you would like to place under blackout using the "Add" button. The blackout will not be star

Name
Oemrep_Database

◢ Related Targets Included in Blackout

Note:

The *Include target downtime in target availability (%) calculation* checkbox should be enabled during the creation of a retroactive blackout/outage, only if Enterprise Manager **did not** detect the outage. Checking this option will include target downtime in the target availability calculation and will result in a ***lower*** availability percentage.

5. Select a reason from the **Reason** drop-down menu and then click **Add** to add the target(s) for which you are creating the retroactive outage.
6. In the **Schedule** region, specify the time period in which the target was down.
7. Click **Submit** to create the retroactive outage. A success confirmation message will be displayed on the Blackouts page.

The following graphic shows target availability history for the repository database before the retroactive outage has been defined. Target availability percentage is 100% with no down time.

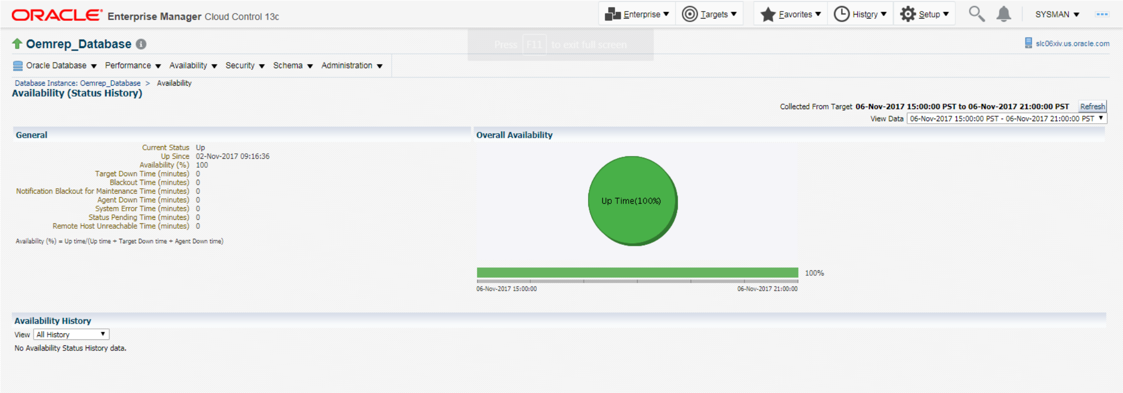

Availability (Status History)

Collected From Target **06-Nov-2017 15:00:00 PST to 06-Nov-2017 21:00:00 PST**

View Data 06-Nov-2017 15:00:00 PST - 06-Nov-2017 21:00:00 PST

General

Current Status	Up
Up Since	02-Nov-2017 09:16:36
Availability (%)	100
Target Down Time (minutes)	0
Blackout Time (minutes)	0
Notification Blackout for Maintenance Time (minutes)	0
Agent Down Time (minutes)	0
System Error Time (minutes)	0
Status Pending Time (minutes)	0
Remote Host Unreachable Time (minutes)	0

Availability (%) = Up time/(Up time + Target Down time + Agent Down time)

View All History

No Availability Status History data.

The next graphic shows the target availability history after a retroactive outage of 56 minutes has been defined. Target availability percentage has been reduced to 84% with 56 minutes of target down time.

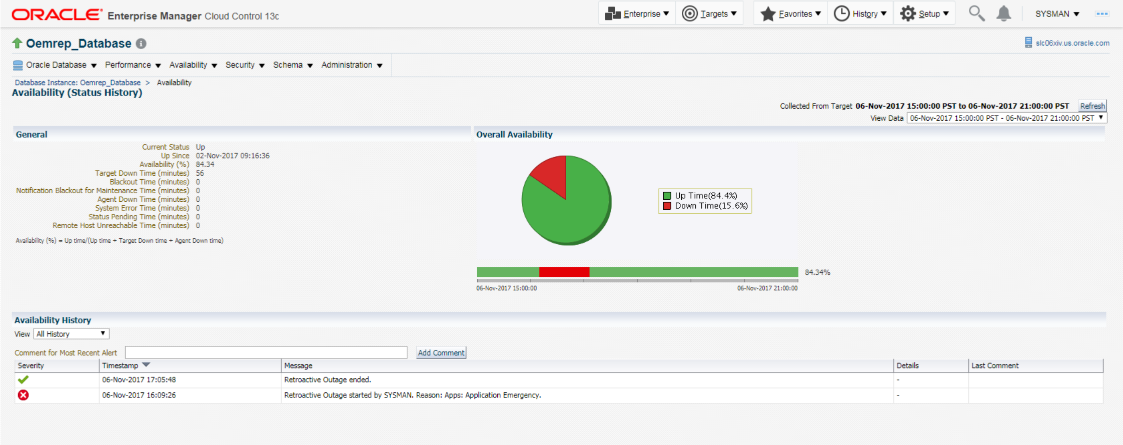

Availability (Status History)

Collected From Target **06-Nov-2017 15:00:00 PST to 06-Nov-2017 21:00:00 PST**

View Data 06-Nov-2017 15:00:00 PST - 06-Nov-2017 21:00:00 PST

General

Current Status	Up
Up Since	02-Nov-2017 09:16:36
Availability (%)	84.34
Target Down Time (minutes)	56
Blackout Time (minutes)	0
Notification Blackout for Maintenance Time (minutes)	0
Agent Down Time (minutes)	0
System Error Time (minutes)	0
Status Pending Time (minutes)	0
Remote Host Unreachable Time (minutes)	0

Availability (%) = Up time/(Up time + Target Down time + Agent Down time)

View All History

Comment for Most Recent Alert

Severity	Timestamp	Message	Details	Last Comment
✓	06-Nov-2017 17:05:48	Retroactive Outage ended.	-	
⊗	06-Nov-2017 16:09:26	Retroactive Outage started by SYSMAN. Reason: Apps: Application Emergency.	-	

Creating a Retroactive Outage using the Command Line Interface (EM CLI)

Alternatively, you can create a retroactive blackout/outage using the EM CLI `create_rbk` verb. The verb syntax is as follows:

```
emcli create_rbk
          -add_targets="name1:type1;name2:type2;..."...
          -reason="reason"
          [-propagate_targets]
          -schedule=
              start_time:<yyyy-MM-dd HH:mm>;
              end_time:<yyyy-MM-dd HH:mm>;
              [tzregion:<...>]
          [-outage]
```

Table 8-1 create_rbk Options

Option	Description
add_targets	Targets to add to the blackout, each specified as `target_name:target_type` The `add_targets` option may be specified more than once.
reason	Reason for the retroactive blackout. This is used for storing in backup tables.
propagate_targets	When this option is specified, a blackout for a target of type *host* applies the blackout to all non-agent targets on that host. Regardless of whether this option is specified, a blackout for a target that is a composite or a group applies the blackout to all members of the composite or group.
schedule	Blackout schedule. **Parameters** • start_time The start date/time of the blackout. The format of the value is `yyyy-MM-dd HH:mm:ss`. Example: `2017-09-20 12:12:12` • end_time The end date/time of the blackout. The format of the value is `yyyy-MM-dd HH:mm:ss`. Example: `2017-09-20 12:15:00` • tzregion The timezone region to use. If not provided *tzregion* is defaulted to UTC.
outage	Use this option with caution as it will lower the target availability (%). This option should be used only if Enterprise Manager did not detect the outage.

The following example shows the command output.

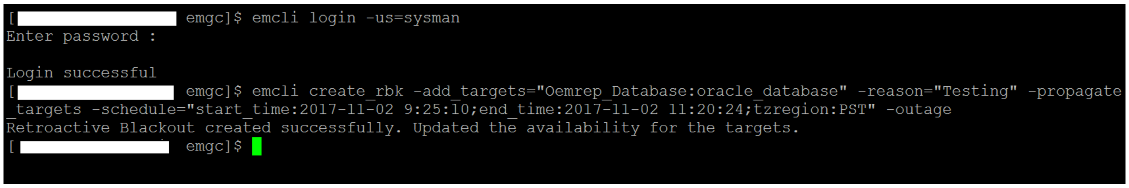

```
[[redacted] emgc]$ emcli login -us=sysman
Enter password :

Login successful
[[redacted] emgc]$ emcli create_rbk -add_targets="Oemrep_Database:oracle_database" -reason="Testing" -propagate
_targets -schedule="start_time:2017-11-02 9:25:10;end_time:2017-11-02 11:20:24;tzregion:PST" -outage
Retroactive Blackout created successfully. Updated the availability for the targets.
[[redacted] emgc]$
```

General Usage Guidelines

- For planned outages, where the administrator forgot to set a blackout, create a retroactive blackout **without enabling** the *Include target downtime in target availability (%) calculation* checkbox. This will **increase** the target's availability %. For example, 84% ==> 100%.
- For unplanned outages, where Enterprise Manager did not detect the outage, create a retroactive blackout and ***enable*** the *Include target downtime in target availability (%) calculation* option. This will **decrease** the target's availability %. For example. 100% ==> 84%.

- For unplanned outages, where Enterprise Manager did detect the outage, nothing needs to be done.

Exclude Targets or Target Types During a Blackout

When creating a blackout on composite target, all members of the composite target, by default, are part of the blackout. You can exclude large numbers of targets or target types from the blackout by using the EMCLI `create_blackout` command to create your blackout.

Under certain circumstances, when blacking out composite targets such as WebLogic domain or eBusiness, where the target type contains associations with its constituent components, such as the database, you may want to exclude specific components from the blackout. This means when the WebLogic domain is blacked out, then the database is included as an indirect member of the blackout. While you can manually exclude indirect members using the UI, this can be impractical for large scale environments. Using EMCLI lets you easily exclude indirect members of composite targets using a command line tool.

For example, if a Weblogic system is blacked out, the associated database is included in the blackout. If that database is used by another system as well, then the implication is that there is a state change in the other system using the database. In this situation, you would want to exclude this database from the blackout.

Using the EMCLI `create_blackout` command, you can exclude composite target components by using the following verb options:

- *exclude_targets*: A list of member targets of the direct blackout members can be specified. These indirect members of the blackout and their members will not be part of the blackout. For example, specifying a database system target will exclude that target and the corresponding database instance from the blackout if it would otherwise be an indirect member of the blackout.
- *exclude_types*: A list of target types can be specified. Indirect members of that type and their members will not be part of the blackout. For example, specifying `oracle_dbsys` will exclude database systems and their members which would be otherwise indirect members of the blackout.
  `exclude_targets` and `exclude_types` can be used in combination.

Example 1: The following example creates a blackout on a WebLogic domain, but excludes the database system and its member targets.

```
emcli create_blackout
      -name="wlblkout"
      -add_targets="Weblogic_domain:weblogic_domain"
      -exclude_types="oracle_dbsys"
      -schedule="duration::30"
      -reason="good reason1"
```

Example 2: The following example creates a blackout on a group which contains hundreds of WebLogic domains. The blackout excludes database systems and its member targets (e.g. Oracle home, Listener, Database instance).

```
emcli create_blackout
      -name=Group_Blackout
      -add_targets="Weblogic_Domain_Group:group"
      -exclude_types=oracle_dbsys
```

```
-schedule="duration:1:30"
-reason="WebLogic Domain Maintenance"
```

For more information about the EMCLI or the `create_blackout` verb, see create_blackout in *Oracle® Enterprise Manager Command Line Interface*.

Controlling Blackouts Using the Command Line Utility

You can control blackouts from the Oracle Enterprise Manager 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 Enterprise Manager console. Similarly, if you are controlling target blackouts from theEnterprise Manager console, do not attempt to control those blackouts from the command line.

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 Enterprise Manager 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_INSTANCE_HOME\bin` directory (Windows).
2. Enter the appropriate command as described in Table 8-2.

> **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 Administering Enterprise Manager Using EMCTL Commands.

Table 8-2 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
  ```

About Blackouts Best Effort

The Blackouts Best Effort feature allows you to create blackouts on aggregate targets, such as groups or systems, for which you do not have Blackout Target (or Higher) privileges on all members of the aggregate target.

Here, an Enterprise Manager administrator has Blackout Target privilege on an aggregate target but do not have OPERATOR privilege on its member/associated targets. You should ideally create a Full Blackout on this aggregate target. When defining the blackout, you are allowed to select any member target, even those member targets for which you have no Blackout Target privileges.

When the blackout actually starts, Enterprise Manager checks privileges on each member target and only blackout those on which you have Blackout Target(or Higher) privileges. This automated privilege check and target blackout selection is Enterprise Manager's "best effort" at blacking out the aggregate target.

When to Use Blackout Best Effort

The Blackout Best Effort functionality is targeted towards the creation of blackouts on targets of any aggregate type, such as Group, Hosts, Application Servers, Web Applications, Redundancy Groups, or Systems.

All targets the blackout creator has Blackout Target (or higher) privilege on will be displayed in the first step of Create/Edit Blackout Wizard. Once the blackout creator selects an aggregate type of target to be included in the Blackout Definition, this Blackout is "Full Blackout" by default.

The creator has the option of choosing the Blackout to run on “All Current" or “Selected" Targets, by selecting the appropriate values from the List box. Only when the "Full Blackout" option is chosen, will Blackout Best Effort affect targets for which the creator does not have Blackout Target (or higher) privileges.

Example Use Case

Consider 3 targets T1,T2 and T3 (all databases). A Group G1 contains all these 3 targets.

User U1 has OPERATOR privilege on T1,T2 and G1. User U1 has VIEW privilege on T3.

User U1 creates a scheduled full blackout on target G1. Scheduled implies that the blackout will start at a later point in time.

At the time of blackout creation, the tip text *Needs Blackout Target privilege, see Tip below the table* would be shown beside target T3.

When this blackout starts, if by that time User U1 has been granted OPERATOR privileges on target T3, then target T3 would also be under blackout. Otherwise only targets T1, T2 and G1 will be under blackout.

9

Managing Groups

This chapter introduces the concept of group management and contains the following sections:

Introduction to Groups

Groups are an efficient way to logically organize, manage, and monitor the targets in your global environments. Each group has its own group home page. The group home page shows the most important information for the group and enables you to drill down for more information. The home page shows the overall status of the group and other information such as current availability, incidents, and patch recommendations for members of the group.

Group Management Tasks

You can use Enterprise Manager to perform the following group management functions:

- Edit the configuration of a selected group, remove groups, and, in the case of an Administration Group, associate or disassociate a Template Collection.
- View the status and health of the group from the System Dashboard
- Drill down from a specific group to collectively monitor and manage its member targets.
- View a roll-up of member statuses and open incidents for members of the group.
- Apply blackouts to all targets in a group.
- Run jobs against a group
- Run a report
- Apply monitoring templates
- Associate compliance standards

In addition to creating groups, you can also create specific types of groups, such as redundancy groups, privilege propagating groups, and dynamic groups. The following sections explain the different types of groups.

Overview of Groups

Groups enable you to collectively monitor and administer many targets as a single logical unit. For example, you can define a group to contain all the databases serving an enterprise application, and define another group to contain all the hosts in a host farm. You can then use these groups to perform administrative operations. To create a group, you can manually select and add the members of the group. If you add an aggregate target, such as a Cluster Database, all of its member targets are automatically added to the group.

A group can include targets of the same type, such as all your production databases, or it could include all the targets on a host which would be comprised of different target types. You can nest static groups inside each other. In the target selector when you are selecting group members, choose Group as the target type, or choose a parent group as part of the process of creating a group.

If a system target is added to a group, it automatically pulls in its member targets. This could be the case of a regular group where a system such as WebLogic Server is added and also pulls in its members, or in a dynamic group where you specify a Target type to be an Oracle WebLogic Server and it also pulls in members of the WebLogic Server even though it does not match the dynamic group criteria. In this scenario, the group operations (for example, runnning jobs, blackouts, and so on) apply to all members of the group.

> **Note:**
>
> Because the Enterprise Manager Repository is a member of the OMS and the Enterprise Manager Repository is a RAC database, the ASMs and listeners are added as group members as well.
>
> You can also check the member relationships of any target by navigating to the target home page. From the <target> menu, select Members—>Topology—>View: System Members

After you configure a group, you can perform various administrative operations, such as:

- View a summary status of the targets within the group.
- View a roll-up of member statuses and open incidents for members of the group.
- View a summary of critical patch advisories.
- View configuration changes during the past 7 days.
- Create jobs and view the status of job executions.
- Create blackouts and view the status of current blackouts.

Overview of Privilege Propagating Groups

Privilege propagating groups enable administrators to propagate privileges to members of a group. You can grant a privilege on a group once to an administrator or a role and have that same privilege automatically propagate to any new member of the group. For example, granting *operator* privilege on a privilege propagating 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. Any aggregate that you add to a privilege propagating group must also be privilege propagating as well. For example, any group that you add to a privilege propagating group must also be privilege propagating.

Privileges on the group can be granted to an Enterprise Manager administrator or a role. Use a role if the privileges you want to grant are to be granted to a group of Enterprise Manager administrators.

For example, suppose you create a 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.

Overview of Dynamic Groups

The membership management for groups is typically manual or static in nature. Manually managing memberships works well for small deployments but not necessarily in large, dynamic environments where new targets come into the system frequently. Groups whose members are added frequently would be easier to maintain if they were to be defined by membership criteria instead of adding targets directly into the group. When the membership criteria is defined once, Enterprise Manager will automatically add targets.

A dynamic group is a group whose membership is determined by membership criteria. The owner of a dynamic group specifies the membership criteria during dynamic group creation (or modification) and membership in the group is determined solely by the criteria specified. Membership in a dynamic group cannot be modified directly because targets cannot be directly added to a dynamic group. Enterprise Manager automatically adds targets that match membership criteria when a dynamic group is created. It also updates group membership as new targets are added or target properties are changed and the target matches the group's membership criteria.

It is important to note that static groups can contain dynamic groups as members but not the other way around. You cannot include a static group as a member of a dynamic group.

Use the Define Membership Criteria function of Dynamic Groups to define the criteria for group membership. Once you have defined criteria, the targets selected by the criteria will be displayed in a read-only table in the Members region of the Groups page. Since dynamic groups are defined by criteria, you can intentionally or unintentionally define criteria that could result in very large groups.

The following requirements apply to dynamic groups:

- Dynamic groups cannot contain static groups, other dynamic groups, or administration groups.
- Administration groups cannot contain dynamic groups, however, a static group can contain dynamic groups as a member.
- OR-based criteria is not supported. All criteria selected on the criteria page are AND-based.
- Supported properties are global properties, user-defined properties, and other attributes specifically supported for groups such as Version, Platform, Target Name, and Type. Other instance properties and config data elements are not supported as membership criteria.
- The View Any Target and Add Any Target privileges are required to create a dynamic group.
- The Full Any Target, Add Any Target, and Create Privilege Propagating Group privileges are required to create a privilege-propagating dynamic group.

Overview of Administration Groups

Administration Groups greatly simplify the process of setting up targets for management in Enterprise Manager by automating the application of management settings such as monitoring settings or compliance standards. Typically, these settings are manually applied to individual target, or perhaps semi-automatically using custom scripts. However, by defining Administration Groups, Enterprise Manager uses specific target properties to direct the target to the appropriate Administration Group and then automatically apply the requisite monitoring and management settings. This level of automation simplifies the target setup process and also enables a datacenter to easily scale as new targets are added to Enterprise Manager for management.

Administration groups are a special type of group used to fully automate application of monitoring and other management settings upon joining the group. When a target is added to the group, Enterprise Manager applies these settings using a Template Collection consisting of Monitoring Templates, compliance standards, and Enterprise Manager policies. This completely eliminates the need for administrator intervention.

To watch Part 1 of a video about using administrative groups and template collections, click here.

To watch Part 2 of the video about using administrative groups and template collections, click here.

Choosing Which Type of Group To Use

There are two major types of groups you can choose to manage targets: Static Groups/ Dynamic groups, which can be one or more groups that you define, and Administration Groups for automating monitoring setup using templates.

You should carefully consider the purpose of your group and the function it serves before determining which type of group to use.

The following table diagrams when you should use Administration Groups or Dynamic Groups.

Table 9-1 When To Use Administration Groups vs. Dynamic Groups

Type of Group	Main Purpose	Membership Based on Criteria	Additional Membership Requirements	Privilege Propagating
Administration Group	Auto-apply monitoring templates	Yes, based on target properties	Target can belong to at most one administration group	Yes (always)
Dynamic Group	Perform any group operation.	Yes, based on target properties	Target can belong to one or more groups	User-specified option

The main purpose of an Administration Group is to automate the application of management settings, such as monitoring settings or compliance standards. When a target is added to the group, Enterprise Manager automatically applies these settings using templates to eliminate the need for administrator action.

Dynamic groups, on the other hand, can be used to manage many targets as a single unit where you can define the group membership by defining the properties that constitute the group. For example, you could use dynamic groups to manage privileges or groups that you create containing the targets that are managed for different support teams.

Managing Groups

By combining targets in a group, Enterprise Manager provides management features that enable you to efficiently manage these targets as one group. Using the Group functionality, you can:

- View a summary status of the targets within the group.
- Monitor 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.

When you choose Groups from the Targets menu in the Enterprise Manager, the Groups page appears. 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 member status summary and rollup of incidents for members in a group.
- Create Groups, Dynamic Groups or the Administration Group hierarchy, edit the configuration of a selected group, remove groups, and, in the case of an Administration Group, associate or disassociate a Template Collection.
- 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.
- Customize the homepage of a specific group

Redundancy systems and special high availability groups are not accessed from this Groups page. You can access them from the All Targets page or you can access Redundancy Systems and other systems from the Systems page.

Creating and Editing Groups

Enterprise Manager Groups enable administrators to logically organize distributed targets for efficient and effective management and monitoring.

To create a group, follow these steps:

1. From the Enterprise Manager Console, choose **Targets** then choose **Groups**. Alternately, you can choose **Add Target** from the **Setup** menu and choose the menu option to add the specific type of group.
2. Click **Create** and choose the type of group you want to create. The Enterprise Manager Console displays a set of Create Group pages that function similarly to a wizard.
3. On the General tab of the Create Group page, enter the **Name** of the Group you want to create. If you want to make this a privilege propagating group, then enable the Privilege Propagation option by clicking **Enabled**. If you enable Privilege Propagation for the group, the target privileges granted on the group to an administrator or a role are propagated to the member targets. As with regular groups with privilege propagation, the Create Privilege Propagating Group privilege is required for creation of privilege propagating dynamic groups. In addition, the Full any Target privilege is required to enable privilege propagation because only targets on which the owner has Full Target privileges can be members, and any target can potentially match the criteria and a system wide Full privilege is required. To create a regular dynamic group, the View any Target system wide privilege is required as the group owner must be able to view any target that can potentially match the membership criteria.
4. Configure each page, then click **OK**. You should configure all the pages before clicking **OK**. For more information about these steps, see the online help.

After you create the group, you always have immediate access to it from the Groups page.

You can edit a group to change the targets that comprise the group, or change the metrics that you want to use to summarize a given target type. To edit a group, follow these steps:

1. From the Enterprise Manager Console, choose **Targets** then choose **Groups**.
2. Click the group **Name** for the group you want to edit.
3. Click **Edit** from the top of the groups table.
4. Change the configuration for a page or pages, then click **OK**.

Creating Dynamic Groups

The owner of a dynamic group specifies the membership criteria during dynamic group creation (or modification) and membership in the group is determined solely by the criteria specified. Membership in a dynamic group cannot be modified directly. Enterprise Manager automatically adds targets that match membership criteria when a dynamic group is created. It also updates group membership as new targets are added or target properties are changed and the targets match the group's membership criteria.

To create a dynamic group, follow these steps:

1. From the Groups page, click **Create** and then select **Dynamic Group** from the drop-down list. Alternately, you can choose **Add Target** from the **Setup** menu and then select **Group**.
2. On the General tab of the Create Dynamic Group page, enter the **Name** of the Dynamic Group you want to create. If you want to make this a privilege propagating dynamic group, then enable the Privilege Propagation option by clicking **Enabled**. If you enable Privilege Propagation for the group, the target privileges granted on the group to an administrator or a role are propagated to the member targets. As with regular groups with privilege propagation, the Create Privilege Propagating Group privilege is required for creation of privilege propagating dynamic groups. In addition, the Full any Target privilege is required to enable privilege propagation because only targets on which the owner has Full Target privileges can be members, and any target can potentially match the criteria and a system wide Full privilege is required. To create a regular dynamic group, the View any Target system wide privilege is required as the group owner must be able to view any target that can potentially match the membership criteria.

   The privilege propagating group feature contains two 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

     Grant this privilege to an administrator or role that enables him to become group administrator for the group. This means he can perform operations on the group, share privileges on the group with other administrators, etc.

     The Group Administration Privilege is available for both Privilege Propagating Groups and conventional groups. If you are granted this privilege, you can grant full privilege access to the group to other Enterprise Manager users without having to be the SuperAdministrator to grant the privilege.

3. In the Define Membership Criteria section, define the criteria for the dynamic group membership by clicking **Define Membership Criteria**.

   The Define Membership Criteria page appears where you can Add or Remove properties of targets to be included in the group. Group members must match one value in each of the populated target properties. Use the Member Preview section to review a list of targets that match the criteria. Click **OK** to return to the General page.

At least one of the criteria on the Define Membership Criteria page must be specified. You cannot create a Dynamic group without at least one of the target types, on hosts or target properties specified. Use the following criteria for dynamic groups:

- Target type(s)
- Department
- On Host
- Target Version
- Lifecycle Status
- Operating System
- Line of Business
- Platform
- Location
- CSI
- Cost Center
- Contact
- Comment

You can add or remove properties using the **Add** or **Remove** Target Properties button on the Define Membership Criteria page.

The **Indirect Members** drop-down menu provides the option to include or exclude any indirect member targets as part of the dynamic group. Indirect members are targets whose target properties DO NOT match dynamic group criteria. However, they have been added to the dynamic group because their parent target is a direct member of the dynamic group. Whenever aggregate targets, such as targets that have member targets, are added to a dynamic group, all members of the aggregate target are, by default, also added to the group. An example of an aggregate target is Oracle WebLogic Server. If that target is added to a group, then all Application Deployment targets on it are also pulled into the group, even if their properties do not satisfy the group criteria.

4. Enter the **Time Zone**. The time zone you select is used for scheduling operations such as jobs and blackouts on this group. The groups statistics charts will also use this time zone.
5. Click the **Charts** tab. Specify the charts that will be shown in the Dynamic Group Charts page. By default, the commonly used charts for the target types contained in the Dynamic Group are added.
6. Click the **Columns** tab to add columns and abbreviations that will be seen in the Members page and also in the Dashboard.
7. Click the **Dashboard** tab to specify the parameters for the System Dashboard. The System Dashboard displays the current status and incidents and compliance violations associated with the members of the Dynamic Group in graphical format.
8. Click the **Access** tab. Use the Access page to administer access privileges for the group. On the Access page you can grant target access to Enterprise Manager roles and grant target access to Enterprise Manager administrators.
9. Click **OK** to create the Dynamic Group.

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-aggregate target as the member of a privilege propagating group. For aggregate targets in Cloud Control version 12*c*, cluster and RAC databases and other propagating groups can be added as members. Cloud Control version 12*c* supports more aggregate target types, such as redundancy systems, systems and services.

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.

Converting Conventional Groups to Privilege Propagating Groups

In Enterprise Manager release 12*c* you can convert conventional groups to privilege propagating groups (and vice-versa) through the use of the specified EM CLI verb. Two new parameters have been added in the *modify_group* EM CLI 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 EM CLI verbs: *modify_system*, *modify_redundancy_group*, and *modify_aggregrate_service*.

The EM CLI 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 EM CLI verbs, see the *EM CLI Reference Manual*.

Viewing and Managing Groups

Enterprise Manager 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 Group Home page provides the following sections:

- A General section that displays the general information about the group, such as the Owner, Group Type, and whether the group is privilege propagating. You can drill down to the Edit Group page to enable or disable privilege propagating by clicking on the Privilege Propagating field.
- 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.
- An Overview of Incidents and Problems section that displays the summary of incidents on members of the group that have been updated in the recent period of time. It also shows a count of open problems as well as problems updated in recent period of time.

  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 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 the compliance of members of the group against the compliance standards defined for the group. This section also shows a rollup of violations by severity (critical, warning, minor warning) as well as the average compliance score(%).
- A Job Activity section that displays a summary of jobs for the targets in the group whose start date is within the last 7 days. You can click Show to see the latest run or all runs. Click View to select and reorder the columns that appear in the table or to adjust scrolling and expanding the table.
- A Blackouts section that displays information about current or pending blackouts. You can also create a blackout from this section.
- A Patch Recommendations section that displays the Oracle patch recommendations that are applicable to your enterprise. You can view patch recommendations by classification or target type.

  You can navigate to My Oracle Support to view all recommendations by clicking the All Recommendations link.
- 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 also view inventory summary information in the context of different dimensions. From here you can click See Details to display the Inventory and Usage page.
- A Configuration Changes section that displays the number of configuration changes to the group in the previous 7 days. You can click the number to display a page that displays detailed information about the changes. Enterprise Manager automatically collects configuration information for group targets and changes to configurations are recorded and may be viewed from that page.

Viewing a Group

To view a group, follow these steps:

1. From the Enterprise Manager Console, choose **Targets** then choose **Groups**. A summary table lists all defined groups.
2. Click the desired group to go to the Home page of that group.

You can use View By filters (located in the upper right corner of the home page) to change the view of the homepage to members of targets of a specific type. When you do this, the Group homepage refreshes to only show information for targets of that type. Additional regions of interest might display. For example, DBAs might switch to the Database filter to view information specifically on Database targets in the group.

You can also personalize the home page by clicking the Actions icon in the upper right corner of each region on the home page to move that region up or down on the page. You can also expand or contract a region by clicking the arrow icon in the upper left corner of each region.

You can also navigate to other management operations on the group using the Group menu. For example, you can view all the members in a group by choosing **Member** from the Group menu. Likewise you can view the **Membership History** of the group by choosing Membership History from the Group menu.

Overview of Group Charts

Group Charts enable you to monitor the collective performance of a 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.

Overview of Group Members

Enterprise Manager allows you to summarize information about the member targets in a group. It provides information on their current availability status, roll-up of open incidents and compliance violations, 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 rank members by a certain criterion (for example, database targets in order of decreasing wait time percentage). You can display default key performance metrics based on the targets you select, but you can customize these to include additional metrics that are important for managing your group.

You can view the members of a group by choosing **Members** from the Group menu. Enterprise Manager displays the Members page where you can view the table of members filtered by All Members, Direct Members, or Indirect Members. Direct members are targets directly added to the group. Indirect members are targets that are members of a direct member target, and are automatically included into the group because their parent target was added to the group. The page provides the option to **Export** or **Edit** the group.

You can also access information about membership history by choosing **Membership History** from the Group menu. The Membership History page displays changes in the group membership over time.

Viewing Group Status History

You can view Status History for a group to see the historical availability of a member during a specified time period or view the current status of all group members. You can access the

Status History page by choosing **Monitoring** from the Group menu and then selecting **Status History**.

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 of a member, you can click a Status icon on the View Group Status History page 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.

About the System Dashboard

The System Dashboard enables you to proactively monitor the status, incidents and compliance violations in the group as they occur. The color-coded interface is designed to highlight problem areas — 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 determine the problem areas for any target and drill down for details as needed. An incident table is also included to provide a summary for all open incidents in the group. The incidents in the table are presented in reverse chronological order to show the most recent incidents first, but you can also click any column in the table to change the sort order. The colors in top bar of the Member Targets table change based on the incident's critical level. The priority progresses from warning to critical to fatal. If the group has at least one fatal incident (irrespective of critical or warning incidents), the top bar becomes dark red. If the group has at least one critical incident (irrespective of warning incidents), the top bar becomes faint red. If the group has only warning incidents, the top bar turns yellow. If the group has no incidents, the top bar remains colorless.

The Dashboard auto-refreshes based on the Refresh Frequency you set on the Customize Dashboard page.

The Dashboard allows you to drill down for more detailed information. You can click 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
- Status icon for a metric with key values to access the metric page with a list of all key values
- Dashboard header to access the group home page
- Incidents and Problems table to view summary information about all incidents or specific categories of incidents.

Click **Customize** to access the Customize Dashboard page. This page allows you to change the refresh frequency and display options for the Member Targets table at the top of the dashboard. You can either show all individual targets or show by target type. There is also the option to expand or contract the Incidents and Problems table at the bottom. To change the columns shown in the Member Targets table, go to the Columns tab of the Edit Group page which you can access by choosing Target Setup from the Group menu.

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 incidents displayed are rolled up for the targets in that specific target type.

If you minimize the dashboard window, pertinent alert information associated with the group or system is still displayed in the Microsoft Windows toolbar.

You can use Information Publisher reports to make the System Dashboard available to non-Enterprise Manager users. First, create a report and include the System Monitoring Dashboard reporting element. In the report definition, choose the option, Allow viewing without logging in to Enterprise Manager. Once this is done, you can view it from the Enterprise Manager Information Publisher Reports website.

Using 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:**
>
> Using Information Publisher

10
Using Administration Groups

Administration groups greatly simplify the process of setting up targets for management in Enterprise Manager by automating the application of management settings such as monitoring settings or compliance standards. Typically, these settings are manually applied to individual target, or perhaps semi-automatically using custom scripts. However, by defining administration groups, Enterprise Manager uses specific target properties to direct the target to the appropriate administration group and then automatically apply the requisite monitoring and management settings. Any change to the monitoring setting will be automatically applied to the appropriate targets in the administration group. This level of automation simplifies the target setup process and also enables a datacenter to easily scale as new targets are added to Enterprise Manager for management.

This chapter covers the following topics:

What is an Administration Group?

Administration groups are a special type of group used to fully automate application of monitoring and other management settings targets upon joining the group. When a target is added to the group, Enterprise Manager applies these settings using a template collection consisting of monitoring templates, compliance standards, and Enterprise Manager policies. This completely eliminates the need for administrator intervention. The following illustration demonstrates the typical administration group workflow:

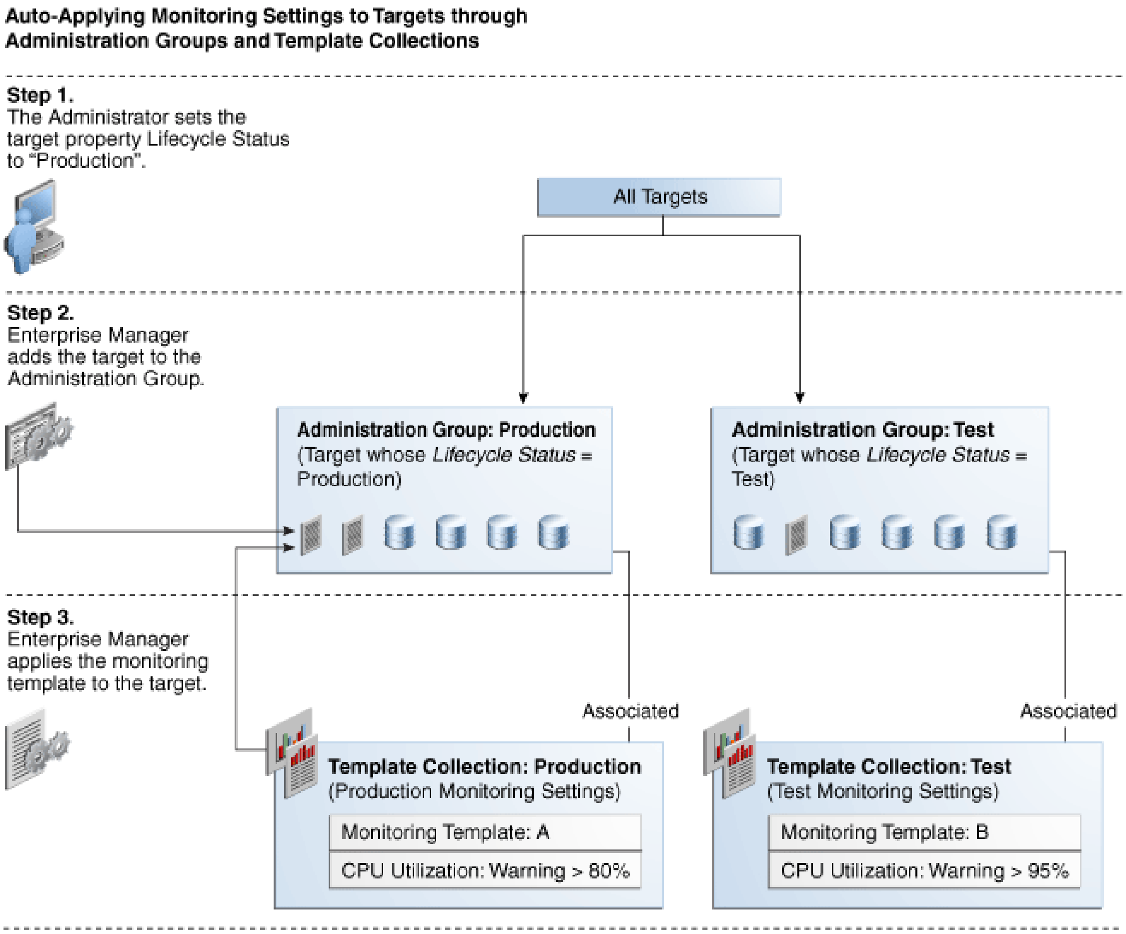
Auto-Applying Monitoring Settings to Targets through Administration Groups and Template Collections
Step 1.
The Administrator sets the target property Lifecycle Status to "Production".
All Targets
Step 2.
Enterprise Manager adds the target to the Administration Group.
Administration Group: Production
(Target whose Lifecycle Status = Production)
Administration Group: Test
(Target whose Lifecycle Status = Test)
Step 3.
Enterprise Manager applies the monitoring template to the target.
Associated
Associated
Template Collection: Production
(Production Monitoring Settings)
Monitoring Template: A
CPU Utilization: Warning > 80%
Template Collection: Test
(Test Monitoring Settings)
Monitoring Template: B
CPU Utilization: Warning > 95%

The first step involves setting a target's Lifecycle Status property when a target is first 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 prioritization consists of the following levels:

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

As shown in step two of the illustration, once Lifecycle Status is set, Enterprise Manger uses it to determine which administration group the target belongs.

In order to prevent different monitoring settings to be applied to the same target, administration groups were designed to be mutually exclusive with other administration groups in terms of group membership. Administration groups can also be used for hierarchically classifying targets in an organization, meaning a target can belong to at most one administration group. This also means you can only have one administration group hierarchy in your Enterprise Manager deployment.

For example, in the previous illustration, you have an administration group hierarchy consisting of two subgroups: *Production* targets and *Test* targets, with each subgroup having its own template collections. In this example, the Production group inherits monitoring settings from monitoring template A while targets in the Test subgroup inherit monitoring settings from monitoring template B.

Developing an Administration Group

In order to create an administration group, you must have both *Full Any Target* and *Create Privilege Propagating Group* target privileges.

Developing an administration group is performed in two phases:

- **Planning**
  - Plan your administration group hierarchy by creating a group hierarchy in a way reflects how you monitor your targets.
  - Plan the management settings associated with the administration groups in the hierarchy.
    - Management settings: Monitoring settings, Compliance standard settings, Enterprise Manager policy settings
    - For Monitoring settings, you can have additional metric settings or override metric settings lower in your hierarchy
    - For Compliance standards or Enterprise Manager 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, Enterprise Manager 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.

Planning an Administrative Group

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

You can only have one administration group hierarchy in your Enterprise Manager deployment, thus ensuring that administration group member targets can only directly belong to one administration group. This prevents monitoring conflicts from occurring as a result of having a target join multiple administration groups with different associated monitoring settings.

To define the hierarchy, you want to think about the highest (root) level as consisting of all targets that have been added to Enterprise Manager. Next, think about how you want to divide your targets along the lines of how they are monitored, where targets that are monitored in one way are in one group, and targets that are monitored in another way are part of another group. For example, Production targets might be monitored one way and Test targets might be monitored in another way. You can further divide individual groups if there are further differences in monitoring. For example, your Production targets might be further divided based on the line of business they support because they might have additional metrics that need to be monitored for that line of business. Eventually, you will end up with a hierarchy of groups under a root node.

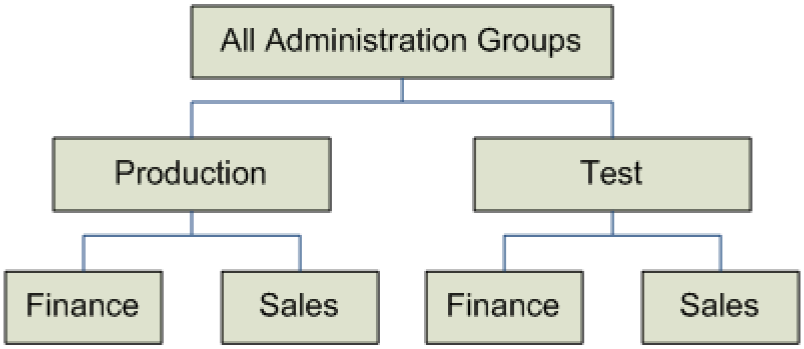

The attributes used to define each level of grouping and thus the administration group membership criteria are based on *global target properties* as well as *user-defined target properties*. These target properties are attributes of every target and specify operational information within the organization. For example, *location*, *line of business* to which it belongs, and *lifecycle status*. The global target properties that can be used in the definition of administration groups are:

- Lifecycle Status

> **Note:**
>
> Lifecycle Status target property is of particular importance because it denotes a target's operational status. Lifecycle Status can be any of the following: Mission Critical, Production, Staging, Test, or Development.

- Location
- Line of Business
- Department
- Cost Center
- Contact
- Platform
- Operating System
- Target Version
- Customer Support Identifier
- High Availability Role
- Target Type (Allowed but not a global target property.)

You can create custom user-defined target properties using the EM CLI verbs *add_target_property* and *set_target_property_value*. See Verb Reference in the *Oracle Enterprise Manager Command Line Interface Guide* for more information.

Enterprise Manager (EM) 13.5 Release Update 13 and later supports a user-defined ‘High Availability Role’ target property. This property can be used to distinguish a database's role: if it is a primary or a standby database. When a database role change occurs, Enterprise Manager automatically updates this user-defined target property value, if defined, with the new database role (primary or standby). This property can then be applied towards your Dynamic Groups or Administration Groups.

For example, you may have an Administration Group strictly for primary databases and a separate Administration Group strictly for standby databases. This is to ensure that you can send separate notifications for each of these groups to respective database administrators. Once a database role change occurs, if the ‘High Availability Role’ target property is included in the Administration Group hierarchy definition then EM automatically updates the value of this property to the new database role. Hence the respective database will automatically become a new member of the respective Administration Group matching the new criteria. The administrators of the respective Administration Group will thus be notified of any new corresponding alerts.

How to Use ‘High Availability Role’ Target Property:

1. Create a user-defined target property and label it ‘High Availability Role’.

2. Set the OMS property ‘oracle.sysman.db.utilizeHARoleTargetProperty’ to ‘true’ to enable the feature.

3. Update your master list of values for this property to explicitly include the values 'Primary' and 'Standby'.

4. Assign target property values, ‘Primary’ or ‘Standby’, to your desired database targets.

5. Create or update your Administration Group hierarchy definition to include the ‘High Availability Role’."

Example:

1. The user-defined target property, ‘High Availability Role’ has been created.
   a. It has the target property values of 'Primary' and 'Standby'.
2. There are two databases: DB1 (primary database) and DB2 (standby database) that are assigned the target property value that corresponds to their current role.
   a. DB1 (primary) ==> ‘High Availability Role = Primary’
   b. DB2 (standby) ==> ‘High Availability Role = Standby’
3. Database role change occurs (automatic)
   a. EM updates the target property values of the databases if they are now Primary or Standby
      i. DB1 (formerly primary, now standby) ==> 'High Availability Role' property changes from ‘Primary' to ‘Standby’
      ii. DB2 (formerly standby, now primary) ==> 'High Availability Role' property changes from ‘Standby' to ‘Primary’
4. EM automatically puts the new primary and new standby database in the appropriate groups

You cannot manually add targets to an administration group. Instead, you set the target properties of the target (prospective group member) to match the membership criteria defined for the administration group. Once the target properties are set, Enterprise Manager automatically adds the target to the appropriate administration group.

Target Properties Master List

To be used with administration group (and dynamic groups), target properties must be specified in a uniformly consistent way by all users in your managed environment. In addition, you may want to limit the list of target property values that can be defined for a given target property value. To accomplish this, Enterprise Manager lets you define a target properties master list. When a master list is defined for a specific target type, a drop-down menu containing the predefined property values appears in place of a text entry field on a target's target properties page. You use the following EM CLI verbs to manage the master properties list:

- **use_master_list**: Enable or disable a master list used for a specified target a property.
- **add_to_property_master_list**: Add target property values to the master list for a specified property.
- **delete_from_property_master_list**: Delete values from the master list for specified property.
- **list_property_values**: List the values for a property's master list
- **list_targets_having_property_value**: Lists all targets with the specified property value for this specified property name.
- **rename_targets_property_value**: Changes the value of a property for all targets.

For more information about the master list verbs, see the Verb Reference in the *Oracle Enterprise Manager Command Line Interface Guide* for more information.

> **Note:**
>
> You must have Super Administrator privileges in order to define/maintain the target properties master list.

Enterprise Manager Administrators and Target Properties

When creating an Enterprise Manager administrator, you can associate properties such as Contact, Location, and Description. However, there are additional resource allocation properties that can be associated with their profile. These properties are:

- Department
- Cost Center
- Line of Business

It is important to note that these properties are persistent--when associated with an administrator, the properties (which mirror, in part, the target properties listed above) are automatically passed to any targets that are discovered or created by the administrator.

Example

In the following administration group hierarchy, two administration groups are created under the node *Root Administration Group*, *Production* and *Test*, because monitoring settings for production targets will differ from the monitoring settings for test targets.

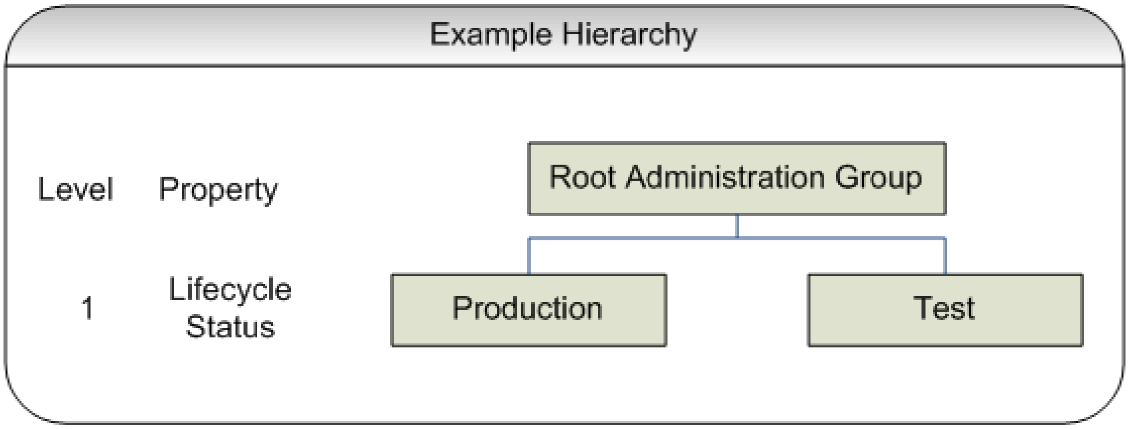

In this example, the group membership criteria are based on the *Lifecycle Status* target property. Targets whose *Lifecycle Status* is 'Production' join the Production group and targets whose *Lifecycle Status* is 'Test' join the Test group. For this reason, *Lifecycle Status* is the target property that determines the first level in the administration group hierarchy. The values of Lifecycle Status property determine the membership criteria of the administration groups in the first level: Production group has membership criteria of "Lifecycle Status = Production" and Test group has membership criteria of "Lifecycle Status = Test' membership criteria.

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.

The end result of this hierarchy planning exercise is summarized in the following table.

Root Level (First Row)	Level 1 target property (second row) Lifecycle Status	Level 2 target property (third row) Line of Business
Root Administration Group	Production or Mission Critical	Finance
_	_	Sales
_	Staging or Test or Development	Finance
_	_	Sales

Each cell of the table represents a group. The values in each cell represent the values of the target property that define membership criteria for the group.

It is possible to have the group membership criteria be based on more than one target property value. In that case, any target whose target property matches any of the values will be added to the group. For example, in the case of the Production group, if the *Lifecycle Status* of a target is either *Production* or *Mission Critical*, then it will be added to the Production group.

It is also important to remember that group membership criteria is cumulative. For example, for the *Finance* group under *Production or Mission Critical* group, a target must have its *Lifecycle Status* set to *Production or Mission Critical* **AND** its *Line of Business* set to *Finance* before it can join the group. If the target has its *Lifecycle Status* set to *Production* but does not have its *Line of Business* set to *Finance* or *Sales*, then it does not join any administration group.

For this planning example, the resulting administration group hierarchy would appear as shown in the following graphic.

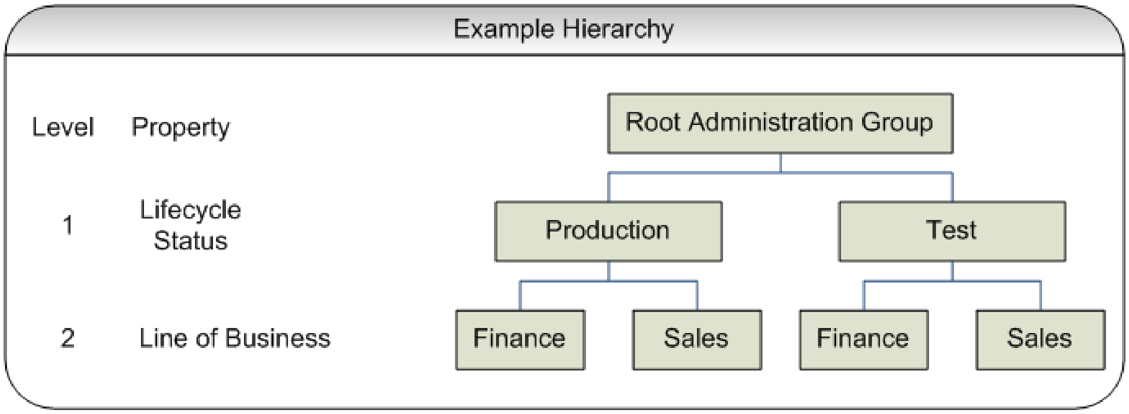

It is important to note that a target can become part of hierarchy if and only if its property values match criteria at both the levels. A target possessing matching values for *lifecycle status* cannot become member of the administration group at the first level. Also, all targets in the administration group hierarchy will belong to the lowest level groups.

Step 2: Assign Target Properties

After establishing the desired administration group 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. For targets that have already been added to Enterprise Manager, you can also set the target properties via the console or using the EM CLI verb *set_target_property_value*, See the *Oracle Enterprise Manager Command Line Interface Guide* for more information. Note that when running *set_target_property_value*, any prior values of the target property are overwritten. If you set target properties before hierarchy creation, it will join the group after it is created. The targets whose properties are set using EM CLI will automatically join their appropriate administration groups. Target properties can, however, be set after the administration group hierarchy is created.

For small numbers of targets, you can change target properties directly from the Enterprise Manager console.

1. From an Enterprise Manager target's option menu, select **Target Setup**, then select **Properties**.

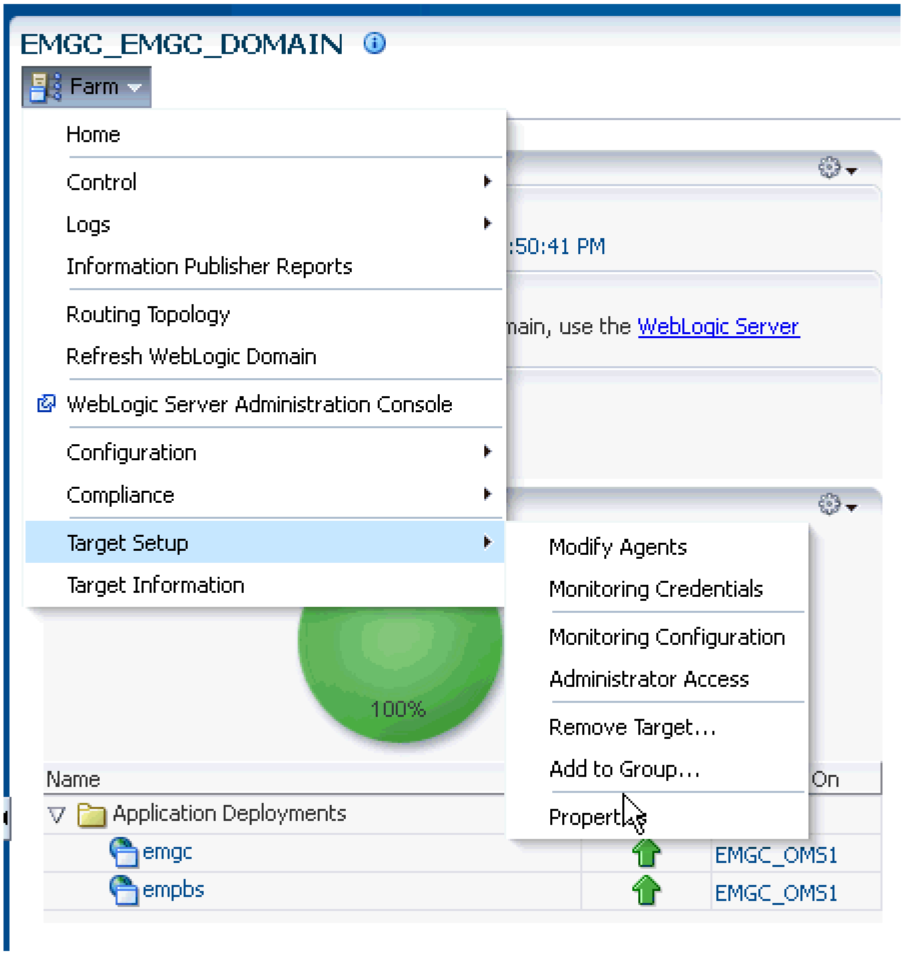

2. On the **Target Properties** page, click **Edit** to change the property values.

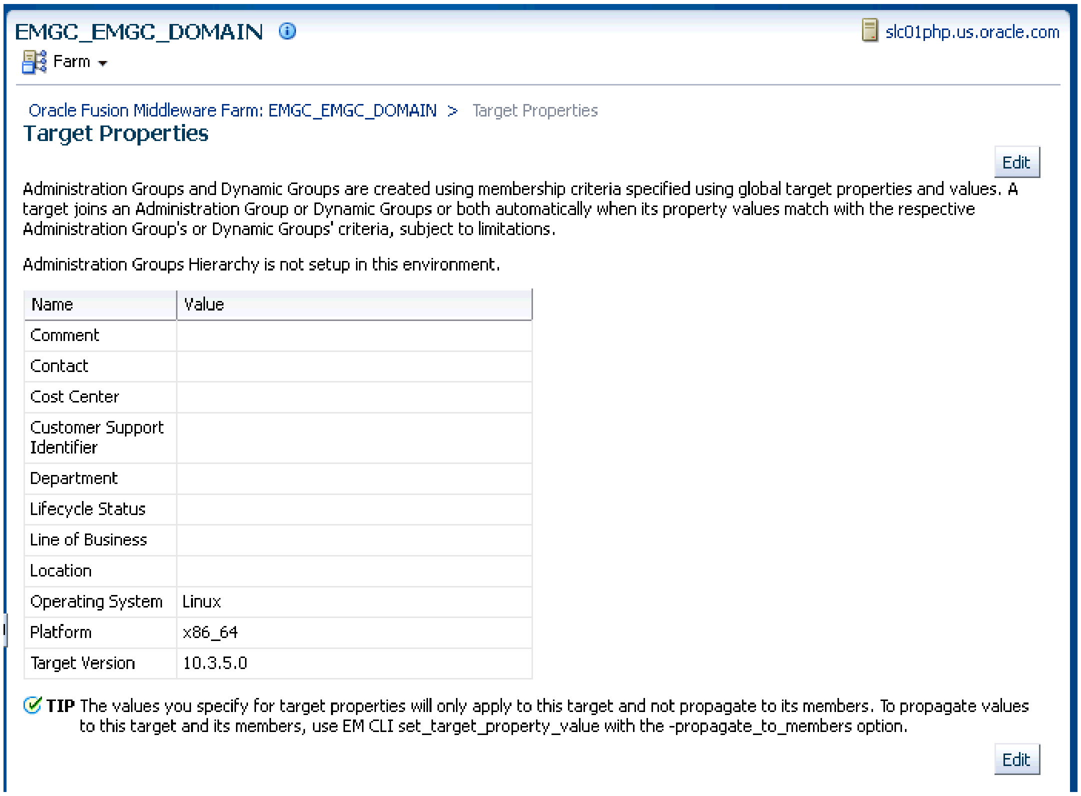
EMGC_EMGC_DOMAIN
Farm
slc01php.us.oracle.com
Oracle Fusion Middleware Farm: EMGC_EMGC_DOMAIN > Target Properties
Edit

Target Properties

Administration Groups and Dynamic Groups are created using membership criteria specified using global target properties and values. A target joins an Administration Group or Dynamic Groups or both automatically when its property values match with the respective Administration Group's or Dynamic Groups' criteria, subject to limitations.

Administration Groups Hierarchy is not setup in this environment.

Name	Value
Comment	
Contact	
Cost Center	
Customer Support Identifier	
Department	
Lifecycle Status	
Line of Business	
Location	
Operating System	Linux
Platform	x86_64
Target Version	10.3.5.0

✓ **TIP** The values you specify for target properties will only apply to this target and not propagate to its members. To propagate values to this target and its members, use EM CLI set_target_property_value with the -propagate_to_members option.

To help you specify the appropriate target property values used as administration group criteria, pay attention to the instructional verbiage at the top of the page.

3. Once you have set the target properties, click **OK**.

For large numbers of targets, it is best to use the Enterprise Manager Command Line Interface (EM CLI) `set_target_property_value` verb to perform a mass update. For more information about this EM CLI verb, see the Enterprise Manager Command Line Interface guide.

Administration groups are privilege-propagating: Any privilege that you grant on the administration group to a user (or role) automatically applies to all members of the administration group. For example, if you grant Operator privilege on the Production administration group to a user or role, then the user or role automatically has Operator privileges on all targets in the administration group. Because administration groups are always privilege propagating, any aggregate target that is added to an administration group must also be privilege propagating.

> **Note:**
>
> An aggregate target is a target containing other member targets. For example, a Cluster Database (RAC) is an aggregate target has RAC instances.

A good example of aggregate target is the Privilege Propagating Group. See "Managing Groups" for more information.

At any time, you can use the **All Targets** page to view properties across all targets. To view target properties:

1. From the **Targets** menu, select **All Targets** to display the All Targets page.
2. From the **View** menu, select **Columns**, then select **Show All**.
3. Alternatively, if you are interested in specific target properties, choose **Columns** and then select **Show More Columns**.

Step 3: Prepare for Creating Template Collections

Template collections contain the monitoring settings and other management settings that are meant to be applied to targets as they join the administration group. Monitoring settings for targets are defined in monitoring templates. Monitoring templates are defined on a per target type basis, so you will need to create monitoring templates for each of the different target types in your administration group. You will most likely create multiple monitoring templates to define the appropriate monitoring settings for an administration group. For example, you might create a database Monitoring template containing the metric settings for your production databases and a separate monitoring template containing the settings for your non-production databases. Other management settings that can be added to a template collection include Compliance Standards and Enterprise Manager Policies. Ensure all of these entities that you want to add to your template collection are correctly defined in Enterprise Manager before adding them to template collections.

If you have an administration group hierarchy defined with more than two levels, such as the hierarchy shown in the following figure, it is important to understand how management settings are applied to the targets in the administration group.

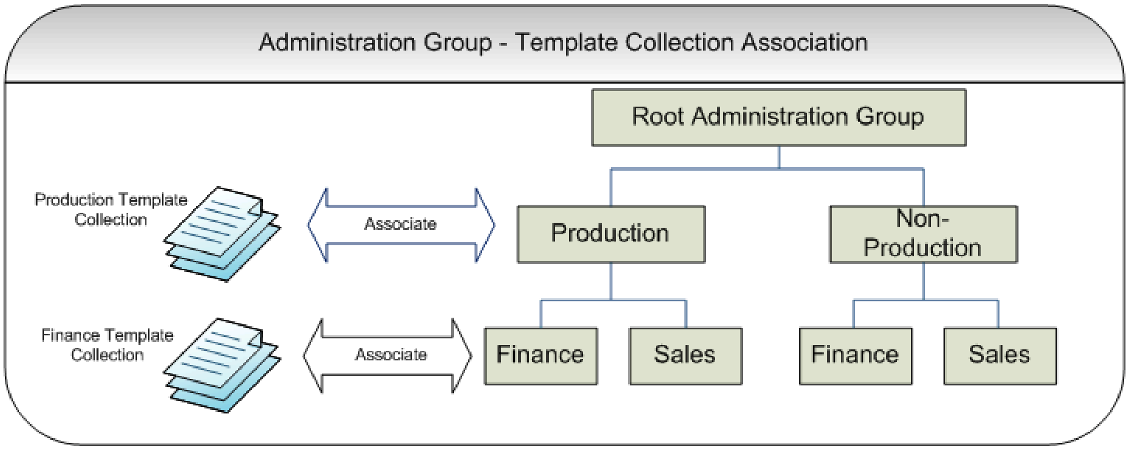
Administration Group - Template Collection Association
Root Administration Group
Production Template Collection
Associate
Production
Non-Production
Finance Template Collection
Associate
Finance
Sales
Finance
Sales

Each group in the administration group hierarchy can be associated with a template collection (containing monitoring templates, compliance standards, and Enterprise Manager policies). If you associate a template collection containing monitoring settings with the *Production* group, then the monitoring settings will apply to the *Finance* and *Sales* subgroup under *Production*. If the *Finance* group under *Production* has additional monitoring settings, then you can create a monitoring template with only those additional monitoring settings. (Later, this monitoring template should be added to another template collection and associated with the *Finance* group). The monitoring settings from the *Finance Template Collection* will be logically combined with the monitoring settings from the *Production Template Collection*. In case there are duplicate metric settings in both template collections, then the metric settings from the *Finance Template Collection* takes precedence and will be applied to the targets in the *Finance* group. This precedence rule only applies to the case of metric settings. In the case of compliance standard rules and Enterprise Manager policies, even if there are duplicate compliance standard rules and Enterprise Manager policies in both template collections, they will be all applied to the targets in the *Finance* group.

Once you have completed all the planning and preparation steps, you are ready to begin creating an administration group.

Implementing Administration Groups and Template Collections

With the preparatory work complete, you are ready to begin the four step process of creating an administration group hierarchy 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 perform each step.

This process involves:

1. Creating the administration group hierarchy.
2. Create monitoring templates.
3. Creating template collections.
4. Associating template collections to administration group.
5. 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.

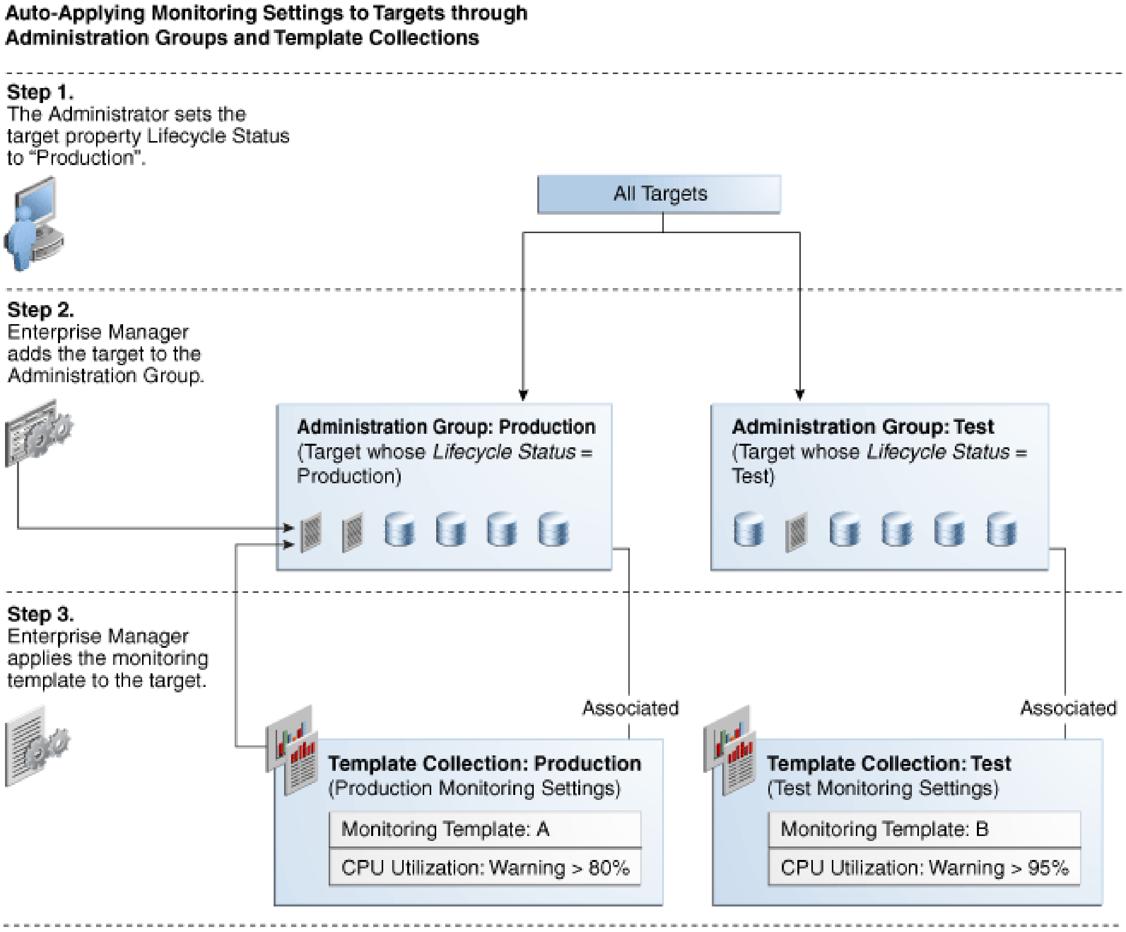
Auto-Applying Monitoring Settings to Targets through Administration Groups and Template Collections
Step 1.
The Administrator sets the target property Lifecycle Status to "Production".
All Targets
Step 2.
Enterprise Manager adds the target to the Administration Group.
Administration Group: Production
(Target whose *Lifecycle Status* = Production)
Administration Group: Test
(Target whose *Lifecycle Status* = Test)
Step 3.
Enterprise Manager applies the monitoring template to the target.
Associated
Associated
Template Collection: Production
(Production Monitoring Settings)
Monitoring Template: A
CPU Utilization: Warning > 80%
Template Collection: Test
(Test Monitoring Settings)
Monitoring Template: B
CPU Utilization: Warning > 95%

Creating the Administration Group Hierarchy

The following four primary tasks summarize the administration group creation process. These tasks are conveniently arranged in sequence via tabbed pages.

> **Note:**
>
> In order to create the administration group hierarchy, you must have both **Full Any Target** and **Create Privilege Propagating Group** target privileges.

Task 1: Access the Administration Group and Template Collections page.

Task 2: Define the hierarchy.

From the **Hierarchy** tab, you define the administration group hierarchy that matches the way you manage your targets. See Defining the Hierarchy.

Task 3: Define the Template Collections.

From the **Template Collections** tab, you define the monitoring and management settings you want applied to targets. See Defining Template Collections.

Task 4: Associate the Template Collections with the Administration Group.

From the **Associations** tab, you tie the monitoring and management settings to the appropriate administration group. See Associating Template Collections with Administration Groups.

Accessing the Administration Group Home Page

All administration group operations are performed from the Administration Groups home page.

From the **Setup** menu, select **Add Target** and then select **Administration Groups**. The Administration Groups home page displays.

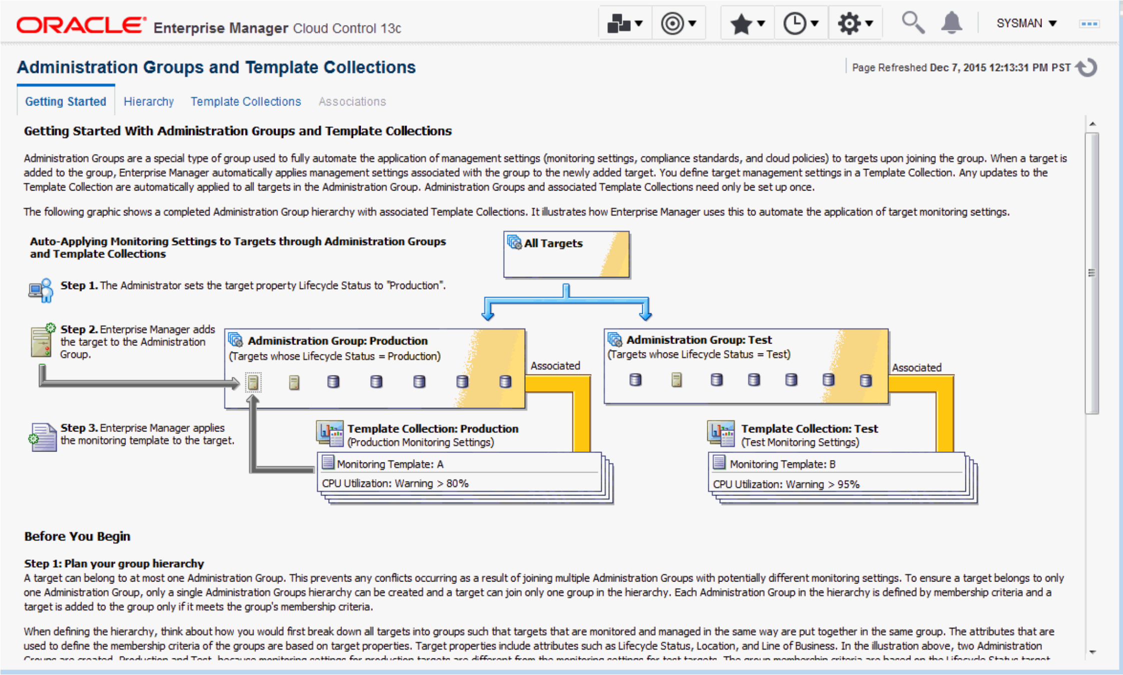
ORACLE Enterprise Manager Cloud Control 13c
SYSMAN
Page Refreshed Dec 7, 2015 12:13:31 PM PST
Getting Started
Hierarchy
Template Collections
Associations

Administration Groups and Template Collections

Getting Started With Administration Groups and Template Collections

Administration Groups are a special type of group used to fully automate the application of management settings (monitoring settings, compliance standards, and cloud policies) to targets upon joining the group. When a target is added to the group, Enterprise Manager automatically applies management settings associated with the group to the newly added target. You define target management settings in a Template Collection. Any updates to the Template Collection are automatically applied to all targets in the Administration Group. Administration Groups and associated Template Collections need only be set up once.

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.

Auto-Applying Monitoring Settings to Targets through Administration Groups and Template Collections

Step 1. The Administrator sets the target property Lifecycle Status to "Production".

Step 2. Enterprise Manager adds the target to the Administration Group.

Step 3. Enterprise Manager applies the monitoring template to the target.

All Targets
Administration Group: Production
(Targets whose Lifecycle Status = Production)
Associated
Administration Group: Test
(Targets whose Lifecycle Status = Test)
Associated
Template Collection: Production
(Production Monitoring Settings)
Monitoring Template: A
CPU Utilization: Warning > 80%
Template Collection: Test
(Test Monitoring Settings)
Monitoring Template: B
CPU Utilization: Warning > 95%

Before You Begin

Step 1: Plan your group hierarchy
A target can belong to at most one 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, only a single Administration Groups hierarchy can be created and a target can join only one group in the 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.

When defining the hierarchy, think about how you would first break down all targets into groups such that targets that are monitored and managed in the same way are put together in the same group. The attributes that are used to define the membership criteria of the groups are based on target properties. Target properties include attributes such as Lifecycle Status, Location, and Line of Business. In the illustration above, two Administration Groups are created, Production and Test, because monitoring settings for production targets are different from the monitoring settings for test targets. The group membership criteria are based on the Lifecycle Status target

Read the relevant information on the **Getting Started** page. The information contained in this page summarizes the steps outlined in this chapter. For your convenience, links are provided that take you to appropriate administration group functions, as well as the Enterprise Manager **All Targets** page where you can view target properties.

Defining the Hierarchy

On this page you define the administration group hierarchy that reflects the organizational hierarchy you planned earlier and which target properties are associated with a particular hierarchy level.

On the left side of the page are two tables: Hierarchy Levels and Hierarchy Nodes. There is also a drop-down menu to include/exclude indirect members.

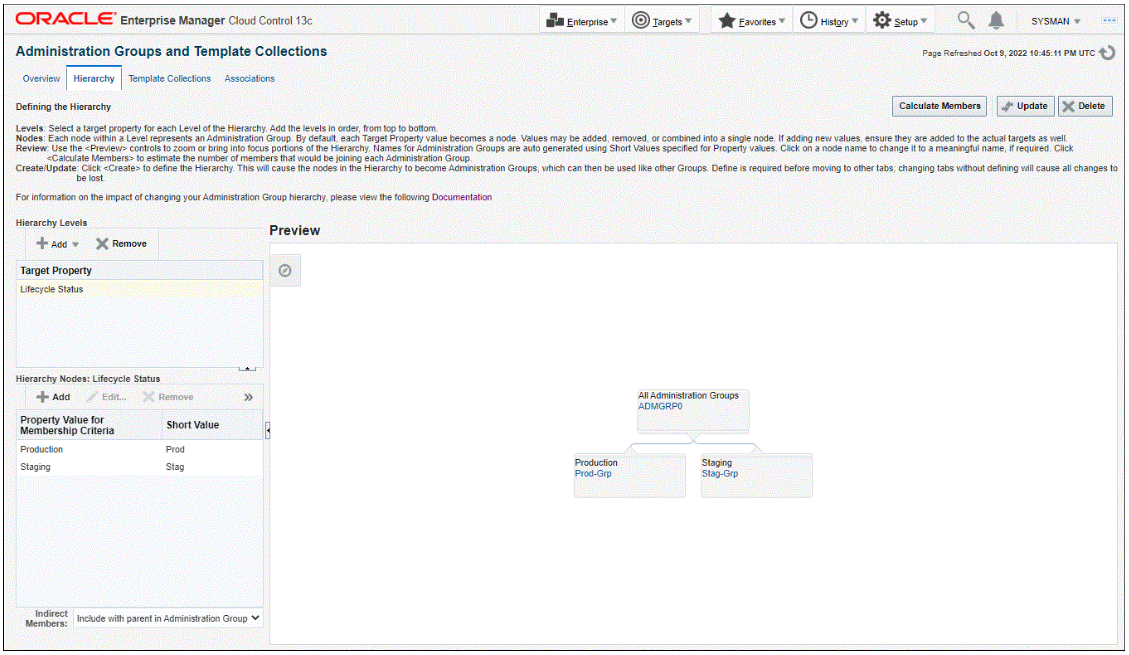

Administration Groups and Template Collections

Page Refreshed Oct 9, 2022 10:45:11 PM UTC

Defining the Hierarchy

Levels: Select a target property for each Level of the Hierarchy. Add the levels in order, from top to bottom.
Nodes: Each node within a Level represents an Administration Group. By default, each Target Property value becomes a node. Values may be added, removed, or combined into a single node. If adding new values, ensure they are added to the actual targets as well.
Review: Use the <Preview> controls to zoom or bring into focus portions of the Hierarchy. Names for Administration Groups are auto generated using Short Values specified for Property values. Click on a node name to change it to a meaningful name, if required. Click <Calculate Members> to estimate the number of members that would be joining each Administration Group.
Create/Update: Click <Create> to define the Hierarchy. This will cause the nodes in the Hierarchy to become Administration Groups, which can then be used like other Groups. Define is required before moving to other tabs; changing tabs without defining will cause all changes to be lost.

For information on the impact of changing your Administration Group hierarchy, please view the following Documentation

Hierarchy Levels

Target Property
Lifecycle Status

Hierarchy Nodes: Lifecycle Status

Property Value for Membership Criteria	Short Value
Production	Prod
Staging	Stag

Indirect Members: Include with parent in Administration Group

Preview

The **Hierarchy Levels** table allows you to add the target properties that define administration group hierarchy. The **Hierarchy Nodes** table allows you to define the values associated with the target properties in the **Hierarchy Levels** table. When you select a target property, the related property values are made available in the **Hierarchy Nodes** table, where you can add/remove/merge/split the values. In the **Hierarchy Nodes** table, each row corresponds to a single administration group. The Short Value column displays abbreviated value names that are used to auto-generate group names.

The **Hierarchy Levels** table allows you to add the target properties that define each level in the administration group hierarchy. The **Hierarchy Nodes** table allows you to define the values associated with the target properties in the **Hierarchy Levels** table. Each row in the **Hierarchy Nodes** table will correspond to a node or group in the administration group hierarchy for that level. When you select a target property in the **Hierarchy Levels** table, the related property values are made available in the **Hierarchy Nodes** table, where you can add/remove/merge/split the values. Merge two or more values if either value should be used as membership criteria for the corresponding administration group. The Short Value column displays abbreviated value names that are used to auto-generate group names.

The **Indirect Members** drop-down menu provides you with the option to include or exclude any indirect member targets as part of the administration group hierarchy. For more information on indirect members, see Group Member Type and Synchronization (Indirect Members).

Adding a Hierarchy Level

1. On the **Administration Group** page, click the **Hierarchy** tab.
2. From the **Hierarchy Levels** table, click **Add** and choose one of the available target properties. You should add one property/level at a time instead of all properties at once.
3. With the target property selected in the **Hierarchy Levels** table, review the list of values shown in the **Hierarchy Nodes** table. The values of the target property in the **Hierarchy Nodes** table.

   Enterprise Manager finds all existing values of the target property across all targets and displays them in the **Hierarchy Nodes** table. For some target properties, such as Lifecycle Status, predefined property values already exist and are automatically displayed in the

Hierarchy Nodes table. You can select and remove target property values that will not be used as membership criteria in any administration group. However, property values that are not yet available but will be used as administration group membership criteria, will need to be added.

The next step shows you how to add property values.

4. From the **Hierarchy Nodes** table, click **Add**. The associated property value add dialog containing existing values from various targets displays. Add the requisite value(s). Multiple values can be specified using a comma separated list. For example, to add multiple locations such as San Francisco and Zurich, add the **Location** target property to the **Hierarchy Level** table. Select **Location** and then click **Add** in the **Hierarchy Nodes** table. The **Values for Hierarchy Nodes** dialog displays. Enter "San Francisco,Zurich" as shown in the following graphic.

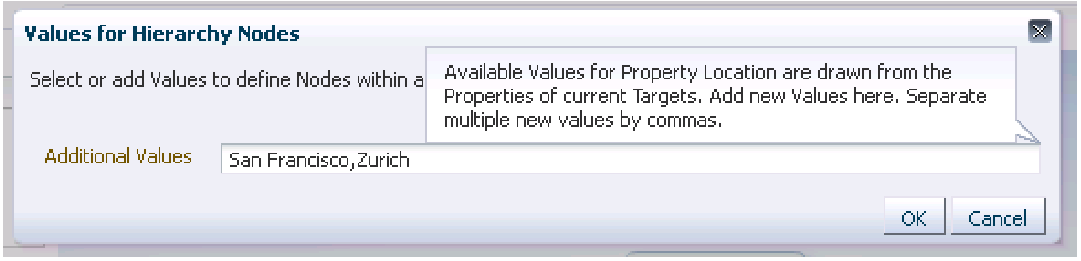

Values for Hierarchy Nodes

Select or add Values to define Nodes within a

Available Values for Property Location are drawn from the Properties of current Targets. Add new Values here. Separate multiple new values by commas.

Additional Values: San Francisco,Zurich

Extending Administration Group Hierarchy Maximum Limits

There is a default maximum for the number of values that can be supported for a target property as administration group criteria. If you see a warning message indicating that you have reached this maximum value, you can extend it using the OMS property *admin_groups_width_limit*. Specify the maximum number of values that should be supported for a target property. For example, to support up to 30 values for a target property that will be used in administration group criteria, set the admin_groups_width_limit as follows (using the OMS `emctl` utility):

```
emctl set property -name admin_groups_width_limit -value 30 -module emoms
```

You can also add up to four levels after the root node of an administration group hierarchy. If there is a need to add additional level, you will first need to change the OMS *admin_groups_height_limit* property to the maximum height limit. For example, if you want to create to administration group hierarchy consisting of five levels after the root node, set the *admin_groups_height_limit* property as follows (using the OMS emctl utility):

```
emctl set property -name admin_groups_height_limit -value 5 -module emoms
```

This is a global property and only needs to be set once using the emctl utility of any OMS. This is also a dynamic property and does not require a stop/restart of the OMS in order to take effect.

Click **Create**. The two locations "San Francisco" and "Zurich" appear as nodes in the **Preview** pane as shown in the following graphic.

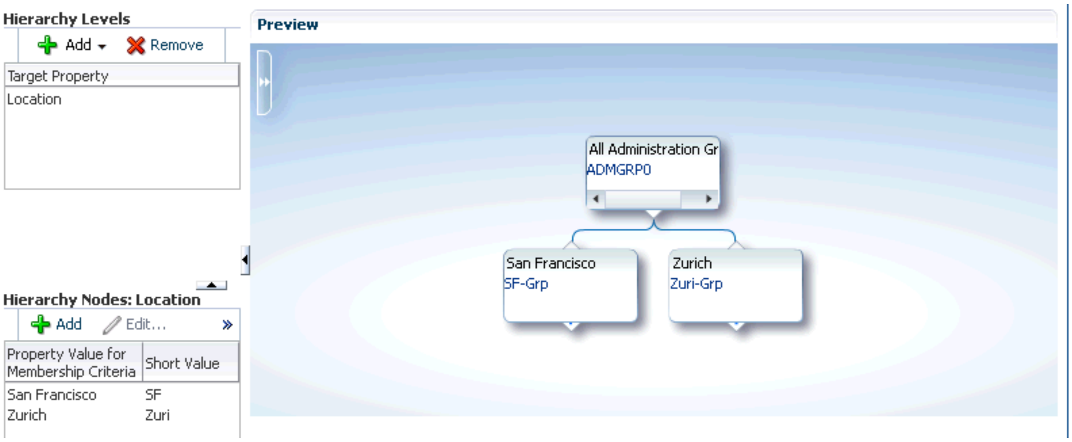

Under certain circumstances, it may be useful to treat multiple property values as one: Targets may have different target property values, but should belong to the same administration group because they have the same monitoring profile/settings. For example, a combination of values is needed for the *Lifecycle Status* property where you have the following:

- Production Group: Criteria is based on *Lifecycle Status = Mission Critical or Production*
- Non-Production Group: Criteria is based on *Lifecycle Status = Development, Test, Staging*

In this situation, the two *Lifecycle Status* properties should be merged (combined into a single node).

To merge property values:

a. Select a target property from the list of chosen properties in the **Hierarchy Levels** table. The associated property values are displayed.

b. Select two or more property values by holding down the *Shift* key and clicking on the desired values.

c. Click **Merge**.

5. From the **Indirect Members** drop-down menu, choose whether you want to include or exclude any indirect member targets as part of the administration group. For more information on indirect group members, see Group Member Type and Synchronization (Indirect Members).

6. Continue adding hierarchy levels until the group hierarchy is complete. The **Preview** pane dynamically displays any changes you make to your administration group hierarchy.

7. Set the time zone for the group.

   a. Click on the group name. The Administration Group Details dialog displays allowing you to select the appropriate time zone.

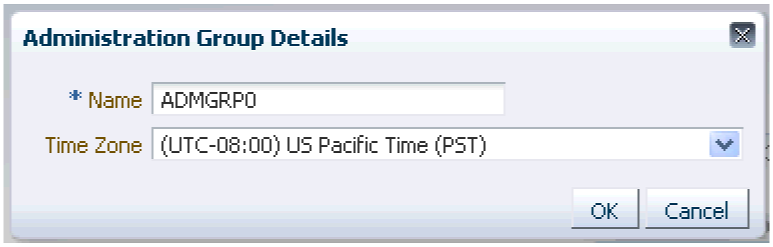
Administration Group Details

* Name: ADMGRP0

Time Zone: (UTC-08:00) US Pacific Time (PST)

The administration group time zone is used for displaying group charts and also for scheduling operations on the group. Because this is also the default time zone for all subgroups that may be created under this group, you should specify the time zone at the highest level group in the administration group hierarchy before the subgroups are created. Note that the parent group time zone will be used when creating any child subgroups, but user can always select a child subgroup and change its time zone.

The auto-generated name can also be changed.

8. Click **Create** to define the hierarchy.

> **Note:**
>
> Review and define the complete hierarchy before clicking **Create**.

Even after your administration group hierarchy has been created, you can always make future updates if organizational needs change. For example, adding/removing group membership criteria property values, which equates to creating/deleting additional administration groups for a given level. Using the previous example, if in addition to San Francisco and Zurich you add more locations, say New York and Bangalore, you can click **Add** in the **Hierarchy Node** table to add additional locations, as shown in the following graphic. For more information about changing the administration group hierarchy, see "Changing the Administration Group Hierarchy".

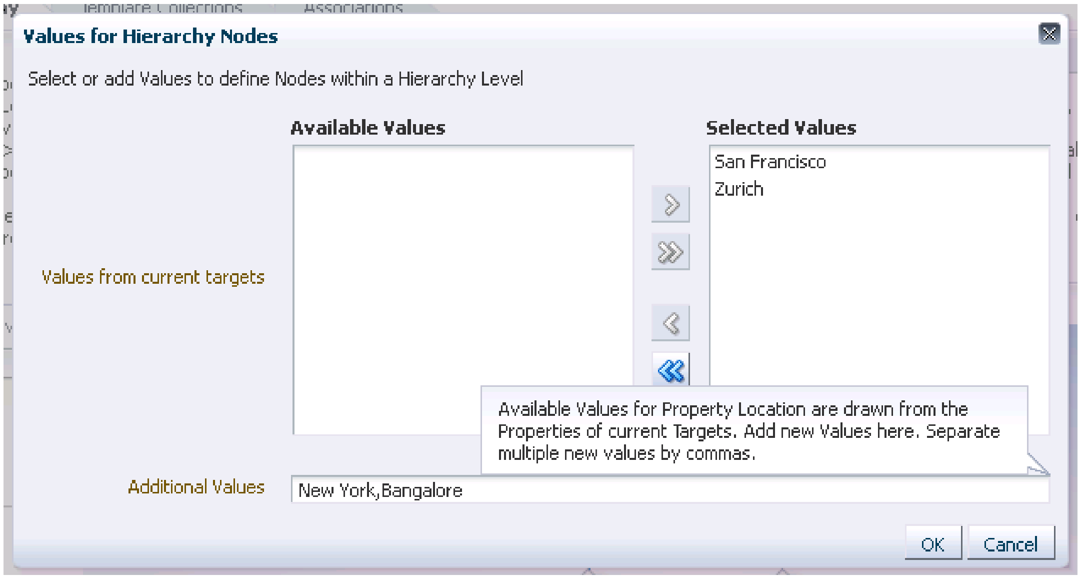

Values for Hierarchy Nodes

Select or add Values to define Nodes within a Hierarchy Level

Values from current targets — Available Values: (none); Selected Values: San Francisco, Zurich

Available Values for Property Location are drawn from the Properties of current Targets. Add new Values here. Separate multiple new values by commas.

Additional Values: New York,Bangalore

Click **Update** to save your changes.

Defining Template Collections

A template collection is an assemblage of monitoring/management settings to be applied to targets in the administration group. Multiple monitoring templates can be added to a template collection that in turn is associated with an administration group. However, you can only have one monitoring template of a particular target type in the template collection. The monitoring template should contain the complete set of metric settings for the target in the administration group. You should create one monitoring template for each type of target in the administration group. 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 two templates for databases. When direct members targets are added to an administration group, the template monitoring and management settings are automatically applied. A template will completely replace all metric settings in the target. This means applying the template copies over metric settings (thresholds, corrective actions, collection schedule) to the target, removes the thresholds of the metrics that are present in the target, but not included in the template. Removing of thresholds disables alert functionality for these metrics. Metric data will continue to be collected.

Template collections may consist of three types of monitoring/management setting categories:

- Monitoring Templates (monitoring settings)
- Compliance Standards (compliance policy rules)
- Enterprise Manager Policies (Enterprise Manager policies such as determining when to start virtual machines or scale out clusters).

> **Note:**
>
> These monitoring/management settings do not apply to indirect members. For more information about direct and indirect members, see Group Member Type and Synchronization.

When creating a template collection, you can use the default monitoring templates, compliance standards, or Enterprise Manager templates supplied with Enterprise Manager or you can create your own. For more information, see Using Monitoring Templates .

To create a template collection:

1. Click the **Template Collections** tab. The Template Collection page displays.
2. Click **Create**.
3. In the **Name** field, specify the template collection name.
4. Click the template collection member type you want to add (Monitoring Template, Compliance Standard, Enterprise Manager Polices). The requisite definition page appears.
5. Click **Add**. A list of available template entities appears.

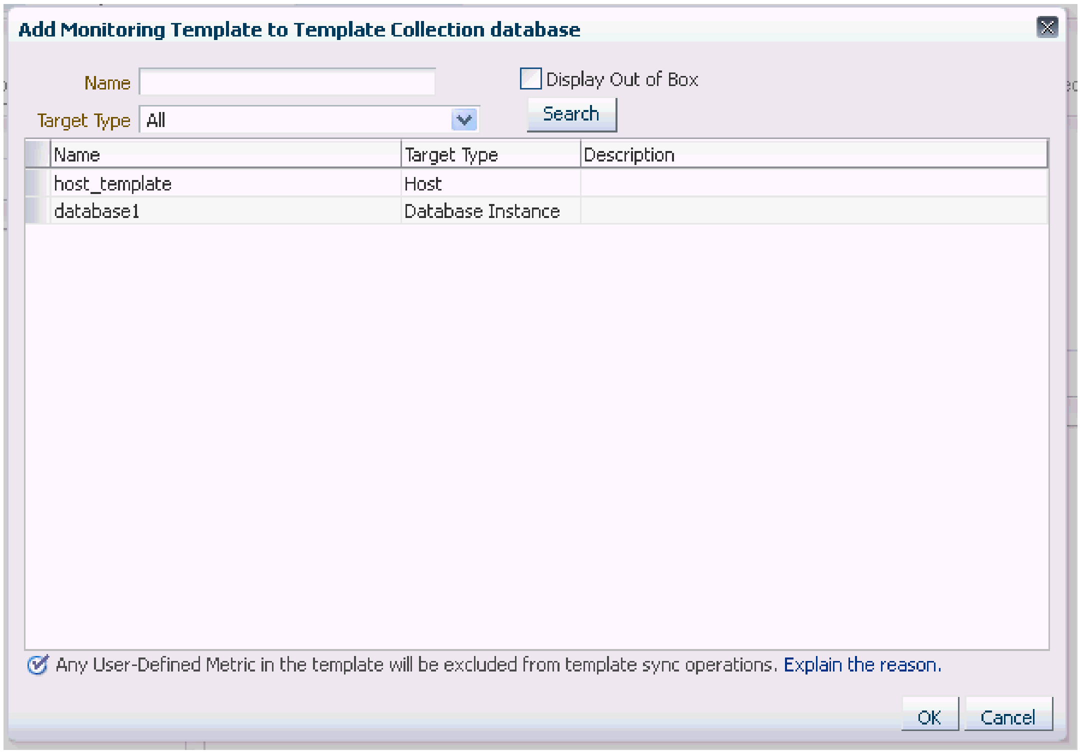

Add Monitoring Template to Template Collection database

Name

Target Type All

□ Display Out of Box

Name	Target Type	Description
host_template	Host	
database1	Database Instance	

Any User-Defined Metric in the template will be excluded from template sync operations. Explain the reason.

6. Select the desired template entities you want added to the template collection.
7. Click **OK**.
8. Continue adding template entities (Monitoring Template, Compliance Standard, Enterprise Manager Policies) as required.
9. Click **Save**. The newly defined collection appears in the **Template Collections Library**.
10. To create another template collection, click **Create** 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 **Cancel**. This restores the template collection to its state when it was last saved.

Required Privileges

To create a template collection, you must have the *Create Template Collection* resource privilege. To include a monitoring template into a template collection, you need at least *View* privilege on the specific monitoring template or *View Any Monitoring Template* privilege, which allows you to view any monitoring template and add it to the template collection. The following table summarizes privilege requirements for all Enterprise Manager operations related to template collection creation.

Enterprise Manager Operation	Minimum Privilege Requirement
Create administration group hierarchy.	Full Any Target Create Privilege Propagating Group

Enterprise Manager Operation	Minimum Privilege Requirement
Create monitoring templates.	Create Monitoring Template
Create template collection.	Create template collection (resource privilege). VIEW on the monitoring template to be added to the template collection or View any monitoring template (resource privilege).
Create compliance standards.	Create Compliance Entity No privileges are required to view compliance standards.
Create Enterprise Manager policies.	Create Any Policy View Enterprise Manager Policy
Associate template collection with administration group.	VIEW on the specific template collection. Manage Target Metrics on the group.
Perform on-demand synchronization.	OPERATOR on the group or Manage Target Metrics.
Define global synchronization schedule.	Enterprise Manager Super Administrator privileges.
Set the value of target properties for a target (allows the target to "join" an administration group).	Configure Target on the specific target
Delete an administration group hierarchy.	Full Any Target

Corrective Action Credentials

A corrective action is an automated task that is executed in response to a metric alert. When a corrective action is part of a monitoring template/template collection, the credentials required to execute the corrective action will vary depending on how the template is applied.

The two situations below illustrate the different credential requirements.

- *The corrective action is part of a monitoring template that is manually applied to a target.*

  When the corrective action runs, it can use one of the following:

  - The preferred credentials of the user who is applying the template

    or

  - The user-specified named credentials.

  The user selects the desired credential option during the template apply operation.

- *The corrective action is part of a monitoring template within a template collection that is associated with an administration group.*

  When the corrective action runs, the preferred credentials of the user who is associating the template collection with the administration group is used.

Associating Template Collections with Administration Groups

Once you have defined one or more template collections, you need to associate them to administration groups in the hierarchy. You can associate a template collection with one or more administration groups. As a rule, you should associate the template collection with the applicable administration group residing at the highest level in the hierarchy as the template collection will also be applied to targets joining any subgroup.

The **Associations** page displays the current administration 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

> **Note:**
>
> For users that do not have View privilege on all administration groups, you can also perform the association/disassociation operation from the Groups page (from the **Targets** menu, select **Groups**).

1. Click the **Associations** tab. The Associations page displays.

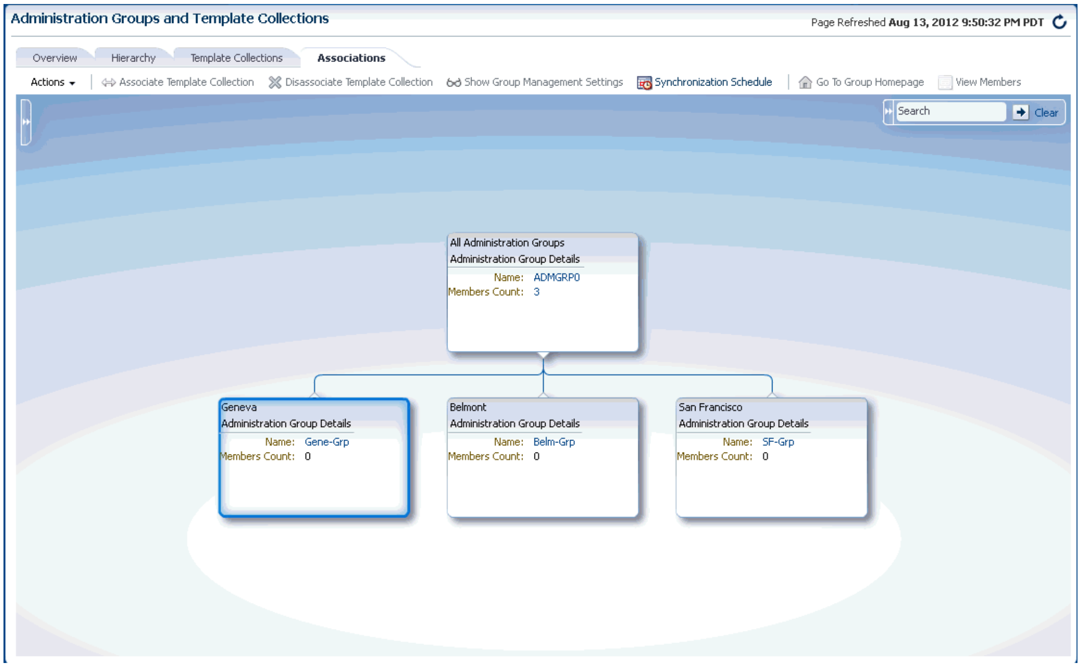
Administration Groups and Template Collections
Page Refreshed Aug 13, 2012 9:50:32 PM PDT
Overview
Hierarchy
Template Collections
Associations
Actions
Associate Template Collection
Disassociate Template Collection
Show Group Management Settings
Synchronization Schedule
Go To Group Homepage
View Members
Search
Clear

All Administration Groups
Administration Group Details
Name: ADMGRP0
Members Count: 3

Geneva
Administration Group Details
Name: Gene-Grp
Members Count: 0

Belmont
Administration Group Details
Name: Belm-Grp
Members Count: 0

San Francisco
Administration Group Details
Name: SF-Grp
Members Count: 0

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 list of targets affected by this operation is displayed. Confirm or discard the operation.

   > **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. The administrator needs at least Manage Target Metrics privileges on the group.

Note:

Settings from monitoring templates applied at lower levels in the hierarchy override settings inherited from higher levels. This does not apply to compliance standards or Enterprise Manager policies.

Compliance Standard Rules

When compliance standard rules are defined in different template collections at different locations in the Administration Group hierarchy, a union of compliance standard rules from all template collections across all levels of the hierarchy will be used.

Searching for Administration Groups

While the administration group UI is easy to navigate, there may be cases where the administration group hierarchy is inordinately large, thus making it difficult to find individual groups. At the upper right corner of the Associations page is a search function that greatly simplifies finding groups in a large hierarchy.

Figure 10-1 Administration Group Search Dialog

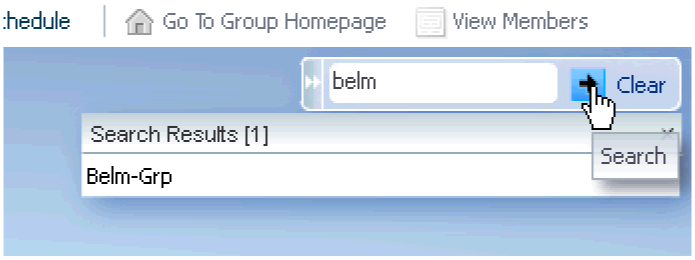

To search for a specific administration group:

1. If not already displayed, expand the Search interface.

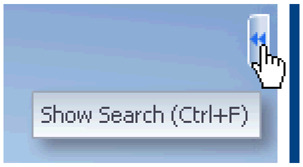

2. Enter either a full or partial group name and click **Search**.

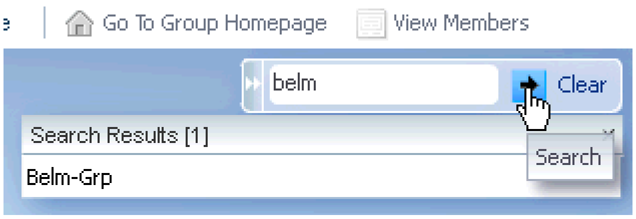

As shown in the graphic, the search results display a list of administration groups that match the search criteria. You can then choose an administration group from the list by double-clicking on the entry. The administration group hierarchy will then display a vertical slice (subset) of the administration group hierarchy from the root node to the group you selected.

Figure 10-2 Administration Group Search: Graphical Display

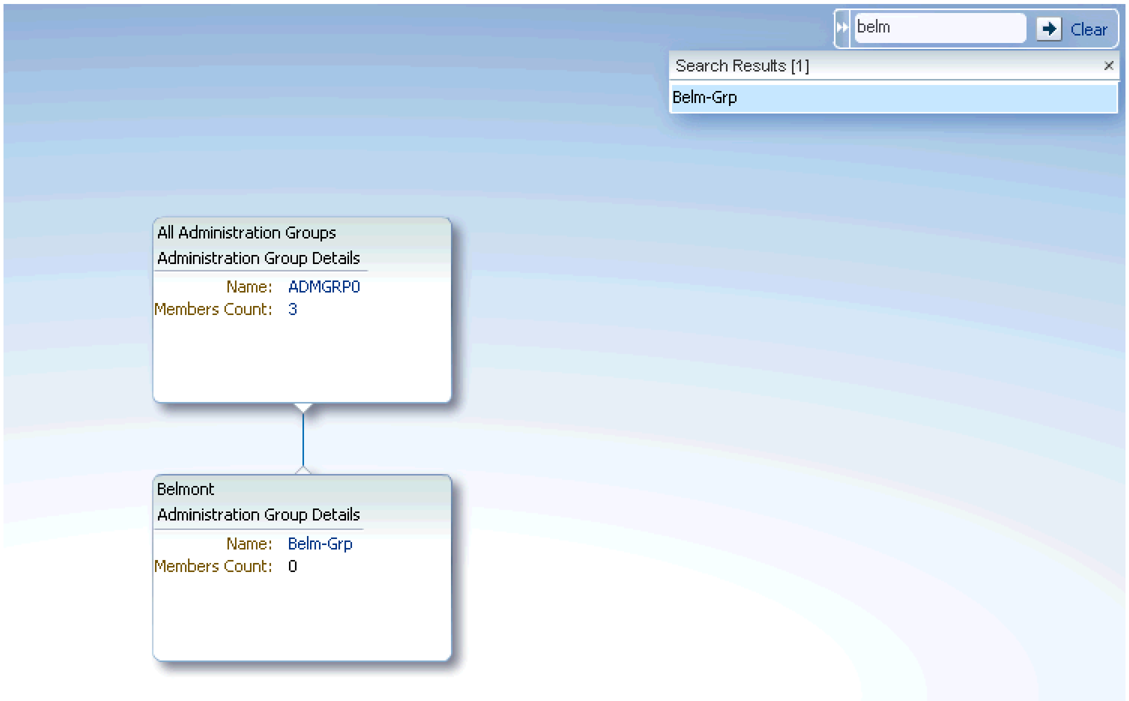

To restore the full administration group hierarchy, click **Clear**.

Group Names and Searches

In order to perform effective searches for specific administration groups, it is helpful to know how Enterprise Manager constructs an administration group name: Enterprise Manger uses the administration group criteria to generate names. For example, you have an administration group with the following criteria:

- Lifecycle Status: Development or Mission Critical
- Department: DEV
- Line of Business: Finance or HR
- Location: Bangalore

Enterprise Manager assembles a group name based on truncated abbreviations. In this example, the generated administration group name is *DC-DEV-FH-Bang-Grp*

As you are building the hierarchy, you can change the abbreviation associated with each value (this is the Short Value column next to the property value in the Hierarchy Nodes table. Hence, you can specify a short value and Enterprise Manager will use that value when constructing new names for any subgroups created.

During the design phase of an administration group, you have the option of specifying a custom name. However, if there is large number of groups, it is easier to allow Enterprise Manager to generate unique names.

Setting the Global Synchronization Schedule

In order to apply the template collection/administration group association, you must set up a global synchronization schedule. This schedule is used to perform synchronization operations, such as applying templates to targets in administration groups. If no synchronization schedule is set up, when a target joins an administration group, Enterprise Manager will auto-apply the associated template. However, if there are changes to the template later on, then Enterprise Manager will only apply these based on synchronization schedule, otherwise these operations are pending.When there are any pending synchronization operations, they will be scheduled on the next available date based on the synchronization schedule.

> **Note:**
>
> You **must** set the synchronization schedule as there is no default setting. You can specify a non-peak time such as weekends.

To set up the synchronization schedule:

1. Click **Synchronization Schedule**. The Synchronization Schedule dialog displays.

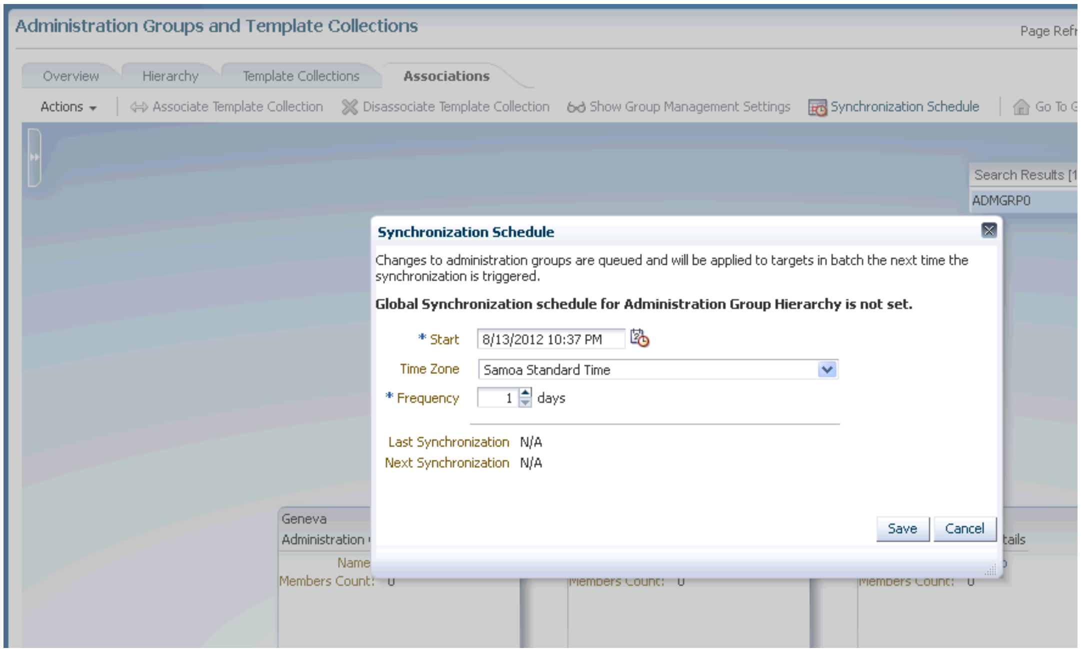

Synchronization Schedule

Changes to administration groups are queued and will be applied to targets in batch the next time the synchronization is triggered.

Global Synchronization schedule for Administration Group Hierarchy is not set.

* Start	8/13/2012 10:37 PM
Time Zone	Samoa Standard Time
* Frequency	1 days
Last Synchronization	N/A
Next Synchronization	N/A

2. Click **Edit** and then choose a date and time you want any pending sync operations (For example, template apply operations) to occur. By default, the current date and time is shown.

Note:

You can specify a start date for synchronization operations and interval in days. Whenever there are any pending sync operations, then they will be scheduled on the next available date based on this schedule.

3. Click **Save**.

When Template Collection Synchronization Occurs

The following table summarizes when Template Collection Synchronization operations (such as apply operations) occur on targets in administration groups.

Action	When Synchronization Occurs
Target is added to an administration group (by setting its target properties)	Immediate upon joining the administration group.
Template collection is associated with the administration group.	Targets in an administration group will be synchronized based on next scheduled date in global Synchronization Schedule.
Changes are made to any of the templates in the template collection.	Targets in an administration group will be synchronized based on the next scheduled date in global Synchronization Schedule.
Target is removed from an administration group (by changing its target properties).	No change in target's monitoring settings. Compliance Standards and Enterprise Manager Policies will be disassociated with the target. Immediate synchronization operation occurs.
Template collection is disassociated with administration group.	No change in target's monitoring settings for all the targets in the administration group. Compliance Standards and Enterprise Manager Policies will be disassociated with the target. Targets under the administration group will be synchronized based on next schedule date in Global Synchronization Schedule.
User performs an on-demand synchronization by clicking on the **Start Synchronization** button in the **Synchronization Status** region in the administration group's homepage.	Immediate synchronization operation occurs.

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** region, you can view the status of the monitoring template, compliance standard, and/or Enterprise Manager policies synchronization (In Sync, Pending, or Failed).

   You can initiate an immediate synchronization by clicking **Start Synchronization**.

Group Member Type and Synchronization

There are two types of administration group member targets:

- *Direct Members*: Group members whose target properties match the administration group criteria. Monitoring settings, compliance standards, Enterprise Manager policies from the associated template collection are applied to direct members.

  Only direct members are represented in the targets count in the Synchronization Status region.

- *Indirect Members*: Indirect members are targets whose target properties DO NOT match administration group criteria. However, they have been added to the administration group because their parent target is a direct member of the administration group. Whenever aggregate targets, i.e., targets that have member targets, are added to an administration (or dynamic) group, all members of the aggregate target are, by default, also added to the group. An example of an aggregate target is Oracle WebLogic Server. If that target is added to a group, then all Application Deployment targets on it are also pulled into the group, even if their properties do not satisfy the group criteria. **Indirect group members will NOT be part of any template apply/sync operations.**

  When defining the administration group hierarchy, you have the option of including or excluding any indirect member targets using the Indirect Members drop-down menu. This feature is available with Oracle Enterprise Manager 13c Release 5 Update 10 (13.5.0.10) and later.

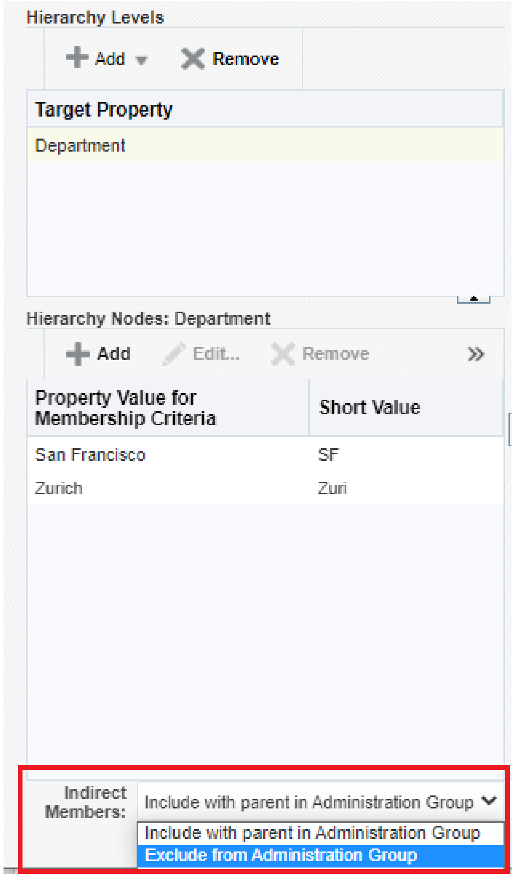
Hierarchy Levels
Add
Remove
Hierarchy Nodes: Department
Add
Edit...
Remove
Include with parent in Administration Group
Exclude from Administration Group

Target Property
Department

Property Value for Membership Criteria	Short Value
San Francisco	SF
Zurich	Zuri

Indirect Members: Include with parent in Administration Group

Only direct members are represented in the targets count in the Synchronization Status region.

1. From the hierarchy diagram, click on a group name to access the group's home page. You can also access this information from **All Targets** groups page.
2. From the **Group** menu, select **Members** and then one of the display menu options.

System Targets and Administration Groups

If a system target gets added to an administration group because it matches group criteria, then the system target and its constituent members are also added. However, for template apply purposes, it will only operate on the direct members that also match the administration group criteria. Template apply operations will not occur on member targets whose target properties do not match administration group criteria. All other group operations, such as jobs and blackouts, will apply on all members, both direct and indirect.

Disassociating a Template Collection from a Group

To disassociate a template collection from an administration group.

1. From the **Setup** menu, select **Add Target** and then **Administration Groups**. The Administration Group home page displays.
2. Click on the **Associations** tab to view the administration group hierarchy diagram.
3. From the hierarchy diagram, select the administration group with the template collection you wish to remove. If necessary, use the **Search** option to locate the administration group.
4. Click **Disassociate Template Collection**. The number of targets affected by this operation is displayed. Click **Continue** or **Cancel**.

The template collection is immediately removed. See "When Template Collection Synchronization Occurs" in Defining the Hierarchy for more information.

Viewing Aggregate (Group Management) Settings

For any administration group, you can easily view what template collection components (monitoring templates, compliance standards, and/or Enterprise Manager policies) are associated with individual group members.

> **Note:**
>
> For monitoring templates, the settings for a target could be a union of two or more monitoring templates from different template collections.

1. From the **Setup** menu, select **Add Target** and then **Administration Groups**. The Administration Group home page displays.
2. Click on the **Associations** tab to view the administration group hierarchy diagram.
3. From the hierarchy diagram, select the desired administration group.
4. Click **Show Group Management Settings**.

   The **Administration Group Details** page displays.

   This page displays all aggregate settings for monitoring templates, compliance standards and Enterprise Manager policies that will be applied to members of the selected administration group (listed by target type). The page also displays the synchronization status of group members.

To can change the display to show a different branch of the administration group hierarchy, click **Select Branch** at the upper-right area of the page. This function lets you display hierarchy branches by choosing different target property values

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, details of the Associated Template Collection for the group selected in hierarchy viewer, job activity, or critical patch advisories. To view administration group home pages:

1. From the hierarchy diagram, select an administration group.
2. Click **Goto Group Homepage.** The homepage for that particular administration group displays.

Alternatively, from the Enterprise Manger **Targets** menu, choose **Groups**. From the table, you can expand the group hierarchy.

Identifying Targets Not Part of Any Administration Group

From the **Associations** page, you can determine which targets do not belong to any administration group by generating an *Unassigned Targets Report*.

1. From the **Actions** menu, select **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.

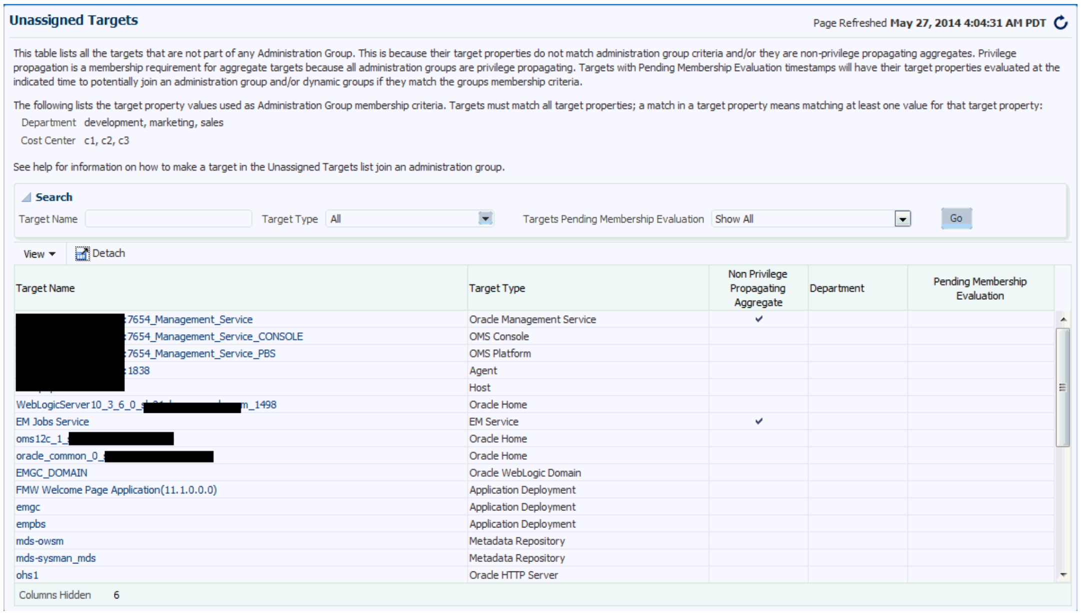

Unassigned Targets

Page Refreshed **May 27, 2014 4:04:31 AM PDT**

This table lists all the targets that are not part of any Administration Group. This is because their target properties do not match administration group criteria and/or they are non-privilege propagating aggregates. Privilege propagation is a membership requirement for aggregate targets because all administration groups are privilege propagating. Targets with Pending Membership Evaluation timestamps will have their target properties evaluated at the indicated time to potentially join an administration group and/or dynamic groups if they match the groups membership criteria.

The following lists the target property values used as Administration Group membership criteria. Targets must match all target properties; a match in a target property means matching at least one value for that target property:

Department	development, marketing, sales
Cost Center	c1, c2, c3

See help for information on how to make a target in the Unassigned Targets list join an administration group.

Search

Target Name:
Target Type: All
Targets Pending Membership Evaluation: Show All

Target Name	Target Type	Non Privilege Propagating Aggregate	Department	Pending Membership Evaluation
[redacted]7654_Management_Service	Oracle Management Service	✓		
[redacted]7654_Management_Service_CONSOLE	OMS Console			
[redacted]7654_Management_Service_PBS	OMS Platform			
[redacted]1838	Agent			
[redacted]	Host			
WebLogicServer10_3_6_0_[redacted]m_1498	Oracle Home			
EM Jobs Service	EM Service	✓		
oms12c_1_[redacted]	Oracle Home			
oracle_common_0_[redacted]	Oracle Home			
EMGC_DOMAIN	Oracle WebLogic Domain			
FMW Welcome Page Application(11.1.0.0.0)	Application Deployment			
emgc	Application Deployment			
empbs	Application Deployment			
mds-owsm	Metadata Repository			
mds-sysman_mds	Metadata Repository			
ohs1	Oracle HTTP Server			

Columns Hidden 6

2. From the **View** menu, choose the customization options to display only the desired information.

> **Note:**
>
> The **Non-Privilege Propagating Aggregate** column indicates whether a target is a non-privilege propagating aggregate. This type of target cannot be added to an administration group, which are by design privilege propagating. For this reason, any aggregate target added to administration group must also be privilege propagating. To make an aggregate target privilege propagating, use the EM CLI verb *modify_system* with *-privilege_propagation=true option*.

On this page, you can review the list to see if there any targets that need to be added to the administration group. Click on the target names shown in this page to access the target's **Edit Target Properties** page where you can change the target property values. After making the requisite changes and clicking OK, you are returned to the **Unassigned Targets** page.

For information on changing target properties, see "Planning an Administrative Group".

3. Click your browser *back* button to return to the **Administration Groups and Template Collections** homepage.

Changing the Administration Group Hierarchy

Organizations are rarely static--new lines of business may be added or perhaps groups are reorganized due to organizational expansion. To accommodate these changes, you may need to make changes to the existing administration group hierarchy.

Beginning with Enterprise Manager 12*c* Release 12.1.0.3, you can change the administration group hierarchy without having to rebuild the entire hierarchy. You can easily perform administration group alterations such as adding more groups to each hierarchy level, merging two or more groups, or adding/deleting entire hierarchy levels. All of these operations can be performed from the Hierarchy page.

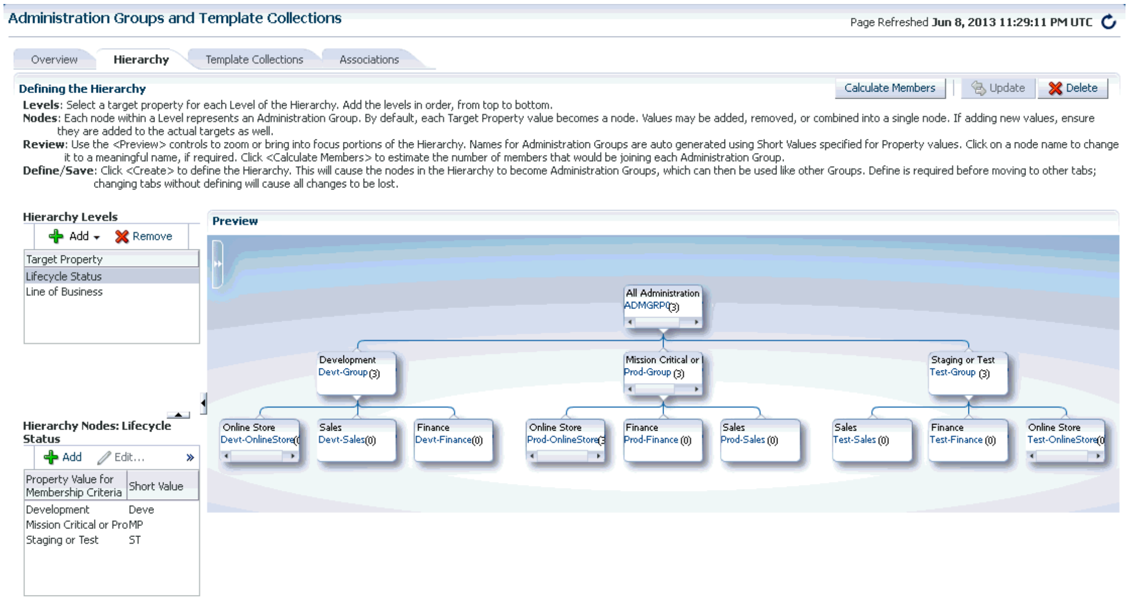

Levels: Select a target property for each Level of the Hierarchy. Add the levels in order, from top to bottom.

Nodes: Each node within a Level represents an Administration Group. By default, each Target Property value becomes a node. Values may be added, removed, or combined into a single node. If adding new values, ensure they are added to the actual targets as well.

Review: Use the <Preview> controls to zoom or bring into focus portions of the Hierarchy. Names for Administration Groups are auto generated using Short Values specified for Property values. Click on a node name to change it to a meaningful name, if required. Click <Calculate Members> to estimate the number of members that would be joining each Administration Group.

Define/Save: Click <Create> to define the Hierarchy. This will cause the nodes in the Hierarchy to become Administration Groups, which can then be used like other Groups. Define is required before moving to other tabs; changing tabs without defining will cause all changes to be lost.

Target Property
Lifecycle Status
Line of Business

Property Value for Membership Criteria	Short Value
Development	Deve
Mission Critical or Pro	MP
Staging or Test	ST

> **Note:**
>
> After making any change to the administration group hierarchy, click **Update** to save your changes.

Adding a New Hierarchy Level

Adding a new hierarchy level equates to adding a new target property to the administration group criteria. For this reason, you must set the value of this target property for all your targets in order for them to continue to be part of the administration group hierarchy. Any new target property added/hierarchy level added will always be added as the bottom-most level of the hierarchy. You cannot insert a new level between levels.

To insert a hierarchy level, you must remove a hierarchy level, then add the levels you want. Think carefully before removing a hierarchy level as removing a level will result in the deletion of groups corresponding to that hierarchy level.

See "Adding a Hierarchy Level" in Defining the Hierarchy for step-by-step instructions on adding a new level.

Removing a Hierarchy Level

Removing a hierarchy level equates to deleting a target property, which in turn causes groups at that level to be deleted. For this reason, think carefully about the groups that will be removed when you remove the hierarchy level, especially if those groups are used in other functional areas of Enterprise Manager.

To remove a hierarchy level:

1. On the **Administration Group** page, click the **Hierarchy** tab.
2. From the **Hierarchy Levels** table, select a hierarchy level and click **Remove**
3. Click **Update** to save your changes.

If any of those groups have an associated template collection, then the monitoring settings of the subgroups of the deleted group will be impacted since the subgroups obtained monitoring settings from the associated template collection. You may need to review the remaining template collections and re-associate the template collection with the appropriate administration group.

Merging Administration Groups

If you want to merge two or more administration groups, you merge their corresponding target property criteria in the administration group hierarchy definition. The group merge operation consists of retaining one of the groups to be merged and then moving over the targets from the other groups into the group that is retained. Once the targets have been moved, the other groups will be deleted.

You choose which group is retained by choosing its corresponding target property value. The group(s) containing the selected target property value as part of its criteria is retained. If the retained target property criteria corresponds to multiple groups, i.e. group containing subgroups, the movement of targets will actually occur at the lowest level administration groups since the targets only reside in the lowest level administration groups. The upper-level administration groups' criteria will be updated to include the criteria of the other groups that have been merged into it.

To merge groups:

1. Select a target property from the list of chosen properties in the **Hierarchy Levels** table. You choose the target property corresponding to the groups you want to merge.

   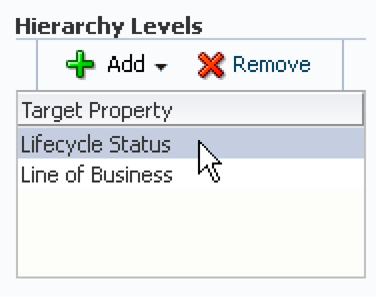

Target Property
Lifecycle Status
Line of Business

   For example, let us assume you want to merge <Devt-Group> with <Test or Stage Group>. In the hierarchy, this corresponds to target property Lifecycle Status. The associated property values are displayed.

2. Select two or more property values corresponding to the groups you would like to merge by holding down the *Shift* or *CTRL* key and clicking on the desired values.

   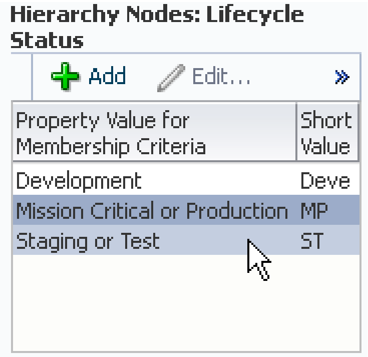

Property Value for Membership Criteria	Short Value
Development	Deve
Mission Critical or Production	MP
Staging or Test	ST

3. Click **Merge**.

   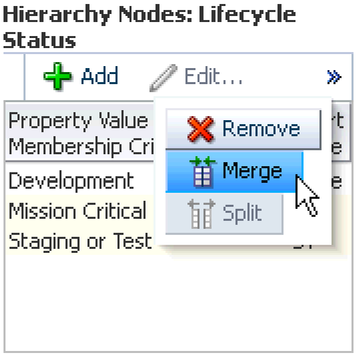

   The Merge Values dialog displays.

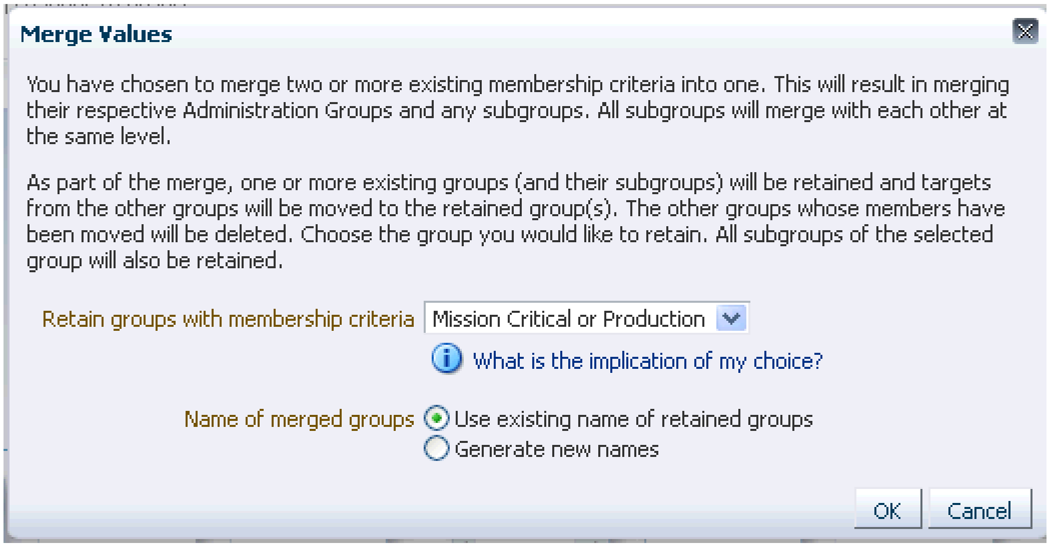

Merge Values

You have chosen to merge two or more existing membership criteria into one. This will result in merging their respective Administration Groups and any subgroups. All subgroups will merge with each other at the same level.

As part of the merge, one or more existing groups (and their subgroups) will be retained and targets from the other groups will be moved to the retained group(s). The other groups whose members have been moved will be deleted. Choose the group you would like to retain. All subgroups of the selected group will also be retained.

Retain groups with membership criteria: Mission Critical or Production

What is the implication of my choice?

Name of merged groups: ● Use existing name of retained groups ○ Generate new names

Again, by merging membership criteria (target properties), you are merging administration groups and their respective subgroups. You choose the administration group to be retained. The other groups will be merged into that group.

4. Choose the group you wish to retain and specify whether you want to use the existing name of the retained group or specify a new name.

> **Note:**
>
> When deciding which group to retain, consider choosing the group that is used in most group operations such as incident rule sets, system dashboard, or roles. These groups will be retained and the members of the other merged groups will join the retained groups. After the merge, group operations on the retained groups will also now apply to the members from the other merged groups. Doing so minimizes the impact of the merge.

5. Click **OK** to merge the groups.
6. Click **Update** to save the new hierarchy.

Example

Your administration group consists of the following:

Hierarchy Levels

- Lifecycle Status
- Line of Business

Hierarchy Nodes

- *Lifecycle Status*
  - Development
  - Mission Critical or Production
  - Staging or Test
- *Line of Business*

- Online Store
- Sales
- Finance

The following graphic shows the administration group hierarchy.

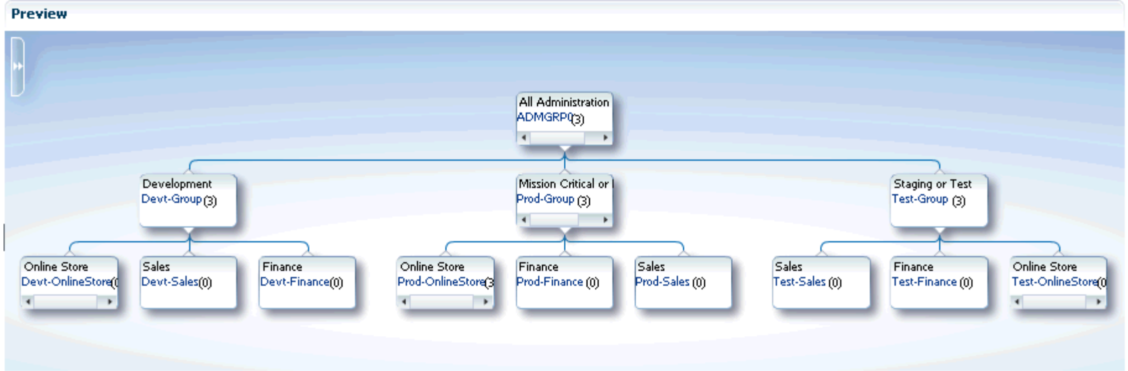

You decide that you want to merge the *Mission Critical or Production* group with the *Staging or Test* group because they have the same monitoring settings.

Choose Lifecycle Status from the **Hierarchy Levels** table.

From the **Hierarchy Nodes** table, choose both *Mission Critical or Production* and *Staging or Test*.

Select **Merge** from the **Hierarchy Nodes** menu.

The **Merge Values** dialog displays. In this case, you want to keep original name (*Mission Critical or Production*) of the retained group.

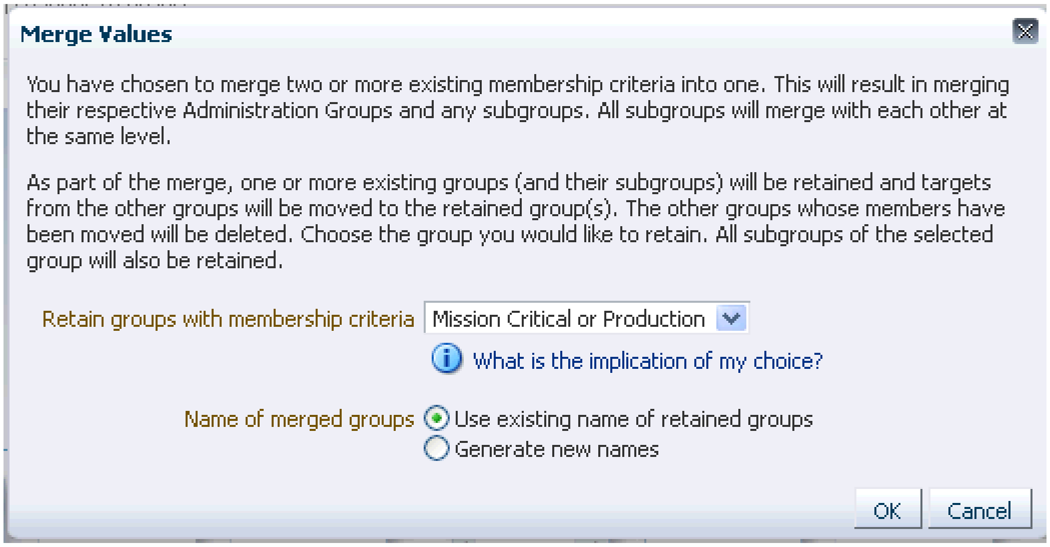

Merge Values

You have chosen to merge two or more existing membership criteria into one. This will result in merging their respective Administration Groups and any subgroups. All subgroups will merge with each other at the same level.

As part of the merge, one or more existing groups (and their subgroups) will be retained and targets from the other groups will be moved to the retained group(s). The other groups whose members have been moved will be deleted. Choose the group you would like to retain. All subgroups of the selected group will also be retained.

Retain groups with membership criteria: Mission Critical or Production

What is the implication of my choice?

Name of merged groups: ● Use existing name of retained groups
○ Generate new names

After clicking **OK** to complete the merge, the resulting administration group hierarchy is displayed. All targets from *Test-Sales* group moved to the *Prod-Sales* group. The *Test-Sales* group was deleted. All targets from the *Test-Finance* group moved to the *Prod-Finance* group. The *Test-Finance* group got deleted.

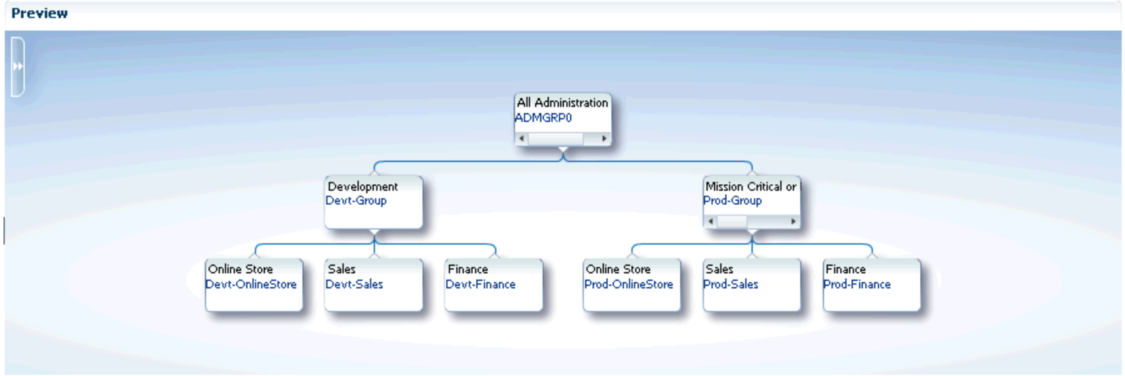

Click **Update** to save the changes.

Removing Administration Groups

You can completely remove an administration group hierarchy or just individual administration groups from the hierarchy. Deleting an administration group will not delete targets or template collections, but it will remove associations. Any stored membership criteria is removed. When you delete an administration group, any stored membership criteria is removed.

To remove the entire administration group hierarchy:

1. From the **Setup** menu, select **Add Target**, then select **Administration Groups**.
2. Click on the **Hierarchy** tab.
3. Click **Delete**.

To remove individual administration groups from the hierarchy:

1. From the **Setup** menu, choose **Add Target**, then select **Administration Groups**.
2. Click on the **Hierarchy** tab.
3. From the **Hierarchy Levels** table, choose the target property that corresponds to the hierarchy level containing the administration group to be removed.
4. From the **Hierarchy Nodes** table, select the administration group (**Property Value for Membership Criteria**) to be removed.
5. Choose **Remove** from the drop-down menu.
6. Click **Update**.

11 Using Dashboards

Enterprise Manager dashboards are data visualization tools that gather real-time data from the enterprise and display them in easy-to-interpret widgets. Dashboards enable you to monitor your IT infrastructure or specific areas or targets at a glance, and quickly identify outliers and take corrective action.

To access **Dashboards** in Enterprise Manager, open the navigation menu (upper-left corner), select **Enterprise**, and then select **Dashboards**. On the **Dashboards** page, the list of out-of-the-box dashboards and custom dashboards, if any, is displayed. On this page, you can:

- Click **Create dashboard** to create a custom dashboard. For information, see Creating a Custom Dashboard.
- Click **Import dashboards** to import a dashboard. For information, see Exporting and Importing Dashboards.
- Click **Show widgets** to view the list of available out-of-the-box and custom widgets. On the **Widgets** page, you can click the ⋮ icon for a custom widget and click the available options to perform tasks such as duplicating or deleting the widget. Note that these options are currently disabled for out-of-the-box widgets.
- Click **Show filters** to view the list of available out-of-the-box and custom filters. On the **Filters** page, you can click the ⋮ icon for an out-of-the-box filter and click **Duplicate** to duplicate the filter. For custom filters, you can click the ⋮ icon and click the available options to perform tasks such as duplicating or deleting the filter.
- Select an option in the **Created by** or **Updated by** drop-down lists or enter text in the **Search** field to filter the list of dashboards. For example, select **Oracle** in the **Created by** drop-down list to only view out-of-the-box dashboards.
- Sort the list of dashboards by clicking the **Sort ascending** or **Sort descending** icon (▾ or ▴) in the column headings.
- Enter text in the **Search...** field to view the dashboards that have the text in the name or description.
- Click the name of a dashboard to open it.
- Click the ⋮ icon for a dashboard and click the available options to perform tasks such as opening the dashboard in the view or edit mode, duplicating the dashboard, or deleting the dashboard. For out-of-the-box dashboards, only the **Open** and **Duplicate** options are enabled in this menu.

For Dashboards-related new features available in Enterprise Manager 24ai Release 1 (24.1), see Enterprise Monitoring New Features in *Enterprise Manager Introduction*.

Note:

In addition to Enterprise Manager dashboards, you can create dashboards using Grafana, which is an open-source technology used for metric analytics and visualization. In the Oracle Enterprise Manager App for Grafana, you can extract OMS repository metric data and display it graphically for fast, intuitive access to performance and metric information. For information on enabling the Oracle Enterprise Manager App for Grafana, see Enable the Oracle Enterprise Manager App for Grafana.

Out-of-the-Box Enterprise Manager Dashboards

The out-of-the-box Enterprise Manager dashboards provide a summary of the health, compliance, and other details of the assets in your enterprise or the targets that are a part of a group. The data in these dashboards is organized in widgets and presented using visualization options such as graphs, thereby facilitating quick and easy interpretation.

The following out-of-the-box dashboards are listed on the **Dashboards** page and you can click the name of the dashboard to access it and monitor the data available in the widgets:

- **Target Summary**: Provides an overview of the health and performance of a target. In this dashboard, you can select a single target and monitor its availability, a summary of open incidents, and key metrics over the last 24 hours. In addition, the tabs below the widgets display historical metric data, the list of incidents, compliance standards, and jobs for the target. In this dashboard, you can click the titles and contents of the widgets to go to the corresponding Enterprise Manager page to obtain additional inputs. For example, if monitoring a database instance in the **Target Summary** dashboard, click the title of the **Availability** widget to go to the **Availability (Status History)** page and obtain detailed availability information for the database.
- **Enterprise Summary**: Provides an overview of the assets in the enterprise, their health, and compliance. The widgets in this dashboard are organized within three major categories, **Overview**, **Infrastructure**, and **Compliance**. The widgets display information such as the status of the targets, inventory data such as the number of hosts and databases, compliance information, number of incidents and jobs, patch recommendations, and version details.
- **Group Summary**: Provides an overview of the health, top metrics, and compliance information of the targets in a particular group. As a prerequisite task to using this dashboard, an administrator user must create a group. For information on creating a group, see Creating and Editing Groups.
  On accessing this dashboard and selecting a group in the **Group** drop-down list, you can monitor target status, target types, incident trend, compliance information, and inventory data. In addition, you can also monitor the top metrics for hosts, databases, and WebLogic server.
- **Federation Summary – By Sites**: Provides a consolidated view of targets, incidents, jobs, and compliance data across a list of federated Enterprise Manager sites. It can be further filtered by a group that exists on multiple sites.
- **Federation Summary – By Groups**: Provides a consolidated view of targets, incidents, jobs, and compliance data for selected groups across a list of federated Enterprise Manager sites.

Note that before you use the out-of-the-box **Federation Summary** dashboards, you must configure Enterprise Manager Federation. For information, see Configuring Enterprise

Manager Federation in *Oracle Enterprise Manager Advanced Installation and Configuration Guide*.

In addition to the out-of-the-box dashboards listed above, you can use the out-of-the-box Exadata dashboards to monitor Exadata fleet configuration and capacity. These dashboards are a part of the Exadata Management Pack and Zero Data Loss Recovery Appliance Management Pack. For more information, see Exadata Fleet Dashboards in *Oracle Enterprise Manager Advanced Management and Monitoring for Engineered Systems*.

In the out-of-the-box dashboards, click **Actions** and click one of the following options:

- **Duplicate**: Create a duplicate of the out-of-the-box dashboard and customize it to meet your requirements.
- **Print**: Print the dashboard.
- **Auto-refresh**: Specify the frequency at which the data in the dashboard must be refreshed automatically.
- **Refresh**: Refresh the data in the dashboard.

In the out-of-the-box widgets, click the data or value in a widget to navigate to the corresponding Enterprise Manager page. For example, if you click the number in the **Databases** widget, the **All Targets** page, which is refined to only list database targets, is displayed.

Customizing an Out-of-the-Box Dashboard

You can customize out-of-the-box dashboards in Enterprise Manager to meet your specific requirements.

By customizing an out-of-the-box dashboard, you can add filters and widgets, remove the widgets that you do not need, alter the size of the widgets and change the layout to ensure that the widgets that have data of importance are prominently displayed.

To customize a dashboard and add a widget:

1. Navigate to the **Dashboards** page.
2. For the out-of-the-box dashboard that you want to customize, click the ⋮ icon and click **Duplicate**.

   > **Note:**
   >
   > To duplicate the out-of-the-box **Target Summary** dashboard, you must first open it, specify **Target Type** and **Target Name**, click **Actions** in the upper-right corner, and then click **Duplicate**.

3. In the **Duplicate dashboard** dialog, specify a name for the duplicate dashboard and click **OK**.
4. Click the name of the duplicate dashboard on the **Dashboards** page to open it.
5. Click **Actions** and then click **Edit** to go to the edit mode.

   In the edit mode, click the ⋮ icon in the upper-right corner of each widget to delete the widget, alter its size, or change the layout by moving the widget.
6. Optionally, click **+ Add time selector** to add a time selector that will allow you to select a duration of time for which you want to display data in the widgets in the dashboard.

7. Optionally, in the **Filters** tab, click or drag and drop the filter you want to add to the dashboard and select an option in the drop-down list to narrow down and view data within the context of the selected option.
8. Optionally, in the **Widgets** tab, create a custom widget or click or drag and drop an out-of the box widget to add it to the dashboard.

   For information on:

   - Adding an out-of the box widget, see step *7* in Creating a Custom Dashboard.
   - Creating a custom widget and adding it to the dashboard, see Creating a Custom Query-based Widget.
9. Click **Save changes** to save the dashboard.

The dashboard is in view mode. In the view mode, click the ⋮ icon on a widget and then click **Export to CSV** to export widget data to a CSV file. Note that the **Export to CSV** option is disabled if the widget does not support this feature.
If you want to make changes to the contents of the dashboard, click **Actions** in the upper-right corner of the dashboard and then click **Edit** to go to the edit mode.

Creating a Custom Dashboard

You can create a custom dashboard by adding the out-of-the-box widgets in the widget library and creating and adding custom query-based widgets.

Before you create a dashboard, you must ensure that you have the `CREATE_ANY_DASHBOARD` privilege. An Enterprise Manager Super Administrator can grant the `CREATE_ANY_DASHBOARD` privilege to a user or role using the `emcli grant_privs` command. Alternatively, a Super Administrator can grant the *Create Dashboards* privilege using the Dashboards resource type in the **Security** user interface.

To create a custom dashboard:

1. Navigate to the **Dashboards** page.
2. Click **Create Dashboard**.

   A blank untitled dashboard is displayed with the **About**, **Widgets**, and **Filters** tabs on the right.
3. Optionally, in the upper-right side of the dashboard, select the duration of time for which you want to display data in the widgets in the dashboard. You can click the icons adjacent to the time selector to edit or delete it.
4. Click the **About** tab, specify a name for the dashboard, and optionally add a description for the dashboard.
5. Optionally, click the **Filters** tab, click or drag and drop the filter you want to add to the dashboard, and select an option in the drop-down list to narrow down and view data within the context of the selected option.
6. Click the **Widgets** tab.

   The list of widgets in the widget library is displayed under **Add widgets**. You can enter text in the **Search...** field to view the widgets that have the text in the name or description.

   Optionally, click the **Launch widgets table view** icon (▦) to view the widget library in a tabular format. On the **Widget library table view** page, select an option in the **Created by** drop-down list or enter text in the **Search...** field to filter the list of widgets and click **Apply changes to widget library** to view the filtered list in the **Widgets** tab. For example, in the **Widget library table view** page, select a user name in the **Created by** drop-down list, and click **Apply changes to widget library** to only view the widgets created by that user in the **Widgets** tab.

7. Add widgets to the dashboard. To do so:

   - Drag and drop an out-of-the-box widget into the dashboard.
   - Click the **Add widget group and widgets** icon (+) in the **Widgets** tab and click **Create query-based widget** to create a custom query-based widget. For information, see Creating a Custom Query-based Widget.

   If you are adding an out-of-the-box widget, then you may have to manually configure inputs for relevant data to be displayed. In such cases, the **Configure <*Name*> Input for <*Widget Name*> (Required)** dialog is displayed when you add the widget to the dashboard. This dialog has the following options:

   - **Link with an existing filter**: Select to link the input to an existing filter in the dashboard. Note that this option is only displayed if a filter related to the input has already been added to the dashboard.
   - **Add a new filter**: Select to link the input with a filter available in the filter library.
   - **Specify the input**: Select to specify a fixed value of the input that will be used only in the context of this widget.

   > **Note:**
   >
   > Some of the required widget inputs such as **Aggregate Target Name** are automatically linked to the **Aggregate Target Name** filter and the **Configure <*Name*> Input for <*Widget Name*> (Required)** dialog is not displayed. For example, when adding the **Group - Target Types - Pie** widget, you are not prompted to configure inputs for the widget and the **Aggregate Target Name** filter is automatically added to the filter panel. If you want to make changes to the widget input that is automatically linked to a filter, you can do so in the **Edit Widgets** tab.

   The **Edit Widgets** tab on the right displays the details of the widget you have added. This includes the list of required inputs, which are denoted by an asterisk. Some of the inputs are configured when adding the widget and the others may have to be configured by clicking the **Edit** icon (✎). On clicking the **Edit** icon, the **Configure <*Name*> Input for <*Widget Name*>...** dialog is displayed and you have the option of linking the input to an existing filter in the dashboard, adding a new filter, or specifying a fixed value as input.

   If you have added a custom widget, then you have the additional option of clicking **Edit Widget** to go to the widget builder and edit the input parameters or visualization options.

8. Click **Save changes** to save the dashboard.

After you click **Save changes**, the dashboard is in view mode. In the view mode, click the ⋮ icon on a widget and then click **Export to CSV** to export widget data to a CSV file. Note that the **Export to CSV** option is disabled if the widget does not support this feature.

If you want to make changes to the contents of your dashboard, click **Actions** in the upper-right corner of the dashboard and then click **Edit** to go to the edit mode.

You can allow other Enterprise Manager users to view or edit the dashboard you have created. For information, see Editing Dashboard Privileges.

Using Out-of-the-Box Metric Widgets

You can add out-of-the-box metric widgets to a dashboard and visualize and display the metrics stored in Enterprise Manager.

The out-of-the-box metric widgets access and obtain the metric data uploaded to the Enterprise Manager Management Repository and display it in a visual format. The visualization of metric data in a widget makes it easy to monitor and interpret. Depending on the duration for which you want to display metric data in the widgets, it is fetched from the following sources:

- Metric data for a duration of less than 2 days is fetched from RAW data.
- Metric data for a duration of 2 to 14 days is fetched from the hourly roll-up table.
- Metric data for a duration of more than 14 days is fetched from the daily roll-up table.

Here is a scenario that illustrates how to create a new dashboard and use the **Target Metric - Line** metric widget to display the **Active Sessions: CPU** metric for a specific Pluggable Database (PDB):

1. Navigate to the **Dashboards** page.
2. Click **Create Dashboard** to create a new dashboard.
   A blank untitled dashboard is displayed.
3. In the **About** tab, specify a name for the dashboard, and optionally add a description for the dashboard.
4. In the **Widgets** tab, type **Target Metric - Line** in the search field to view the widget in the list of out-of-the-box widgets.
5. Click or drag and drop the **Target Metric - Line** widget to add it to the new dashboard.
   As you add the widget to the dashboard, you must configure the inputs to display the **Active Sessions : CPU** metric for a specific PDB. Here is more information on the options you can select:
   a. In the **Configure Target Type input for Target Metric - Line (Required)** dialog, select the **Specify the Target Type input** option and select the **Pluggable Database** option in the drop-down list.
   b. In the **Configure Target Name input for Target Metric - Line (Required)** dialog, select the **Specify the Target Name input** option and select the specific PDB in the drop-down list.
   c. In the **Configure Metric Name input for Target Metric - Line (Required)** dialog, select the **Specify the Metric Name input** option and select **Wait Count** in the drop-down list.
   d. In the **Configure Metric Column input for Target Metric - Line (Required)** dialog, select the **Specify the Metric Column input** option and select **Active Sessions : CPU** in the drop-down list.

   On configuring these inputs, the **Target Metric - Line** widget is added to the dashboard and displays the number of active sessions using CPU for the specified PDB in a line chart. The **Edit Widgets** tab on the right displays the details of the widget, and this includes the list of inputs configured for the widget. To make changes to the input configuration, click the **Edit** icon (✎) adjacent to the input. For example, in the **Target Metric - Line** widget, the time input is linked to the dashboard time by default, however, you can edit this input and specify a different duration of time.
6. Click **Save changes** to save the dashboard.

Similar to the **Target Metric - Line** widget that displays metric data for a single target, you can also use the **Group - Metrics - Line** widget that displays historical metric data for a target type. For the **Group - Metrics - Line** widget, in addition to the required inputs that must be configured when adding the widget, you can configure the following optional inputs to further customize the metric data:

- **Aggregate Target Type**: The **Group** option is selected by default, however, you can change the aggregate target type. Any target type that has been defined to have members, is listed as an aggregate target type.
- **Number of Targets**: The default number of top targets for which data is displayed is 5, however, you can change this to display data for more or fewer targets. The top <*n*> targets are determined based on the highest average metric value queried from the hourly roll-up table.
- **Rollup Column**: The pre-aggregated column that must be selected from the historical metric table. Historical metric data is pre-aggregated using standard functions such as average, min, and max. **Average** is the default function, however, you can edit the **Rollup Column** input and select the **Min** or **Max** functions.
- **Key Aggregate Function**: The function used to aggregate multiple key metric values for a target for a given time stamp. **Average** is the default function, however, you can edit the **Key Aggregate Function** input and select the **Min**, **Max**, or **Sum** functions.
  Here is a scenario that illustrates how **Key Aggregate Function** can be used. A Windows host target has multiple drives (keys), such as `C:` and `D:`, and the file system in the `C:` drive uses 1.5 GB and the file system in the `D:` drive uses 2 GB. You can select the **FileSystem Used** metric, which is a key metric for host targets and select **Sum** as the **Key Aggregate Function** to view the total file system usage in the Windows host at a given time: `1.5 + 2 = 3.5 GB`.
- **Time**: The time input is set to the last 7 days by default, however, you can specify a different duration of time for which metric data will be displayed.

Editing Dashboard Privileges

You can grant or revoke the privileges required to view, edit, and delete a custom dashboard on the **Dashboards** page in Enterprise Manager.

After a custom dashboard is created, an Enterprise Manager Super Administrator can edit privileges to grant the following dashboard privileges:

- `VIEW_DASHBOARD`: Allows a user to view the dashboard.
- `EDIT_DASHBOARD`: Allows a user to view and edit the dashboard.
- `FULL_DASHBOARD`: Allows a user to view, edit, and delete the dashboard.

To edit dashboard privileges:

1. Navigate to the **Dashboards** page.
2. Click the ⋮ icon for the custom dashboard and click **Edit privileges...**.

   > **Note:**
   >
   > The **Edit privileges...** option is disabled for out-of-the-box dashboards and if you are not the Enterprise Manager Super Administrator.

3. In the **Edit dashboard privileges** panel:

a. Select one or more users and the privileges you want to grant to the users, and click **Add**.
   Note that if assigning the **View Dashboard** privilege to a user, the **View Saved Search** (`VIEW_SAVED_SEARCH`) privilege, which allows the user to view the widgets in the dashboard, is also automatically assigned.

b. Review the selected users and privileges under **Dashboard Privileges**.

   To revoke the privilege granted to a user, click the 🗑 icon in the **Revoke** column.

c. Click **Save**.

d. Review the **Save dashboard privileges** message and click **OK**.

The changes to the dashboard privileges are saved.
Note that you can also use the `emcli grant_privs` command and the `DashboardID`, which is a part of the dashboard metadata identified by the `"id"` attribute, to grant dashboard privileges. Here is an example of the EMCLI command:

```
$ emcli grant_privs -name="<USER>" -
privilege="EDIT_DASHBOARD;DASHBOARD_ID=<DashboardID>"
```

Exporting and Importing Dashboards

You can export a custom dashboard along with its widgets and filters, and import it to a different environment.

> **Note:**
>
> You cannot export and import out-of-the-box dashboards.

Exporting a Dashboard

You can export a custom dashboard in JavaScript Object Notation (JSON) format.

1. Navigate to the **Dashboards** page.
2. Click the ⋮ icon for the dashboard that you want to export, and click **Export**.

   Alternatively, click the name of the dashboard that you want to export, click **Actions** in the upper-right corner, and click **Export**.

The JSON file with the configuration information of the dashboard is downloaded.

Importing a Dashboard

You can import a dashboard using the configuration information JSON file that was downloaded when exporting the dashboard.

When importing a dashboard, its widgets and filters are imported too.

To import a dashboard:

1. Navigate to the **Dashboards** page.
2. Click **Import dashboards**.

3. Select the configuration information JSON file of the dashboard you want to import, and click **Open**.

   The **Import dashboards** dialog is displayed.

4. In the **Import dashboards** dialog, review the message and click **Import**.

The dashboard is imported. If the imported dashboard or widget already exists, then it's updated.

Creating a Custom Query-based Widget

You can create a custom query-based widget and add it to a new or existing dashboard.

Before you create a widget, you must ensure that you have the `CREATE_SAVED_SEARCHES` privilege. An Enterprise Manager Super Administrator can grant the `CREATE_SAVED_SEARCHES` privilege to a user or role using the `emcli grant_privs` command. Alternatively, a Super Administrator can grant the *Create Saved Searches* privilege using the Saved Searches resource type in the **Security** user interface.

Other prerequisites for creating query-based widgets include a working knowledge of:

- SQL queries and how to create and edit them
- Oracle Enterprise Manager Management Repository views

The SQL query retrieves the data you want to view in the widget from an Enterprise Manager Management Repository view. These views provide access to target, metric, and monitoring information stored in the Management Repository. For information on Management Repository views, see About Management Repository Views in *Oracle Enterprise Manager Management Repository Views Reference*.

To access the widget builder to create a custom query-based widget:

1. Navigate to the **Dashboards** page.
2. Click **Create Dashboard** or open an existing dashboard in edit mode.
3. In the **Widgets** tab, under **Add widgets**, click the **Add widget group and widgets** icon (+), and then click **Create query-based widget**.
   A blank untitled widget is displayed in the widget builder.

You can use the following types of SQL queries as input to create a query-based widget:

- **Custom SQL**: Use the sample custom SQL query, edit it, or replace it with a new query to create a widget that meets your specific requirements. For more information, see Creating a Widget Using Custom SQL.

  > **Note:**
  >
  > If creating a widget using custom SQL, you can also select **EM Federation** as the data source and the SQL query then retrieves the data from federated Enterprise Manager sites. For information on configuring Enterprise Manager Federation, see Configuring Enterprise Manager Federation in *Oracle Enterprise Manager Advanced Installation and Configuration Guide*.

- **Predefined SQL**: Select from the out-of-the-box named SQL queries, which display the requested data in the widget. For more information, see Creating a Widget Using Predefined SQL.

On specifying the input to the widget in the widget builder, use the available visualization options to display data to support various scenarios.

> **Note:**
>
> - Only custom query-based widgets can be edited in the widget builder.
> - Any changes made to a custom query-based widget will impact its use in all the dashboards of which it is a part. For example, if a custom query-based widget is used in two dashboards and the widget is edited in one of the dashboards and changes are made to the parameter configuration, then this may result in the widget not displaying any data in the other dashboard.

You can allow other Enterprise Manager users to view or edit the widget you have created. To grant the required privileges, use the `emcli grant_privs` command and the `SavedSearchID`, which is a part of the saved search metadata identified by the `"id"` attribute. The privilege levels include:

- `VIEW_SAVED_SEARCH`: Allows a user to view the saved search.
- `EDIT_SAVED_SEARCH`: Allows a user to view and edit the saved search.
- `FULL_SAVED_SEARCH`: Allows a user to view, edit, and delete the saved search.

Here is an example of the EMCLI command:

```
$ emcli grant_privs -name="<USER>" -
privilege="EDIT_SAVED_SEARCH;SAVED_SEARCH_ID=<SavedSearchID>"
```

Creating a Widget Using Custom SQL

1. Navigate to the widget builder.
2. Optionally, click the **Open/close filter panel** icon (▼) in the upper-left corner to view the filters added to the dashboard and make changes, if required.
3. From the **Data Source** drop-down list, select one of the following data sources:
   - **EM Repository**: Select this option to query data from Enterprise Manager Management Repository views and tables. For information on Management Repository views, see About Management Repository Views in *Oracle Enterprise Manager Management Repository Views Reference*.
   - **EM Federation**: Select this option to query data from federated Enterprise Manager sites. For information on Enterprise Manager Federation, see Configuring Enterprise Manager Federation in *Oracle Enterprise Manager Advanced Installation and Configuration Guide*.

   > **Note:**
   >
   > The **EM Federation** data source option is only available when creating a widget using custom SQL. You cannot use this data source when creating a widget using predefined SQLs.

4. From the **SQL Selector** drop-down list, select **Custom SQL**.

   The **Custom SQL** option is selected by default and a sample query is added to the **SQL Query** field and the corresponding widget is displayed. In the sample query, the `SELECT` statement queries column names from the selected data source, **EM Repository** or **EM Federation**. The visualization options are also selected by default to visualize the data retrieved by the sample query.

5. In the **SQL Query** field, replace the sample query with a new query or customize the SQL query to meet specific requirements by adding custom input parameters for the widget.

   Optionally, to add a custom input parameter, you can replace the part of the query that you want to customize with a standard bind parameter marker denoted by a question mark (?). This is similar to JDBC bind parameter support. Custom input parameters are of two types, a **regular input parameter** to specify input such as target types and a **time input parameter**, which is essentially a composite parameter, to specify the duration of time (start and end time) for which data will be displayed in the widget.

   > **Note:**
   >
   > This is an advanced concept and some familiarity with the widget builder and bind parameters is required. If you are creating a widget using custom SQL for the first time, it is recommended that you go to step *6*.

   - **To add a regular input parameter and specify a custom value**:

     a. Replace the part of the query to which you want to add the regular input parameter with the standard bind parameter marker (?). Here is an example in which the target types in the custom SQL query are replaced with a bind parameter marker, so you can add custom input values for the target types you want to view in the widget.
     Replace the target types in bold:

```
SELECT
    target_type as "Target Type",
    type_display_name as "Type Display Name",
    count(target_type) as "Target Type Count"
FROM
    mgmt$target
WHERE
    target_type in (
        'oracle_database', 'weblogic_j2eeserver',

'host','oracle_apache','rac_database','oracle_exadata','exalytics_ma
chine'
    )
GROUP BY
    target_type,
    type_display_name
```

     with the bind parameter marker

```
SELECT
    target_type as "Target Type",
    type_display_name as "Type Display Name",
    count(target_type) as "Target Type Count"
FROM
    mgmt$target
WHERE
    target_type = ?
GROUP BY
    target_type,
    type_display_name
```

b. In the **Settings** tab in the lower-right corner, click **Add input +**.

c. In the **Configure input for Untitled Widget (Required)** dialog:

   i. **Is time input**: Ensure that this check box is *not* selected. Note that this check box should only be selected to specify time input parameters. For information, see To add a time input parameter and specify custom values.

   ii. **Input Name**: Enter a name for the input. For example, **TargetType**.

   iii. **Input Label**: Enter a display name for the input. The input name is displayed in this field and you can make changes to it, if required.

   iv. **Parameter Indexes**: Enter the index value to map the input parameter's value to the bind parameter marker in the SQL query.

   v. **Default settings**: Review and make changes to the default settings, if required.

   > **Note:**
   >
   > If you update the **Default settings** to configure the input parameter, the same settings will be applied whenever this widget is added to a different dashboard. To configure an input parameter whose value can be edited after the widget is created and added to the widget library, use the options in the **Configure widget input** section of this dialog.

   - **Required**: Deselect to indicate that this is not a required widget input. By default, this check box is selected and if you do not configure the input parameter, data will not be displayed in the widget.

   - **Allow multiple values**: Select to allow multiple values to be configured for the input parameter. Here is an example:

     ```
     WHERE
         TARGET_TYPE IN (select /*+ Cardinality(theArray 10)*/
     COLUMN_VALUE from (TABLE (CAST(? AS
     MGMT_MEDIUM_STRING_TABLE)) theArray) )
     ```

   - **Allow widget users to link the input with a filter**: Deselect this check box to only specify a fixed input parameter value. This check box is selected by default and if you deselect it, the fields available to link the input to a filter are not displayed in the dialog.

   - **Allow widget users to specify the input**: Deselect this check box to only link the input to a filter. This check box is selected by default and if you deselect it, the fields available to specify a fixed input parameter value are not displayed in the dialog.

vi. **Configure widget input**: Specify the value for the regular input parameter:
   - **Link the <*name*> input with an existing filter**: Select to link the input parameter to an existing filter in the widget.
   - **Specify the <*name*> input**: Select to specify a fixed input value.
   - **Not configured**: Select this option to leave the input unconfigured.

vii. Click **Save changes**.

- **To add a time input parameter and specify custom values**:

  a. Add bind parameter markers for a time input parameter in the custom SQL query. Here are a couple of examples:
     - `WHERE` clause

       ```
       WHERE mtx.collection_timestamp > CAST(TO_UTC_TIMESTAMP_TZ(?) at
       time zone tgt.timezone_region AS DATE)
       AND mtx.collection_timestamp < CAST(TO_UTC_TIMESTAMP_TZ(?) at
       time zone tgt.timezone_region AS DATE)
       ```

       In this example, `mtx.collection_timestamp` is in the target time zone and the bind variable (for both the start and end time represented by ?) is the UTC time zone string that is cast to the target time zone for comparison with `collection_timestamp`. When constructing the `WHERE` clause for a time input parameter:
       * Avoid using a function around timestamp columns that are indexed or partitioned, as this may nullify the usage of the index or partition. If you are bounding timestamps by an upper or lower value, apply the function to that side of the equality operation.
       * Ensure that the `TIMESTAMP` and `DATE` values contain no built-in time zone information but can be cast from one time zone to another.
     - `SELECT` clause

       ```
       SELECT FROM_TZ(CAST(mmh.rollup_timestamp as TIMESTAMP),
       mmh.timezone_region) AT TIME ZONE 'UTC' from mgmt$metric_hourly
       mmh
       WHERE ...
       ```

       In this example, a column select statement is used to convert time in one particular time zone (`timezone_region`) to the UTC time zone and return a value in the `TIMESTAMP WITH TIME ZONE` format.

       > **Note:**
       >
       > You must ensure that the selected `TIMESTAMP` and `DATE` values are returned in the correct time zone as various columns and tables in Enterprise Manager store `TIMESTAMP` and `DATE` values differently.

  b. In the **Settings** tab in the lower-right corner, click **Add input +**.

  c. In the **Configure input for Untitled Widget (Required)** dialog:

     i. **Is time input**: Select this check box.

ii. **Input Name**: This field is not editable and displays **time**.

iii. **Input Label**: Edit the default display name, **Time**, if required.

iv. **Start Time Indexes**: Enter the index value to map the time input parameter's start time value to the bind parameter marker in the SQL query.

v. **End Time Indexes**: Enter the index value to map the time input parameter's end time value to the bind parameter marker in the SQL query.

vi. **Default settings**: Review and make changes to the default settings, if required.

> **Note:**
>
> If you update the **Default settings** to configure the input parameter, the same settings will be applied whenever this widget is added to a different dashboard. To configure an input parameter whose value can be edited after the widget is created and added to the widget library, use the options in the **Configure widget input** section of this dialog.

For information on the **Default settings**, see To add a regular input parameter and specify a custom value.

vii. **Configure widget input**: Specify the values for the time input parameter:

- **Link with Dashboard time**: Select to link the time input parameter values with the time range selected for the dashboard.
- **Specify the Time input**: Select this option to specify a custom time range.
- **Not configured**: Select this option to leave the input unconfigured.

viii. Click **Save changes**.

On adding the regular or time input parameters to the SQL query, you can view the input parameter and value in **Configured widget inputs** in the **Settings** tab. You can also view the bind index mapped to the input parameter in **Parameter Index Mapping** adjacent to the SQL query.

6. Review the changes made to the SQL query and click **Run**.

   On clicking **Run**, the query is executed and data is displayed in the widget. Below the widget, a table with the various attributes of the raw data for the widget is displayed and you can click **Hide raw data** to hide it.

7. Use the other tabs in the lower-right corner to perform the following tasks:

> **Note:**
>
> You can click **JSON** in the upper-right corner to specify or edit the details on the **About**, **Visualization**, and **Settings** tabs, in a JSON editor. The JSON editor provides you with greater flexibility and more options to visualize data in the widget, however, it's recommended that you use it only if you are familiar with editing widget JSON. For more information on the JSON editor, contact Oracle Support.

- **About**: Add a name and description for the widget.

- **Visualization**: Select a chart type and customize the visualization by specifying or modifying additional options.
  If you select **Table** as the chart type, then you have the option of removing some of the columns or select the **All Columns** check box to display all columns. Chart types that have Y and X axes, such as the **Line** chart have the following additional visualization options:
  - **Y Axis**: Select the data attribute you want projected on the Y axis. Y axis supports numeric data attributes and only those are listed.
  - **X Axis**: Select the data attribute you want projected on the X axis.
  - **Series**: Select the data attribute that you want to plot in separate series in the chart. To use this option, add the `GROUP BY` clause for a certain data attribute in the query and view the categories of data in the widget. This option is also available for a **Pie** chart.
  - **Color By**: Select the data attribute for which you want to view color distinct values in the chart. This option is also available for a **Pie** chart.
  - **Y Axis Title**: Specify a title for the data attribute projected on the Y axis.
  - **Legend**: Specify the location of the legend in the widget. This option is also available for a **Pie** chart.
  - **Stacked**: Select this option if you want to stack the data attributes displayed in the chart, by color.
  - **Use solid color**: When you use an **Area** chart for visualization, the area that represents the presence of data is covered in solid color. This option is selected by default.

8. Click **Save** to save the widget.

After you save the widget, it is added to the dashboard and is also listed in the **Widgets** tab.

Click the ⋮ icon in the upper-right corner of the widget and click **Edit** to view the widget in the **Edit widgets** tab and click **Edit Widget** to go to the widget builder and edit the input parameters or visualization options.

Creating a Widget Using Predefined SQL

1. Navigate to the widget builder.
2. Optionally, click the **Open/close filter panel** icon (▼) in the upper-left corner to view the filters added to the dashboard and make changes, if required.
3. From the **SQL Selector** drop-down list, select an option listed under **Predefined SQLs**.

   The predefined SQLs are Oracle-defined named SQL queries, which display the requested data in the widget. When you select a predefined SQL query, the query is displayed for your reference in the **SQL Query** field. Click the tooltip adjacent to the **SQL Selector** drop-down list to view a description.
4. Optionally, select the **Convert to Custom SQL** check box to convert the predefined SQL query into a custom SQL query and customize it. On selecting this check box, the index values mapped to the input parameters in the SQL query are displayed in **Parameter Index Mapping** adjacent to the SQL query and the configured widget inputs are displayed in the **Settings** tab in the lower-right corner. For information on working with custom SQL queries, see Creating a Widget Using Custom SQL.

5. In the **Settings** tab in the lower-right corner, click the **Edit** icon (✎) adjacent to the widget inputs to configure the input and display relevant data in the widget, if required. Note that the list of required inputs are denoted by an asterisk.

   For example, if you select **Group - Incidents count of last 7 days** in the **SQL Selector** drop-down list, then the **Group Type** and **Group Name** widget inputs are listed in the **Settings** tab and must be configured. Here is how you can configure the widget input parameters for **Group - Incidents count of last 7 days**:

   a. Click the **Edit** icon (✎) adjacent to the **Group Type** input in the **Settings** tab.

   b. In the **Configure Group Type input for Untitled Widget (Required)** dialog, specify the **Group Type** input in one of the following sections:

      - **Default settings**: Review the default configuration settings for the input and select a filter. By default, the **Allow widget users to link the Group Type input with a filter** and **Allow widget users to specify the Group Type input** check boxes are selected and you can select a default filter and a value to link the input with an available filter. This is useful when the widget is added to a dashboard that uses the same filter, and this configuration setting will apply to future widgets.

      - **Configure widget input**: Select one of the following options to configure specific inputs that will only apply to this widget:

        - **Link the Group Type input with an existing filter**: Select this option to link the input with an existing filter. Note that this option is only displayed if a filter is already added to the dashboard in which you are creating the widget.

        - **Specify the Group Type input**: Select this option to specify a fixed value as the input for the widget.

        - **Not configured**: Select this option to leave the input unconfigured.

   c. Click **Save changes** to save the configuration for the **Group Type** input.

   On configuring the **Group Type** input, the **Edit** icon (✎) adjacent to the **Group Name** input is enabled and you have to perform the same set of steps to configure it. After the input parameters are configured, data is displayed in the widget. Below the widget, a table with the various attributes of the raw data for the widget is displayed and you can click **Hide raw data** to hide it.

6. Use the other tabs in the lower-right corner to perform the following tasks:

   > **Note:**
   >
   > You can click **JSON** in the upper-right corner to specify or edit the details on the **About**, **Visualization**, and **Settings** tabs, in a JSON editor. The JSON editor provides you with greater flexibility and more options to visualize data in the widget, however, it's recommended that you use it only if you are familiar with editing widget JSON. For more information on the JSON editor, contact Oracle Support.

   - **About**: Add a name and description for the widget.

   - **Visualization**: Select a chart type and customize the visualization by specifying or modifying additional options.
     If you select **Table** as the chart type, then you have the option of removing some of the columns or select the **All Columns** check box to display all columns. Chart types that

have Y and X axes, such as the **Line** chart have the following additional visualization options:

- **Y Axis**: Select the data attribute you want projected on the Y axis. Y axis supports numeric data attributes and only those are listed.
- **X Axis**: Select the attribute of the data you want projected on the X axis.
- **Series**: Select the data attribute that you want to plot in separate series in the chart. Typically, the `GROUP BY` clause in a custom SQL query can be used for series. This option is also available for a **Pie** chart.
- **Color By**: Select the data attribute for which you want to view color distinct values in the chart. This option is also available for a **Pie** chart.
- **Y Axis Title**: Specify a title for the data attribute projected on the Y axis.
- **Legend**: Specify the location of the legend in the widget. This option is also available for a **Pie** chart.
- **Stacked**: Select this option if you want to stack the data attributes displayed in the chart, by color.
- **Use solid color**: When you use an **Area** chart for visualization, the area that represents the presence of data is covered in solid color. This option is selected by default.

7. Click **Save** to save the widget.

After you save the widget, it is added to the dashboard and is also listed in the **Widgets** tab.

Click the ⋮ icon in the upper-right corner of the widget and click **Edit** to view the widget in the **Edit widgets** tab and click **Edit Widget** to go to the widget builder and edit the input parameters or visualization options.

Adding a Widget Group

You can add a widget group container to group widgets in a dashboard.

The widget group can be used to group related widgets in a new or existing dashboard. For example, you can add a widget group that contains all the compliance-related widgets in the dashboard.

To add a widget group to a dashboard:

1. Navigate to the **Dashboards** page.
2. Click **Create Dashboard** or open an existing dashboard in edit mode.
3. In the **Widgets** tab, under **Add widgets**, click the **Add widget group and widgets** icon (+), and then click **Add widget group**.
   A blank widget group container is added to the dashboard.
4. Drag and drop the widgets you want to group together into the widget group container.
5. Click **Edit widgets** on the **Widgets** tab to view and specify a name for the widget group, and review and configure inputs for the widgets in the group, if required.

   Optionally, click the ⋮ icon on the widget group to perform tasks such as altering its size or placement, deleting or duplicating it, and ungrouping all the widgets in the group. Click the ⋮ icon on a single widget in the widget group and click **Ungroup** to remove that widget from the group.
6. Click **Save changes** to save the dashboard.

12 Using Monitoring Templates

Monitoring templates simplify the task of setting up monitoring for large numbers of targets by allowing you to specify the monitoring and Metric and Collection Settings once and applying them to many groups of targets as often as needed.

This chapter covers the following topics:

About Monitoring Templates

Monitoring templates let you standardize monitoring settings across your enterprise by allowing you to specify the monitoring settings once and apply them to your monitored targets. You can save, edit, and apply these templates across one or more targets or groups. A monitoring template is specified for a particular target type and can only be applied to targets of the same type. For example, you can define one monitoring template for test databases and another monitoring template for production databases.

A monitoring template defines all Enterprise Manager parameters you would normally set to monitor a target, such as:

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

Once a monitoring template is defined, it can be applied to your targets. This can be done either manually through the Enterprise Manager console, via the command line interface (EM CLI), or automatically using template collections. See "Defining Template Collections" for more information. For any target, you can preserve custom monitoring settings by specifying metric settings that can never be overwritten by a template.

Oracle-Certified Templates

In addition to templates that you create, there are also Oracle-certified templates. These templates contain a specific set of metrics for a specific purpose. The purpose of the template is indicated in the description associated with the template.

Example: The template called *Oracle Certified - Enable AQ Metrics for SI Database* contains metrics related to Advanced Queueing for single instance databases. You can use this Oracle-certified template if you want to use the AQ metrics. Or you can copy the metric settings into your own template.

Definition of a Monitoring Template

A monitoring template defines all Enterprise Manager parameters you would normally set to monitor a target. A template specifies:

- **Name**: A unique identifier for the template. The template name must be globally unique across all templates defined within Enterprise Manager.
- **Description**: Optional text describing the purpose of the template.
- **Target Type**: Target type to which the template applies.
- **Owner**: Enterprise Manager administrator who created the template.
- **Metrics**: Metrics for the target type. A monitoring template allows you to specify a subset of all metrics for a target type. With these metrics, you can specify thresholds, collection schedules and corrective actions.
- **Other Collected Items**: Additional collected information (non-metric) about your environment.

Default Templates (Auto Apply Templates)

Under certain circumstances, Oracle's out-of-box monitoring settings may not be appropriate for targets in your monitored environment. Incompatible Metric and Collection Settings for specific target types can result in unwanted/unintended alert notifications. Enterprise Manager allows you to set default monitoring templates that are automatically applied to newly added targets, thus allowing you to apply monitoring settings that are appropriate for your monitored environment.

> **Note:**
>
> Super Administrator privileges are required to define default monitoring templates.

Viewing a List of Monitoring Templates

To view a list of all Monitoring Templates, from the **Enterprise** menu, select **Monitoring** and then **Monitoring Templates**.

The Monitoring Templates page displays all the out-of-box templates and the templates for which you have has at least VIEW privilege on. Enterprise Manager Super Administrators can view all templates.

Figure 12-1 Monitoring Templates

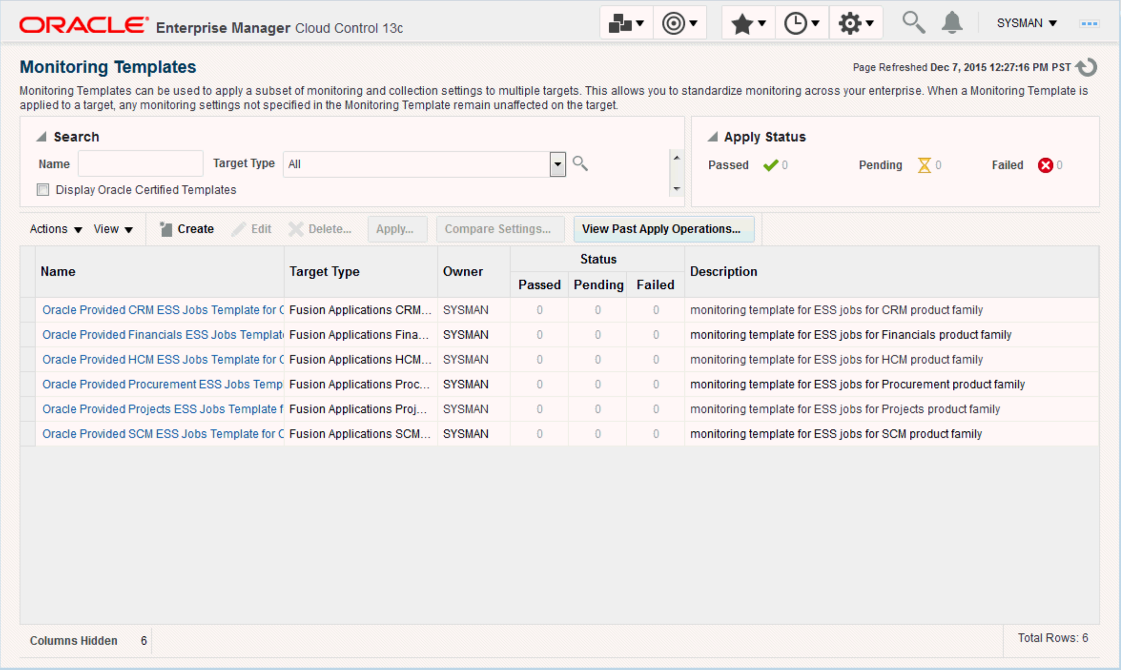

Monitoring Templates

Page Refreshed Dec 7, 2015 12:27:16 PM PST

Monitoring Templates can be used to apply a subset of monitoring and collection settings to multiple targets. This allows you to standardize monitoring across your enterprise. When a Monitoring Template is applied to a target, any monitoring settings not specified in the Monitoring Template remain unaffected on the target.

Search

Name   Target Type All

☐ Display Oracle Certified Templates

Apply Status

Passed ✓ 0   Pending ⌛ 0   Failed ✖ 0

Actions ▼ View ▼ Create Edit Delete... Apply... Compare Settings... View Past Apply Operations...

Name	Target Type	Owner	Status			Description
			Passed	Pending	Failed	
Oracle Provided CRM ESS Jobs Template for (	Fusion Applications CRM...	SYSMAN	0	0	0	monitoring template for ESS jobs for CRM product family
Oracle Provided Financials ESS Jobs Templat	Fusion Applications Fina...	SYSMAN	0	0	0	monitoring template for ESS jobs for Financials product family
Oracle Provided HCM ESS Jobs Template for (	Fusion Applications HCM...	SYSMAN	0	0	0	monitoring template for ESS jobs for HCM product family
Oracle Provided Procurement ESS Jobs Temp	Fusion Applications Proc...	SYSMAN	0	0	0	monitoring template for ESS jobs for Procurement product family
Oracle Provided Projects ESS Jobs Template f	Fusion Applications Proj...	SYSMAN	0	0	0	monitoring template for ESS jobs for Projects product family
Oracle Provided SCM ESS Jobs Template for (	Fusion Applications SCM...	SYSMAN	0	0	0	monitoring template for ESS jobs for SCM product family

Columns Hidden 6   Total Rows: 6

You can begin the monitoring template creation process from this page.

Creating a Monitoring Template

Monitoring templates allow you to define and save monitoring settings for specific target types. As such, specific Enterprise Manager privileges are required in order to create monitoring templates.

There are two resource privileges that can be granted to a user/role that allows you to create and/or view monitoring templates:

- *Create Monitoring Template*

  This privilege allows you to create a monitoring template.

- *View Any Monitoring Template*

  This privilege allows you to view any monitoring template.

These privileges can be granted from the Resource Privilege page of an Enterprise Manager user, or when creating a role.

Monitoring templates adhere to a typical access model: You can grant either FULL or VIEW access on a template to other users or roles. VIEW access allows you to see and use the monitoring template. FULL access allows you to see, use, edit and delete a monitoring template. The template owner can change access to a template.

By default, Enterprise Manager Super Administrators have FULL access on all monitoring templates.

Monitoring template allow you to define and save monitoring settings for specific target types. To define a new template:

1. From the **Enterprise** menu, select **Monitoring** and then **Monitoring Templates.**

2. Click **Create**. Enterprise Manager gives you the option of selecting either a specific target or a target type. Template monitoring settings are populated according to the selected target or target type. Click **Continue**.

> **Note:**
>
> If the selected target type is either Web Application or Service, you will only be able to select those targets for which you have Operator privilege.

3. Enter requisite template information on the General, Metric Thresholds, and Other Collected Items tabs.

   On the Metric Thresholds tab, you can delete or add monitoring template metrics. To delete existing metrics, select one or more metrics and click **Remove Metrics from Template**.

   To add metrics, click **Add Metrics to Template**. The Add Metrics to Template page displays as shown in the following graphic.

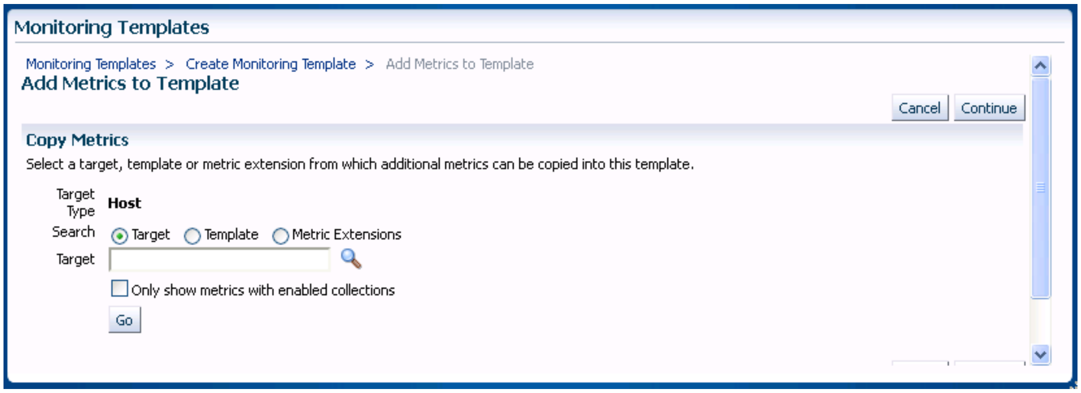
Monitoring Templates
Monitoring Templates > Create Monitoring Template > Add Metrics to Template
Add Metrics to Template
Cancel
Continue
Copy Metrics
Go

Select a target, template or metric extension from which additional metrics can be copied into this template.

Target Type: Host

Search: ● Target ○ Template ○ Metric Extensions

Target:

□ Only show metrics with enabled collections

   On this page, you can select a source from which you can copy metrics to the template. Sources include specific targets, other monitoring templates, or metric extensions. When adding metric to a template from another template, if the metric has adaptive settings, then the adaptive settings will not be copied: Only the metric will be copied along with existing thresholds.

   Note: You must define the metric extension thresholds in order to add it to a monitoring template.

   Click **Continue** once your have finished modifying the template metrics.

4. Once you have finished entering requisite information, click **OK**.

Editing a Monitoring Template

The Monitoring Templates page lists all viewable templates. To edit a template, you must have FULL access privileges.

To edit a Monitoring Template:

1. From the **Enterprise** menu, select **Monitoring**, and then **Monitoring Templates**.
2. Choose the desired template from the table.

3. Click **Edit**.
4. Once you have finished making changes, click **OK**.

Sharing Access with Other Users

By default, template owners (creators) have FULL access privileges on the template and Enterprise Manager Super Administrators have FULL access privileges on all templates. Only the template owner can change access to the template. You, as owner, can grant VIEW (view the template) or FULL (edit or delete the template) on the template to a user or role.

Applying Monitoring Templates to Targets

As mentioned earlier, a monitoring template can be applied to one or more targets of the same target type, or to composite targets such as groups. For composite targets, the template is applied to all member targets that are of the appropriate type. If you applied the template manually or via EM CLI, once a template is applied, future changes made to the template will not be automatically propagated to the targets: You must reapply the template to all affected targets

Administration Groups and Template Collections: Applying Monitoring Templates Automatically

Monitoring templates can be automatically applied whenever a new targets are added to your Enterprise Manager environment. Automation is carried out through Administration Groups and Template Collections Administration Groups are a special type of group used to automate application of monitoring settings to targets upon joining the group. When a target is added to the administration group Enterprise Manager applies monitoring settings from the associated template collection consisting of monitoring templates, compliance standards, and Enterprise Manager policies. If changes are later made to the monitoring template, Enterprise Manager automatically applies the changes to the relevant targets based on the synchronization schedule. For more information, see "Using Administration Groups ".

Applying a Monitoring Template

To apply a template, you must have at least *Manage Target Metrics* target privileges on the destination target(s).

1. From the **Enterprise** menu, select **Monitoring** and then **Monitoring Templates**.
2. Select the desired template from the table.
3. Click **Apply**.
4. Select the desired apply options and the target(s) to which you want the templates applied. See Monitoring Template Application Options for additional information.
5. Click **OK**.

Monitoring Template Application Options

You can choose aggregate targets such groups, systems or clusters as destination targets. The templates will apply to the appropriate members of the group/system/cluster as they currently exist. If new members are later added to the group, you will need to re-apply the template to those new members. Template application is performed in the background as asynchronous jobs, so after the apply operation is performed, you can click on the link under the Pending Apply Operations column in the main templates table to see any apply operations that still are pending.

When applying a Monitoring Template, metric settings such as thresholds, comparison operators, and corrective actions are copied to the destination target. In addition, metric collection schedules including collection frequency and upload interval are also copied to the target. You determine how Enterprise Manager applies the metric settings from the template to the target by choosing an apply option.

Apply Options

Template apply options control how template metric and policy settings are applied to a target. Two template apply options are available:

- **Template will completely replace all metric settings in the target**: When the template is applied, all metrics and policies defined in the template will be applied to the target. Pre-existing target monitoring settings not defined in the template will be disabled: Metric thresholds will be set to NULL or blank. Policies will be disabled. This effectively eliminates alerts from these metrics and policies by clearing current severities and violations.
- **Template will only override metrics that are common to both template and target**: When the template is applied, only metrics and policies common to both the template and target are updated. Existing target metric and policies that do not exist in the template will remain unaffected. When this option is selected, additional template apply options are made available for metrics with key value settings.

Metrics with Key Value Settings

A metric with key value settings is one that can monitor multiple objects at different thresholds. For example, the Filesystem Space Available(%) metric can monitor different mount points using different warning and critical thresholds for each mount point.

When the template contains a metric that has key value settings, you can choose one of three options when applying this template to a target. As an example, consider the case where the template has the following metric:

Filesystem Space Available(%)

Mount Point	Operator	Warning Threshold	Critical Threshold
/		40	20
/private		30	20
/private2		20	20
/u1		30	20
All Others		25	15

And a host target has the same metric at different settings:

Mount Point	Operator	Warning Threshold	Critical Threshold
/		30	10
/private		25	15
/private2		20	20
All Others		25	15

These are the results for each option:

1) All key value settings in the template will be applied to the target, any additional key values settings on the target will not be removed

When the template is applied to the target using this copy option, all the template settings for the mount points, /, /private, and /U1 will be applied. Existing target settings for mount points not covered by the template remain unaffected. Thus, the resulting settings on the target for this metric will be:

Mount Point	Operator	Warning Threshold	Critical Threshold
/		40	20
/private		30	20
/u1		30	20

2) All key value settings in the template will be applied to target, any additional key value settings on the target will be removed.

When the template is applied to the target using this copy option, all template settings will be applied to the target. Any object-specific threshold settings that exist only on the target will be removed, any object-specific thresholds that are only in the template will be added to the target. Thus, the final settings on the target will be:

Mount Point	Operator	Warning Threshold	Critical Threshold
/		40	20
/private		30	20
/u1		30	20
All Others		25	15

3) Only settings for key values common to both template and target will be applied to the target

When the template is applied to the target using this copy option, only the settings for the common mount points, / and /private will be applied. Thus, the resulting settings on the target for this metric will be:

Mount Point	Operator	Warning Threshold	Critical Threshold
/		40	20
/private		30	20
/private2		20	20
All Others		25	15

Comparing Monitoring Templates with Targets

The intended effect of applying Monitoring Templates to destination targets is not always clear. Deciding how and when to apply a template is simplified by using the Compare Monitoring Template feature of Enterprise Manager. This allows you to see at a glance how metric and collection settings defined in the template differ from those defined on the destination target. You can easily determine whether your targets are still compliant with the monitoring settings you have applied in the past. This template comparison capability is especially useful when used with aggregate targets such as groups and systems. For example, you can quickly compare the metric and collection settings of group members with those of a template, and then apply the template as appropriate.

Performing a Monitoring Template-Target comparison:

1. From the **Enterprise** menu, select **Monitoring** and then **Monitoring Templates**.
2. Choose the desired template from the table.
3. Click **Compare Settings**. The Compare Monitoring Template page displays.
4. Click **Add** to add one or more destination targets. The Search and Select dialog displays.
5. Select a one or more destination targets and then click **Select**. The selected targets are added to the list of destination targets.
6. Select the newly added destination targets and then click **Continue**. A confirmation message displays indicating the *Compare Template Settings* job was successfully submitted.
7. Click **OK** to view the job results. Note: Depending on the complexity of the job run, it may take time for the job to complete.

When is a metric between a template and a target considered "different"?

The metric is said to be different when any or all the following conditions are true (provided the target does not have "Template Override" set for that metric):

- The Warning Threshold settings are different
- The Critical Threshold settings are different.
- The Collection Schedules are different.
- The Upload Intervals are different.
- The number of occurrences (for which the metric has to remain at a value above the threshold before an alert is raised) are different.)
- For metric extensions, in addition to the above, the OS Command/SQL statement used to evaluate the metric extension is different. Note that this applies only if the name and the return type are the same.
- The metric extension marked for delete will be shown as "different" on the destination target and the template only if:
  - A metric extension with the same name exists on both the destination target and template.
  - The return type (String, Numeric) of the metric extension is the same on both the destination target and template.
  - The metric type is the same on both the destination target and the template.

Comparing Metric Settings Using Information Publisher

In addition to viewing metric differences between Monitoring Templates and destination targets using the Compare Monitoring Template user-interface, you can also use Information Publisher to generate reports containing the target-template differences. Using Information Publisher's reporting capabilities gives you more flexibility for displaying and distributing metric comparison data. For more information, see "Using Information Publisher".

Create a Report Definition

1. From the **Enterprise** menu, select **Reports** and then **Information Publisher Reports**.
2. Click **Create**. The Create Report Definition user interface is displayed.
3. On the General page, specify the report name, how targets should be included, target privileges, report time period, and display options.
4. On the Elements page, click **Add** to access the Add Element page.
5. Select the Monitoring Template Comparison element and click **Continue** to return to the Element page.
6. Once you have added the report element, click the **Set Parameter** icon to specify requisite operational parameters. On this page, you specify a report header, select a monitoring template, destination targets, and template application settings for multiple threshold metrics. Click **Continue** to return to the Elements page.
7. Click **Layout** to specify how information should be arranged in the report.
8. Click **Preview** to validate that you are satisfied with the data and presentation of the report.
9. On the Schedule page, define when reports should be generated, and whether copies should be saved and/or sent via e-mail, and how stored copies should be purged.
10. On the Access page, click **Add** to specify which Enterprise Manager administrators and/or roles will be permitted to view this generated report. Additionally, if you have GRANT_ANY_REPORT_VIEWER system privilege, you can make this report definition accessible to non-credential users via the Enterprise Manager Reports Website
11. Click **OK** when you are finished.
12. Validate the report definition. If the parameters provided conflict, validation errors or warnings will appear and let you know what needs attention.
13. Once the report definition has been saved successfully, it appears in the Report Definition list under the Category and Subcategory you specified on the General page.

Viewing the Report

1. Find the template comparison report definition in the Report Definition list. You can use the Search function to find or filter the list of report definitions.
2. Click on the report definition title. If the report has a specified target, the report will be generated immediately. If the report does not have a specified target, you will be prompted to select a target.

Scheduling Reports for Automatic Generation

1. Create or edit a report definition.
2. On the Schedule page, choose the **Schedule Report** option.
3. Specify a schedule type. The schedule parameters on this page change according to the selected schedule type.

When reports are scheduled for automatic generation, you have the option of saving copies to the Management Repository and/or sending an e-mail version of the report to designated recipients.

If a report has been scheduled to save copies, a copy of the report is saved each time a scheduled report completes. When a user views a report with saved copies by clicking on the report title, the most recently saved copy of the report is rendered. To see the complete list of saved copies click on the Saved Copies link at the top of the report. Enterprise Manager administrators can generate a copy of the report on-demand by clicking on the Refresh icon on the report.

Exporting and Importing Monitoring Templates

For portability, monitoring templates can be exported to an XML file and then imported into another Enterprise Manager installation as an active template.

> **Note:**
>
> You can export templates from Enterprise Manager 10g release 2 or higher and import them into Enterprise Manager 13c.

Exporting a Monitoring Template

To export a template to an XML file:

1. From the **Enterprise** menu, select **Monitoring** and then **Monitoring Templates**.
2. Select the desired monitoring template from the table.
3. Click **Export**.

   Note: If you are running an Enterprise Manager 11*g* or earlier release, use the EM CLI *export_template* verb to perform the export operation. Note that if the monitoring template contains *policy rules* from earlier Enterprise Manager releases (pre-12*c)* , these will not be imported into Enterprise Manager 12*c* as policy rules no longer exist in this release.

Importing a Monitoring Template

To import a template from an XML file:

1. From the **Enterprise** menu, select **Monitoring** and then **Monitoring Templates**.
2. Click Import. The Import Template page displays.
3. Specify the monitoring template XML file you want to import.
4. Click **Import**.

Upgrading Enterprise Manager: Comparing Monitoring Templates

When upgrading from one Enterprise Manager release to the next, you will accumulate monitoring templates that have been created for different releases. Beginning with Enterprise Manager release 12.1.0.4, you can generate a post-upgrade Monitoring Template Difference Report that allows you to view what templates had been created for various Enterprise Manager releases. To generate the Monitoring Template Difference Report, from the **Setup** menu, select **Manage the Manager**, and then **Post Upgrade Tasks**.

Changing the Monitoring Template Apply History Retention Period

You can view monitoring template apply history using the predefined report in Information Publisher. From the **Enterprise** menu, select **Reports** and then **Information Publisher.** On the Information Publisher page, you can enter "template" in the **Title** text entry field and click **Go**. The predefined report *Monitoring Template Apply History (last 7 days)* appears in the report list.

By default, Enterprise Manager retains the monitoring template apply history for a period of 31 days. If required, you can change the retention period to a value suitable for your monitoring needs. Although the retention period cannot be indefinite, it can be set to an extremely long period of time. Enterprise Manager provides the following PL/SQL API to change the retention period.

```
mgmt_template_ui.modify_purge_policy(p_retention_days=><num_days>)
```

This procedure takes a NUMBER as input (*num_days*).

13
Using Metric Extensions

Metric extensions provide 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:

- What are Metric Extensions?
- Metric Extension Lifecycle
- Working with Metric Extensions
- Adapters
- Metric Extension Command Line Verbs

What are Metric Extensions?

Metric extensions allow you to create metrics on any target type. 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.

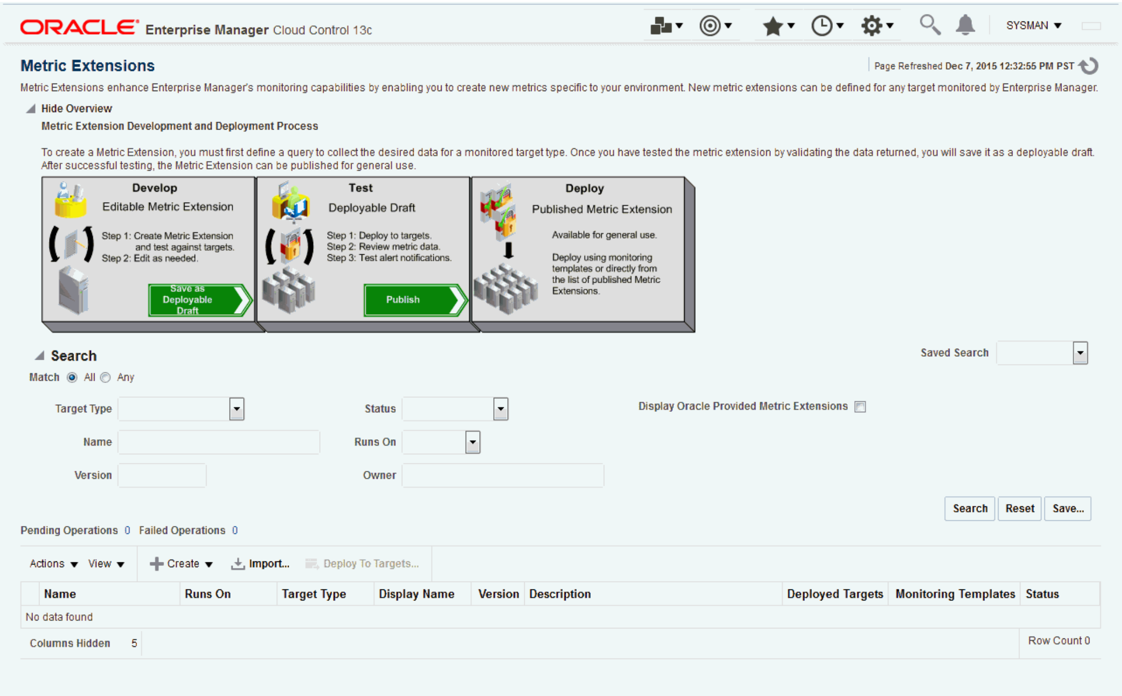

Metric Extensions

Page Refreshed Dec 7, 2015 12:32:55 PM PST

Metric Extensions enhance Enterprise Manager's monitoring capabilities by enabling you to create new metrics specific to your environment. New metric extensions can be defined for any target monitored by Enterprise Manager.

Hide Overview

Metric Extension Development and Deployment Process

To create a Metric Extension, you must first define a query to collect the desired data for a monitored target type. Once you have tested the metric extension by validating the data returned, you will save it as a deployable draft. After successful testing, the Metric Extension can be published for general use.

Develop
Editable Metric Extension
Step 1: Create Metric Extension and test against targets.
Step 2: Edit as needed.

Test
Deployable Draft
Step 1: Deploy to targets.
Step 2: Review metric data.
Step 3: Test alert notifications.

Deploy
Published Metric Extension
Available for general use.
Deploy using monitoring templates or directly from the list of published Metric Extensions.

Search

Match ● All ○ Any

Saved Search

Target Type

Status

Display Oracle Provided Metric Extensions □

Name

Runs On

Version

Owner

Pending Operations 0 Failed Operations 0

Name	Runs On	Target Type	Display Name	Version	Description	Deployed Targets	Monitoring Templates	Status
No data found								

Columns Hidden 5

Row Count 0

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 atonal 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. However, in the case of user-defined metrics, multiple pieces of data must be collected by creating multiple user-defined metrics. Because the data is being collected separately, it is not possible to refer to the associated data when alerts are generated.
- **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**: User-defined metrics only allow OS user-defined metrics against host targets and SQL user-defined metrics against database targets. No other target types are permitted. If, for example, you want to deploy a user-defined metric against WebLogic instances in your environment, you will not be able to do so since it is neither a host or database target type.

Most importantly, the primary difference between metric extensions and user-defined metrics is that, unlike user-defined metrics, metric extensions 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.

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

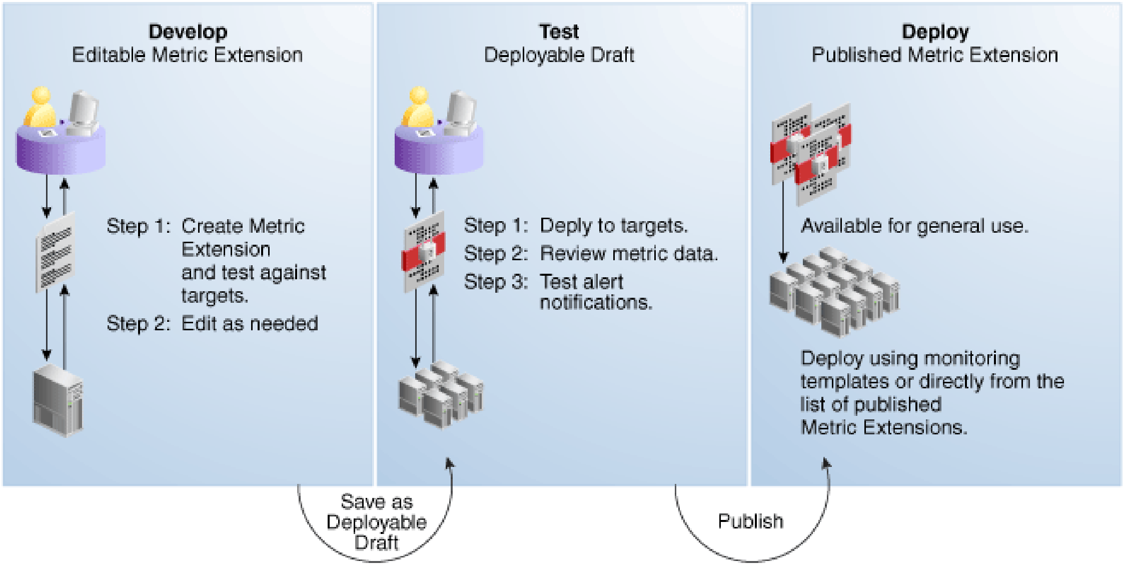
Develop
Editable Metric Extension
Step 1: Create Metric Extension and test against targets.
Step 2: Edit as needed
Test
Deployable Draft
Step 1: Deply to targets.
Step 2: Review metric data.
Step 3: Test alert notifications.
Deploy
Published Metric Extension
Available for general use.
Deploy using monitoring templates or directly from the list of published Metric Extensions.
Save as Deployable Draft
Publish

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.

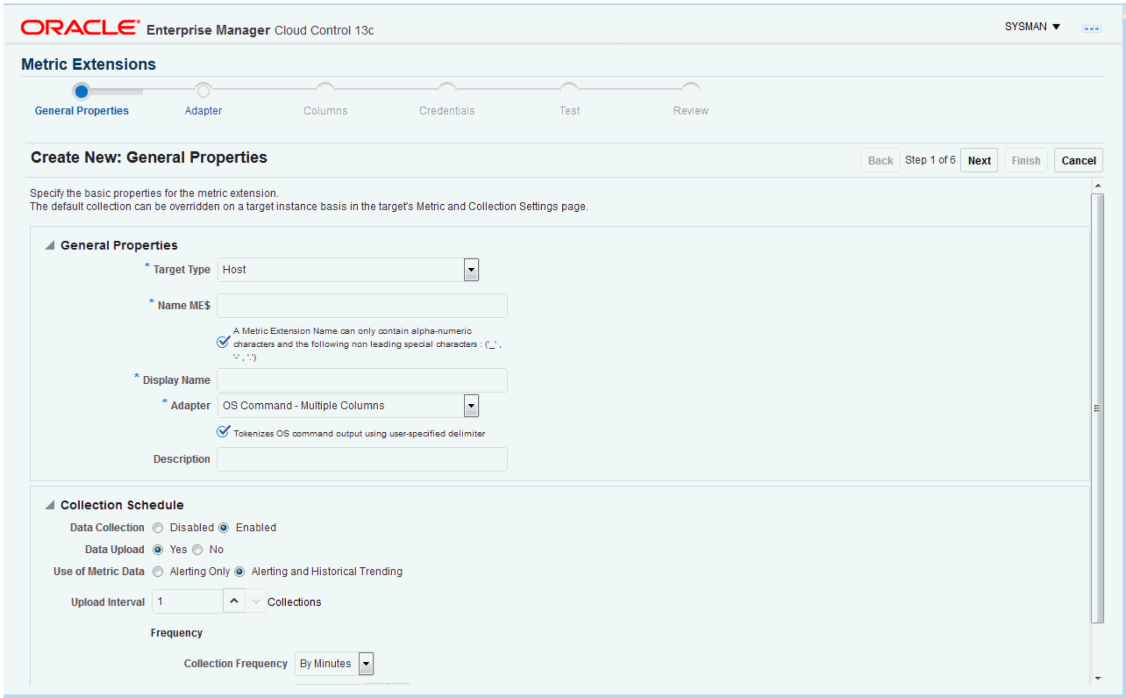

Create New: General Properties

Specify the basic properties for the metric extension.
The default collection can be overridden on a target instance basis in the target's Metric and Collection Settings page.

General Properties

* Target Type: Host

* Name ME$:

A Metric Extension Name can only contain alpha-numeric characters and the following non leading special characters : ('_' , '-' , '.')

* Display Name:

* Adapter: OS Command - Multiple Columns

Tokenizes OS command output using user-specified delimiter

Description:

Collection Schedule

Data Collection ○ Disabled ● Enabled

Data Upload ● Yes ○ No

Use of Metric Data ○ Alerting Only ● Alerting and Historical Trending

Upload Interval 1 Collections

Frequency

Collection Frequency By Minutes

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 it can be accessed 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 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 multiple 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. From the Actions menu, select **Publish Metric Extension**.

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 select 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.

> **Note:**
>
> In the case of a metric extension deployment missing from the agent, or a target's metric extension attachment missing from agent, try using the EM CLI verb `resync_target`. For more information, see resync_target.

Updating Metric Extensions

Beginning with Enterprise Manager Release 12.1.0.4, metric extensions can be updated using the Enterprise Manager Self-update feature. See Updating Cloud Control for more information.

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 EM CLI 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.

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" privilege edit and create next versions of a particular metric extensions. The metric extension creator has this privilege by default. This privilege must be granted on a per-metric extension basis.
- **Full Metric Extension**: Allows users with 'Create Metric Extension' privilege to edit and create new versions 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.

Granting Create Metric Extension Privilege

To grant create metric extension privileges to another administrator:

1. From the **Setup** menu, select **Security**, then select **Administrators**.
2. Choose the Administrator you would like to grant the privilege 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 Privileges, click the **Create Metric Extension** check box.
6. Click **Continue**, review changes, and click **Finish** in the Review tab.

Managing Administrator Privileges

Before an Enterprise Manager administrator can edit or delete a metric extension created by another administrator, that administrator must have been granted requisite access privileges. *Edit* privilege allows editing and creating next versions of the extension. *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, select **Security**, then select **Administrators**.
2. Choose the Administrator you would like to grant access to.
3. Click **Edit**.
4. Go to **Resource Privileges** and click Manage Privilege Grants (pencil icon) for the Metric Extensions resource type.
5. Under **Resource Privileges**, 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.

   On this page, you can also grant an administrator the *Create Metric Extension* privilege, which will allow them to manage metric extension access. See "Managing Administrator Access to Metric Extensions" for more information.
6. If you would additionally like to allow delete operations, then click the pencil icon in the **Manage Resource Privilege 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.

Managing Administrator Access to Metric Extensions

Administrators commonly share the responsibility of monitoring and managing targets within the IT environment. Consequently, creating and maintaining metric extensions becomes a collaborative effort involving multiple administrators. Metric extension owners can control access directly from the metric extension UI.

Granting Full/Edit Privileges on a Metric Extension

As metric extension owner or Super Administrator, perform the following actions to assign full/edit privileges on a metric extension to another administrator:

1. From the **Enterprise** menu, select **Monitoring,** then select **Metric Extensions.**
2. Choose a metric extension requiring update.
3. From the **Actions** menu, select **Manage Access**.
4. Click **Add**. The administrator selection dialog box appears. You can filter the list by administrator, role, or both.
5. Choose one or more administrators/roles from the list.
6. Click **Select**. The chosen administrators/roles appear in the access list.

   In the **Privilege** column, **Edit** is set by default. Choose **Full** from the drop-down menu to assign **Full** privileges on the metric extension.

   Edit Privilege: Allows an administrator to make changes to the metric extension but not delete it.

   Full Privilege: Allows an administrator to edit and also delete the metric extension. The privilege granted to a user or role applies to all versions of the metric extension.
7. Click **OK**.

Revoking Access Privileges on a Metric Extension

As metric extension owner or Super Administrator, perform the following actions to revoke metric extension privileges assigned to another administrator:

1. From the **Enterprise** menu, select **Monitoring,** then select **Metric Extensions.**
2. Choose a metric extension requiring update.
3. From the **Actions** menu, select **Manage Access**.
4. Choose one or more administrators/roles from the list.
5. Click **Remove**. The chosen administrators/roles is deleted from the access list.
6. Click **OK**.

Enterprise Manager allows metric extension ownership to be transferred from the current owner of the metric extension to another administrator as long as that administrator has been granted the *Create Metric Extension* privilege.

> **Note:**
>
> The Enterprise Manager Super Administrator has full managerial access to all metric extensions (view, edit, and ownership transfer).

As mentioned above, *manage access* is only enabled for the owner of the extension or an Enterprise Manager Super User. Once the ownership is transferred, the previous owner does not have any management privileges on the metric extension unless explicitly granted before ownership transfer. The **Change Owner** option is only available to users and not roles.

Manage access allows the metric extension owner or Super Administrator to grant other Enterprise Manager users or roles the ability to edit, modify, or delete metric extensions.

Transferring Metric Extension Ownership

Enterprise Manager allows metric extension ownership to be transferred from the current owner of the metric extension to another administrator as long as that administrator has been granted the *Create Metric Extension* privilege.

> **Note:**
>
> The Enterprise Manager Super Administrator has full managerial access to all metric extensions (view, edit, and ownership transfer).

As mentioned above, *manage access* is only enabled for the owner of the extension or an Enterprise Manager Super User. Once the ownership is transferred, the previous owner does not have any management privileges on the metric extension unless explicitly granted before ownership transfer. The **Change Owner** option is only available to users and not roles.

Manage access allows the metric extension owner or Super Administrator to grant other Enterprise Manager users or roles the ability to edit, modify, or delete metric extensions.

Creating a New Metric Extension

To create a new metric extension:

1. From the **Enterprise** menu, select **Monitoring,** then select **Metric Extensions.**
2. From the **Create** menu, select **Metric Extension**. 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.

Collection Schedule

You defined the frequency with which metric data is collected and how it is used (Alerting Only or Alerting and Historical Trending) by specifying collection schedule properties.

Depending on the target type selected, an *Advanced* option region may appear. This region may (depending on the selected target type) contain one or two options that determine whether metric data continues to be collected under certain target availability/ alert conditions. The options are:

- **Option 1**: Continue metric data collection even if the target is down. This option is visible for all target types except for *Host* target types as it is not possible to collect metric data when the host is down.
- **Option 2**: Continue metric data collection when an alert severity is raised for a specific target metric. This metric is defined in such a way (AltSkipCondition element is defined on this metric) that when a severity is generated on this metric, the metric collections for other target metrics are stopped. The explanatory text above the checkbox for this option varies depending on the selected target type.

  The Management Agent has logic to skip evaluation of metrics for targets that are known to be down to reduce generation of metric errors due to connection failures. If the AltSkipCondition element is defined for that target metric, other metrics are skipped whenever there is an error in evaluating the Response metric or there is a non-clear severity on the Response:Status metric. There are two situations where a metric collection will be skipped or not happen:

- When a target is down (option 1). This is same as the Severity on Response/ Status metric.
- When a target is UP, but there is a severity on any other metric. Such conditions are called Alt Skip (Alternate Skip) conditions.

Option 2 is only visible if an AltSkipCondition defined for one of the target's metrics. For example, this option will not be visible if the selected target type is *Oracle Weblogic Domain*, but will be visible if the selected target type is *Database Instance*.

The following graphic shows the Advanced collection schedule options.

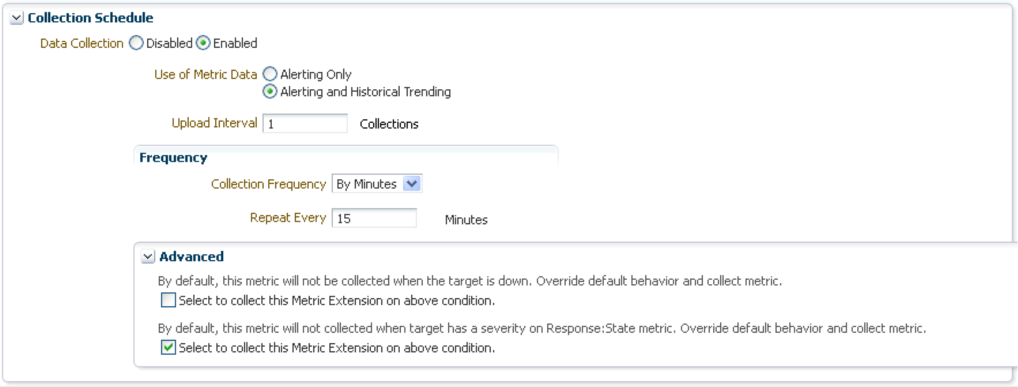

Collection Schedule

Data Collection ○ Disabled ● Enabled

Use of Metric Data ○ Alerting Only
● Alerting and Historical Trending

Upload Interval 1 Collections

Frequency

Collection Frequency By Minutes

Repeat Every 15 Minutes

Advanced

By default, this metric will not be collected when the target is down. Override default behavior and collect metric.
□ Select to collect this Metric Extension on above condition.

By default, this metric will not collected when target has a severity on Response:State metric. Override default behavior and collect metric.
☑ Select to collect this Metric Extension on above condition.

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

   - **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. You can also create rate and delta metric columns based on an existing data column. See *Rate and Delta Metric Columns* below.

     > **Note:**
     >
     > Key columns might contribute to the creation of a large number of unique keys, such as `metric_key_id`, causing load on the repository.
     >
     > For example, *Timestamp* should not be a Key column, as it will create unique metric_key_id for each collection, and the existing metric_key_id and metric_item_id will no longer be reused.

   - **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 distinct CPU. The syntax is to specify the key column values in a comma separated list, the "=" symbol, followed by the alert threshold. Multiple thresholds for different rows can be separated by the semi-colon symbol ";". For example, if the key columns 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**

> **Note:**
>
> You must expand the Advanced region in order to view the Manually Clearable Alert option.

  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 Occurrences 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 an 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 click **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.

Rate and Delta Metric Columns

You can create additional metric columns based on an existing data column that measures the rate at which data changes or the difference in value (delta) since the last metric collection. The rate/delta metric definition will be allowed when a metric's collection frequency is periodic. For example, collected every 10 minutes. Converseley, a metric that is computed every Monday and Tuesday only cannot have a rate/delta metric as data sampling is too infrequent.

After at least one data column has been created, three additional options appear in the **Add** menu as shown in the following graphic.

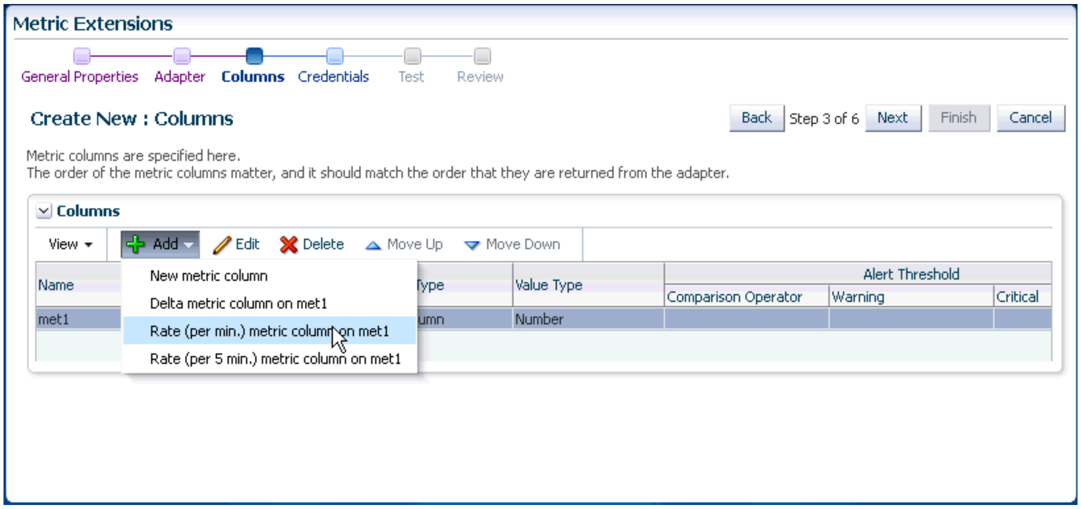

Create New : Columns

Metric columns are specified here.
The order of the metric columns matter, and it should match the order that they are returned from the adapter.

Name	ype	Value Type	Alert Threshold		
			Comparison Operator	Warning	Critical
met1	umn	Number			

- Add Delta metric columns based on another metric column

  Example: You want to know the difference in the table space used since the last collection.

  Delta Calculation:

  current metric value - previous metric value

- Add Rate Per Minute metric column based on another metric column

  Example: You want to know the average table space usage per minute based on the table space column metric which is collected every 1 hr.

  Rate Per Minute Calculation:

  (current metric value - previous metric value)/ collection schedule

  where the *collection schedule* is in minutes.

- Add Rate Per Five Minutes metric column based on another metric column

  Example: You want to know the average table space usage every five minutes based on the table space column which is collected say every 1 hour]

  Rate Per Five Minute Calculation:

  *[(current metric value - previous metric value)/ collection schedule] * 5*

  where the *collection schedule* is in minutes.

To create a rate/delta metric column, click on an existing data column in the table and then select one of the rate/delta column options from the **Add** menu.

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**.

Creating a New Metric Extension (Create Like)

To create a new metric extension based on an existing metric extension:

1. From the **Enterprise** menu, select **Monitoring**, then select **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, select **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**.

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 deployable draft, it cannot be edited, you can only create a new version.

To edit an existing metric extension:

1. From the **Enterprise** menu, select **Monitoring**, then select **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, select **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**.

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, select **Monitoring**, then select **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, select **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**.

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, select **Monitoring**, then select **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 system 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, select **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**.

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, select **Monitoring**, then select **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, select **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 Production, Deployable Draft and Published metric extension versions.
6. Confirm the export file is downloaded.

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, select **Monitoring**, then select **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, select **Delete**. Enterprise Manager prompts you to confirm the deletion.
5. Confirm the deletion.

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, select **Monitoring**, then select **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, select **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, select **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.

Creating an Incident Rule to Send Email from Metric Extensions

One of the most common tasks administrators want Enterprise Manager to perform is to send an email notification when a metric alert condition occurs. Specifically, Enterprise Manager monitors for alert conditions defined as incidents. For a given incident you create an incident rule set to tell Enterprise Manager what actions to take when an incident occurs. In this case, when an incident consisting of an alert condition defined by a metric extension occurs, you need to create an incident rule to send email to administrators. For instructions on sending email for metric alerts, see "Sending Email for Metric Alerts".

For information incident management see Using Incident Management .

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, select **Monitoring**, then select **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, select **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.

Creating Repository-side Metric Extensions

Beginning with Enterprise Manager Release 12.1.0.4, you can create repository-side metric extensions. This type of metric extension allows you to use SQL scripts to extract information directly from the Enterprise Manager repository and raise alerts for the target against which the repository-side extension is run. For example, you can use repository-side metric extensions to raise an alert if the total number of alerts for a host target is greater than 5. Or perhaps, raise an alert if the CPU utilization on that host is greater than 95% AND the number of process running on that host is greater than 500. Repository-side metrics allows you to monitor your Enterprise Manager infrastructure with greater flexibility.

To create a repository-side metric:

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

   Collection Schedule

   You defined the frequency with which metric data is collected and how it is used (Alerting Only or Alerting and Historical Trending) by specifying collection schedule properties.
5. Create the SQL query to be run against the Enterprise Manager Repository. Explicit instructions for developing the query as well as examples are provide on the SQL Query page.

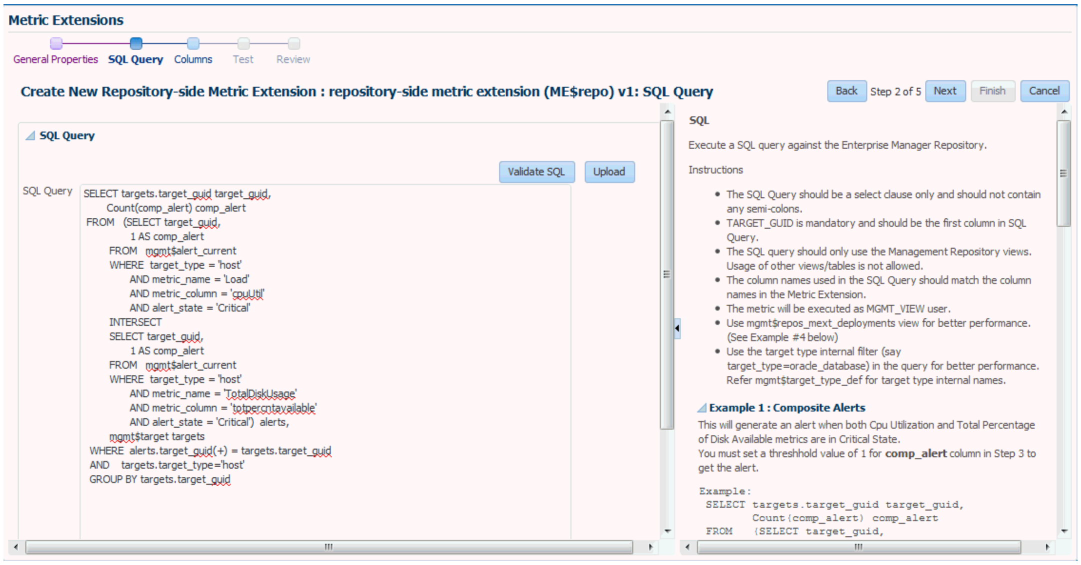

Create New Repository-side Metric Extension : repository-side metric extension (ME$repo) v1: SQL Query

SQL Query

```
SELECT targets.target_guid target_guid,
       Count(comp_alert) comp_alert
 FROM   (SELECT target_guid,
               1 AS comp_alert
        FROM   mgmt$alert_current
        WHERE  target_type = 'host'
               AND metric_name = 'Load'
               AND metric_column = 'cpuUtil'
               AND alert_state = 'Critical'
        INTERSECT
        SELECT target_guid,
               1 AS comp_alert
        FROM   mgmt$alert_current
        WHERE  target_type = 'host'
               AND metric_name = 'TotalDiskUsage'
               AND metric_column = 'totpercntavailable'
               AND alert_state = 'Critical')  alerts,
        mgmt$target targets
  WHERE  alerts.target_guid(+) = targets.target_guid
  AND    targets.target_type='host'
  GROUP BY targets.target_guid
```

SQL

Execute a SQL query against the Enterprise Manager Repository.

Instructions

- The SQL Query should be a select clause only and should not contain any semi-colons.
- TARGET_GUID is mandatory and should be the first column in SQL Query.
- The SQL query should only use the Management Repository views. Usage of other views/tables is not allowed.
- The column names used in the SQL Query should match the column names in the Metric Extension.
- The metric will be executed as MGMT_VIEW user.
- Use mgmt$repos_mext_deployments view for better performance. (See Example #4 below)
- Use the target type internal filter (say target_type=oracle_database) in the query for better performance. Refer mgmt$target_type_def for target type internal names.

Example 1 : Composite Alerts

This will generate an alert when both Cpu Utilization and Total Percentage of Disk Available metrics are in Critical State.
You must set a threshhold value of 1 for **comp_alert** column in Step 3 to get the alert.

```
Example:
 SELECT targets.target_guid target_guid,
        Count(comp_alert) comp_alert
 FROM   (SELECT target_guid,
```

Click **Validate SQL** to test the query.

If you already have a SQL script, you can click **Upload** to load the SQL from an external file.

6. From the Columns page, you can view/edit columns returned by the SQL query. You may edit the columns, however, you cannot add or delete columns from this page.

   - **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. You can also create rate and delta metric columns based on an existing data column. See *Rate and Delta Metric Columns* below.

   - **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 distinct CPU. The syntax is to specify the key column values in a comma separated list, the "=" symbol, followed by the alert threshold. Multiple thresholds for different rows can be separated by the semi-colon symbol ";". For example, if the key columns 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**

> **Note:**
>
> You must expand the Advanced region in order to view the Manually Clearable Alert option.

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 Occurrences 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 an 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 click **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.

- Add Delta metric columns based on another metric column

  Example: You want to know the difference in the table space used since the last collection.

  Delta Calculation:

  current metric value - previous metric value

- Add Rate Per Minute metric column based on another metric column

  Example: You want to know the average table space usage per minute based on the table space column metric which is collected every 1 hr.

  Rate Per Minute Calculation:

(current metric value - previous metric value)/ collection schedule

where the *collection schedule* is in minutes.

- Add Rate Per Five Minutes metric column based on another metric column

  Example: You want to know the average table space usage every five minutes based on the table space column which is collected say every 1 hour]

  Rate Per Five Minute Calculation:

  *[(current metric value - previous metric value)/ collection schedule] * 5*

  where the *collection schedule* is in minutes.

To create a rate/delta metric column, click on an existing data column in the table and then select one of the rate/delta column options from the **Add** menu.

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**.

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 available from the UI. For example, when creating a metric extension for an Automatic Storage Management target type, only three adapters (OS Command-Single Column, OS Command-Multiple Columns, and SQL) are available from the UI.

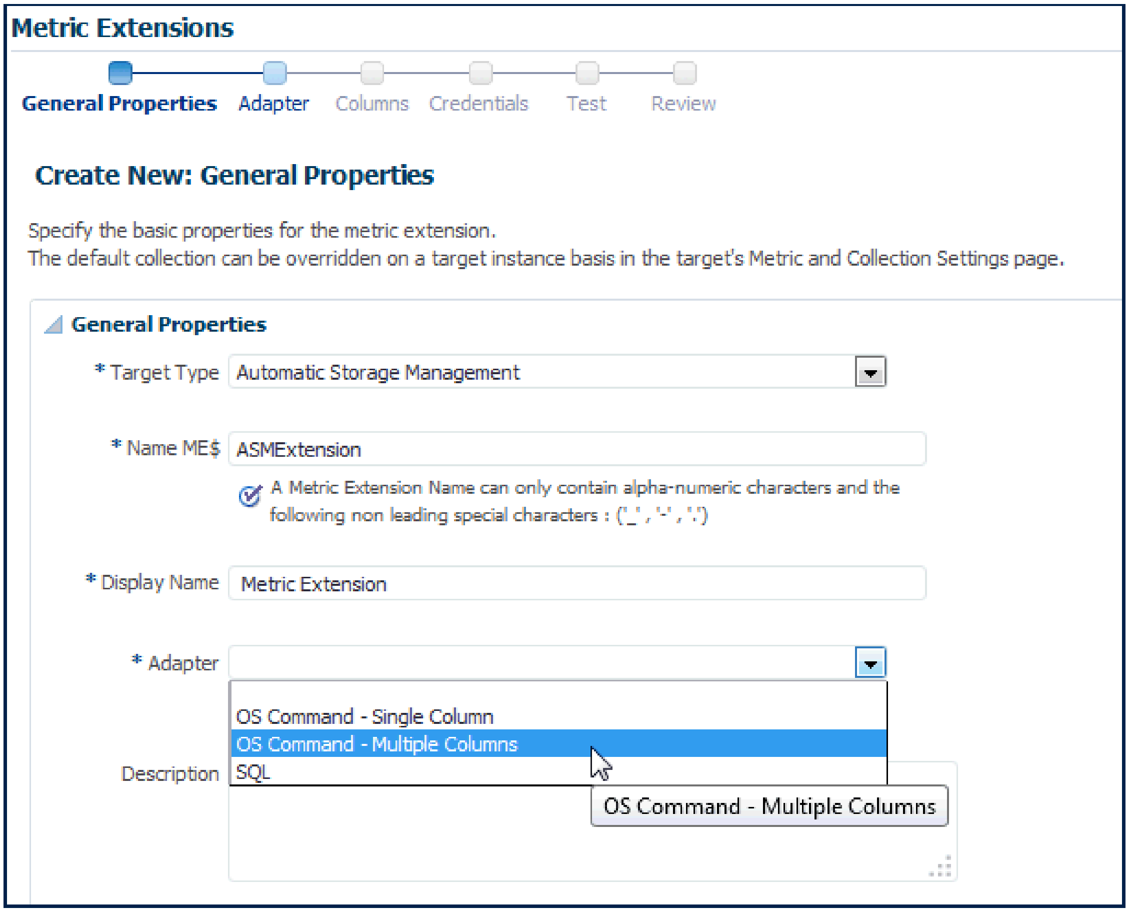

Create New: General Properties

Specify the basic properties for the metric extension.
The default collection can be overridden on a target instance basis in the target's Metric and Collection Settings page.

General Properties

* Target Type: Automatic Storage Management

* Name ME$: ASMExtension

A Metric Extension Name can only contain alpha-numeric characters and the following non leading special characters : ('_' , '-' , '.')

* Display Name: Metric Extension

* Adapter:

OS Command - Single Column
OS Command - Multiple Columns
SQL
Description
OS Command - Multiple Columns

A target type's out-of-box metric definition defines the adapters for which it has native support, and only those adapters will be shown in the UI. No other adapters are supported for that target type.

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

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"*.

> **Note:**
>
> *Value="%TYPE%"* may not be applicable for certain target-types (for example: host, wls). For such target types, use *Value=TYPE* instead.

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

Overriding default monitoring credentials by creating and using a custom monitoring credential set for host target.

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 Enterprise Manager 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:mypwd'
```

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=out_x
em_result=out_y
```

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.

> **Note:**
>
> *Value="%TYPE%"* may not be applicable for certain target-types (for example: host, wls). For such target types, use *Value=TYPE* instead.

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.

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:

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  out_x  out_y
  ```

  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.

> **Note:**
>
> *Value="%TYPE%"* may not be applicable for certain target-types (for example: host, wls). For such target types, use *Value=TYPE* instead.

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.

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 specified.

  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

Overriding default monitoring credentials by creating and using a custom monitoring credential set for database target.

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 Enterprise Manager by going to **Setup**, then selecting **Security**, then selecting **Monitoring Credentials**. Alternatively, this can be performed using the Enterprise Manager Command Line Interface.

```
> 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:mypwd'
```

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 expected 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.

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 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;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 the Monitoring Using Web Services and JMX chapter in the for an in-depth example of creating a JMX-based Metric Extension.

Metric Extension Command Line Verbs

Metric extensions can be manipulated outside the UI via the Enterprise Manager Command Line Interface (EM CLI). 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>
```

14
Advanced Threshold Management

There are monitoring situations in which different workloads for a target occur at regular (expected) intervals. Under these conditions, a static alert threshold would prove to be inaccurate. For example, the accurate alert thresholds for a database performing Online Transaction Process (OLTP) during the day and batch processing at night would be different. Similarly, database workloads can change based purely on different time periods, such as weekday versus weekend. In both these situations, fixed, static values for thresholds might result in false alert reporting.

Advanced Thresholds allow you to define and manage alert thresholds that are either adaptive (self-adjusting) or time-based (static).

- *Adaptive Thresholds* are thresholds based on statistical calculations from the target's observed behavior (metrics).
- *Time-based Thresholds* are user-defined threshold values to be used at different times of the day/week to account for changing target workloads.

This chapter covers the following topics:

- Accessing the Advanced Threshold Management Page
- Adaptive Thresholds
- Time-based Static Thresholds
- Determining What is a Valid Metric Threshold

Accessing the Advanced Threshold Management Page

You manage advanced thresholds from the Enterprise Manager console. The Advanced Threshold Management page allows you to create time-based static thresholds and adaptive thresholds. To access this page:

1. From a target home page (host, for example), navigate to the **Metric Collection and Settings** page.
2. From the Related Links region, click **Advanced Threshold Management**.

   The Advanced Threshold Management page displays.

Adaptive Thresholds

Adaptive thresholds are statistically computed thresholds that adapt to target workload conditions. Adaptive thresholds apply to all targets (both Agent and repository-monitored).

Important Concepts

Creating an adaptive threshold is based on the following key concepts:

- **Baseline periods**

For the purpose of performance evaluation, a baseline period is a period of time used to characterize the typical behavior of the system. You compare system behavior over the baseline period to that observed at some other time.

There are two types of baseline periods:

- **Moving window baseline periods**: Moving window baselines are defined as some number of days prior to the current date. This "window" of days forms a rolling interval that moves with the current time. The number of days that can be used to define moving window baseline in Enterprise Manager are:
  - 7 days
  - 14 days
  - 21 days
  - 30 days

    Example: Suppose you have specified trailing 7 days as a time period while creating moving window baseline. In this situation, the most recent 7-day period becomes the baseline period for all metric observations and comparisons today. Tomorrow, this reference period drops the oldest day and picks up today.

    Moving window baselines allow you to compare current metric values with recently observed history, thus allowing the baseline to incorporate changes to the system over time. Moving window baselines are suitable for systems with predictable workload cycles.

> **Note:**
>
> Enterprise Manager computes moving window statistics every day rather than sampling.

Registering Adaptive Threshold Metrics

Adaptive threshold metrics are not immediately available by default; they must be defined and added to the system (registered) in order for them to become available for use by Enterprise Manager. Not all metrics can have adaptive thresholds: Adaptive Threshold metrics must fall into one of the following categories:

- Load
- LoadType
- Utilization
- Response

You can register adaptive threshold metrics from the Advanced Threshold Management page.

1. From a target menu (Host is used in this example), select **Monitoring** and then **Metric and Collection Settings**.
2. In the Related Links area, click **Advanced Threshold Management**.The Advanced Threshold Management page displays.

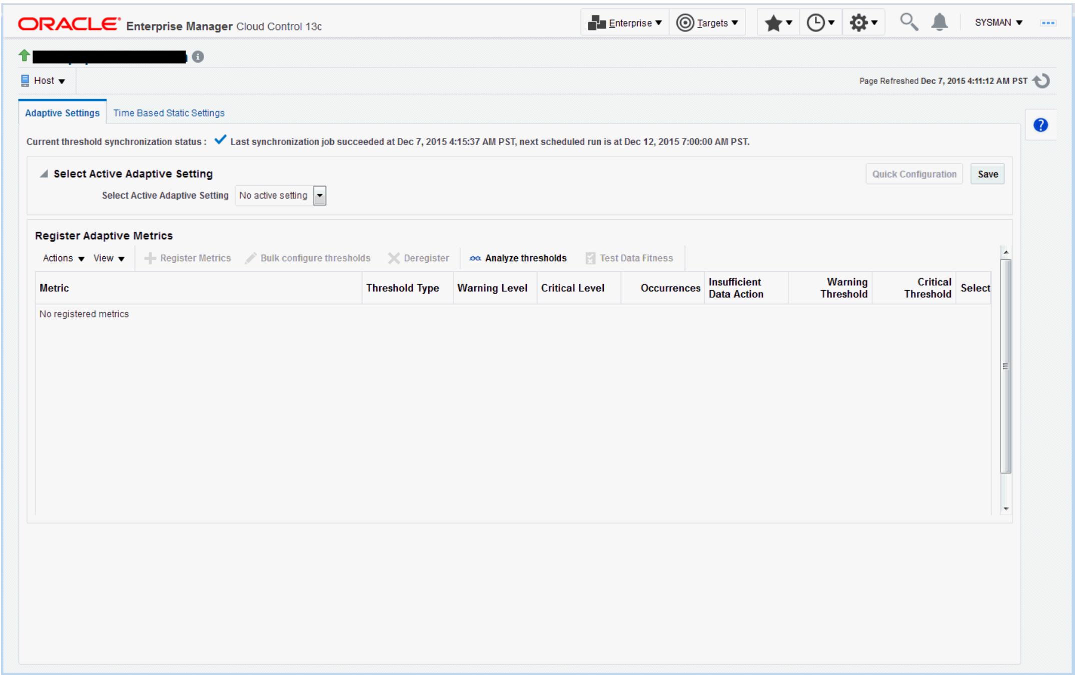

[redacted]

Page Refreshed Dec 7, 2015 4:11:12 AM PST

Current threshold synchronization status : ✓ Last synchronization job succeeded at Dec 7, 2015 4:15:37 AM PST, next scheduled run is at Dec 12, 2015 7:00:00 AM PST.

Select Active Adaptive Setting

Select Active Adaptive Setting: No active setting

Register Adaptive Metrics

Metric	Threshold Type	Warning Level	Critical Level	Occurrences	Insufficient Data Action	Warning Threshold	Critical Threshold	Select
No registered metrics								

3. From the **Select Active Adaptive Setting** menu, select **Moving Window**, additional controls are displayed allowing you to define the moving window's *Threshold Change Frequency* and the *Accumulated Trailing Data* that will be used to compute the adaptive thresholds.

   - *Threshold Change Frequency* (The target timezone is used.)

     – **None**: One set of thresholds will be calculated using past data. This set of thresholds will be valid for the entire week.

       None should be used when there is no usage pattern between daytime versus nighttime or within hours of a day.

     – **By Day and Night**: Two sets of thresholds (day and night) will be calculated using past data. Day thresholds will be calculated using previous day's daytime data, Night thresholds will be calculated using previous day's nighttime data. Thresholds will be changed every day and night.

       By Day and Night should be used when there are distinct performance and usage variations between day hours and night hours.

     – **By Weekdays and Weekend**: Two sets of thresholds (weekdays and weekend) will be calculated using past data. Weekdays thresholds will be calculated using the previous weekdays data. Weekend thresholds will be calculated using the previous weekend data. Thresholds will be changed at start of the weekdays and start of the weekend.

     – **By Day and Night, over Weekdays and Weekend**: Four sets of thresholds will be calculated using past data. *Weekdays Day* thresholds will be calculated using the previous weekday's daytime data, *Weekdays Night* thresholds will be calculated using the previous weekday's nighttime data. *Weekends Day* thresholds will be calculated using previous weekend's daytime data, *Weekends Night* thresholds will be calculated using previous weekend's nighttime data. Thresholds will be changed each day and night.

       Weekday day hours (7a.m. to 7p.m)

       Weekend day hours (7am to 7pm)

    Weekday night hours (7pm to 7am)

    Weekend night hours (7pm to 7am)

    - **By Day of Week**: Seven sets of thresholds will be calculated, one for each day of the week. Thresholds will be calculated using the previous week's same-day data. Thresholds will be changed every day.

      By Day of Week should be used when there is significant daily variation in usage for each day of week.

    - **By Day and Night, per Day of Week**: Fourteen sets of thresholds will be calculated, one for each day of the week and one for each night of the week. *Day* thresholds will be calculated using previous weeks same-day daytime data, *Night* thresholds will be calculated using the previous week's same-day nighttime data. Thresholds will be changed every day and night each day of the week.

- *Accumulated Trailing Data*

  Total time period for which metric data will be collected. Options are 7, 14, 21, and 28 days. In general, you should select the larger value as the additional data helps in computing more accurate thresholds.

4. Click **Save** and accept the confirmation. The **Register Metrics** button becomes active in the Register Adaptive Metrics region. Click R**egister Metrics.** The Metric Selector dialog displays.

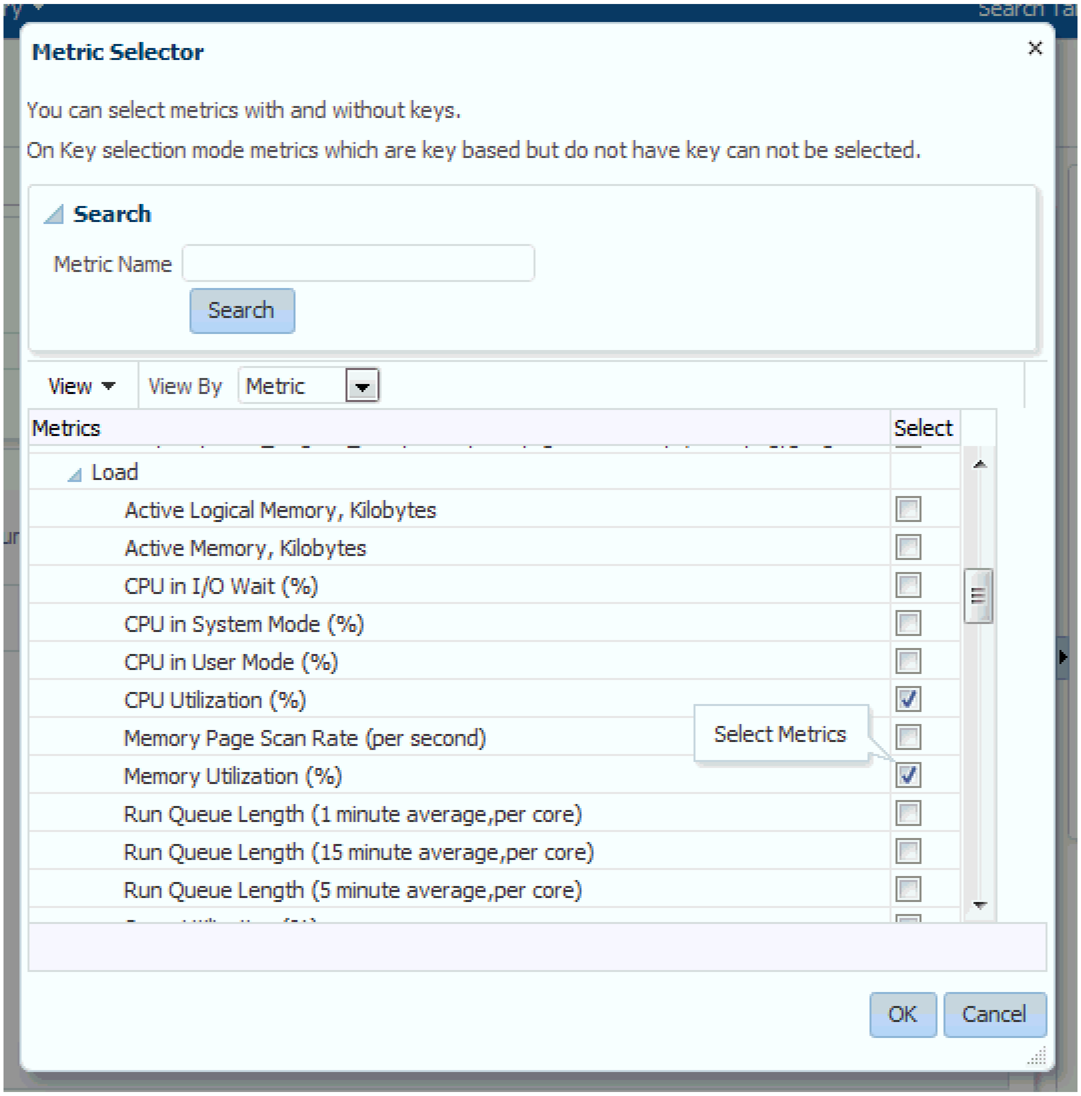

Metric Selector

You can select metrics with and without keys.

On Key selection mode metrics which are key based but do not have key can not be selected.

Metric Name

Metrics	Select
Load	
Active Logical Memory, Kilobytes	□
Active Memory, Kilobytes	□
CPU in I/O Wait (%)	□
CPU in System Mode (%)	□
CPU in User Mode (%)	□
CPU Utilization (%)	☑
Memory Page Scan Rate (per second)	□
Memory Utilization (%)	☑
Run Queue Length (1 minute average,per core)	□
Run Queue Length (15 minute average,per core)	□
Run Queue Length (5 minute average,per core)	□

5. Select the desired metric(s) and then click **OK**. A confirmation dialog displays stating that the selected metric(s) will be added to this target's Adaptive Setting. Click **Yes** to confirm the action. The selected metrics appear in the Register Adaptive Metrics region.

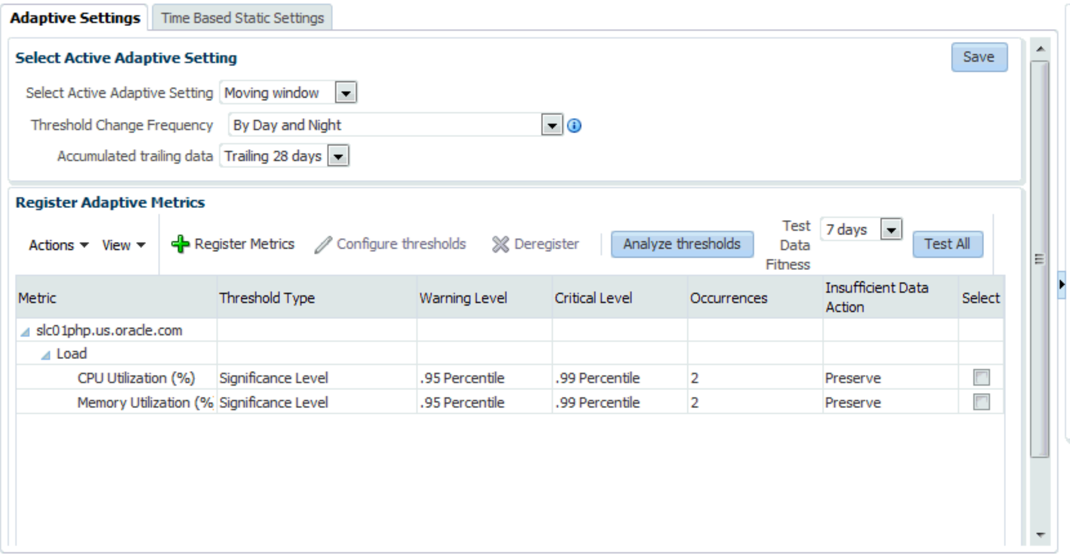

Select Active Adaptive Setting

Select Active Adaptive Setting: Moving window

Threshold Change Frequency: By Day and Night

Accumulated trailing data: Trailing 28 days

Register Adaptive Metrics

Test Data Fitness: 7 days

Metric	Threshold Type	Warning Level	Critical Level	Occurrences	Insufficient Data Action	Select
slc01php.us.oracle.com						
Load						
CPU Utilization (%)	Significance Level	.95 Percentile	.99 Percentile	2	Preserve	□
Memory Utilization (%	Significance Level	.95 Percentile	.99 Percentile	2	Preserve	□

6. Once registered as adaptive metrics, you can then select individual metrics to configure thresholds. When metrics are first registered, by default, Enterprise Manager enables *Significance Level* and sets the warning and critical thresholds at 95 and 99 percentile respectively.

Configuring Adaptive Thresholds

Once you have registered the adaptive metrics, you now have the option of configuring the thresholds if the predefined thresholds do not meet your monitoring requirements.

To configure adaptive thresholds:

1. From the Register Adaptive Metrics region, select the metric(s) you wish to configure and click **Bulk Configure Thresholds**. The Configure Thresholds dialog displays.

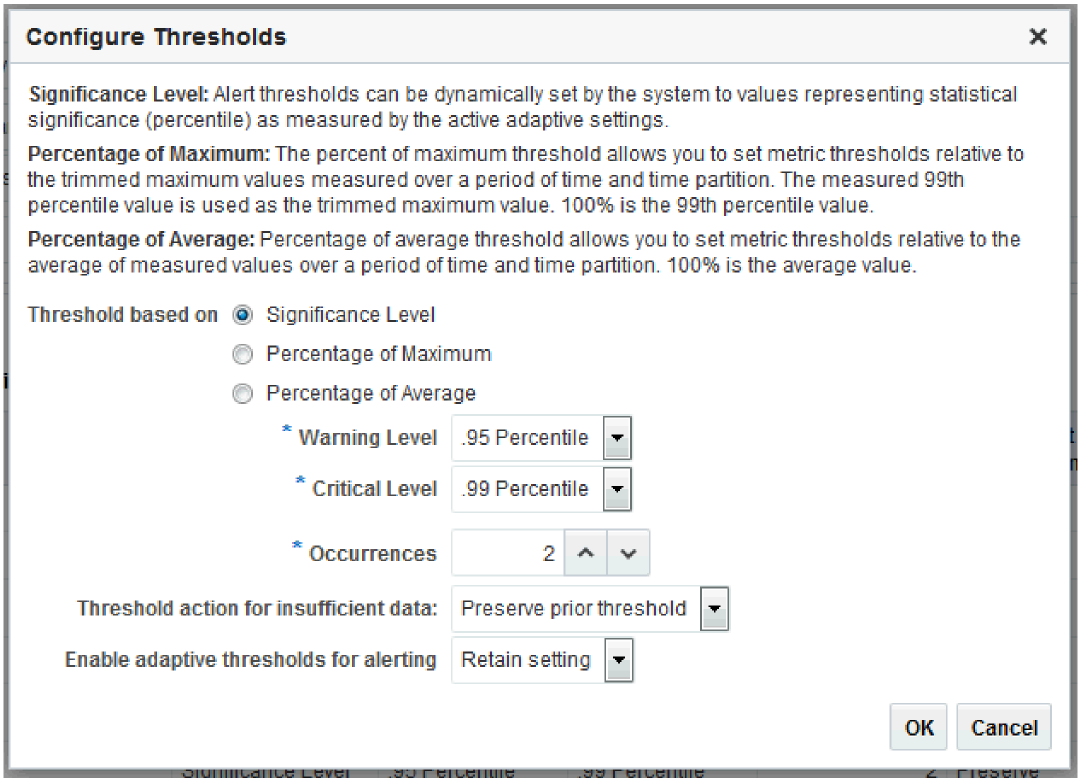

Configure Thresholds

Significance Level: Alert thresholds can be dynamically set by the system to values representing statistical significance (percentile) as measured by the active adaptive settings.

Percentage of Maximum: The percent of maximum threshold allows you to set metric thresholds relative to the trimmed maximum values measured over a period of time and time partition. The measured 99th percentile value is used as the trimmed maximum value. 100% is the 99th percentile value.

Percentage of Average: Percentage of average threshold allows you to set metric thresholds relative to the average of measured values over a period of time and time partition. 100% is the average value.

Threshold based on ● Significance Level
○ Percentage of Maximum
○ Percentage of Average

* Warning Level: .95 Percentile

* Critical Level: .99 Percentile

* Occurrences: 2

Threshold action for insufficient data: Preserve prior threshold

Enable adaptive thresholds for alerting: Retain setting

2. Choose whether you want your threshold to be based on:

   Significance Level: Thresholds based on significance level use statistical relevance to determine which current values are statistical outliers. The primary reason to use Significance Level for alerting is that you are trying to detect statistical outliers in metric values as opposed to simply setting a threshold value. Hence, thresholds are percentile based. For example, if the significance level is set to .95 for a warning threshold, the metric threshold is set where 5% of the collected metric values fall outside this value and any current values that exceed this value trigger an alert. A higher significance level of .98 or .99 will cause fewer alerts to be triggered.

   Percentage of Maximum: These types of thresholds compute the threshold values based on specified percentages of the maximum observed over the period of time you selected. Percentage-of-maximum-based alerts are generated if the current value is at or above the percentage of maximum you specify. For example, if a maximum value of 1000 is encountered during a time group, and if 105 is specified as the Warning level, then values above 1050 (105% of 1000 = 1050) will raise an alert.

   Percentage of Average. Based on the time grouping and bucketing, the threshold average is computed allowing you to set metric thresholds relative to the average of measured values over a period of time and time partition. 100% is the average value.

   For all types of alerts you can set the **Occurrences** parameter, which is the number of times the metric crosses a threshold value before an alert is generated.

   Clear Threshold: Thresholds for the selected metrics will be cleared. No Alert will be generated. Use this option when you do not want any thresholds set for the metrics but you do not want to remove historical data. **Important**: Deregistering metrics will remove the historical data.

   Depending on the option selected, the Warning, Critical, and Occurrence setting options will change.

   You can set the deviation over the computed average value.

The **Threshold action for insufficient data** menu allows you set the appropriate action for Enterprise Manager to take if there is not enough data to calculate a valid metric threshold. There are two actions available: *Preserve the prior threshold* and *Suppress Alerts*.

Maximum Allowed Threshold allows you to set a threshold limit that, if crossed, will raise a critical alert. A warning alert will automatically be raised at 25% less the maximum value.

3. Click **OK** to set the changes.

Determining whether Adaptive Thresholds are Correct

Even though Enterprise Manager will use the adaptive threshold settings to determine an accurate target workload-metric threshold match, it is still be necessary to match the metric sampling schedule with the actual target workload. For example, your moving window baseline period (see *Moving Window Baseline Periods*) should match the target workloads. In some situations, you may not know the actual target workloads, in which case setting adaptive thresholds may be problematic.

To help you determine the validity of your adaptive thresholds, Enterprise Manager allows you to analyze threshold using various adaptive settings to determine whether the settings are correct.

To analyze existing adaptive thresholds:

1. From the Register Adaptive Metrics region, click **Analyze Thresholds**.

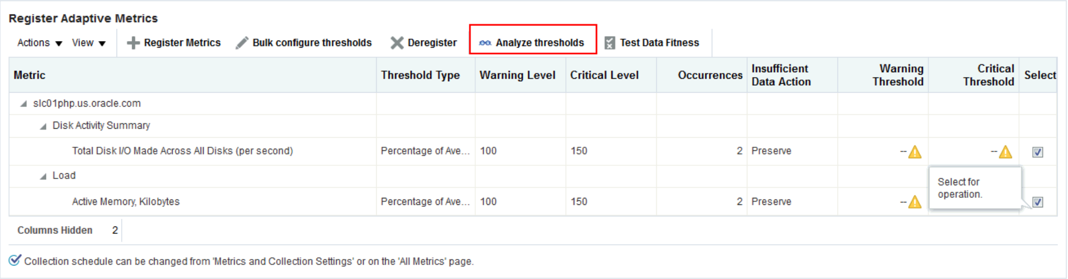

Metric	Threshold Type	Warning Level	Critical Level	Occurrences	Insufficient Data Action	Warning Threshold	Critical Threshold	Select
slc01php.us.oracle.com								
Disk Activity Summary								
Total Disk I/O Made Across All Disks (per second)	Percentage of Ave...	100	150	2	Preserve	-- ⚠	-- ⚠	☑
Load								
Active Memory, Kilobytes	Percentage of Ave...	100	150	2	Preserve	-- ⚠		☑

Select for operation.

Columns Hidden 2

Collection schedule can be changed from 'Metrics and Collection Settings' or on the 'All Metrics' page.

The Analyze Threshold page displays containing historical metric data charts (one for each metric).

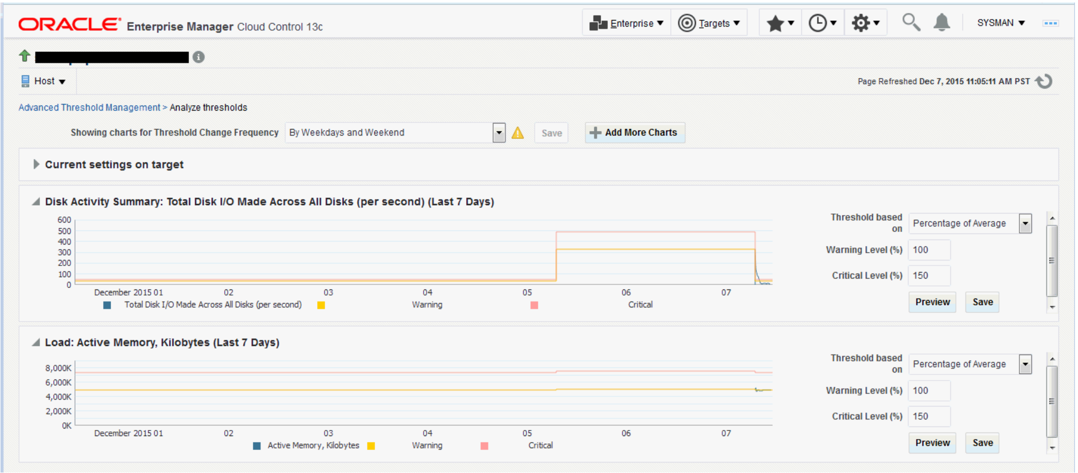

[redacted]

Page Refreshed Dec 7, 2015 11:05:11 AM PST

Advanced Threshold Management > Analyze thresholds

Showing charts for Threshold Change Frequency: By Weekdays and Weekend

Disk Activity Summary: Total Disk I/O Made Across All Disks (per second) (Last 7 Days)

Threshold based on: Percentage of Average

Warning Level (%): 100

Critical Level (%): 150

Load: Active Memory, Kilobytes (Last 7 Days)

Threshold based on: Percentage of Average

Warning Level (%): 100

Critical Level (%): 150

2. Modify the adaptive threshold parameters to closely match metric threshold settings with the target workload. You can experiment with the following adaptive metric parameters:

   Threshold Change Frequency

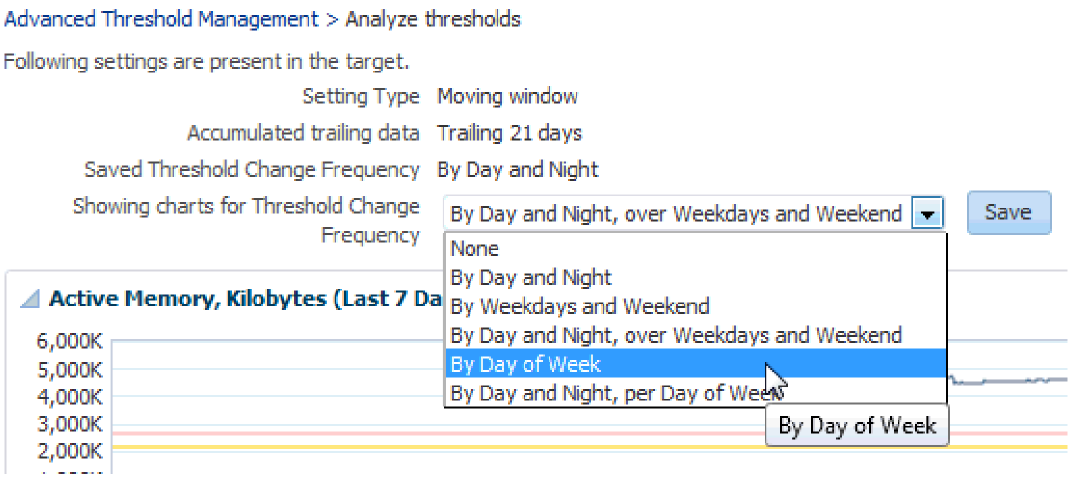

Advanced Threshold Management > Analyze thresholds

Following settings are present in the target.

Setting Type	Moving window
Accumulated trailing data	Trailing 21 days
Saved Threshold Change Frequency	By Day and Night
Showing charts for Threshold Change Frequency	By Day and Night, over Weekdays and Weekend

Save
None
By Day and Night
By Weekdays and Weekend
By Day and Night, over Weekdays and Weekend
By Day of Week
By Day and Night, per Day of Week
By Day of Week
Active Memory, Kilobytes (Last 7 Da
6,000K
5,000K
4,000K
3,000K
2,000K

   Threshold Based On

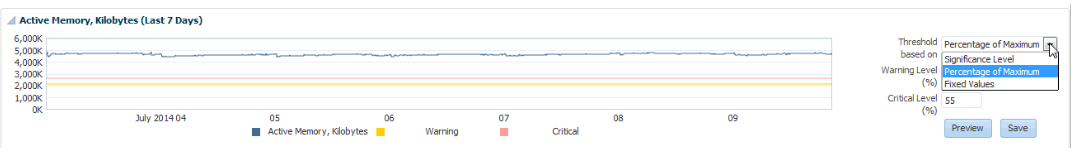

Threshold based on	Percentage of Maximum
Warning Level (%)	
Critical Level (%)	55

   Metric Warning and Critical Thresholds

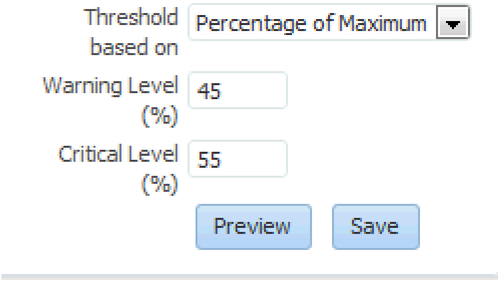

Threshold based on	Percentage of Maximum
Warning Level (%)	45
Critical Level (%)	55

3. Once you are satisfied with the modifications for the Threshold Change Frequency or any of the individual metrics, click **Save** to set the new parameters.

Testing Adaptive Metric Thresholds

Because adaptive metric thresholds utilize statistical sampling of data over time, the accuracy of the thresholds will rely on the quantity and quality of the data collected. Hence, a sufficient amount of metric data needs to have been collected in order for the thresholds to be valid. To verify whether enough data has been collected for metrics registered with adaptive thresholds, use the **Test Data Fitness** function.

1. From the Registered Adaptive Metrics region, click **Test Data Fitness**.

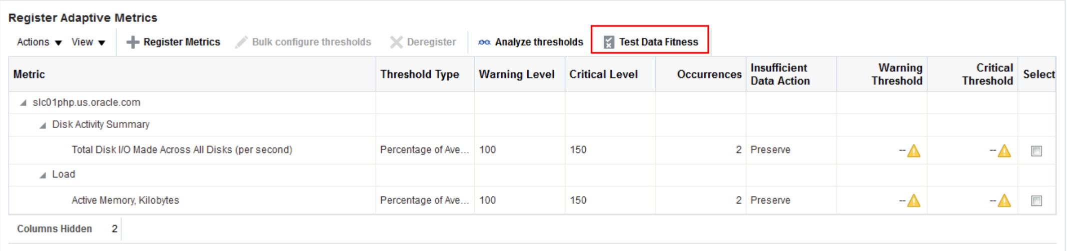

Register Adaptive Metrics

Metric	Threshold Type	Warning Level	Critical Level	Occurrences	Insufficient Data Action	Warning Threshold	Critical Threshold	Select
slc01php.us.oracle.com								
Disk Activity Summary								
Total Disk I/O Made Across All Disks (per second)	Percentage of Ave...	100	150	2	Preserve	--⚠	--⚠	□
Load								
Active Memory, Kilobytes	Percentage of Ave...	100	150	2	Preserve	--⚠	--⚠	□

Columns Hidden 2

Enterprise Manager evaluates the adaptive threshold metrics and then displays the results in the Test Metrics window.

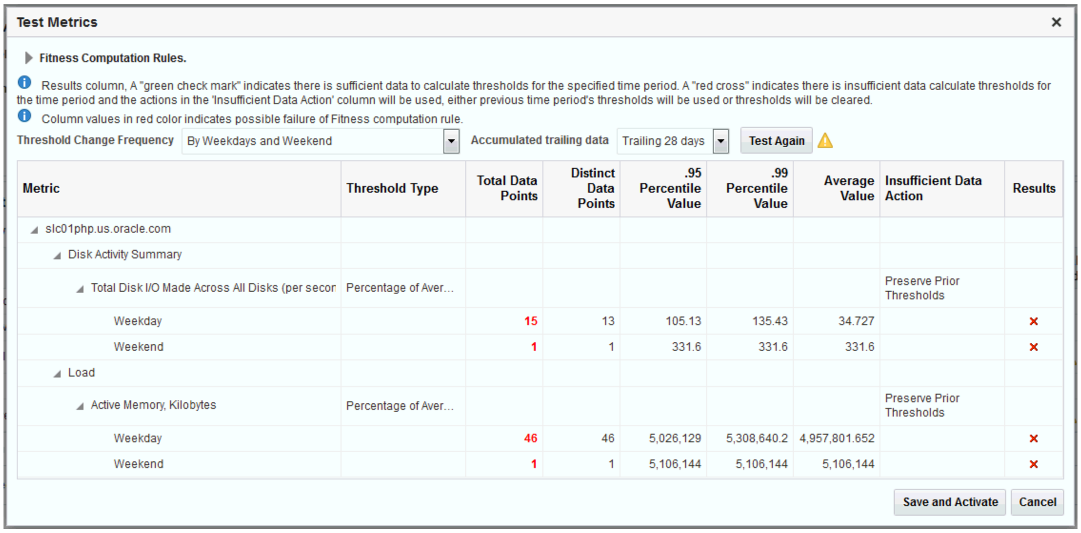

Test Metrics

Fitness Computation Rules.

Results column, A "green check mark" indicates there is sufficient data to calculate thresholds for the specified time period. A "red cross" indicates there is insufficient data calculate thresholds for the time period and the actions in the 'Insufficient Data Action' column will be used, either previous time period's thresholds will be used or thresholds will be cleared.

Column values in red color indicates possible failure of Fitness computation rule.

Threshold Change Frequency: By Weekdays and Weekend   Accumulated trailing data: Trailing 28 days

Metric	Threshold Type	Total Data Points	Distinct Data Points	.95 Percentile Value	.99 Percentile Value	Average Value	Insufficient Data Action	Results
slc01php.us.oracle.com								
Disk Activity Summary								
Total Disk I/O Made Across All Disks (per secor	Percentage of Aver...						Preserve Prior Thresholds	
Weekday		15	13	105.13	135.43	34.727		✗
Weekend		1	1	331.6	331.6	331.6		✗
Load								
Active Memory, Kilobytes	Percentage of Aver...						Preserve Prior Thresholds	
Weekday		46	46	5,026,129	5,308,640.2	4,957,801.652		✗
Weekend		1	1	5,106,144	5,106,144	5,106,144		✗

If there is sufficient data collected to compute the adaptive threshold, a green check appears in the results column. A red 'x' appears in the results column if there is insufficient data collected. To resolve this situation, you can use a longer accumulating trailing data window. Additionally, you can have metric data collected more frequently.

2. Click **Save and Activate** once you are finished viewing the results.

Deregistering Adaptive Threshold Metrics

If you no longer want specific metrics to be adaptive, you can deregister them at any time. To deregister an adaptive threshold metric:

1. From the Register Adaptive Metrics regions, select the metric(s) you wish to deregister.
2. Click **Deregister**. A confirmation displays asking if you want the metric removed from the target's adaptive setting.
3. Click **Yes**.

Setting Adaptive Thresholds using Monitoring Templates

You can use monitoring templates to apply adaptive thresholds broadly across targets within your environment. For example, using a monitoring template, you can apply adaptive threshold setting for the CPU Utilization metric for all Host targets.

To apply adaptive thresholds using monitoring templates:

1. Create a template out of a target that already has adaptive threshold settings enabled.

   From the **Enterprise** menu, select **Monitoring** and then **Monitoring Templates**.The Monitoring Templates page displays.
2. Click **Create**. The **Create Monitoring Template: Copy Monitoring Settings** page displays.
3. Choose a target on which adaptive thresholds have already been set and click **Continue**.
4. Enter a template **Name** and a brief **Description**. Click **OK**.

Once the monitoring template has been created, you can view or edit the template as you would any other template. To modify, add, or delete adaptive metrics in the template:

1. From the **Enterprise** menu, select **Monitoring** and then **Monitoring Templates**.The **Monitoring Templates** page displays.
2. On the **Monitoring Templates** page, select the monitoring template from the list.
3. From the **Actions** menu, select **Edit Advanced Monitoring Settings**. The **Edit Advanced Monitoring Settings** page displays with the **Adaptive Settings** tab selected.
4. Modify the adaptive metrics as required.

Time-based Static Thresholds

Time-based static thresholds allow you to define specific threshold values to be used at different times to account for changing workloads over time. Using time-based static thresholds can be used whenever the workload schedule for a specific target is well known or if you know what thresholds you want to specify.

Registering Time-based Static Thresholds

To register metrics with time-based static thresholds:

1. From the target menu (Host is used in this example), select **Monitoring** and then **Metric and Collection Settings**.
2. In the Related Links area, click **Advanced Threshold Management**.The Advanced Threshold Management page displays.

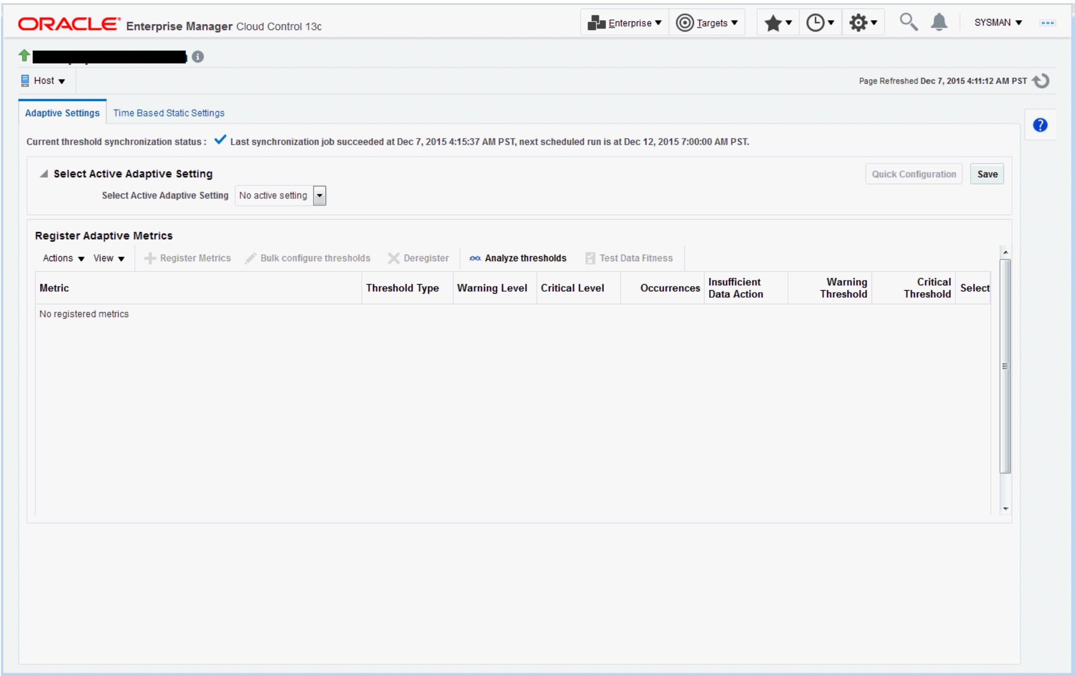

[redacted]

Page Refreshed Dec 7, 2015 4:11:12 AM PST

Adaptive Settings | Time Based Static Settings

Current threshold synchronization status : ✓ Last synchronization job succeeded at Dec 7, 2015 4:15:37 AM PST, next scheduled run is at Dec 12, 2015 7:00:00 AM PST.

Select Active Adaptive Setting

Select Active Adaptive Setting: No active setting

Quick Configuration | Save

Register Adaptive Metrics

Actions | View | Register Metrics | Bulk configure thresholds | Deregister | Analyze thresholds | Test Data Fitness

Metric	Threshold Type	Warning Level	Critical Level	Occurrences	Insufficient Data Action	Warning Threshold	Critical Threshold	Select
No registered metrics								

3. Click on the **Time Based Static Settings** tab.
4. Select the **Threshold Change Frequency**.
5. Click **Register Metrics**.

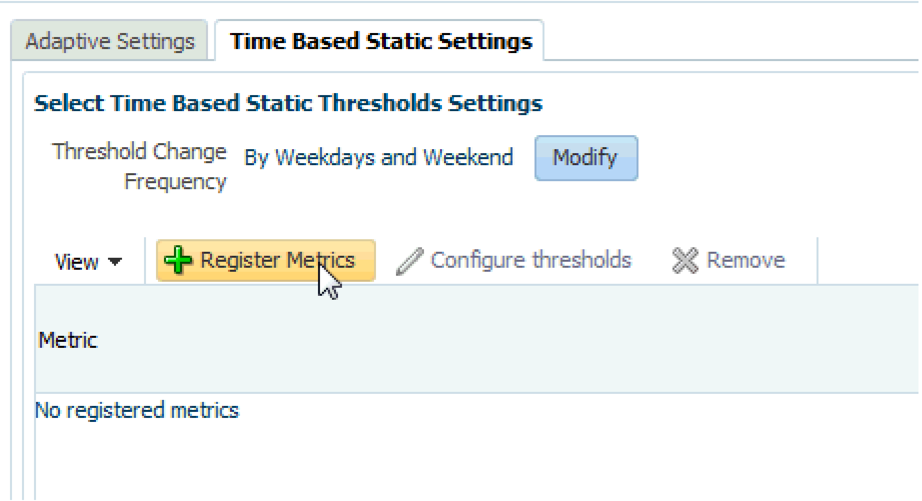

Adaptive Settings | Time Based Static Settings

Select Time Based Static Thresholds Settings

Threshold Change Frequency: By Weekdays and Weekend

Modify

View | Register Metrics | Configure thresholds | Remove

Metric
No registered metrics

The Metric Selector dialog displays.

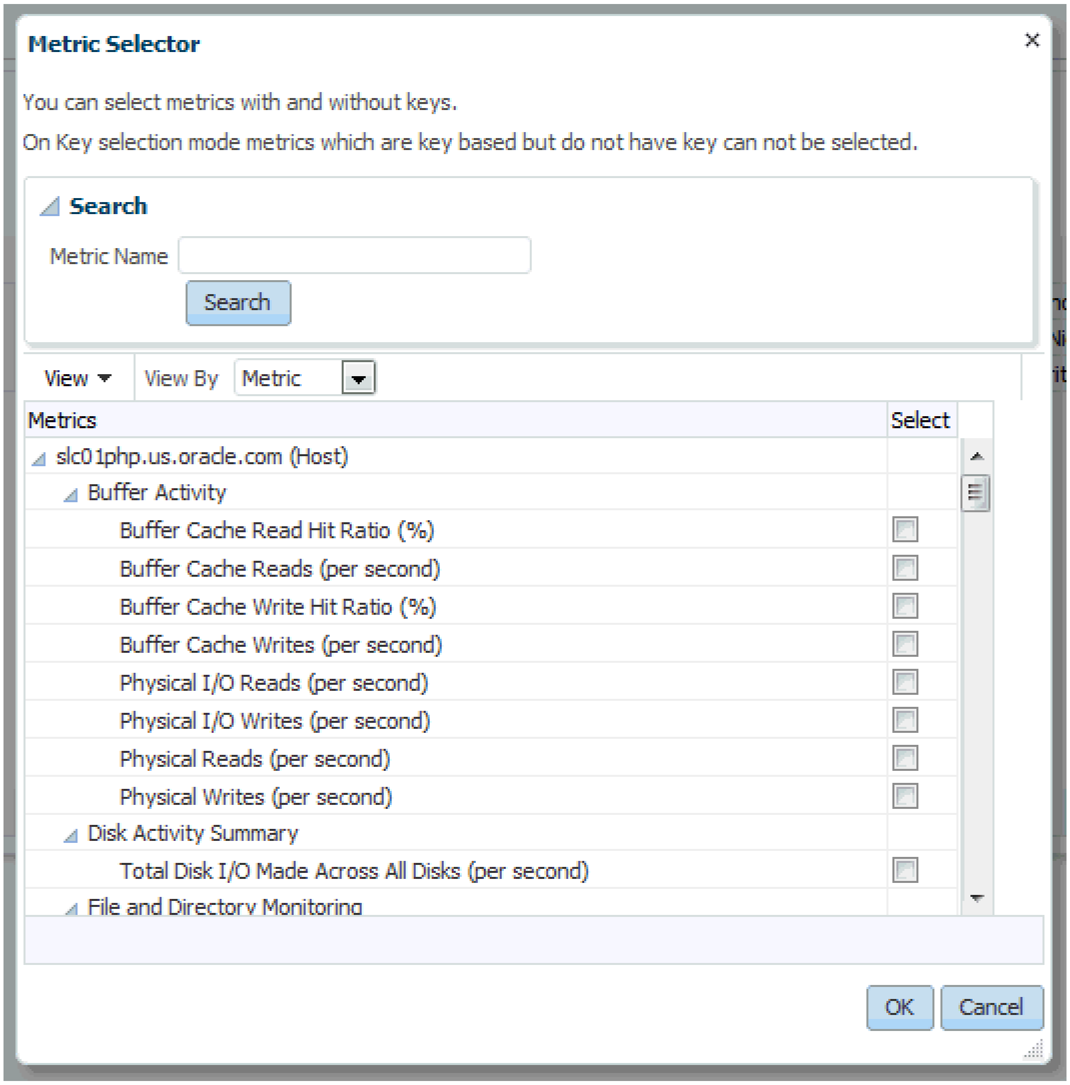

Metric Selector

You can select metrics with and without keys.

On Key selection mode metrics which are key based but do not have key can not be selected.

Search

Metric Name

Metrics	Select
slc01php.us.oracle.com (Host)	
Buffer Activity	
Buffer Cache Read Hit Ratio (%)	□
Buffer Cache Reads (per second)	□
Buffer Cache Write Hit Ratio (%)	□
Buffer Cache Writes (per second)	□
Physical I/O Reads (per second)	□
Physical I/O Writes (per second)	□
Physical Reads (per second)	□
Physical Writes (per second)	□
Disk Activity Summary	
Total Disk I/O Made Across All Disks (per second)	□
File and Directory Monitoring	

6. Select the desired metric(s) and click **OK**.

   The selected metrics appear in the Registered Metrics table.

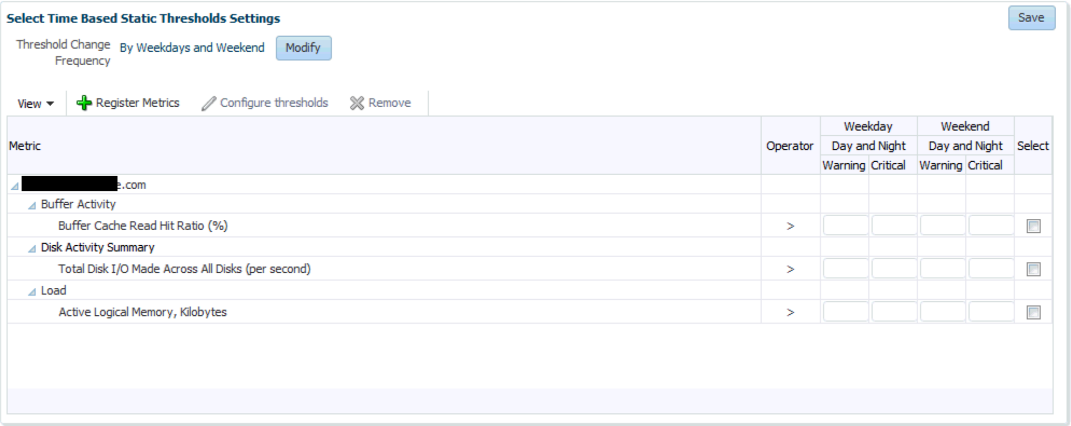

Select Time Based Static Thresholds Settings

Threshold Change Frequency: By Weekdays and Weekend

Metric	Operator	Weekday Day and Night Warning	Weekday Day and Night Critical	Weekend Day and Night Warning	Weekend Day and Night Critical	Select
[redacted].com						
Buffer Activity						
Buffer Cache Read Hit Ratio (%)	>					□
Disk Activity Summary						
Total Disk I/O Made Across All Disks (per second)	>					□
Load						
Active Logical Memory, Kilobytes	>					□

7. Enter the desired metric thresholds and click **Save** once you are done.

If you want to set the thresholds for multiple metrics simultaneously, check the *Select* box for the metrics you want to update and click **Bulk Configure Thresholds**. The Configure Thresholds dialog displays.

Enter the revised Warning and Critical threshold values and click **OK**. A confirmation dialog displays stating that existing metric threshold values will be overwritten. Click **Yes**.

8. Optionally, you can change the Threshold Change Frequency. To do so, from the Time Based Static Thresholds Settings page, click **Modify**. The Modify Threshold Change Frequency dialog displays allowing you to select a new change frequency. Select a new frequency and click **OK**.

   A confirmation dialog displays stating that changing Threshold Change Frequency will affect all the registered metrics and whether you want to continue. Click **Yes** to proceed.

9. Click **Save** to ensure all changes have been saved to the Enterprise Manager repository.

Deregistering Time-based Static Thresholds

If you no longer require time-based static threshold metrics, you can deregister them from the target.

To deregister time-based static metric thresholds:

1. From the Time Based Static Thresholds tab, select the metric(s) you want to deregister.

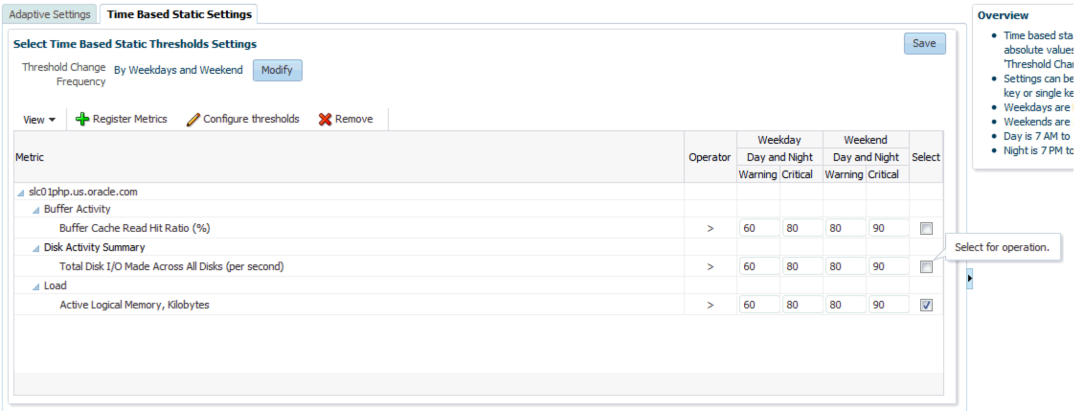

Metric	Operator	Weekday		Weekend		Select
		Day and Night		Day and Night		
		Warning	Critical	Warning	Critical	
slc01php.us.oracle.com						
Buffer Activity						
Buffer Cache Read Hit Ratio (%)	>	60	80	80	90	□
Disk Activity Summary						
Total Disk I/O Made Across All Disks (per second)	>	60	80	80	90	□
Load						
Active Logical Memory, Kilobytes	>	60	80	80	90	☑

2. Click **Remove**. The metric entry is removed from list of

3. Click **Save** to save the changes to the Enterprise Manager repository.

Determining What is a Valid Metric Threshold

As previously discussed, static thresholds do not account for expected performance variation due to increased/decreased workloads encountered by the target, such as the workload encountered by a warehouse database target against which OLTP transactions are performed. Workloads can also change based on different time periods, such as weekday versus weekend, or day versus night. These types of workload variations present conditions where fixed static metric threshold values may cause monitoring issues, such as the generation of false and/or excessive metric alerts. Ultimately, your monitoring needs dictate how to best go about obtaining accurate metric thresholds.

15

Utilizing the Job System and Corrective Actions

The Enterprise Manager 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:

- Diagnosing Job System Issues with the Job System Console
- Job System Purpose and Overview
- Preliminary Considerations
- Creating Jobs
- Viewing and Analyzing Job Status
- Generating Job Event Criteria
- Creating Event Rules For Job Status Change
- Using Diagnostic Tools
- Creating Corrective Actions

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
- Centralizes environment job scheduling into one robust tool

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 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 15-1 Job Activity Page

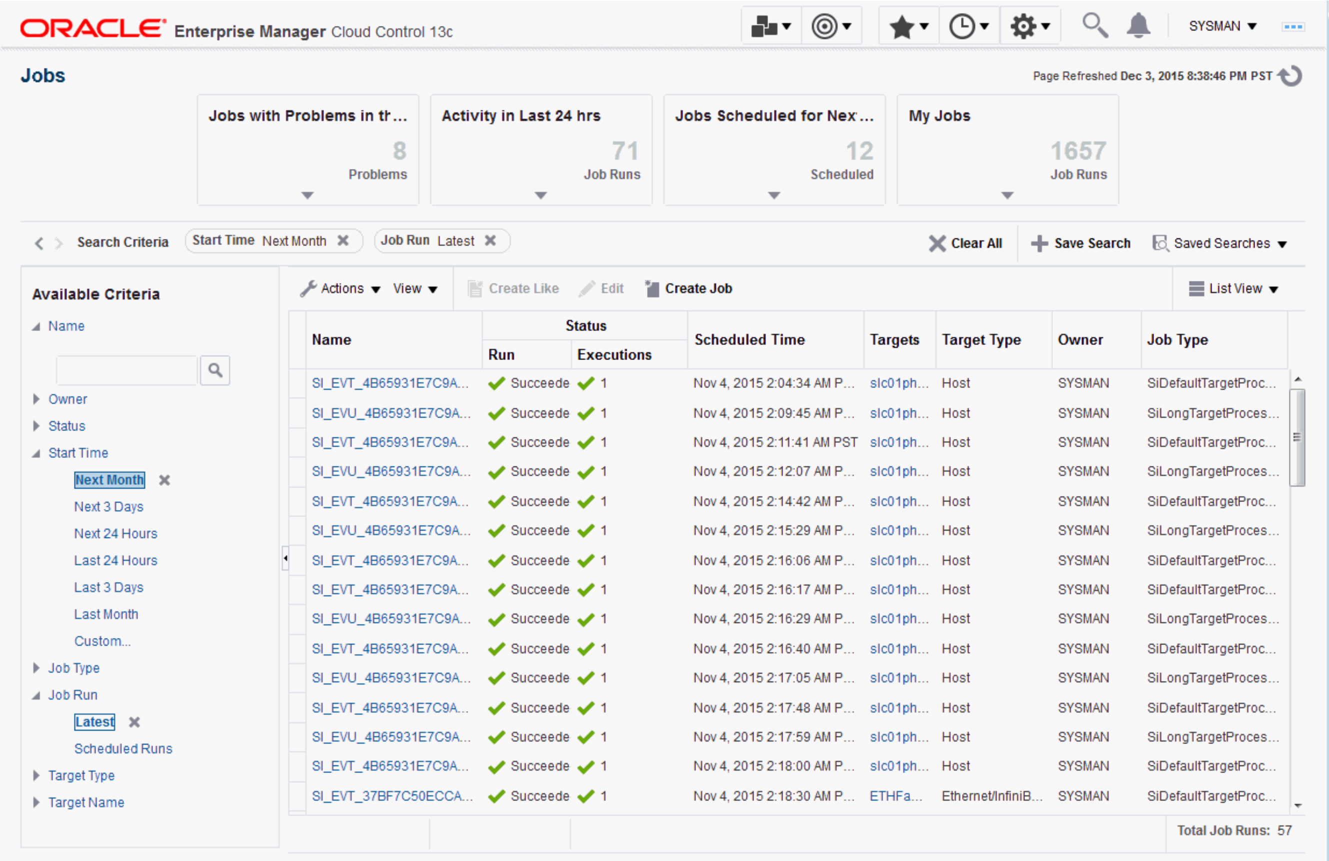

Name	Status		Scheduled Time	Targets	Target Type	Owner	Job Type
	Run	Executions					
SI_EVT_4B65931E7C9A...	✓ Succeede	✓ 1	Nov 4, 2015 2:04:34 AM P...	slc01ph...	Host	SYSMAN	SiDefaultTargetProc...
SI_EVU_4B65931E7C9A...	✓ Succeede	✓ 1	Nov 4, 2015 2:09:45 AM P...	slc01ph...	Host	SYSMAN	SiLongTargetProces...
SI_EVT_4B65931E7C9A...	✓ Succeede	✓ 1	Nov 4, 2015 2:11:41 AM PST	slc01ph...	Host	SYSMAN	SiDefaultTargetProc...
SI_EVU_4B65931E7C9A...	✓ Succeede	✓ 1	Nov 4, 2015 2:12:07 AM P...	slc01ph...	Host	SYSMAN	SiLongTargetProces...
SI_EVT_4B65931E7C9A...	✓ Succeede	✓ 1	Nov 4, 2015 2:14:42 AM P...	slc01ph...	Host	SYSMAN	SiDefaultTargetProc...
SI_EVU_4B65931E7C9A...	✓ Succeede	✓ 1	Nov 4, 2015 2:15:29 AM P...	slc01ph...	Host	SYSMAN	SiLongTargetProces...
SI_EVT_4B65931E7C9A...	✓ Succeede	✓ 1	Nov 4, 2015 2:16:06 AM P...	slc01ph...	Host	SYSMAN	SiDefaultTargetProc...
SI_EVT_4B65931E7C9A...	✓ Succeede	✓ 1	Nov 4, 2015 2:16:17 AM P...	slc01ph...	Host	SYSMAN	SiDefaultTargetProc...
SI_EVU_4B65931E7C9A...	✓ Succeede	✓ 1	Nov 4, 2015 2:16:29 AM P...	slc01ph...	Host	SYSMAN	SiLongTargetProces...
SI_EVT_4B65931E7C9A...	✓ Succeede	✓ 1	Nov 4, 2015 2:16:40 AM P...	slc01ph...	Host	SYSMAN	SiDefaultTargetProc...
SI_EVU_4B65931E7C9A...	✓ Succeede	✓ 1	Nov 4, 2015 2:17:05 AM P...	slc01ph...	Host	SYSMAN	SiLongTargetProces...
SI_EVT_4B65931E7C9A...	✓ Succeede	✓ 1	Nov 4, 2015 2:17:48 AM P...	slc01ph...	Host	SYSMAN	SiDefaultTargetProc...
SI_EVU_4B65931E7C9A...	✓ Succeede	✓ 1	Nov 4, 2015 2:17:59 AM P...	slc01ph...	Host	SYSMAN	SiLongTargetProces...
SI_EVT_4B65931E7C9A...	✓ Succeede	✓ 1	Nov 4, 2015 2:18:00 AM P...	slc01ph...	Host	SYSMAN	SiDefaultTargetProc...
SI_EVT_37BF7C50ECCA...	✓ Succeede	✓ 1	Nov 4, 2015 2:18:30 AM P...	ETHFa...	Ethernet/InfiniB...	SYSMAN	SiDefaultTargetProc...

Total Job Runs: 57

Besides accessing the Job Activity page from the Enterprise menu, you can also access this page from any target-specific menu for all target types by selecting Job Activity from the target type's 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.

Changing Job Activity Summary Table Views

In addition to listing job executions in a conventional tabular format, Enterprise Manager also allows you to display the job executions using in using alternate views that can make job execution data more meaningful. For example, seeing which jobs are using the most resource or are taking the longest to run.

There are three job activity display options:

- **List**: Conventional tabular display of job information.
- **Summary**: Displays a graphical rollup of job runs based on selected attributes of attributes. Clicking on any cell in the rollup view further refines the search and takes you to the tabular view where you can analyze job runs in more granular detail.

Job Searches

For convenience, you can define job searches that allow you to view/access specific jobs that are of frequent interest. By default, there are predefined job searches at the top of the Job Activity page.

- Jobs with problems in the Last 24 Hours
- Job Run activity in the last 24 hours
- Jobs scheduled to run in the next 24 hours
- Jobs belonging to the currently logged in user

Saving Job Searches

Saving commonly used job searches allows you to view pertinent job information quickly.

To create a saved search:

1. From the Available Criteria region, choose the requisite parameters for your search. The results display immediately in the Job Activity table.
2. Click the **Save Search** button. The Create Saved Search dialog displays.
3. Enter a name for your search.
4. Choose how want your saved search displayed:

   Show this search when I come to this page (available from the Saved Searches drop-down list)

   OR

   Show on top of the page (available as one of the summary boxes at the top of the page)
5. Click **OK** to save your search.
6. Click **Run** to run your search and view the results.

Editing Saved Job Searches

To edit a saved job search:

1. From the **Saved Searches** menu, select **Manage Saved Searches**. The Manage Saved Searches dialog displays.
2. Click the **Edit** icon (pencil). The Edit Saved Search dialog displays. Note: In addition to editing search criteria, you can set whether you wan the job to be a **Default** (Choosing **Default** runs the search whenever you access the Job Activity page.) or want the job to appear at the top of the Job Activity page.
3. Click **OK** to save your Job search changes.
4. Click **Done** to close the Manage Saved Searches dialog.

Importing/Exporting Saved Job Searches

By default, you can only see searches that you have created. To use searches other administrators have created, you can import them. Conversely, you can make job searches you have created available to other administrators by exporting your searches. Once the job searches are accessible to you, they will appear in your list of saved searches.

What Are Job Executions and Job Runs?

The following sections explain the characteristics of each of these.

Job Executions

Job executions are usually associated with one target, such as a patch job on a particular database. These are called single-target jobs because each execution has only one target. However, job executions are not always a one-to-one mapping to a target. Some executions have multiple targets, such as comparing hosts. These jobs are called single-execution jobs, since there is only one execution for all the targets. When a job is run against multiple targets, it runs in one or many executions depending on whether it is a single-execution or single-target job. A few jobs have no target. These jobs are called targetless jobs and run in one execution.

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.

Job Runs

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.

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, or a maximum of one day, the execution of this job would be marked "Skipped." A job is also skipped when the scheduled time plus the grace period has passed.

- **Resume jobs** —

  After you suspend a job, any scheduled executions do not occur until you decide to resume the job.

- **Retry all failed executions in a job run** —

  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 all failed executions in a job run**. Successful executions are never retried.

> **Note:**
>
> For more information on job executions and runs, refer to Enterprise Manager online help.

Preliminary Considerations

Before proceeding to the procedural information presented in Creating Jobs, it is suggested that you read the topics presented in the sections below:

- Administrator Roles
- Creating Scripts
- Sharing Job Responsibilities
- Submitting Jobs for Groups

Administrator Roles

Enterprise Manager provides the following administrator types:

- **Administrator** — Most jobs and other activities should be initiated using this "normal" user type
- **Super Administrator** — There may be limited use cases for a super administrator to run jobs, create blackouts, or own targets.
- **Repository owner (SYSMAN)** — The special repository owner user SYSMAN should almost never own or do any of the tasks listed for the other two types above. This user should only be reserved for top-level actions, such as setting up the site and so forth.

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.

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.

Submitting Jobs for Groups

Rather than listing a large number of targets individually, you can use a group as the target of a job. All member targets in the group that match the selected target type of the job are selected as actual targets of the job when it runs. If the membership of the group changes, the actual target list of the job changes with it. If the job repeats, each iteration (or "run") of the job executes on the matching targets in the group at the time of the run.

Overriding the Target Type Selection

To override the target type selection for a group, set targetType=<override_target_type> in the input file for the create_job verb. For example, the default target type for OSCommand jobs is "host". To submit a job against a group of databases, specify:

```
target_list=my_db_group:composite
targetType=oracle_database
```

Note that any targets in the group that do not match the target type selected are ignored.

> **See Also:**
>
> Managing Groups

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

Selecting a Job Type

Using the Job System, you can create a job by clicking Create Job in the Job Activity page and selecting the job type from the Select Job Type dialog. You can find a specifc job type by either searching for the name of a job type or by specifying a target type.

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.

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 **Job**, then **Activity.**
2. Click **Create Job**. The Select Job Type dialog displays.
3. Choose the **OS Command** job type and click **Select**.

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. You can also include clusters in the target list if they are of the same base target type. For example, a host cluster would be selected if the target type is "host" and a RAC database would be selected if the target type is database.
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 Specifying a Single Operation or Specifying a Script.
3. Click the **Credentials** property page link.

> **Note:**
>
> The OS Command relies on the target host's shell to execute the command/interpreter specified. On *nix systems, it is /bin/sh -c and on Windows systems, it is cmd /c. The command line specified is interpreted by the corresponding shell.

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 host-based credentials for the selected targets, 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.

> **Tip:**
>
> preferred credentials are useful when a job is submitted on multiple targets and each target needs to use different credentials for authentication.

- **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 for the host target on the 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. This applies even for named credentials.

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 or PowerBroker, refer to the product guides on their respective product documentation pages.

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. The job system sets the job status to Skipped, if it cannot start the execution between the scheduled time and the time equal to the scheduled time plus the grace period, or within the grace period from the scheduled time.

   - **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:

    - Indefinite: The job will run at the defined repeat interval until it is manually unscheduled.
    - Specified Date: The job will run at the defined repeat interval until the Specified Date is reached.

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. For private roles, it is possible to grant Full access privileges.

   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 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 Enterprise Manager 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 Using 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.

Specifying a Single Operation

> **Note:**
>
> The following information applies to step Creating an OS Command Job in Creating an OS Command Job .

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
  ```

- The job status depends on the exit code returned by the command. If the command execution returns 0, the job returns a status of Succeeded. If it returns any other value, it returns a job status of Failed.

Specifying a Script

> **Note:**
>
> The following information applies to step Creating an OS Command Job in Specifying a Script .

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.

The job status depends on the exit code returned by the last command in the script. If the last command execution returns 0, the job returns a status of Succeeded. If it returns any other value, it returns a job status of Failed. You should implement proper exception handling in the script and return non-zero exit codes when appropriate. This will avoid situations in which the script failed, but the job reports the status as Succeeded.

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
  ```

You 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
```

Access Level Rules

> **Note:**
>
> The following rules apply to Task 6, "Specify Who Can Access the Job - (optional)".

- 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.

Creating a SQL Script Job

The basic process for creating a SQL script job is the same as described in Creating an OS Command Job. The following sections provide supplemental information specific to script jobs:

- Specifying Targets
- Specifying Options for the Parameters Page
- Specifying Host and Database Credentials
- Returning Error Codes from SQL Script Jobs

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.

Specifying 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.

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.

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;
```

Creating a Multi-task Job

The basic process for creating a multi-task job is the same as described in Creating an OS Command Job. The following sections provide supplemental information specific to multi-task jobs:

- Job Capabilities
- Specifying Targets for a Multi-task Job
- Adding Tasks to the Job

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.

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.

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.

Viewing and Analyzing Job Status

Viewing the Aggregate Status of All Jobs

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 15-2 shows the Jobs section at the bottom of the Enterprise Summary page.

Figure 15-2 Summary of Target Jobs on the Enterprise Summary Page

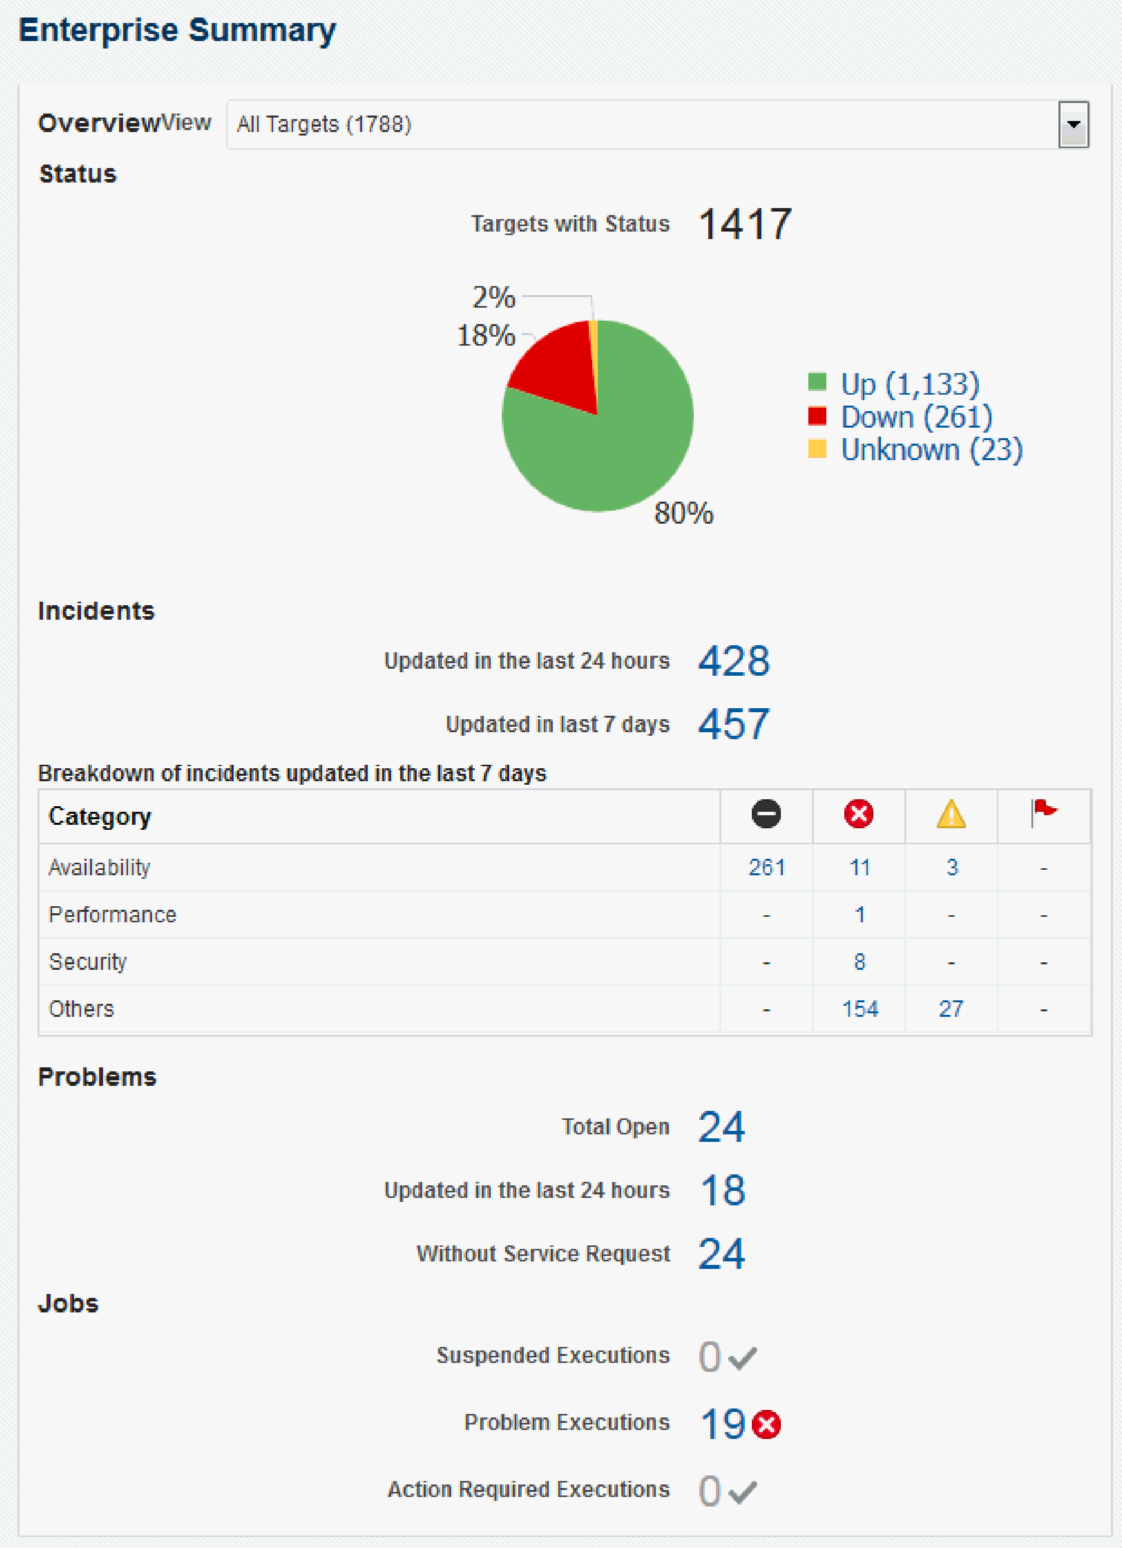

Enterprise Summary

Overview View: All Targets (1788)

Status

Targets with Status 1417

- 🟩 Up (1,133)
- 🟥 Down (261)
- 🟨 Unknown (23)

Incidents

Updated in the last 24 hours 428

Updated in last 7 days 457

Breakdown of incidents updated in the last 7 days

Category	⛔	✖	⚠	🚩
Availability	261	11	3	-
Performance	-	1	-	-
Security	-	8	-	-
Others	-	154	27	-

Problems

Total Open 24

Updated in the last 24 hours 18

Without Service Request 24

Jobs

Suspended Executions 0 ✓

Problem Executions 19 ✖

Action Required Executions 0 ✓

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.

Viewing the General Status of a Particular Job

To find out general status information for a particular job or jobs you have submitted, search for them in the Job Activity page, shown in Figure 15-1.

Viewing the Status of Job Executions

You can view detailed information about a single execution or multiple executions. A single execution can have a single step or multiple steps.

To view the status of executions:

1. From the Job Activity page, click the **Name** link for the job of interest. The Job Execution page displays.
2. In addition to displaying the job execution summary, you can drill down on specific tasks for further information.

Switching to Enhanced View

Beginning with Cloud Control version 12.1.0.4, you can optionally invoke a view of job runs that combines the views of several drill-downs on one page. To enable the enhanced view, execute the following command:

```
emctl set property -name oracle.sysman.core.jobs.ui.useAdfExecutionUi -value true
```

To revert to the standard view, execute the following command:

```
emctl set property -name oracle.sysman.core.jobs.ui.useAdfExecutionUi -value false
```

> **Note:**
>
> These commands do not require you to restart the OMS.

Viewing the Enhanced Status of a Job Run with a Single Execution

A single execution can have a single step or multiple steps. To view the status of a single execution:

Viewing the Enhanced Status of a Job Run with Multiple Executions

To view the status of multiple executions:

1. From the Job Activity page, click the **Name** link or **Status** link for a job containing multiple executions.

   The Job Execution page appears.
2. Click on an execution of interest in the left table.

   The details for the particular execution appears on the right side of the page.

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 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.

To access this page, from the Setup menu, select **Incidents**, then **Job Events**.

Figure 15-3   Job Event Generation Criteria Page

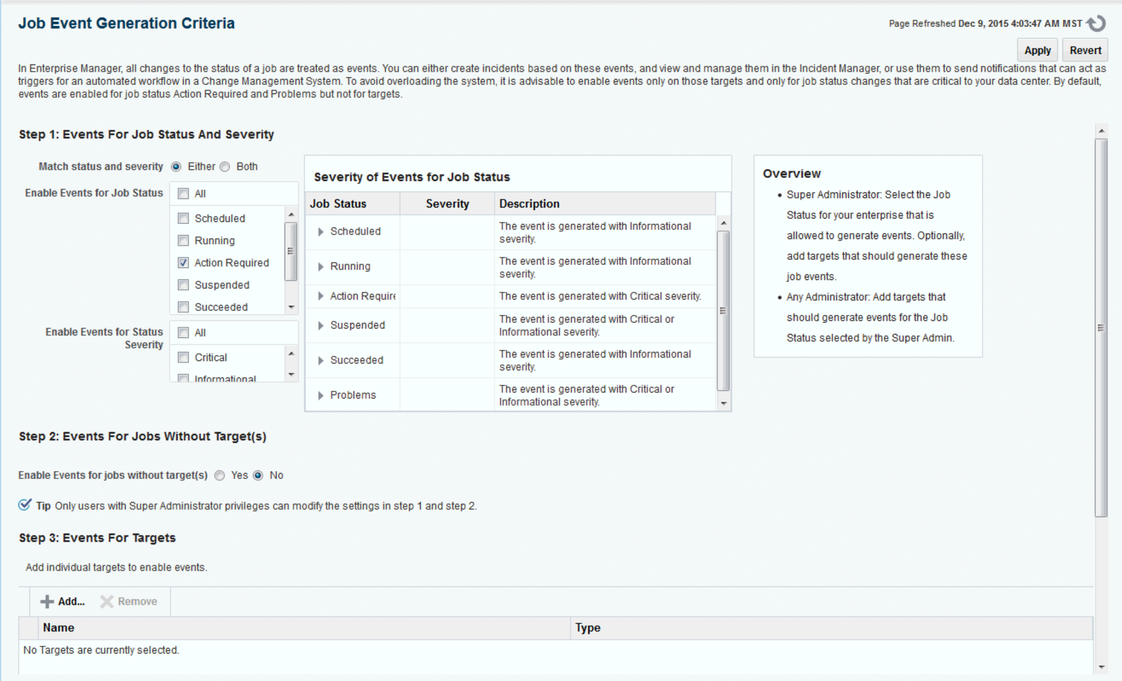

Job Event Generation Criteria

Page Refreshed Dec 9, 2015 4:03:47 AM MST

In Enterprise Manager, all changes to the status of a job are treated as events. You can either create incidents based on these events, and view and manage them in the Incident Manager, or use them to send notifications that can act as triggers for an automated workflow in a Change Management System. To avoid overloading the system, it is advisable to enable events only on those targets and only for job status changes that are critical to your data center. By default, events are enabled for job status Action Required and Problems but not for targets.

Step 1: Events For Job Status And Severity

Match status and severity ● Either ○ Both

Enable Events for Job Status
□ All
□ Scheduled
□ Running
☑ Action Required
□ Suspended
□ Succeeded

Enable Events for Status Severity
□ All
□ Critical
□ Informational

Severity of Events for Job Status

Job Status	Severity	Description
Scheduled		The event is generated with Informational severity.
Running		The event is generated with Informational severity.
Action Require		The event is generated with Critical severity.
Suspended		The event is generated with Critical or Informational severity.
Succeeded		The event is generated with Informational severity.
Problems		The event is generated with Critical or Informational severity.

Overview

- Super Administrator: Select the Job Status for your enterprise that is allowed to generate events. Optionally, add targets that should generate these job events.
- Any Administrator: Add targets that should generate events for the Job Status selected by the Super Admin.

Step 2: Events For Jobs Without Target(s)

Enable Events for jobs without target(s) ○ Yes ● No

✓ Tip Only users with Super Administrator privileges can modify the settings in step 1 and step 2.

Step 3: Events For Targets

Add individual targets to enable events.

Name	Type
No Targets are currently selected.	

Enabling Events For Job Status, Status Severity, 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 Enterprise Manager 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 "Enable Events for Job Status"region, select the statuses for which you want to publish events.
   b. In the "Enable Events for Status Severity" region, select whether you want to enable events for a critical status, informational status, for both.
   c. In the **"Enable Events for Jobs Without Target(s)"** section, select **Yes** if you want to create events for jobs that are not associated with any target.

   d. In the "Events for Targets" section, click **Add** to add targets for which you want the job events to be enabled.

6. Click **Apply**.

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 Enterprise Manager 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**.

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

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 Enterprise Manager as an administrator.

4. From the **Setup** menu, select **Incidents** and then **Incident Rules**. The Incident Rules Page appears.
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**.

   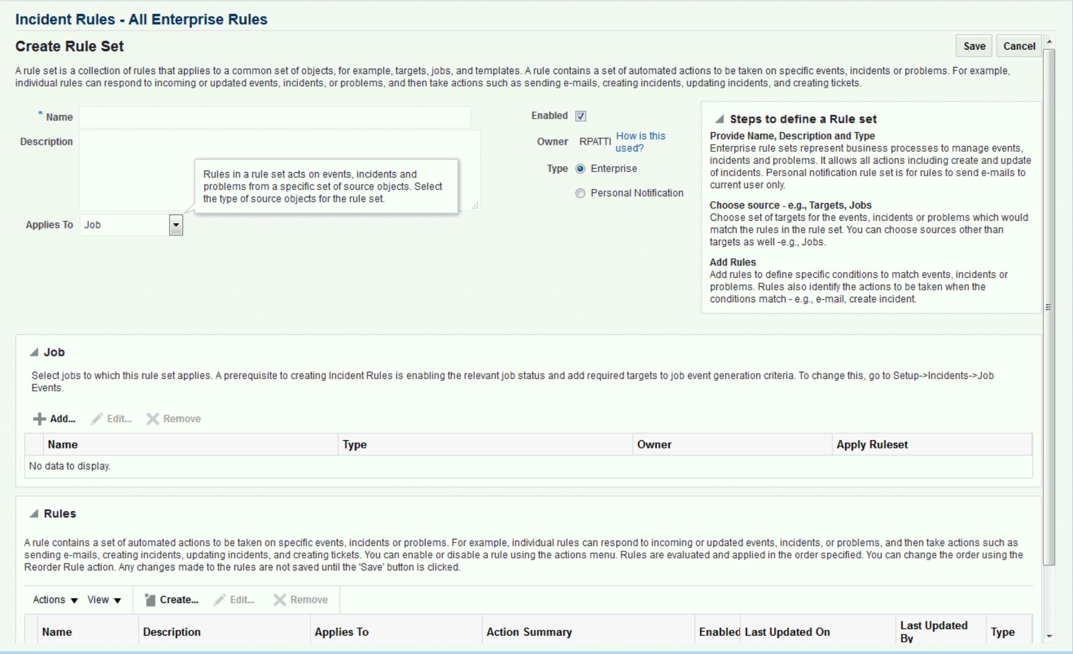

   Incident Rules - All Enterprise Rules

   Create Rule Set

   A rule set is a collection of rules that applies to a common set of objects, for example, targets, jobs, and templates. A rule contains a set of automated actions to be taken on specific events, incidents or problems. For example, individual rules can respond to incoming or updated events, incidents, or problems, and then take actions such as sending e-mails, creating incidents, updating incidents, and creating tickets.

   * Name

   Description

   Rules in a rule set acts on events, incidents and problems from a specific set of source objects. Select the type of source objects for the rule set.

   Applies To: Job

   Enabled: ☑

   Owner: RPATTI

   Type: ● Enterprise ○ Personal Notification

   Steps to define a Rule set

   Provide Name, Description and Type
   Enterprise rule sets represent business processes to manage events, incidents and problems. It allows all actions including create and update of incidents. Personal notification rule set is for rules to send e-mails to current user only.

   Choose source - e.g., Targets, Jobs
   Choose set of targets for the events, incidents or problems which would match the rules in the rule set. You can choose sources other than targets as well -e.g., Jobs.

   Add Rules
   Add rules to define specific conditions to match events, incidents or problems. Rules also identify the actions to be taken when the conditions match - e.g., e-mail, create incident.

   Job

   Select jobs to which this rule set applies. A prerequisite to creating Incident Rules is enabling the relevant job status and add required targets to job event generation criteria. To change this, go to Setup->Incidents->Job Events.

Name	Type	Owner	Apply Ruleset
No data to display.			

   Rules

   A rule contains a set of automated actions to be taken on specific events, incidents or problems. For example, individual rules can respond to incoming or updated events, incidents, or problems, and then take actions such as sending e-mails, creating incidents, updating incidents, and creating tickets. You can enable or disable a rule using the actions menu. Rules are evaluated and applied in the order specified. You can change the order using the Reorder Rule action. Any changes made to the rules are not saved until the 'Save' button is clicked.

Name	Description	Applies To	Action Summary	Enabled	Last Updated On	Last Updated By	Type

   In the Job tab, click **Add** to add jobs for which you want to create event rules.
7. In the Add Jobs dialog box, if you select the job **By Pattern**, provide **Job name like** and select the **Job Type**. Specify **Job owner like**. For the **Specific jobs** choice, select the job. Click **OK**.
8. In the **Rules** tab, click **Create**.

   In the Select Type of Rule to Create dialog box that appears, you can 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.
9. Select **Incoming events and updates to events**, and in the Create New Rule: Select Events page, do the following:

a. **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.

b. Set the other criteria for which you want to set the rule as displayed in the graphic below.

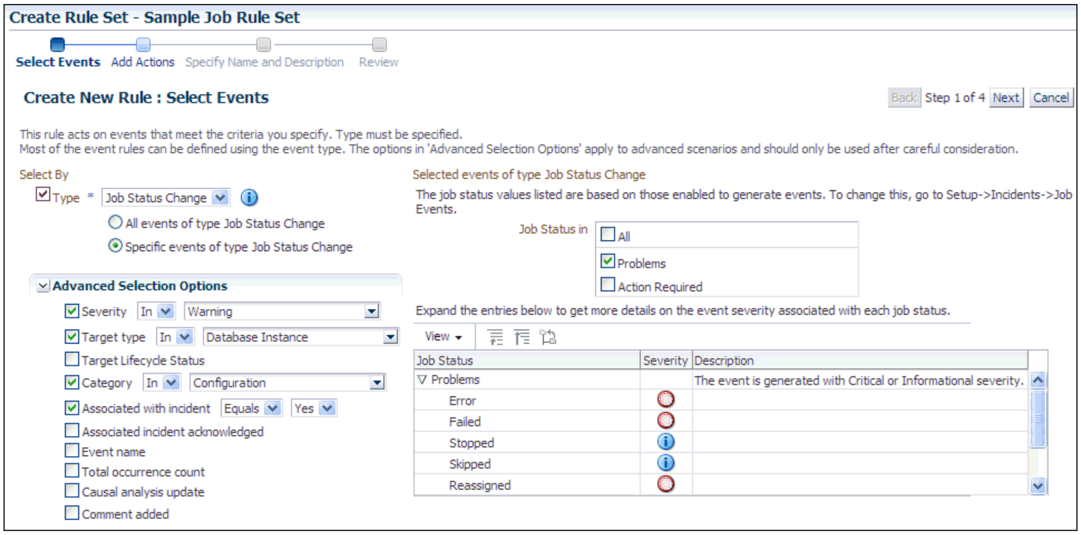

Create Rule Set - Sample Job Rule Set

Create New Rule : Select Events

This rule acts on events that meet the criteria you specify. Type must be specified.
Most of the event rules can be defined using the event type. The options in 'Advanced Selection Options' apply to advanced scenarios and should only be used after careful consideration.

Select By

☑ Type * Job Status Change

○ All events of type Job Status Change
● Specific events of type Job Status Change

Advanced Selection Options

☑ Severity In Warning
☑ Target type In Database Instance
☐ Target Lifecycle Status
☑ Category In Configuration
☑ Associated with incident Equals Yes
☐ Associated incident acknowledged
☐ Event name
☐ Total occurrence count
☐ Causal analysis update
☐ Comment added

Selected events of type Job Status Change

The job status values listed are based on those enabled to generate events. To change this, go to Setup->Incidents->Job Events.

Job Status in ☐ All
☑ Problems
☐ Action Required

Expand the entries below to get more details on the event severity associated with each job status.

Job Status	Severity	Description
▽ Problems		The event is generated with Critical or Informational severity.
Error	●	
Failed	●	
Stopped	ⓘ	
Skipped	ⓘ	
Reassigned	●	

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

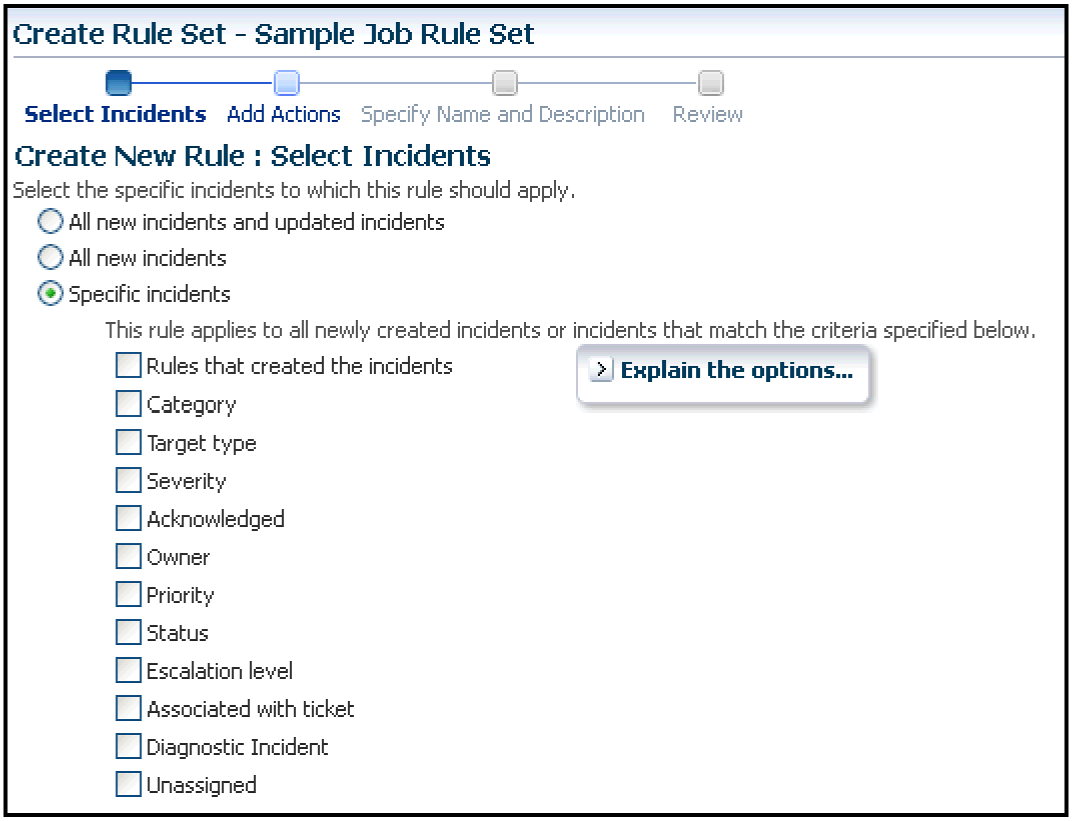

Create Rule Set - Sample Job Rule Set

Create New Rule : Select Incidents

Select the specific incidents to which this rule should apply.

○ All new incidents and updated incidents
○ All new incidents
● Specific incidents

This rule applies to all newly created incidents or incidents that match the criteria specified below.

□ Rules that created the incidents
□ Category
□ Target type
□ Severity
□ Acknowledged
□ Owner
□ Priority
□ Status
□ Escalation level
□ Associated with ticket
□ Diagnostic Incident
□ Unassigned

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 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.

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**.

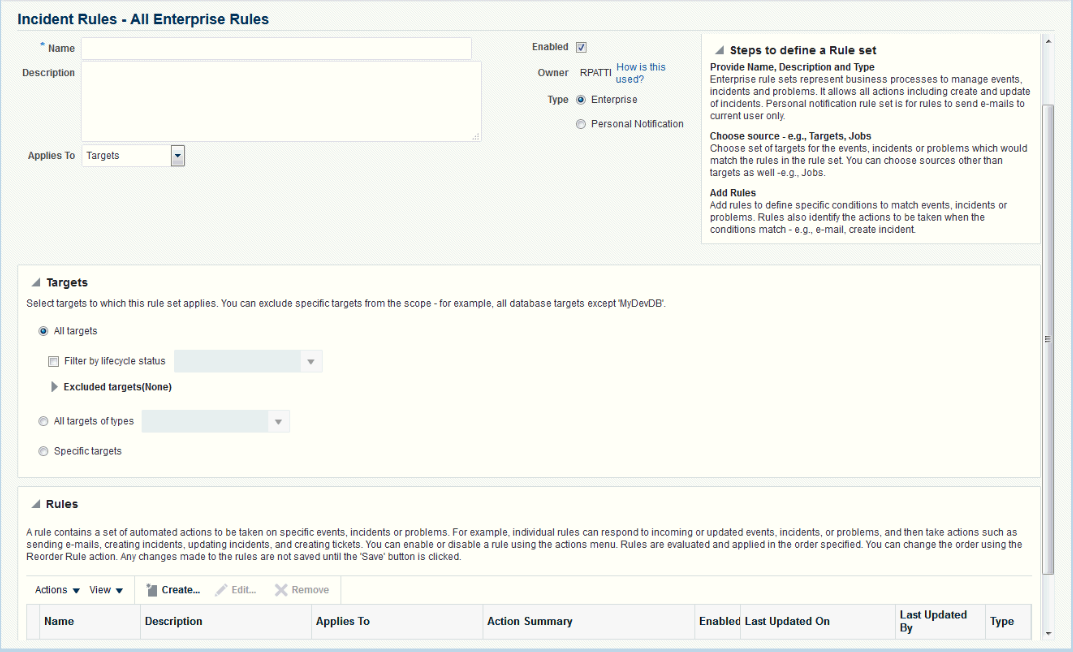

Incident Rules - All Enterprise Rules

* Name

Description

Applies To: Targets

Enabled ☑

Owner: RPATTI

Type: ● Enterprise ○ Personal Notification

Steps to define a Rule set

Provide Name, Description and Type
Enterprise rule sets represent business processes to manage events, incidents and problems. It allows all actions including create and update of incidents. Personal notification rule set is for rules to send e-mails to current user only.

Choose source - e.g., Targets, Jobs
Choose set of targets for the events, incidents or problems which would match the rules in the rule set. You can choose sources other than targets as well -e.g., Jobs.

Add Rules
Add rules to define specific conditions to match events, incidents or problems. Rules also identify the actions to be taken when the conditions match - e.g., e-mail, create incident.

Targets

Select targets to which this rule set applies. You can exclude specific targets from the scope - for example, all database targets except 'MyDevDB'.

● All targets

□ Filter by lifecycle status

Excluded targets(None)

○ All targets of types

○ Specific targets

Rules

A rule contains a set of automated actions to be taken on specific events, incidents or problems. For example, individual rules can respond to incoming or updated events, incidents, or problems, and then take actions such as sending e-mails, creating incidents, updating incidents, and creating tickets. You can enable or disable a rule using the actions menu. Rules are evaluated and applied in the order specified. You can change the order using the Reorder Rule action. Any changes made to the rules are not saved until the 'Save' button is clicked.

Name	Description	Applies To	Action Summary	Enabled	Last Updated On	Last Updated By	Type

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** area, 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.

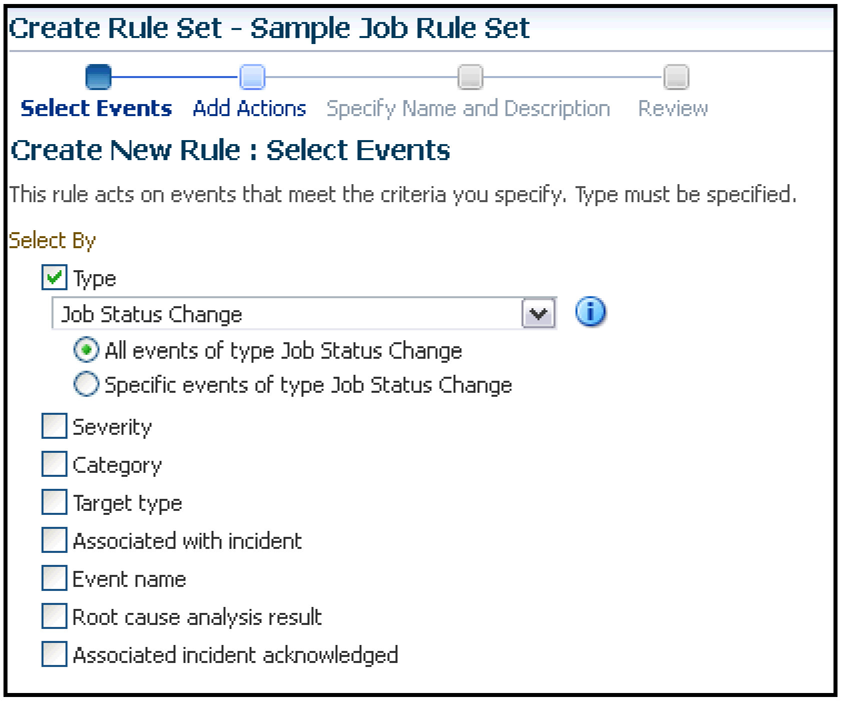

Create Rule Set - Sample Job Rule Set

Create New Rule : Select Events

This rule acts on events that meet the criteria you specify. Type must be specified.

Select By

☑ Type

Job Status Change

● All events of type Job Status Change

○ Specific events of type Job Status Change

□ Severity

□ Category

□ Target type

□ Associated with incident

□ Event name

□ Root cause analysis result

□ Associated incident acknowledged

- 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.

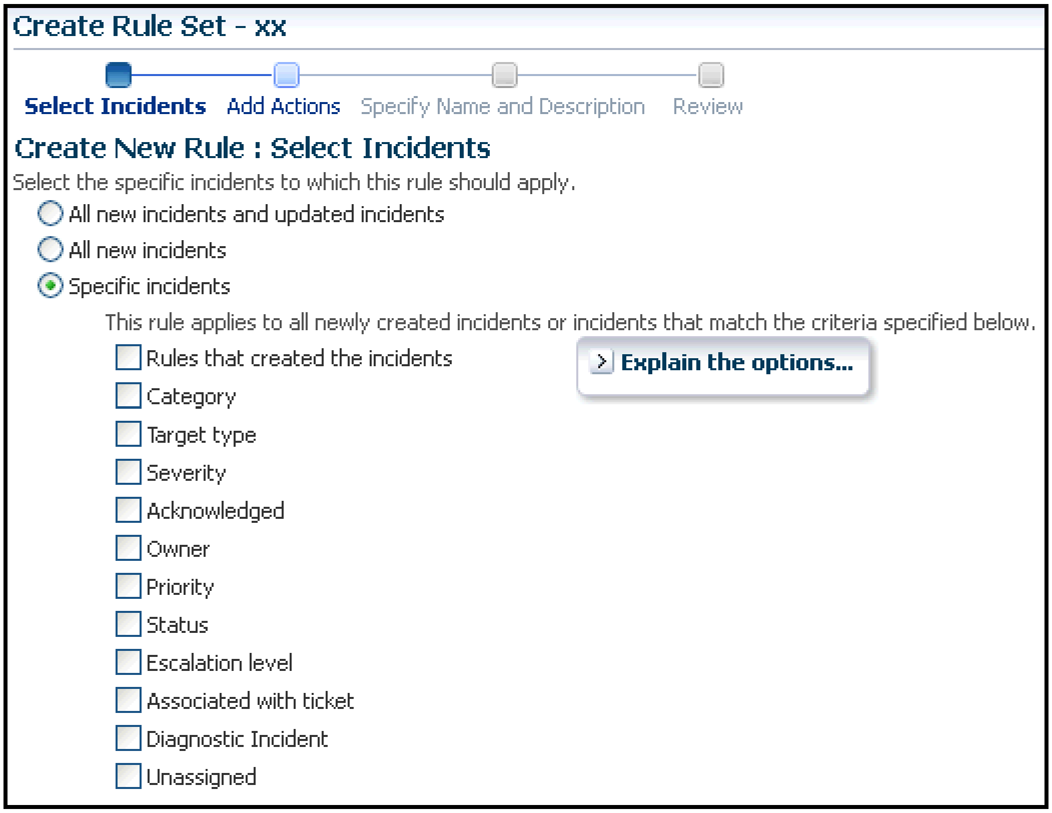

Create Rule Set - xx

Create New Rule : Select Incidents

Select the specific incidents to which this rule should apply.

○ All new incidents and updated incidents
○ All new incidents
● Specific incidents

This rule applies to all newly created incidents or incidents that match the criteria specified below.

□ Rules that created the incidents
□ Category
□ Target type
□ Severity
□ Acknowledged
□ Owner
□ Priority
□ Status
□ Escalation level
□ Associated with ticket
□ Diagnostic Incident
□ Unassigned

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.

Creating Corrective Actions

Corrective actions enable you to specify automated responses to metric alerts and events. 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 is 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.

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.

You define corrective actions for individual metrics for monitored targets. The following sections provide instructions on setting up corrective actions and viewing the details of a corrective action execution:

- Privilege and Access Requirements for Corrective Actions
- Creating Corrective Actions for Metrics
- Creating Corrective Action Libraries
- Which Credentials Will Be Used When a Corrective Action Runs
- Setting Up Notifications for Corrective Actions
- Providing Agent-side Response Actions
- Viewing the Details of a Corrective Action Execution

Privilege and Access Requirements for Corrective Actions

In order to create, edit, delete, or associate a Corrective Action with a specific entity (such as a target, monitoring template, or event rule), you must have the requisite privileges, as shown in the following table.

You want to:	Required Privileges
Create, edit, or delete a Corrective Action	User creating, editing, or delting the Corrective action must have the CREATE_CA privilege.
Associate a Corrective Action to a target/monitoring template.	User associating the Corrective Action with a target/monitoring template must have the CREATE_CA + Any privileges required by the target/monitoring template.
Apply a monitoring template (with an associated Corrective Action) to a target.	User applying the monitoring template must have CREATE_CA + Any privileges required by the template and target.
Associate a Corrective Action to an event rule.	Rule owner associating a Corrective Action to an event rule must have the CREATE_CA privilege + any privileges required by the Event Rule framework.

Note:

No additional privileges are required to view a Corrective Action.

Sharing Access to Corrective Actions

After you create a Corrective Action, you need to determine the access to corrective actions by other users. You do not need to provide input on the Corrective Action *Access* page if you do not want to share the corrective action.

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.

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.

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.myco.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 Creating Corrective Action Libraries.
- If you want to provide an Agent-side response action, select Agent Response Action as the task type. See 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.

Enabling re-evaluation of metric alerts after Corrective Action execution

Starting with Oracle Enterprise Manager 13c Release 5 Update 19 (13.5.0.19), Enterprise Manager can be enabled to automatically re-evaluate the metric alert after the Corrective Action has been executed, rather than waiting for next collection schedule.

To enable automatic re-evaluation of Metric Alerts, set OMS property for the following from `false` to `true`:

- `oracle.sysman.core.events.reevaluateAlertsOnCASuccess`
  Upon successful Corrective Action completion, this will reevaluate the value of the metric in which the alert was created on in order to clear the alert (assuming the metric is below the threshold).
- `oracle.sysman.core.events.reevaluateAlertsOnCAStop`
  When a Corrective Action is stopped, this will reevaluate the value of the metric in which the alert was created on in order to clear the alert (assuming the metric is below the threshold).

> **Note:**
>
> The OMS property can be set by using emcli set_oms_property, emctl set property, or directly from the Enterprise Manager console.

Creating Corrective Action Libraries

For corrective actions that you use repeatedly, you can define a library corrective actions. After a corrective action is in the library, you can reuse its 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 Creating Jobs for *General* and *Parameters*. 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 Sharing 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. At this point, the recently created Corrective Action is in a draft state, with the following options:
   - Create Like - allow the creation of another Corrective Action with exactly the same, editable configuration as the original one
   - Delete - removes the Corrective Action
   - Edit - allows to edit configuration of the Corrective Action, except the event type and the target type
   - Publish - Make Corrective Action available to be used with target metrics. Corrective Action cannot be used before publishing

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 Creating Corrective Actions for Events in Creating Corrective Actions for Metrics, for more information.

Specifying Preferred Credential Type for Corrective Actions

Preferred credentials are used to simplify access to managed targets by storing the login information for those targets in the Management Repository.

When creating a Corrective Action (CA) definition, you can specify which type of preferred credential should be used when running the CA based on the functional nature of the CA. There are two types of preferred credentials you can select when creating the CA definition:

- Normal
- Privileged

Important: Preferred credentials need to be global named credentials.

For more information about preferred credentials, see Preferred Credentials and Global Preferred Credentials.

Which Credentials Will Be Used When a Corrective Action Runs

In order to run a Corrective Action, it must be run by a user with the CREATE_CA privilege. The following table lists scenarios under which Corrective Actions are run and which user credentials are used when the Corrective Action is executed.

Scenario	Credentials Used
Corrective Action is run when directly associated with metric settings for a target.	The Corrective Action is executed using the privileges of the user who associates the Corrective Action to the target metric threshold.
Corrective Action is run when associated with a monitoring template applied to a target.	Corrective Action is executed using the privileges of the user who associates underlying template to the target.
Corrective Action is associated with a monitoring template that is part of an Administration Group or Template Collections Administration Group.	If the Corrective Action is part of a monitoring template/template collection that is automatically applied to an administration group, the preferred credentials of the user who associated the template with the administration group will be used for the corrective action.
Corrective Action is associated with an Event Rule. **Note**: This applies to target availability events, metric alerts, compliance events	Corrective Action is executed using the privileges of the Event Rule owner.

Using Kerberos Credentials with Corrective Actions

The *Add Space to Tablespace (Advanced)* corrective action can be used with Kerberos credentials. This corrective action is available beginning with Enterprise Manager 13c Release 5 Update 11.

> **Note:**
>
> The older *Add Space to Tablespace* corrective action cannot be used with Kerberos credentials.

Once Enterprise Manager 13c Release 5 Update 11 or later has been installed, you will need to run the following steps in order for *Add Space to Tablespace (Advanced)* to become available as a selectable option from the corrective action drop-down menu.

1. Run the following EMCTL command:

   ```
   emctl register oms metadata -service jobTypes -file dbAddSpaceTSKerb.xml-
   pluginId oracle.sysman.db -sysman_pwd <sysman password>
   ```

2. Bounce the OMS.
3. Update your tablespace metric collection settings to use the new *Add Space to Tablespace (Advanced)* corrective action.

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.

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 Creating Corrective Actions for Metrics.

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>
```

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.

Using Advanced Capabilities

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>
```

Viewing the Details of a Corrective Action Execution

There are two methods of displaying the outcome of a corrective action execution.

- Incident Manager method
  1. From the Enterprise Manager Enterprise Manager console Enterprise menu, select **Monitoring**, then **Incident Manager**.
  2. Click the **Search** icon, select **Events** from the Type drop-down, then click **Get Results**.
  3. Double-click the message of interest in the search results table.

     The Corrective Action History table now appears at the bottom of the page.
  4. Select the desired message in the history table, then click the glasses icon.

     The Corrective Action Execution page now appears, which displays the output of the corrective action, status, start time, end time, and so forth.
- All Metrics method
  1. From the target's home page, select **Monitoring**, then **All Metrics**.
  2. From the tree panel on the left, click the desired metric name.

     A row for the metric alert now appears in the Metric Alert History table.
  3. Click the glasses icon in the Details column as shown below.

     The Incident Manager Event Details page now appears.
  4. In the Corrective Action History table at the bottom of the page, select the message in the history table, then click the glasses icon.

     The Corrective Action Execution page now appears, which displays the output of the corrective action, status, start time, end time, and so forth.

Using Diagnostic Tools

The following sections provided procedures for these diagnostic topics:

- Enabling Job Logging
- Viewing and Downloading Job Logging
- Debugging a Failed Job
- Checking for Incidents Related to a Failed Job
- Packaging an Incident Generated by a Job Step
- Viewing Remote Log Files
- Diagnosing Problems with Enterprise Manager Management Tools

Enabling Job Logging

You can enable and disable object logging for diagnostic purposes. By default, only warning level and above are captured.

To enable job logging for a scheduled job:

1. From the Enterprise menu of the Enterprise Manager console, select **Job**, then **Activity**.
2. In the Top Activity table, click the **Name** link for the job you want to log.
3. In the Execution page that appears, click **Debug** from the Actions menu. Debug logging occurs for the selected job while the job is running.

   A confirmation message appears that states "Successfully enabled logging at DEBUG level."

After the job execution completes, the 'Debug' option under the 'Actions' menu is automatically disabled.

Viewing and Downloading Job Logging

If there is user-visible logging for a particular job, you can view the job execution log by doing the following:

1. In the Top Activity page, click the Name link of the job for which you want to view the log.
   The Job Execution page displays.
2. The job output log is displayed

You can also access job logging from the Execution page by clicking the link that appears in the Status column of the Job Run page, then clicking **Log Report** on the Execution page as shown below.

To view the log for a job step, do the following:

1. In the Top Activity page, click the link of the job for which you want to view the log.
2. In the Job Run table, click the link that appears in the Status column for the step you want to examine.
   The Output Log appears for the step.

To download all outputs of all executions to one file, click **Download Logs** from the **Summary** section at the top of the page.

Debugging a Failed Job

If an execution fails and you had not previously set debug as the logging level, you can choose to set the debug level when you retry the execution. For new job executions, you can set logging at the debug level in advance by clicking the Debug button. The Object Logging field indicates whether logging is enabled at the debug level.

Perform the following procedure if you encounter a job that fails.

1. View the job steps that failed.
2. Check the output for the failed step(s). Aggregated job output is displayed for all steps, and also for specific steps.
3. If the output does not contain the reason for the failure, view the logging output. You may also want to check for any incidents that have occurred while the job was running.
4. Determine the cause of the failure and fix the problem.
5. Enable debug mode, then resubmit the job.

   Note that the checkbox for Debug mode only appears on the confirmation page if the earlier execution was not in Debug mode. If the earlier execution was already in Debug mode, the retried execution is automatically in Debug mode.

Checking for Incidents Related to a Failed Job

It is possible for a job to fail because of an internal code error, a severe scaling issue, or other Enterprise Manager issue for which you may be able to investigate an incident or event trail. For example, if the OMS bounced because all Job Workers were stuck, this would cause many jobs to fail. If the loader were failing, that could also cause some jobs to fail.

1. Check for incidents or alerts in the time-frame of the job.
    - To check for incidents, select **Summary** from the Enterprise menu of the Enterprise Manager console, then view the Incidents section of the Enterprise Summary page.

      For more detailed information, select **Monitoring** from the Enterprise menu, then select **Support Workbench**.
    - To check for alerts,
2. Submit a service request with the related incident(s) or event data.

   All step output, error output, logging, remote log files, and incident dump files for a given job are captured for an incident.
    - To submit a service request, select **My Oracle Support** from the Enterprise menu of the console, then select **Service Requests**.
    - To create a technical SR, click **Create SR** on the Service Requests Home page. To create a contact us SR, click **Create "Contact Us" SR** at the top of the Contact Us Service Requests region, or click **Contact Us** at the top of any My Oracle Support page. If you are creating a technical service request, depending on the Support IDs registered in your profile, you can create hardware or software SRs, or both.

      The Create Service Request wizard guides you through the process of specifying product information and attaching configuration information to the SR when it is filed with Oracle Support. To ensure that Oracle Support has the most accurate target information, select the Configuration tab in the What is the Problem? section of Step 1: Problem, then select a target.
3. Apply a patch that support provides.
    - Select **Provisioning and Patching** from the Enterprise menu of the console, then select **Patches & Updates**.
    - Provide login credentials, then click **Go**.
    - Access the online help for assistance with this page.
4. Try to submit the job again after applying the patch.

To package an incident or manually trigger an incident:

1. Access Support Workbench.
2. Gather all job-related dumps and log files, as well as other data from the same time, and package it for Support.
3. Review the incident-related data in Support Workbench, searching for relevant errors.
4. If you determine the root cause without support intervention, fix the job and resubmit it.

Packaging an Incident Generated by a Job Step

Incidents (and the problems that contain them) are not packaged by default. You will need to package the problems of interest or concern in Support Workbench. You can choose whether to package all problems or only a portion thereof.

> **Note:**
>
> If a job with remote log files is involved in an incident, the remote files are automatically included in the incident as part of packaging. For more information on remote log files, see Viewing Remote Log Files.

To package an incident generated by a job step:

1. On the Log Report page, click the **Incident ID** link.
2. Click the **Problem Key** link on the Support Workbench Incident Details page.
3. In the Support Workbench Problem Details page, either click the **Package the Problem** link or the **Quick Package** button as shown below, then follow the instructions in the Quick Packaging wizard and online help.

Viewing Remote Log Files

Some jobs or Provisioning Adviser Framework (PAF) procedures run external commands, such as DBCA or the installer. These commands generate their own log files local to the system and Oracle home where they ran.

To view remote log files:

1. In the Top Activity page, click the link of the job for which you want to view the log.
2. In the Job Run table, click **Log Report**.
3. In the Log Report page, click the **Remote Log Files** link.
4. Specify host credentials, then click **OK**.

   The Remote File Viewer appears and displays the file contents.

Job System Data Warehouse

The job system Data Warehouse enables you to preserve historical job execution summary data. This data is stored in a view called **`MGMT$JOB_FINISHED_EXECS`**. For more information on the data displayed in the `MGMT$JOB_FINISHED_EXECS` view, see MGMT$JOB_FINISHED_EXECS.

Job System Data Warehouse option is available and is enabled by default starting with Oracle Enterprise Manager 13c Release 5 Release Update 15.

Configuration

Job System Data Warehouse feature is enabled by default, with a default 90 days retention policy.

Enable

Job System Data Warehouse can be enabled by executing the following:

```
begin
    EM_JOB_RECOVERY.enable_finished_exec_warehouse();
end;
```

Disable

Job System Data Warehouse can be disabled by executing the following:

```
begin
    EM_JOB_RECOVERY.enable_finished_exec_warehouse(p_enable => FALSE);
end;
```

Retention policy

Purge time for the Job System Data Warehouse data is set, to 90 days by default, but can be controlled with the command below:

```
begin
EM_JOB_RECOVERY.enable_finished_exec_warehouse(p_retention_days => 30);
end;
```

Diagnosing Problems with Enterprise Manager Management Tools

The Enterprise Manager management portion of the console provides several tools that can assist you in assessing the current state of the job system and determining a proper course of action for optimum performance. All of these tools are accessible from the Setup menu of the Enterprise Manager console.

The following sections provide information on each of the available tools.

Health Overview

1. From the Setup menu of the Enterprise Manager console, select **Manage the Manager**.
2. Select the **Health Overview** sub-menu.

The Job System Status region of the Health Overview page displays the following information:

- **Step Scheduler Status**

  The job step scheduler processes the job steps that are ready to run. If the status indicates that job step scheduler is running in warning or error mode, the job system is not functioning normally. In this case, the job system may run in fail-over mode, where the job dispatcher process may also run the task performed by the job step scheduler periodically. However, the job system may be running below its potential capacity, so resolving this situation would be beneficial.

  Several possible messages can appear:

  – DBMS_SCHEDULER job for step-scheduler not found

    This message is very rare and usually indicates a potentially serious issue. The job was likely removed inadvertently or due to some special processing (patch installation, for example, that requires recycling all DBMS_SCHEDULER jobs). No automatic resolution is possible here, and this would need to be addressed on a case by case basis.

- Failure in checking status

  This is a rare occurrence. The error message is usually shown. The error may disappear on its own as this error indicates that the status could not be calculated.

- DBMS_SCHEDULER is disabled

  All of the DBMS_SCHEDULER jobs are disabled in the environment. This should not occur unless a type of installation is in progress. Resolve this by starting DBMS_SCHEDULER processes.

- All job queue processes are in use

  The DBMS_SCHEDULER processes have been expended. Increase the parameter job_queue_processes in the repository RDBMS.

- All slave processes are in use

  The cause is similar to the above case. In this situation, you need to increase MAX_JOB_SLAVE_PROCESSES of the DBMS_SCHEDULER.

- All sessions are in use

  No RDBMS sessions were available for the DBMS_SCHEDULER. Increase the PROCESSES for the RDBMS.

- Reason for delay could not be established

  This usually appears because none of the above criteria were met, and is the most common warning. The dispatcher may just be overloaded because there is more work than available workers. Check the backlog in this case. The situation should resolve automatically, but if it persists, the number of workers available for the job system may be insufficient for the load the site experiences.

- **Job Backlog**

  The job backlog indicates the number of job steps that have passed their scheduled time but have not executed yet. If this number is high and has not decreased for a long period, the job system is not functioning normally. This situation usually arises if job engine resources are unable to meet the inflow of jobs from system or user activity.

  A high backlog can also happen because of the abnormal processing of specific jobs because they are stuck for extended periods. For more information on stuck job worker threads, do the following:

  1. From the OMS and Repository menu of the Health Overview page, select **Monitoring**, then **Diagnostic Metrics**.

  If the jobs system has a backlog for long periods of time, or if you would like to process the backlog faster, set the following parameters with the emctl set property command. These settings assume that sufficient database resources are available to support more load. These parameters are likely to be needed in a Large configuration with 2 OMS nodes.

Table 15-1 Large Job System Backlog Settings

Parameter	Value
oracle.sysman.core.jobs.shortPoolSize	50
oracle.sysman.core.jobs.longPoolSize	24
oracle.sysman.core.jobs.longSystemPoolSize	20

Table 15-1 (Cont.) Large Job System Backlog Settings

Parameter	Value
oracle.sysman.core.jobs.systemPoolSize	50
oracle.sysman.core.conn.maxConnForJobWorkers	144 This setting may require an increase of the processes setting in the database of 144 * number of OMS servers.

Repository Home Page

1. From the Setup menu of the Enterprise Manager console, select **Manage the Manager**.
2. Select the **Repository** sub-menu.

The Repository Scheduler Jobs Status table in the Management Services and Repository page displays the job system purge status and next run schedule.

Management Services and Repository: All Metrics

There are two navigation paths for accessing the All Metrics page:

- From the Setup menu
- From the Targets menu

Setup Menu Navigation

1. From the Setup menu of the Enterprise Manager console, select **Manage the Manager**.
2. Select the **Health Overview** sub-menu.
3. From the Management Services and Repository page that appears, select **Monitoring** from the OMS and Repository menu, then select **All Metrics**.
4. Scroll down to DBMS Job Status in the left pane, then select a metric.
5. Scroll down further and expand Repository Job Scheduler Performance.

   Definitions for the available metrics are as follows:

   - **Average number of steps marked as ready by the scheduler** — Average number of steps processed by the job step scheduler to mark the steps "ready" for execution. This number usually depends on the job system load over a time period.
   - **Estimated time for clearing current Job steps backlog (Mins)** — Estimated time to clear the backlog assuming the current inflow rate of the job system.
   - **Job step backlog** — Number of job steps that have passed their scheduled time but have not executed yet. If this number is high and has not decreased for a long period, the job system is not functioning normally. This situation usually arises if job engine resources are unable to meet the inflow of jobs from system or user activity.
   - **Latency in marking steps as ready by the scheduler** — The job step scheduler moves scheduled steps to ready queue. This metric indicates the average latency in marking the steps to ready queue. High latency means abnormal functioning of the job step scheduler process.

- **Overall job steps per second** — Average number of steps the job system executes per second.
- **Scheduler cycles** — Frequency of dbms scheduler process. Executes a minimum of 5 cycles per min, and may increase depending on the job system load. A low number usually indicates a problem in the job step scheduler process.

6. Scroll down further and expand the Usage Summary entries for Jobs, then select the metric for which you are interested.

Target Menu Navigation

1. From the Targets menu of the Enterprise Manager console, select **All Targets**.
2. In the left pane of the All Targets page, scroll down and expand **Internal**, then select **OMS and Repository**.
3. Click on the **OMS and Repository** table entry in the page that follows.
4. From the OMS and Repository menu in the Health Overview page that appears, select **Monitoring**, then **All Metrics**.

OMS and Repository: Diagnostic Metrics

1. From the Setup menu of the Enterprise Manager console, select **Manage the Manager**.
2. Select the **Health Overview** sub-menu.
3. From the Management Services and Repository page that appears, select **Monitoring** from the OMS and Repository menu, then select **Diagnostic Metrics**.

pbs_* metrics are relevant for diagnosing issues in the job system. This information is useful if you are searching for more information on stuck job system threads, or job threads usage statistics to determine outliers preventing other jobs from running.

OMS and Repository: Charts

1. From the Setup menu of the Enterprise Manager console, select **Manage the Manager**.
2. Select the **Health Overview** sub-menu.
3. From the Management Services and Repository page that appears, select **Monitoring** from the OMS and Repository menu, then select **Charts**.

Assuming that the job had steps, this page shows historical charts for the overall upload backlog, job step backlog, and overall job steps per second.

Management Servers and Job Activity Details Pages

1. From the Targets menu of the Enterprise Manager console, select **All Targets**.
2. On the left pane under Groups, Systems, and Services, click **Management Servers**.
3. Click **Management Servers** in the Target Name table.

   The Job System region displays a snapshot of job system status and details of processed executions. The Recent Job Executions Summary table displays the total user job executions that are expected to run within a specific time period, the completed count, running count, and the count of executions that are neither completed nor running. This helps you to determine if various user jobs are running as expected in the system.
4. Click the **More Details** link below the summary table.
5. Select the desired time frame in the drop-down for when executions are expected to start.

Jobs and their status, if any, appear in the table.

6. Select the **Job Dispatchers** tab.

   If more than one management server is configured, the page displays the job dispatcher and thread pool utilization information for each management servers.

   - **Dispatcher Utilization (%)** — Measures how frequently the job dispatcher picks up the job steps. High utilization indicates a heavy job system load.
   - **Throughput (steps dispatched/min)** — Indicates the average number of steps other than internal steps processed by dispatcher every minute.
   - **Thread Pool Utilization** — Displays the total number of threads configured for each pool, the average steps selected by the thread pool per minute, and the average number of available threads.

Job System Reports

The job system provides both a diagnostic report and usage report.

Diagnostic Report

1. From the Setup Enterprise of the Enterprise Manager console, select **Reports**.
2. Select the **Information Publisher Reports** sub-menu.
3. Search for **job** in the Title field, then click **Go**.
4. Click the **Job System Diagnostic Report** link in the table.

This report provides an overview of the job system's health and displays diagnostic information about executing jobs or jobs that are possibly delayed beyond their scheduled time. This information is usually relevant for an Oracle Support engineer diagnosing problems in the job system.

Usage Report

Follow the steps above to access this report, except click the **Job Usage Report** link in step 4.

This report provides an overview of the job system usage information over the past 7 days.

Job System Console

The Job System Console provides an in-depth administrators view into the Job System and its components.

To access Job System Console:

1. Log in to the Enterprise Manager console as a user with Super Administrator privileges.
2. From the Setup menu, select **Manage the Manager** and then **Job System Console**.

For more information about the Job System Console, see Diagnosing Job System Issues with the Job System Console.

Diagnosing Job System Issues with the Job System Console

The Job System Console gives you an in-depth view into the Job system. Using intuitive dashboards, the console provides an administrator view of the Job system to diagnose problems and resolve Job system performance issues.

Topics:

- Typical Job System Issues
- Job System Components
- Accessing the Job System Console
- Job System Console: Overview
- Job System Health

Typical Job System Issues

Below are some of the top issues that can affect Job System performance:

- Agent is Down, Unknown, or Suspended in Blackout
- Agent is overloaded resulting in excessive job retries (Metric Extensions can often cause this)
- Priority jobs are getting starved due to failing System Retry Jobs
- DB session hang due to repository background process deadlocks
- OMS UI console to PBS communication failure
- Corrective Actions trigger too frequently due to incorrect metric threshold settings
- User-suspended jobs are locking resources
- Long running jobs are blocking common Job System resources, thus preventing new jobs from running
- Jobs backlog due to stuck head of the queue

The job diagnostics dashboard allows administrators to easily identify the above issues, diagnose the root cause and take appropriate action.

Job System Components

The Enterprise Manager Job System is an OMS subsystem and includes a Job Scheduler and Job Workers. In turn, the Job Scheduler consists of two components: the Job Step Scheduler and the Job Dispatcher. In addition to user-submitted jobs, the majority of the background tasks in Enterprise Manager are run via a series of jobs. Typical tasks carried out by these jobs are calculating the availability of composite targets and sending notifications.

Performance of the Job System relies on numerous components to perform optimally. Job Diagnostics consolidates performance information pertaining to these components into intuitive dashboards for easy comparison and analysis. The primary components of the Job System are shown in the following illustration.

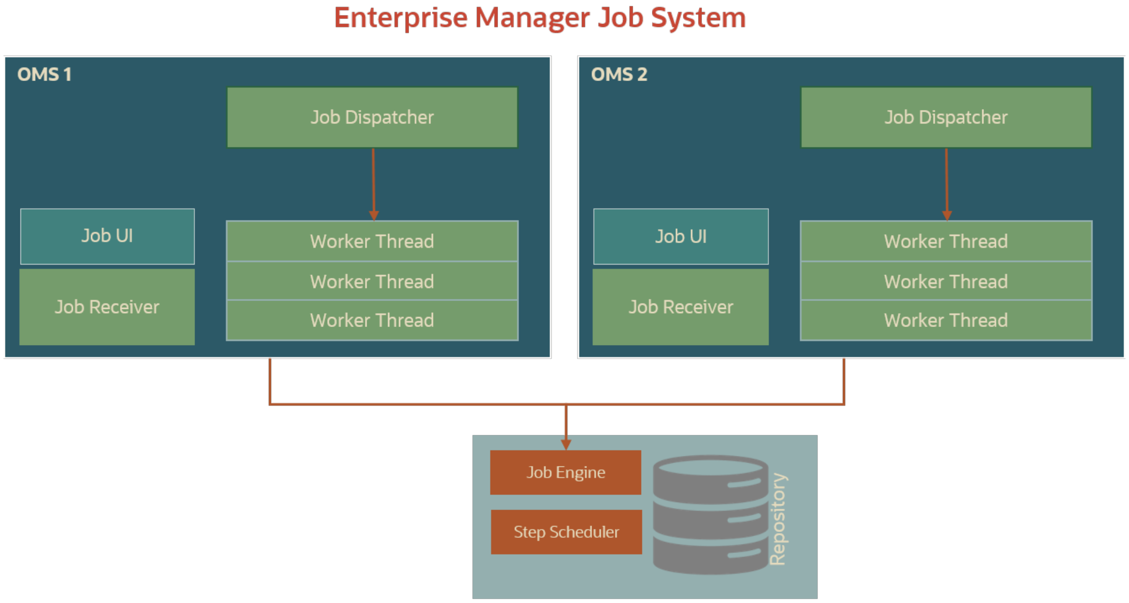

Job System Components Used by Job Diagnostics

- **Job Step Scheduler** – The Job Step Scheduler is a global component so there is only one per Enterprise Manager environment. It is scheduled to run by the DBMS Scheduler. The primary purpose of this component is to mark steps ready for the dispatcher to execute.
- **Job Dispatcher** - The Enterprise Manager Job System also has a notion of a *short jobs* (jobs that complete quickly) and *long jobs* (jobs that run a long time) and has separate worker pools in the OMS (not in the database as with the job workers) to handle those requests. The Job Dispatcher runs locally on each OMS and its purpose is to dispatch the jobs found by the Job Step Scheduler to the Job Workers. If the dispatcher cannot keep up with the work in the queue, the backlog increases. This is not a problem as long as the backlog is temporary. If it is not, then either the dispatcher is not able to keep up with the amount of work which could mean adding another OMS server or there is a problem with the Job Workers and they are not able to accept the work from the dispatcher.
- **Job Workers** – Job Workers take work for a given job step from the Job Dispatcher and process it. This can happen while holding a thread for steps that do processing in java, by contacting the repository for those that use SQL, or by contacting the agent for those that run remotely. If Job Workers are always busy and never free, then capacity needs to be added either via another OMS server or by increasing the number of Job Workers and potentially increasing the number of DB connections (each Job Worker takes a connection to the database).

Accessing the Job System Console

1. Log in to the Enterprise Manager console as a user with Super Administrator privileges.

> **Note:**
>
> You must have Super Administrator privileges in order to access the Job Diagnostics UI.

2. From the Setup menu, select **Manage the Manager** and then **Job System Console**. The *Job System Console* displays.

Job System Console: Overview

The Job System Console Overview, accessible only by super user job admins, is divided into two regions **Health** and **Performance**.

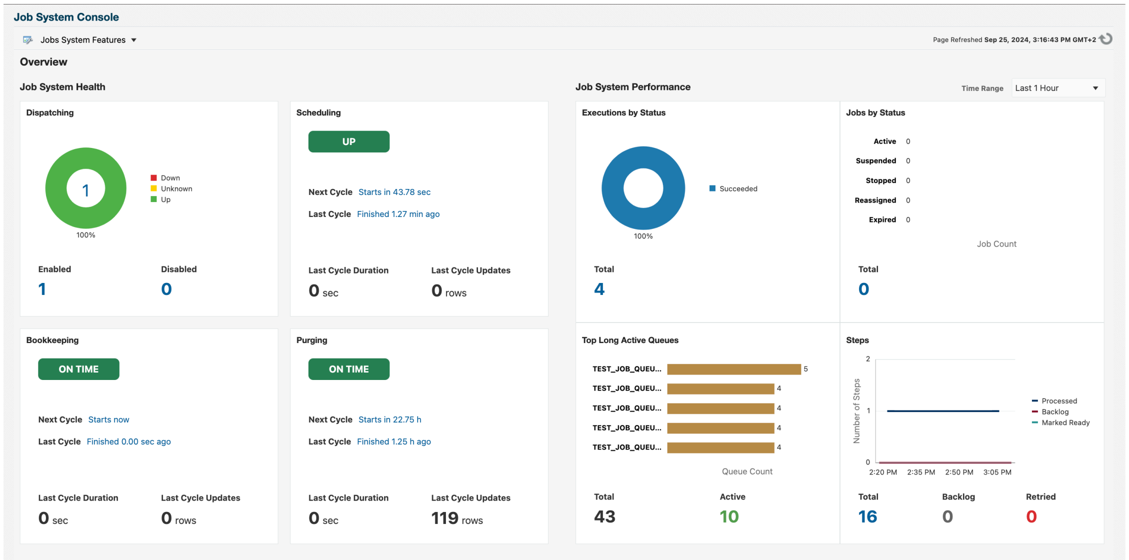

Job System Console

Jobs System Features ▼

Page Refreshed Sep 25, 2024, 3:16:43 PM GMT+2

Overview

Job System Health

Dispatching

Enabled 1

Disabled 0

Scheduling

UP

Next Cycle Starts in 43.78 sec

Last Cycle Finished 1.27 min ago

Last Cycle Duration 0 sec

Last Cycle Updates 0 rows

Bookkeeping

ON TIME

Next Cycle Starts now

Last Cycle Finished 0.00 sec ago

Last Cycle Duration 0 sec

Last Cycle Updates 0 rows

Purging

ON TIME

Next Cycle Starts in 22.75 h

Last Cycle Finished 1.25 h ago

Last Cycle Duration 0 sec

Last Cycle Updates 119 rows

Job System Performance

Time Range: Last 1 Hour ▼

Executions by Status

Total 4

Jobs by Status

Active 0

Suspended 0

Stopped 0

Reassigned 0

Expired 0

Total 0

Top Long Active Queues

Total 43

Active 10

Steps

Total 16

Backlog 0

Retried 0

- **Job System Health**: overall summary internal job system performance and status

  Dispatching - shows information such as status and configuration settings of Dispatchers running in the EM, and allows changing their configuration. Also shows information about usage of worker thread pools and connection pools used by Dispatchers.

  Scheduling - shows information about DBMS jobs responsible for marking steps as ready to be executed.

  Bookkeeping - shows information about the process called bookkeeping which is running frequently and is doing maintenance jobs.

  Purging - shows the information about the process responsible for removing jobs past retention policy.

- **Job System Performance**: overall summary of job system execution performance

  Executions - shows the information about All Job Executions in the EM.

  Jobs - shows the information about All Jobs in the EM.

  Queues - shows the information about All job queues in the EM.

  Steps - shows the information about All job execution steps in the EM.

Table 15-2 Example DBMS jojbs that use Queues

Job Name	Scheduler Job Name	Task
Agent Ping	EM_PING_MARK_NODE_STATUS	Keeps track of the health of the host targets in Enterprise Manager.

Table 15-2 (Cont.) Example DBMS jojbs that use Queues

Job Name	Scheduler Job Name	Task
Daily Maintenance	EM_DAILY_MAINTENANCE	This job does the daily repository maintenance tasks such as partition maintenance, stats updates, etc.
Repository Metrics	MGMT_COLLECTION.COLLECTION_SUBSYSTEM	This job shows the amount of work done for the repository metrics.
Rollup	EM_ROLLUP_SCHED_JOB	This job indicates the amount of data involved in the rollup job.

Job System Health

Dispatching

Job Dispatchers are services that handle dispatching the Job Steps for execution. You can view the current status of all dispatchers in your Enterprise Manager environment as well as change the configuration from the Dispatchers dashboard.

Accessing the Dispatchers dashboard:

1. From the Enterprise Manager Setup menu, select **Manage the Manager** and then **Job System Console**.
2. After clicking **Job System Health** from the drop menu, select the **Dispatching** tab, or by clicking the **Dispatching** module in the **Job System Health** section of the **Overview** page. The Dispatching dashboard displays showing all dispatchers and their current service status, with four metrics at the top indicating status, count, type and maintenance mode of the dispatchers. The metrics can be used to quickly filter the dispatchers.
3. Select a dispatcher from the list.
4. Click **Configuration** to display all configurable parameters for the selected dispatcher.
5. Modify the configuration settings as appropriate.

Configuring Dispatchers

From the Configuration tab, you can modify a dispatcher's configuration parameters. See Sizing Your Enterprise Manager Deployment for more information about optimizing Enterprise Manager configuration.

General

You can switch *Maintenance Mode* for the dispatcher on or off. When in *Maintenance Mode* all steps which are not system steps will fail. System steps are processed normally by the dispatcher.

Intervals

Sleep Interval

- **Sleep Interval Increment**: Maximum number of seconds to sleep between each dispatcher sweep when dispatcher is unable to pick up any steps.
- **Dispatcher Cycle Interval** (minimum): Minimum number of seconds to sleep between each dispatcher sweep (value infractions)

Dispatcher Cycle Interval (maximum): Maximum number of seconds to sleep between each dispatcher sweep when dispatcher is unable to pick up any steps.

Steps Scheduler

- **Health Check Interval**: Time in seconds to check for health of DBMS_SCHEDULER processes (step_dispatcher)
- **Step Scheduler Interval**: The interval (in secs) with which the dispatcher should attempt to run the step scheduler. The step scheduler interval applies when the DBMS_SCHEDULER processes are not running or the dispatcher also needs to handle step scheduling.

Thread Pool

Thread pools provide a way to scope the resources used by the Job System. For example, the user short pool defaults to 25 threads. This allows each OMS to run up to 25 different user steps marked *short running* concurrently.

Job steps can be categorized into 5 broad categories:

- **User Short** -- End user (short running)
- **User Long** -- End user (long running)
- **System Normal**—Steps run by system jobs
- **System Critical**—Steps run by system jobs
- **Internal** -- Steps created by the Job System for performing low-level actions like step time outs, grace period timeouts and bookkeeping steps.

Connection Pool (maximum number of connections allowed for the Dispatcher)

There are three categories of connections:

- **Job Worker**--Worker threads of the Job System that execute particular steps.
- **Job Receiver**—Pool of threads to accept asynchronous status and updates from the agent.
- **Job Dispatcher**—Takes care of the dispatching the steps to various workers for execution.

If the Job Worker percent usage is high, then it means the dispatcher cannot dispatch to all the workers in a timely fashion. In this situation, there could be a resource problem, and the environment could probably benefit from more worker threads. However, do not go beyond doubling the size of the threads. If doubling the number of threads does not seem high enough, contact Oracle as it might be better to add an additional OMS.

Viewing High-Level Details

Click on the *Details* tab to view a list of all dispatchers. Selecting a dispatcher from the list graphically displays the total number of steps executed as well as the number of connections made at a specific point in time.

Scheduling

The **Scheduling** tab shows four information panels at the top:

- **Cycle Count By Status**: shows the number of running, completed ans stuck cycles.
- **Last Cycle**: shows when the last cycle was finished and its status.
- **Cycle Spikes**: shows the maximum duration of spikes in milliseconds and the maximum number of row updates.

- **Next Cycle**: when the next scheduled cycle will start and its status.

Bookkeeping

The **Bookkeeping** tab shows four information panels at the top:

- **Cycle Count By Status**: shows the number of running, completed ans stuck cycles.
- **Last Cycle**: shows when the last cycle was finished and its status.
- **Cycle Spikes**: shows the maximum duration of spikes in milliseconds and the maximum number of row updates.
- **Next Cycle**: when the next scheduled cycle will start and its status.

Purging

The **Purging** tab shows four information panels at the top:

- **Cycle Count By Status**: shows the number of running, completed ans stuck cycles.
- **Last Cycle**: shows when the last cycle was finished and its status.
- **Cycle Spikes**: shows the maximum duration of spikes in milliseconds and the maximum number of row updates.
- **Next Cycle**: when the next scheduled cycle will start and its status.

Job System Performance

Queues

The **Queues** tab shows two information panels at the top:

- **Queue Count**: shows the total number of queues, and the number of active queues.
- **Top 5 Long Queues**: shows the top five queues by number of items. Clicking on a queue will highlight it in the **Activity** table and its details will be shown below.

The **Activity** table lists the available queues, and by clicking a queue, its details will be shown below.To find a specific queue, use the **Search Loaded Records** search in the top-right corner of the table.

Jobs

The **Jobs** tab displays the jobs with the creation date in the time range selected in the top-right corner of the tab. The two information panels at the top:

- **Job Count**: shows the number of total, user and system jobs. Clicking on a counter will filter the jobs listed in the **Activity** table.
- **Jobs by Status**: shows a pie chart displaying the percentage number of jobs with status Active, Expired, Reassigned, Stopped or Suspended. Clicking on a status will filter the jobs listed in the **Activity** table.

The **Activity** table can display a list of queues filtered by **Job Status** and/or **Job Type**, and by clicking a job, its details will be shown below .To find a specific job, use the **Search Job Name** search in the top-right corner of the table.

Executions

The **Executions** tab displays the executions with the creation date in the time range selected in the top-right corner of the tab. The two information panels at the top:

- **Number of Executions**: shows the number of Total, Active, Succeeded and with Problems executions. Clicking on a counter will filter the executions listed in the **Activity** table.
- **Executions by Status**: shows a pie chart displaying the percentage number of executions with status Completed, Aborted, Failed. Clicking on a status will filter the jobs listed in the **Activity** table.

  Possible Statuses:

  Scheduled, Running, Error, Failed, Succeeded, Suspended by User, Suspended: Agent is not Ready, Stopped, Suspended on Lock, Suspended on Event, Suspended on Blackout, Stop Pending, Suspend Pending, Inactive, Queued, Waiting, Skipped, Delete Pending, Reassigned, Credentials Missing, Action Required, Running Time Limit Exceeded, Processing Stopped, Non Stoppable, Suspended: Target is not Ready, Ignored, Corrective Action marked Broken.

The **Activity** table can display a list of executions filtered by **Executions Status** and/or **Job Type**, and by clicking a job, its details will be shown below .To find a specific exectution, use the **Search Job Name** search box in the top-right corner of the table.

To see details of an execution click the link from the **Job Name** column, which will open the page with the details of the execution.

Steps

The **Steps** tab displays the steps with the creation date in the time range selected in the top-right corner of the tab. The two information panels at the top:

- **Step Count**: shows the number of Total, Backlog, Active and Retried steps.
- **Executions by Status**: shows a pie chart displaying the percentage number of steps with status Completed, Aborted, Failed. Clicking on a status will filter the jobs listed in the **Activity** table.
- **Statistics**: shows the number of Backlog, Processed and Marked Ready steps over time.

The **Activity** table can display a list of steps filtered by **Step Status** and/or **Job Type**, and by clicking a step, its details will be shown below .To find a specific step, use the **Search Job Name** search box in the top-right corner of the table.

16
Monitoring Access Points Configured for a Target

Access points are channels of communication using which a software or a hardware deployment can be monitored in Enterprise Manager. They are upload points through which data is collected and uploaded to Oracle Management Service (OMS). While most targets have a single access point, some targets have multiple access points.

This chapter introduces you to the concept of access points, and describes how you can monitor the access points that are configured for a multi-access point target (for example, System Infrastructure Server, System Infrastructure Cisco Switch, and so on). In particular, this chapter covers the following:

Introduction to Monitoring Access Points

Access points are channels of communication using which a software or a hardware deployment can be monitored in Enterprise Manager Cloud Control. They are upload points through which metric data is collected and uploaded to Oracle Management Service.

Typically, to discover a software deployment and monitor it as a target in Enterprise Manager, you must first install an Oracle Management Agent (Management Agent) on the host where the software is deployed. In this case, the Management Agent collects data about the software deployment and acts as the only channel of communication with the software deployment on that host.

However, for some targets (for example, System Infrastructure Server, System Infrastructure Cisco Switch, System Infrastructure Oracle InfiniBand Switch, and System Infrastructure Rack, and so on) you can have multiple access points to collect metric data about the target. The access points can be a chip that is integrated with the hardware, a plug-in that is deployed onto the hardware, and so on.

Different access point types are used to distinguish between the different management interfaces that are used to interact with a given target type, and each access point type typically has its own separate implementation of metrics and actions for the target.

In Figure 16-1, you can see how each of the access points implement a subset of the metrics for a given multi-access point target. The access points with metric names in red indicate that they are the best implementers of those metrics. The access points with metric names in blue indicate that they are the second best implementers of those metrics.

Figure 16-1 Implementation of Access Points

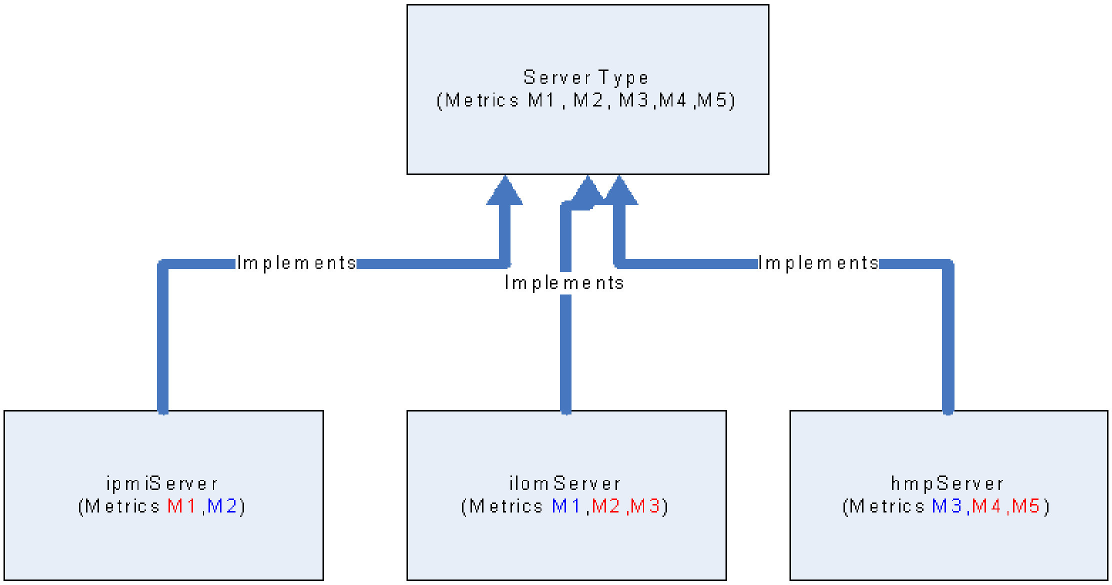

Enterprise Manager enables you to view a list of such access points that are configured for a multi-access point target. In addition, you can view the historical status of access points, the capability maps of access points, and the best implementers for collecting certain types of metric data.

Note:

When a multi-access point target is blacked out, either explicitly or when a group containing such targets is blacked out, the multi-access point target and its access points are blacked out and they go into the *Blackout* status in the Enterprise Manager Console.

On the other hand, if a multi-access point target is not blacked out but all its access points are blacked out by virtue of full host blackouts on all the Management Agents that are monitoring the access points, then the multi-access point target goes into the *Pending* status because none of the access points are available to collect metric data for that target. This is particularly in the case of multi-access point targets that have only one access point. If only one access point is blacked out and all other access points are still available, then the multi-access point target continues to appear as up and running because it has other available access points that can collect metric data.

Viewing a List of Access Points Configured for a Target

Enterprise Manager enables you to view a list of access points that are configured for a multi-access point target. Access points are channels of communication using which a software or a hardware deployment can be monitored in Enterprise Manager. They are upload points through which data is collected and uploaded to Oracle Management Service.

To view a list of access points configured for a multi-access point target, follow these steps:

1. From the **Targets** menu, select **All Targets.**
2. On the All Targets page, select a multi-access point target (for example, System Infrastructure Server, System Infrastructure Switch, and so on.)
3. On the Home page, from the target-specific menu, select **Monitoring,** then select **Access Points - Overview.**
4. On the Access Points - Overview page, you can view a list of access points configured for the selected multi-access point target.

Deleting Access Points Configured for a Target

Enterprise Manager enables you to delete the access points that are not required for monitoring. Access points are channels of communication using which a software or a hardware deployment can be monitored in Enterprise Manager. They are upload points through which data is collected and uploaded to Oracle Management Service.

To delete the access points configured for a multi-access point target, follow these steps:

1. From the **Targets** menu, select **All Targets.**
2. On the All Targets page, select a multi-access point target (for example, System Infrastructure Server, System Infrastructure Switch, and so on.)
3. On the Home page, from the target-specific menu, select **Monitoring,** then select **Access Points - Overview.**

   > **WARNING:**
   >
   > Removing the last access point will remove the multi-access point target from Enterprise Manager.

4. On the Access Points - Overview page, select the access point you want to delete, and click **Remove.**

Viewing the Capability Metric Map for a Target

Enterprise Manager enables you to view the capability map of the access points that are configured for a multi-access point target. Access points are channels of communication using which a software or a hardware deployment can be monitored in Enterprise Manager. They are upload points through which data is collected and uploaded to Oracle Management Service.

An access point can have the capability to either collect metric data or take action against the target. Enterprise Manager provides a tabulated view of the capabilities of the access points. You can view this capability map and determine what each access point is efficient at.

To view the capability metric map for a multi-access point target, follow these steps:

1. From the **Targets** menu, select **All Targets.**
2. On the All Targets page, select a multi-access point target (for example, System Infrastructure Server, System Infrastructure Switch, and so on.)

3. On the Home page, from the target-specific menu, select **Monitoring,** then select **Access Points - Current Capability Map.**
4. On the Access Points - Current Capability Map page, view a list of access points and their capabilities. See Table 16-1.

   You can also identify which access point is the best implementer for collecting the required metric data by selecting **Show only best implementers.**

   If you want to filter the table based on the capability type and the capability name, use the **Capability Type** and **Capability Name** drop-down lists, and click **Search.**

   Table 16-1 describes the columns in the Current Capability Map table.

Table 16-1 Description of the Columns in the Capability Map Table

Column Name	Description
Access Point	Name of the access point configured for the multi-access point target.
Priority	Measure of the capability of the access point to collect metric data. The access point with the highest priority is the first one to collect and upload metric data. If and when this access point becomes unavailable, the access point with the next highest priority takes over and collects the metric data.
Capability Type	Whether the access point is capable of collecting metric data or taking action against the target. The possible values are Action and Collection.
Capability Name	Name of the metric collection capability or the action capability that the access point has.
Is Current Best Implementer	Whether or not the access point is currently the best implementer. The possible values are Yes or No. An access point can have the capability to either collect metric data or take action against the target, and sometimes, two or more access points can have overlapping set of capabilities. Access points are termed best implementers when they are best suited for a certain metric or action.
Status	Stratus of the access point and whether the collection capability is enabled or disabled for that access point.

Viewing the Best Access Point Implementers (and their History) for Various Operations Supported for a Target

Enterprise Manager not only provides a tabulated view of the capabilities of the access points but also highlights the ones that are considered best implementers or best suited for a certain metric or action.

Access points are channels of communication using which a software or a hardware deployment can be monitored in Enterprise Manager. They are upload points through which data is collected and uploaded to Oracle Management Service. An access point can have the capability to either collect metric data or take action against the target, and sometimes, two or more access points can have overlapping set of capabilities. Access points are termed best implementers when they are best suited for a certain metric or action.

The best implementer is determined when an event that changes the capability map for the target occurs. Events that affect the capability map for the target include agent going down, access points being down or broken, capability metric for an access point reporting a different priority, and capability metric reporting a new or removed capabilities. Thus, when an access point defined as the best implementer is unavailable to collect the metric, Enterprise Manager automatically switches over the next best implementer to collect the same metric data.

To view the best access point implementers for a multi-access point target, follow these steps:

1. From the **Targets** menu, select **All Targets.**
2. On the All Targets page, select a multi-access point target (for example, System Infrastructure Server, System Infrastructure Switch, and so on.)
3. On the Home page, from the target-specific menu, select **Monitoring,** then select **Access Points - Current Capability Map.**
4. On the Access Points - Current Capability Map page, view a list of access points and their capabilities. Select **Show Only the Best Implementers.** To further filter the list and view the history of the best implementers, select the start date and end date, and click **Search.**

Modifying or Reconfiguring the Monitoring Properties of the Access Points Configured for a Target

Enterprise Manager enables you to view a list of access points that are configured for a multi-access point target. Access points are channels of communication using which a software or a hardware deployment can be monitored in Enterprise Manager. They are upload points through which data is collected and uploaded to Oracle Management Service.

While discovering a multi-access point target, the access points configured for that target are automatically added to Enterprise Manager for monitoring. These access points are added with default monitoring configuration properties. However, after they are added to Enterprise Manager, you can choose to modify or reconfigure their monitoring properties if you want.

To modify or reconfigure the monitoring properties of the access points configured for a multi-access point target, follow these steps:

1. From the **Targets** menu, select **All Targets.**
2. On the All Targets page, select a multi-access point target (for example, System Infrastructure Server, System Infrastructure Switch, and so on.)
3. On the Home page, do one of the following:
   - From the target-specific menu, select **Monitoring,** then select **Access Points - Overview.** On the Access Points - Overview page, select the access point whose monitoring properties you want to change, and click **Configure.**
   - From the target-specific menu, select **Target Setup,** then **Monitoring Configuration.**
4. On the Monitoring Configuration page, in the **Access Point** drop-down list, ensure that the correct access point is selected. If not, select the correct access point.
5. In the Monitoring Properties section, edit the properties, and click **Save.**

EM CLI Verbs for Managing the Access Points Configured for a Target

Table 16-2 lists the EM CLI verbs for managing the access points configured for a target. For more information about these verbs, see the *Oracle Enterprise Manager Command Line Interface Guide*.

Table 16-2   EM CLI Verbs for Managing the Access Points Configured for a Target

EM CLI Verb	Description
`add_target`	Adds a multi-access point target to the Enterprise Manager system.
`modify_target`	Modifies a multi-access point target that is already existing in the Enterprise Manager system.
`delete_target`	Deletes a multi-access point target.
`get_accesspoints`	Lists the access points configured for a target.
`get_best_implementer`	Lists the best access point implementers for various operations supported for a multi-access point target.
`set credentials`	Modifies or reconfigures the monitoring properties of the access points configured for a multi-access point target.

17 Always-On Monitoring

The Enterprise Manager Always-On Monitoring application provides the ability to monitor critical target status and metric alerts when the Oracle Management Service is unavailable. The service continuously monitors critical targets through the Enterprise Manager Agent and can be easily configured to send email notifications for these events to administrators.

Always-On Monitoring can be configured to send notifications at any time, but is particularly useful when experiencing downtime of your central Enterprise Manager site for maintenance operations such as upgrade and patching.

The Always-On Monitoring is synchronized with Enterprise Manager to reuse the configuration of monitored targets as well as requisite notification data such as notification contacts and email gateway configuration. Once properly configured and synchronized, the service will receive alerts from Enterprise Manager Agents and send email notifications to the appropriate administrators.

This chapter covers the following topics:

- Prerequisites
- Best Practices
- Installing Always-On Monitoring
- Configuring Always-On Monitoring
- Controlling the Service
- Updating Always-On Monitoring
- Data Maintenance
- Controlling Always-On Monitoring Configuration Settings
- Getting Performance Information
- Modifiable Always-On Monitoring Properties
- Creating an SSO Wallet and JKS for CA Certificates
- Diagnosing Problems
- High Availability and Disaster Recovery
- Uninstalling Always-On Monitoring
- Configuring the Always-On Monitoring Application for Secure Communication Using the TLSv1.2 Protocol

Prerequisites

Before installing Always-On Monitoring, ensure that the following prerequisite tasks have been performed:

Environment Setup

- Install or identify an Oracle Database instance to hold the Always-On Monitoring repository.
- Enterprise Manager must be upgraded to 13.4. In addition, the Management Agent should also be upgraded to 13.4.
- If you do not have database administrator access, you need to ask the DB administrator to add you to the Always-On Monitoring repository. You can use the script found in Granting Required Privileges to the Always-On Monitoring Schema Owner.

  Save the EM key to the Enterprise Manager repository. See Saving the Em Key for more information.

Java Requirements

- JDK8 must be installed in the environment where the user will run Always-On Monitoring (including *emsca*). The JAVA_HOME environment variable must point to the JDK8 location.

System Resource Requirements

- The number of open file descriptors used by Always-On Monitoring is *http.queueSize* + *http.maxThreads*. Always-On Monitoring also uses the database connection pool. The number of open database connections ranges between the number of active threads up to a maximum of 100 connections.
- Maximum memory requirement is estimated at 10 Megabytes per thread. Memory demand peaks when Always-On Monitoring is first started because the Agents must communicate the complete status of all their monitored targets.

Installing the Always-On Monitoring Repository Database

You must install and configure an Oracle database instance to house the Always-On Monitoring repository. Because you must create the schema for the repository separately, it is critical to have an idea of the space that needs to be allocated to the schema, and how the space is broken down.

Because the purpose of Always-On Monitoring is to continuously monitor and send notifications when Enterprise Manager is down, the Always-On Monitoring Repository database should be installed as a separate instance of Oracle with the Always-On Monitoring Schema on another machine so that Always-On Monitoring continues to run when the Enterprise Manager Repository is either being upgraded or unexpectedly goes down.

Like the Enterprise Manager Repository, the Always-On Monitoring Repository needs to be installed on an Oracle database. The database version supported depends on the Always-On Monitoring version as shown in the following table:

Always-On Monitoring Version	Oracle Database Version
All Versions	Oracle Database 12.1.0.2.0 with Bundle Patch 10
Always-On Monitoring 13.3.0.0.0 and Greater	Oracle Database 12.2.0.1.0, 18.0.0.0.0, 18.3.0.0.0, and 19.0.0.0.0.

Database Sizing

The factors to consider when sizing and configuring the database are as follows:

- **Undo Tablespace**: The Undo tablespace is utilized when running the SYNC processes, especially during the initial SYNC when potentially larger numbers of rows are pulled from the Enterprise Manager instance to Always-On Monitoring.

- **Temp Tablespace**: The SYNC process will also utilize TEMP tablespace, especially when indexes are created or sorting occurs during the data movement process.
- **Always-On Monitoring Tablespace**: Table data and indexes specific to Always-On Monitoring are stored in this tablespace.
- **Redo Logs**: Redo logs should be large enough to minimize the number of checkpoints that occur during times when the SYNC is occurring. It is recommended to configure 3 x 1 GB REDO log files.
- **Special Oracle Parameter Settings**: For Always-On Monitoring, it is desirable to maintain the parameters used in the Enterprise Manager repository, especially if the Always-On Monitoring schema is to be populated in the Enterprise Manager repository database. To ensure correctness of the Enterprise Manager parameter settings that are passed in as part of the Always-On Monitoring install, you can run the SYNC and other Always-On Monitoring functions that are part of the emsctl verbs. This is important to consider when configuring a separate database instance. As is the case with the Enterprise Manager Repository, the following parameter should be set for the Always-On Monitoring Repository to avoid unforeseen optimizer issues:

  ```
  ALTER SYSTEM SET OPTIMIZER_ADAPTIVE_FEATURES=FALSE SCOPE=BOTH SID='*';
  ```

  The following table represents an example of the sizing of the above components for an Always-On Monitoring schema that is being used for a large enterprise database. Note that these figures represent the sizes of the tablespaces after the initial full SYNC between Always-On Monitoring and Enterprise Manager. Also note that the Redo Log files are not included in these calculations:

Table 17-1 Always-On Monitoring Repository Tablespace Sizing

Tablespace Name	Used Space (MB)	Free Space (MB)	Total Allocation (MB)	Percentage Free (%)
TEMP	0	30,720	30,720	100%
Users	6,430	44,797	51,200	87%
SYSAUX	1,385	85	1,470	6%
SYSTEM	896	4	900	0%
UNDOTBS1	84	30,636	30,720	100%
Totals	8.769	106,246	115,015	92%

For this particular configuration, the absolute minimum for disk space required was 9 GB based on the used Tablespace. To ensure room for growth with this schema, it is recommended to plan on allocating at least 120 GB of space for the Always-On Monitoring database if creating a new Oracle instance. The TEMP and UNDO tablespaces are configured to handle initial and subsequent SYNC operations ensuring there is no issue with TEMP or REDO saturation. In addition, the Always-On Monitoring tablespace is sized to ensure no interruption due to saturation of the tablespace.

In order to accurately size the Always-On Monitoring tablespace for current usage and future growth, it is important to understand the following:

- What data from what tables is being transferred from the Enterprise Manager Repository to the Always-On Monitoring Schema.
- Knowing the number of rows transferred for each table, what is the approximate number of Bytes per Row for the tables and indexes?

The following table shows a subset of the tables and indexes in the Always-On Monitoring schema, highlighting the number of rows, total space consumption, and the number of bytes per row:

Table 17-2 Always-On Monitoring Table and Index Space Allocation

Owner	Segment Name	Segment Type	Partition Name	Size (MB)	Number of Rows
Always-On Monitoring	MGMT_TARGET_PROPERTIES	TABLE	Non-Partitioned	1,472	15,408,952
Always-On Monitoring	MGMT_TARGET_PROPERTIES_IDX_02	INDEX	Non-Partitioned	1,152	15,215,316
Always-On Monitoring	MGMT_TARGET_PROPERTIES_PK	INDEX	Non-Partitioned	856	15,462,104
Always-On Monitoring	MGMT_TARGET_PROPERTIES_IDX_01	INDEX	Non-Partitioned	584	14,871,521
Always-On Monitoring	EM_MANAGEABLE_ENTITIES	TABLE	Non-Partitioned	496	1,177,065
Always-On Monitoring	MGMT_TARGET_PROPERTIES_IDX_03	INDEX	Non-Partitioned	472	16,105,162
Always-On Monitoring	EM_MANAGEABLE_ENTITIES_UK1	INDEX	Non-Partitioned	136	1,201,617
Always-On Monitoring	EM_MANAGEABLE_ENTITIES_UK2	INDEX	Non-Partitioned	120	1,133,804
Always-On Monitoring	EM_MANAGEABLE_ENTITIES_IDX01	INDEX	Non-Partitioned	112	1,184,414
Always-On Monitoring	EM_MANAGEABLE_ENTITIES_IDX07	INDEX	Non-Partitioned	112	1,166,536
Always-On Monitoring	EM_MANAGEABLE_ENTITIES_IDX03	INDEX	Non-Partitioned	96	1,182,817
Always-On Monitoring	EM_MANAGEABLE_ENTITIES_IDX04	INDEX	Non-Partitioned	80	1,218,977
Always-On Monitoring	EM_MANAGEABLE_ENTITIES_IDX05	INDEX	Non-Partitioned	72	1,104,428
Always-On Monitoring	EM_MANAGEABLE_ENTITIES_IDX06	INDEX	Non-Partitioned	59	673,838
Always-On Monitoring	EM_MANAGEABLE_ENTITIES_IDX08	INDEX	Non-Partitioned	44	1,215,411
Always-On Monitoring	EM_MANAGEABLE_ENTITIES_PK	INDEX	Non-Partitioned	41	1,172,101
Always-On Monitoring	EM_MANAGEABLE_ENTITIES_IDX09	INDEX	Non-Partitioned	40	1,192,236
Always-On Monitoring	EM_MANAGEABLE_ENTITIES_IDX02	INDEX	Non-Partitioned	37	1,178,410
Always-On Monitoring	EM_VIOLATIONS	TABLE	Non-Partitioned	28	112,106
Always-On Monitoring	EM_METRIC_COLUMN_VER_PK	INDEX	Non-Partitioned	17	611,561

Database Character Set Definition

To ensure that the data being transferred from the Enterprise Manager Repository to the Always-On Monitoring Repository is stored and represented properly when the Always-On Monitoring Repository is on a separate instance from Enterprise Manager, it is important to ensure that the character set matches between the two databases/instances. For Enterprise Manager, the NLS_CHARACTERSET is defined as AL32UTF8..

If using Oracle Installer to manually set up the database instance that Always-On Monitoring uses, the instance will default to WE8MSWIN1252. Ensure at this point that the option AL32UTF8 is selected from the drop down box so that the correct character set is installed from the start..

To determine the character set on the Enterprise Manager Repository, run the following query:

```
SELECT * from nls_database_parameters where parameter='NLS_CHARACTERSET';
```

Run this query as well on the Always-On Monitoring Repository to ensure the character sets match. If the Always-On Monitoring Repository character set is different, the following is recommended to change the character set on the Always-On Monitoring Repository to match that of the Enterprise Manager Repository.

```
--Shutdown and start the database in a restricted mode.
SHUTDOWN IMMEDIATE;
STARTUP RESTRICT;

--Set JOB_QUEUE_PROCESSES and AQ_TM_PROCESSES to its default value i.e.0
ALTER SYSTEM SET JOB_QUEUE_PROCESSES = 0;
ALTER SYSTEM SET AQ_TM_PROCESSES=0;

-- Set the required character set.
ALTER DATABASE CHARACTER SET AL32UTF8;

 -- if the above fails:
ALTER DATABASE CHARACTER SET INTERNAL_USE AL32UTF8;

--Shutdown and restart the database in normal mode.
SHUTDOWN IMMEDIATE;
STARTUP;
/
```

Creating the Always-On Monitoring Repository User

The Always-On Monitoring user should be created by a database administrator with SYSDBA privileges. If you are running emsca as a user with SYSDBA privileges, you can skip this section as the emsca configuration assistant will automatically create the user for you.

> **Note:**
>
> The Always-On Monitoring Repository should not be housed within the Enterprise Manager Repository database. Housing the Always-On Monitoring Repository in a separate database allows Always-On Monitoring to function even if the Enterprise Manager Repository goes down.

Connect to the database to be used as the Always-On Monitoring Repository, create the Always-On Monitoring user and then grant it the required privileges.

> **Note:**
>
> The Always-On Monitoring user should be created by a database administrator with SYSDBA privileges. Alternatively, you can supply SYSDBA credentials when running the emsca configuration utility. emsca will then create the Always-On Monitoring repository user automatically.

If you have SYSDBA access, you can skip this step since it will be done automatically by the emsca configuration assistant. See Using the Always-On Monitoring Configuration Assistant (EMSCA) for more information.

Granting Required Privileges to the Always-On Monitoring Schema Owner

To ensure proper functionality of Always-On Monitoring, the Always-On Monitoring schema owner needs to have the correct privileges. The following sample script details all grants required for the Always-On Monitoring Schema:

```
create user ems identified by ems;
grant CREATE SESSION,
      ALTER SESSION,
      CREATE DATABASE LINK,
      CREATE MATERIALIZED VIEW,
      CREATE PROCEDURE,
      CREATE PUBLIC SYNONYM,
      CREATE ROLE,
      CREATE SEQUENCE,
      CREATE SYNONYM,
      CREATE TABLE,
      CREATE TRIGGER,
      CREATE TYPE,
      CREATE VIEW,
      UNLIMITED TABLESPACE,
      SELECT ANY DICTIONARY to ems;
grant EXECUTE ON SYS.DBMS_CRYPTO to ems;
grant EXECUTE ON SYS.DBMS_AQADM to ems;
grant EXECUTE ON SYS.DBMS_AQ to ems;
grant EXECUTE ON SYS.DBMS_AQIN to ems;
grant EXECUTE on SYS.DBMS_LOCK to ems;
grant EXECUTE ON SYS.DBMS_SCHEDULER to ems;
grant create job to ems;
```

Under certain circumstances you may need to grant privileges directly to an Always-On Monitoring user. However, directly granting privileges to any user may violate security policy and compliance rules. In this situation, you can can assign most of the privileges to a role and then assign that role to the user.

For multitenant environments, when creating a user in a multitenant container database (CDB) and not in one of its constituent pluggable databases (PDB), then the user has to be prefixed by "C##".

Best Practices

Oracle recommends keeping an Always-On Monitoring instance running at all times regardless of whether Enterprise Manager is up or down. By telling the Always-On Monitoring to turn off notifications, you can have it run in the background while Enterprise Manager is up. During Enterprise Manager downtime, whether planned or unplanned, you can issue an *emsctl* command to tell Always-On Monitoring to restart notifications. The primary advantage to having Always-On Monitoring running at all times is that it will stay up to date with any changes made to your Enterprise Manager installation. Always-On Monitoring keeps itself in sync with Enterprise Manager by periodically running a synchronization job. This job runs every 24 hours by default, however the job execution time can be changed to meet the needs of your monitored environment.

You may also run Always-On Monitoring only when needed instead of having it run constantly in the background. However, doing so will require Always-On Monitoring to synchronize with the Enterprise Manager OMS. Depending on the size of your Enterprise Manager deployment, synchronization may take 10-15 minutes in order to obtain current information on all Enterprise Manager targets and metrics.

> **Note:**
>
> Always-On Monitoring and Enterprise Manager must be in SYNC before Enterprise Manager goes down.

Installing Always-On Monitoring

Always-On Monitoring is a self-contained application that is supplied with the Enterprise Manager software distribution.

Steps to install Always-On Monitoring:

1. Uninstall any previous installations of Always-On Monitoring. For more information, see Uninstalling Always-On Monitoring.
2. Locate the latest Always-On Monitoring ZIP file. The ZIP file can be found in the `/sysman/ems` directory of the OMS.
3. Copy the ZIP file to the Always-On Monitoring host and unzip the contents to the location where Always-On Monitoring is to be installed.

Installing Always-On Monitoring from an Enterprise Manager Software Distribution

The Always-On Monitoring installation zip file is shipped with Enterprise Manager. From the `sysman/ems` directory, find the Always-On Monitoring installation ZIP file (e.g. ems.zip). Unzip the file to the location where you want to install Always-On Monitoring.

Installing Multiple Always-On Monitoring Instances

You can set up multiple Always-On Monitoring (AOM) instances to ensure Always-On Monitoring is available in the event of outages. See Setting Up Multiple Always-On Monitoring Instances for more information.

Configuring Always-On Monitoring

Once Always-On Monitoring has been installed, the Always-On Monitoring configuration assistant script can be used to configure the service to communicate with the Enterprise Manager repository.

Always-On Monitoring configuration involves the following steps:

1. Saving the Em Key
2. Using the Always-On Monitoring Configuration Assistant (EMSCA)
3. Removing the Em Key
4. Configuring Email Servers in Enterprise Manager
5. Configuring Downtime Contacts in Enterprise Manager
6. Synchronizing Always-On Monitoring with Enterprise Manager for the First Time
7. Configuring Enterprise Manager to Work with Always-On Monitoring
8. Starting Always-On Monitoring
9. Enabling Notifications
10. Verifying the Always-On Monitoring Upload URL on Enterprise Manager

Saving the Em Key

Always-On Monitoring requires the *Em key* in order to be able to synchronize with the Enterprise Manager repository and reuse the SMTP (email) gateway configuration information from Enterprise Manager.

> **Note:**
>
> The *Em key* needs to be copied to the Enterprise Manager repository BEFORE configuring Always-On Monitoring. Not doing so will result in errors when running *emsca*.

Once Always-On Monitoring configuration is complete (*emsca* has successfully run through completion), immediately remove the *Em key* from the Enterprise Manager repository. Note: Always-On Monitoring does not have to be restarted. To store the key in the repository, run the following *emctl* commands from the `$MW_HOME/bin` directory, where $MW_HOME is the Enterprise Manager Middleware Home directory.

```
% emctl config emkey -copy_to_repos
```

Once Always-On Monitoring configuration is complete, execute the following command to remove the key.

```
% emctl config emkey -remove_from_repos
```

Using the Always-On Monitoring Configuration Assistant (EMSCA)

The Always-On Monitoring configuration assistant, emsca, is a script located under the Always-On Monitoring installation scripts directory. Running emsca requires the bash shell to be installed.

If you have database administrator credentials, you can pass them to the emsca script and it will create the Always-On Monitoring user for you.

If you do not have database administrator credentials, you need to:

1. Ask the database administrator to run the script. See Creating the Always-On Monitoring Repository User.
2. Run `emsca -createEmsDbUser=false`

The following example usage scenarios assume Always-On Monitoring is installed in a location referred to using the environment variable AOM_HOME. Run the configuration assistant located in $AOM_HOME/scripts. The EMSCA may be invoked with no parameters and will prompt for the necessary information. Once the configuration has completed, record the Always-On Monitoring Upload URL as it will be used later to configure Enterprise Manager.

Example EMSCA Installation Scenarios

The user installing Always-On Monitoring has SYSDBA credentials.

```
Copyright (c) 2017, 2020 Oracle Corporation.  All rights reserved.
---------------------------------------------------------------
Always-On Monitoring Repository Connection String : myserver.myco.com:25059:s307480
 Always-On Monitoring Repository Username [ems] :
 Always-On Monitoring Repository Password [ems] :

User "ems" cannot be found in the database.
In order to create this user, SYSDBA credentials are required. If you do not want to
continue, answer "n" to the question below.
Create the Always-On Monitoring Repository user [y] :
 Always-On Monitoring Repository SYSDBA Username : sys
 Always-On Monitoring Repository SYSDBA Password :

 Enterprise Manager Repository Connection String : myserver.myco.com:25059:s307480
 Enterprise Manager Repository Username : sysman
 Enterprise Manager Repository Password :

Creating Always-On Monitoring repository user ems
 Agent Registration Password :

 Keystore for host myserver.myco.com created successfully.
 Connecting to Always-On Monitoring Repository.
 Creating Always-On Monitoring Repository schema
 Creating repository storage for Targets data.
 Creating repository storage for Alerts and Availability data.
 Creating repository storage for Notification Metadata data.
 Creating repository storage for Target Metric Metadata data.
 Registering Always-On Monitoring instance
 Always-On Monitoring Upload URL: https://myserver.myco.com:8081/upload
```

The user installing Always-On Monitoring does not have SYSDBA credentials.

Example 1: SYSDBA creates an Always-On Monitoring user and grants user privileges shown in the following example.

aomuser is the name of the Always-On Monitoring user in the database.

SYSDBA privileges are required to run the following script

```
create user aomuser identified by <password>;
grant CREATE JOB,
         CREATE SESSION,
         ALTER SESSION,
         CREATE DATABASE LINK,
         CREATE MATERIALIZED VIEW,
         CREATE PROCEDURE,
         CREATE PUBLIC SYNONYM,
         CREATE ROLE,
         CREATE SEQUENCE,
         CREATE SYNONYM,
         CREATE TABLE,
         CREATE TRIGGER,
         CREATE TYPE,
         CREATE VIEW,
         UNLIMITED TABLESPACE,
         SELECT ANY DICTIONARY to aomuser;

grant EXECUTE ON SYS.DBMS_CRYPTO to aomuser;
grant EXECUTE ON SYS.DBMS_AQADM to aomuser;
grant EXECUTE ON SYS.DBMS_AQ to aomuser;
grant EXECUTE ON SYS.DBMS_AQIN to aomuser;
grant EXECUTE ON SYS.DBMS_SCHEDULER to aomuser;
grant EXECUTE ON SYS.DBMS_LOCK to aomuser;
```

Example 2: The Always-On Monitoring user invokes the EMSCA script..

```
Oracle Enterprise Manager Cloud Control 13c Release 4
Copyright (c) 2017, 2020 Oracle Corporation.  All rights reserved.
---------------------------------------------------------------
 Always-On Monitoring Repository Connection String : myserver.myco.com:25059:s307480
 Always-On Monitoring Repository Username [ems] : aomuser
 Always-On Monitoring Repository Password [ems] :

Always-On Monitoring Repository user "aomuser" has already been created
 Enterprise Manager Repository Connection String : myserver.myco.com:25059:s307480
 Enterprise Manager Repository Username : sysman
 Enterprise Manager Repository Password :

 Agent Registration Password :

 Keystore for host myserver.myco.com created successfully.
 Connecting to Always-On Monitoring Repository.
 Creating Always-On Monitoring Repository schema
 Creating repository storage for Targets data.
 Creating repository storage for Alerts and Availability data.
 Creating repository storage for Notification Metadata data.
 Creating repository storage for Target Metric Metadata data.
 Registering Always-On Monitoring instance
 Always-On Monitoring Upload URL: https://myserver.myco.com:8081/upload
```

SYSDBA creates a role and assigns it to Always-On Monitoring user then emsca is executed.

Some sites prefer to create a role and assign that role to the Always-On Monitoring user. In this situation, you must run the scripts shown in the following examples.

Example 1: SYSDBA creates a role and then assigns it to the Always-On Monitoring user.

```
create user <aom_user> identified by <aom_password>;
create role ems_role;
grant CREATE SESSION,
         ALTER SESSION,
         CREATE DATABASE LINK,
         CREATE MATERIALIZED VIEW,
         CREATE PROCEDURE,
         CREATE PUBLIC SYNONYM,
         CREATE ROLE,
         CREATE SEQUENCE,
         CREATE SYNONYM,
         CREATE TABLE,
         CREATE TRIGGER,
         CREATE TYPE,
         CREATE VIEW,
         SELECT ANY DICTIONARY to ems_role;
grant ems_role to <aom_user>;
```

Example 2: SYSDBA grants the following privileges directly to the Always-On Monitoring user.

```
grant CREATE JOB,
         UNLIMITED TABLESPACE to <aom_user>;

grant EXECUTE ON SYS.DBMS_CRYPTO to <aom_user>;
grant EXECUTE ON SYS.DBMS_AQADM to <aom_user>;
grant EXECUTE ON SYS.DBMS_AQ to <aom_user>;
grant EXECUTE ON SYS.DBMS_AQIN to <aom_user>;
grant EXECUTE ON SYS.DBMS_SCHEDULER to <aom_user>;
grant EXECUTE ON SYS.DBMS_LOCK to <aom_user>;
```

Example 3: The Always-On Monitoring user invokes the *emsca* script.

```
Oracle Enterprise Manager Cloud Control 13c Release 4
Copyright (c) 2017, 2020 Oracle Corporation.  All rights reserved.
---------------------------------------------------------------
 Always-On Monitoring Repository Connection String : myserver.myco.com:25059:s307480
 Always-On Monitoring Repository Username [ems] :
 Always-On Monitoring Repository Password [ems] :

Always-On Monitoring Repository user "ems" has already been created
 Enterprise Manager Repository Connection String : myserver.myco.com:25059:s307480
 Enterprise Manager Repository Username : sysman
 Enterprise Manager Repository Password :

 Agent Registration Password :

 Keystore for host myserver.myco.com created successfully.
 Connecting to Always-On Monitoring Repository.
 Creating Always-On Monitoring Repository schema
 Creating repository storage for Targets data.
 Creating repository storage for Alerts and Availability data.
 Creating repository storage for Notification Metadata data.
 Creating repository storage for Target Metric Metadata data.
 Registering Always-On Monitoring instance
 Always-On Monitoring Upload URL: https://myserver.myco.com:8081/upload
```

Running EMSCA with a Response File

You can also use the -responseFile option to provide the parameters to the script using a response input file. For example

```
$ emsca -responseFile=<response_filename>
```

Where *response_filename* is the response file containing the following script parameters:

```
emsRepConnectString=localhost:1521:xe
emsRepUsername=ems
emsRepPassword=ems
emRepConnectString=mymachine.mycompany.com:15044:semgc3
emRepUsername=sysman
emRepPassword=sysman
#
emsPort=8081
http.protocol=http
#
```

The following table lists available EMSCA parameters.

Table 17-3 EMSCA Parameters

EMSCA Prompt	Description
Always-On Monitoring Repository Connection String	The connect string used to locate the Always-On Monitoring repository. This may be any valid service descriptor or may be a `host:port:sid`.
Always-On Monitoring Repository Username	The username associated with the Always-On Monitoring repository. See Creating the Always-On Monitoring Repository Userfor more details.
Always-On Monitoring Repository Password	The password associated with the Always-On Monitoring repository user. See Creating the Always-On Monitoring Repository User for more details.
Enterprise Manager Repository Connection String	The connect string used to locate the Enterprise Manager repository. This may be any valid service descriptor or may be a host:port:sid.
Enterprise Manager Repository Username	The username of the Enterprise Manager SYSMAN user.
Enterprise Manager Repository Password	The password of the Enterprise Manager SYSMAN user.

Removing the Em Key

After Always-On Monitoring configuration is complete (after emsca has been run), execute the following command to remove the key.

```
% emctl config emkey -remove_from_repos
```

Always-On Monitoring does not have to be restarted.

Configuring Email Servers in Enterprise Manager

When Always-On Monitoring sends email notifications, it does so using email server configurations that it obtains from Enterprise Manager during the synchronization step. For Always-On Monitoring to function properly, email servers must be configured in Enterprise Manager before performing an Always-On Monitoring synchronization.

If changes to the email server configurations are made in Enterprise Manager, Always-On Monitoring must be synchronized again to retrieve the updated email server configuration. See Updating Always-On Monitoring “Running an Incremental Synchronization Manually” for additional details.

Configuring Downtime Contacts in Enterprise Manager

Once Always-On Monitoring is synchronized with Enterprise Manager, the monitoring service can send email notifications to the appropriate administrators when the OMS is down via configured downtime contacts. Prior to running Always-On Monitoring for the first time, downtime contacts should be configured in Enterprise Manager. The downtime contacts that Always-On Monitoring sends notification to may be any one of the following forms.

- Global downtime contact
- Per-target downtime contact based on target property
- Per-target downtime contact based on event rules

If email addresses are specified for both Global and Per-target forms, then all email addresses will be used.

Global Downtime Contact

The global downtime contact is a single property set on the Enterprise Manager site. Once set, all target status events and metric alerts across all targets will be sent to the recipients specified in the global downtime contact property.

```
% emcli set_oms_property -property_name='oracle.sysman.core.events.ems.downtimeContact'
-property_value='<email addresses>'
```

To delete the global downtime contact, run the following command:

```
emctl delete property -name 'oracle.sysman.core.events.ems.downtimeContact'
```

Per-Target Downtime Contact from Target Property

The email addresses for these downtime contacts should be specified in the *Downtime Contact* target property for each target. There are three ways to specify the *Downtime Contact* target property:

- For each target, navigate to the Target Properties page, which can be accessed from the target's homepage. From the *target* menu on the target homepage, select **Target Setup** and then **Properties**. The Target Properties page displays.

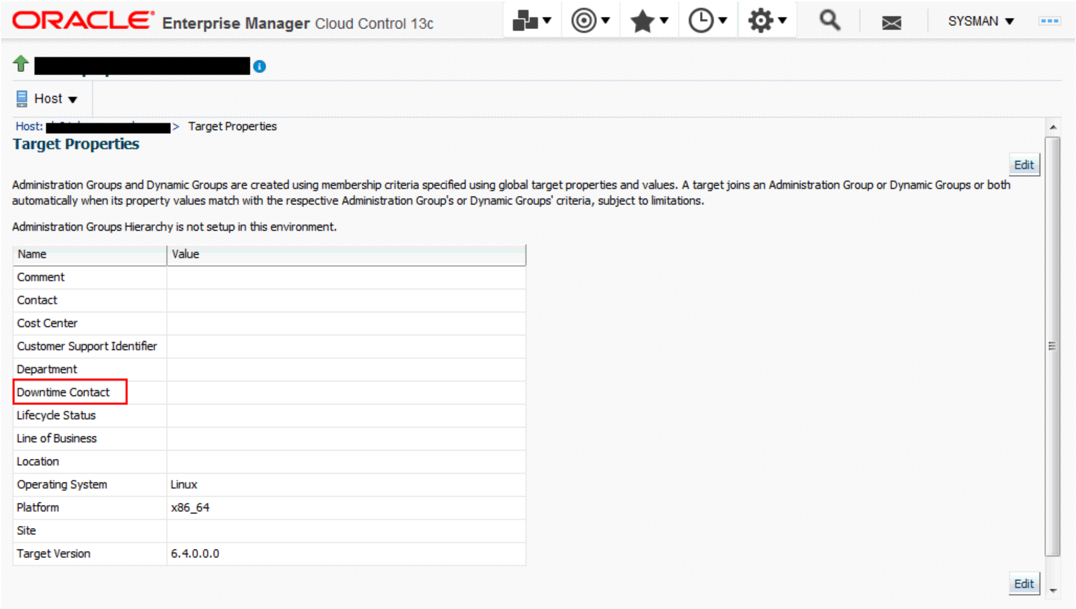

[redacted]

Host: [redacted] > Target Properties

Target Properties

Administration Groups and Dynamic Groups are created using membership criteria specified using global target properties and values. A target joins an Administration Group or Dynamic Groups or both automatically when its property values match with the respective Administration Group's or Dynamic Groups' criteria, subject to limitations.

Administration Groups Hierarchy is not setup in this environment.

Name	Value
Comment	
Contact	
Cost Center	
Customer Support Identifier	
Department	
Downtime Contact	
Lifecycle Status	
Line of Business	
Location	
Operating System	Linux
Platform	x86_64
Site	
Target Version	6.4.0.0.0

  Click on **Edit** and specify the email address in the **Downtime Contact** target property. You can specify multiple email addresses by separating them with commas.

- Use the EM CLI set_target_property_value verb.

```
emcli set_target_property_value -
property_records="target_name:target_type:property_name:property_value"
```

  Example:

```
%./emcli set_target_property_value -property_records="myhost:host:Downtime
Contact:John.Doe@mycompany.com"
```

For more information about the set_target_property_value verb, see the Oracle Enterprise Manager Command Line Interface Guide.

- By leveraging the email recipients in the target down event rules in Enterprise Manager. See the following section for more details.

Per-Target Downtime Contact from Event Rules

Always-On Monitoring may also send email notifications to different users for each target. These contacts are generated in Enterprise Manager based on the event rules for that target. Therefore, as event rules are changed in Enterprise Manager, the contacts must be re-generated and an incremental synchronization performed

By leveraging the event rule setup, the downtime contact will be generated based on the email recipient for the event rule for a *Target Availability* event type where *Down* status has been selected.

> **Note:**
>
> Although downtime contacts are generating using only Target Availability event rules, Always-On Monitoring will send notifications for both target availability and metric threshold alerts.

You can review and update the recipients of your target availability (status down) event rules. Doing so allows you to generate a list of downtime contacts using EM CLI or by submitting the downtime contact generation job.

Generating Downtime Contact using EM CLI

From the EM CLI, use the *generate_downtime_contact* command to produce downtime contacts for a specific target. The command simulates ALL Target Availability rules that would affect that particular target and generates a list of email addresses that includes the output from all rules. The optional *-set* parameter will cause the contact to be saved so that Always-On Monitoring synchronization can retrieve it later.

```
% emcli generate_downtime_contact -target_name=<name> -target_type=<type> -set
```

Generating Downtime Contact Enterprise Manager Job

```
% emcli create_job -input_file=property_file:<job_prop_file> -
name="GenerateDowntimeContacts"
```

Where the *job_prop_file* refers to a properties file with the following content:

```
# job type
type=CoreSetDowntimeContacts
# description
description=You job description
# target names, the list of target names
variable.target_names=host1,database2
# target types, the list of target types corresponding to the target list above
variable.target_types=host,oracle_database
# set on all targets, true to update contacts, false to print output but not save
variable.set_all_targets=true
# frequency to run job, refer to Enterprise Manager documentation for options
schedule.frequency=IMMEDIATE
```

> **Note:**
>
> The process of generating the downtime contacts thru either the EM CLI or the downtime contact generation job, does not immediately make the contacts available to the monitoring service. A separate Always-On Monitoring synchronization must be done each time these contacts are updated. This is done using the emsctl sync command.

Types of Alerts Received by Downtime Contacts

Once configured, Always-On Monitoring will send email notifications for:

- All target status alerts and metric alerts (critical, warning, and clear).
- Metric collection errors (both critical and clear).

Target status alerts for aggregate targets, such as groups and clusters, whose status is based on the status of its component members is not supported. Instead, alerts will be sent for the individual members of these aggregate targets**..**

> **Note:**
>
> Currently, email notification is the only type of notification supported.

The recipients of the email should be email addresses. The recipients are specified using the global property *downtimeContact* or the per-target Downtime Contact. If both are specified, emails will be sent to all recipients specified.

If the global property *downtimeContact* is specified, all alerts on all targets will be sent to the recipients specified in the global property *downtimeContact*. If the per-target property Downtime Contact is specified, all target status and metric alerts for that target will be sent to the recipients specified in that property.

Setting Downtime Contacts for Composite Targets

You can set the Downtime Contacts target property for a composite target so that the target property is propagated to its member targets. You can perform this update using the EMCLI `set_target_property_value` verb. See set_target_property_value in the *Oracle® Enterprise Manager Command Line Interface* for more information about this verb.

Examples

Example 1: The following example sets the AOM Downtime Contact target property for a DB static group.

```
>  ./emcli set_target_property_value -
property_records="db_group:composite:Downtime Contact:aom@test.com" -
propagate_to_members
 Properties updated successfully
```

Example 2: The following example sets the AOM Downtime Contact target property for a DB system target.

```
> ./emcli set_target_property_value -
property_records="eva4.us.oracle.com_sys:oracle_dbsys:Downtime
Contact:aom@test.com" -propagate_to_members
Properties updated successfully
```

Synchronizing Always-On Monitoring with Enterprise Manager for the First Time

Before starting Always-On Monitoring for the first time, you must synchronize Always-On Monitoring with Enterprise Manager. Synchronization copies notification configuration and downtime contacts from Enterprise Manager targets, thus allowing Always-On Monitoring to monitor alerts and send email notifications.

To perform Always-On Monitoring-Enterprise Manager synchronization:

```
% $AOM_HOME/scripts/emsctl sync

Oracle Enterprise Manager Cloud Control 13c Release 5
Copyright (c) 2017, 2021 Oracle Corporation. All rights reserved.
------------------------------------------------------------------
 Connecting to Always-On Monitoring repository.
Starting synchronization with Enterprise Manager.
Synchronizing with Enterprise Manager repository: sysman@myserver.myco.com:35074:semgc4
Synchronizing Targets data.
Synchronizing Alerts and Availability data.
Synchronizing Notification Metadata data.
Synchronizing Target Metric Metadata data.
Synchronization complete at : Thu Oct 22 06:50:56 PDT 2016
```

Note:

Synchronization should be run again whenever there are changes made to Enterprise Manager. For example, after adding new hosts or targets, changing the email server, or changing the downtime contacts, an incremental synchronization is required in order for Always-On Monitoring to pick up the latest configurations.. See "Running an Incremental Synchronization Manually" for additional details.

Configuring Enterprise Manager to Work with Always-On Monitoring

The final step in the configuration of Always-On Monitoring is to update Enterprise Manager to include the upload URL for the service. The upload URL is displayed by Always-On Monitoring Configuration Assistant upon successful completion of the configuration of the service. Always-On Monitoring configuration properties file also includes the protocol and port configured for the service so that the upload URL will be of the form:

https://yourhostname:8081/upload

Once configured, the URL will be sent to all Agents (version 13.1 and greater), both existing Agents and any Agents deployed in future. Hence, configuration only needs to be done once. Agents less than 13.1 will not receive the Always-On Monitoring upload URL.

HTTPS is the default protocol for Always-On Monitoring and the default port is 8081. Please refer to the emsca or Always-On Monitoring configuration file for the configured values.You can find the configuration file at the following location:

```
$AOM_HOME/conf/emsConfig.properties
```

To configure Enterprise Manager with Always-On Monitoring upload location, the following command should be executed from Enterprise Manager.

```
% emctl set property -name "oracle.sysman.core.events.ems.emsURL" -value "https://
yourhostname:8081/upload" -sysman_pwd sysman
```

Starting Always-On Monitoring

Start Always-On Monitoring. This step is important to enable notifications. For details on starting AOM, see Starting Always-On Monitoring.

Enabling Notifications

Always-On Monitoring initial configuration has notifications disabled—the service will not send emails as new alerts are received. You must first enable Always-On Monitoring to send email notifications.The Always-On Monitoring notification enable/disable setting is persistent across starts/stops of Always-On Monitoring. The *enable_notification* command will automatically perform an incremental sync.

```
% $AOM_HOME/scripts/emsctl enable_notification
Oracle Enterprise Manager Cloud Control 13c Release 5
Copyright (c) 2017, 2021 Oracle Corporation. All rights reserved.
------------------------------------------------------------------
Notifications have been enabled. There are downtime contacts configured.
 Connecting to Always-On Monitoring repository.
Starting synchronization with Enterprise Manager.
Synchronizing with Enterprise Manager repository: sysman@myserver.myco.com:35074:semgc4
Synchronizing Targets data.
Synchronizing Alerts and Availability data.
Synchronizing Notification Metadata data.
Synchronizing Target Metric Metadata data.
Synchronization complete at : Thu Oct 22 07:01:20 PDT 2020
```

If you do not want to perform a SYNC then add the -nosync option.

```
$AOM_HOME/scripts/emsctl enable_notification -nosync
Oracle Enterprise Manager Cloud Control 13c Release 5
Copyright (c) 2017, 2021 Oracle Corporation. All rights reserved.
------------------------------------------------------------------
Notifications have been enabled. There are downtime contacts configured.
```

To disable notifications but leave Always-On Monitoring running use the disable_notification command:

```
% $AOM_HOME/scripts/emsctl disable_notification
Oracle Enterprise Manager Cloud Control 13c Release 5
Copyright (c) 2017, 2021 Oracle Corporation.  All rights reserved.
---------------------------------------------------------------
Notifications have been disabled. Filters have been deleted.
```

Filtering Alert Notifications by Event Type and Severity

When only specific alert types are pertinent, you need to have a way of filtering out extraneous alerts so you don't have to scan through all alerts just to find the ones you're interested in.

Always-On Monitoring lets you filter your alert notifications based on either *event type* or *severity*.

Filtering Alert Notifications by Event Type
Alert Notification Filtering by event type lets you exclude alert notifications for certain event types. For example, you want to receive only target availability alerts. You can filter out alerts for the following event types:

- Metric Alert (`metricAlert`)
- Availability (`targetAvail`)
- Metric Evaluation Errors (`metricError`)

Default Settings

- Target availability & metric alert notifications are enabled
- Metric evaluation error alert notifications are disabled

Example 1: To receive only target availability alerts, run the following:

```
$AOM_HOME/scripts/emsctl disable_notification
          -alert_types=targetAvail,metricAlert,metricError
$AOM_HOME/scripts/emsctl enable_notification -alert_types=targetAvail
```

Example 2: To add notification for metric error alerts in addition to target availability alerts:

```
$AOM_HOME/scripts/emsctl enable_notification -alert_types=metricError
Oracle Enterprise Manager Cloud Control 13c Release 5
Copyright (c) 2017, 2021, Oracle Corporation.  All rights reserved.
------------------------------------------------------------------
Notification filters have been enabled. There are downtime contacts
configured.
Notification filters
    Target Availability (targetAvail) : true
    Metric Alerts (metricAlert) : false
    Metric Errors (metricError) : true
```

Filtering Alert Notifications by Severity
Alert Notification Filtering by severity lets you exclude alert notifications for certain severities. For example, you want to receive only Fatal, Critical, and Warning alerts. You can filter out alerts for the following event types:

- Fatal
- Critical
- Warning
- Advisory
- Informational
- Clear

Default Settings

- Fatal, Critical, and Warning alert notifications are enabled
- Advisory, Informational, and Clear alert notifications are disabled

Example 1: To receive only Fatal and Critical severity alerts, run the following:

```
$AOM_HOME/scripts/emsctl disable_notification –
severities=fatal,critical,warning,advisory,informational,clear
$AOM_HOME/scripts/emsctl enable_notification - severities = fatal,critical
```

Example 2: To add notification for Warning alerts in addition to Fatal and Critical severity alerts:

```
$AOM_HOME/scripts/emsctl enable_notification -severities=warning

Oracle Enterprise Manager Cloud Control 13c Release 5
Copyright (c) 2017, 2021, Oracle Corporation.  All rights reserved.
------------------------------------------------------------------
Notification filters have been enabled. There are downtime contacts
configured.
Notification filters - Alert types
    Target Availability (targetAvail) : true
    Metric Alerts (metricAlert) : true
    Metric Errors (metricError) : false
Notification filters - Severity
    Fatal (fatal) : true
    Critical (critical) : true
    Warning (warning) : true
    Advisory (advisory) : false
    Informational (informational) : false
    Clear (clear) : false
```

Verifying the Always-On Monitoring Upload URL on Enterprise Manager

You can determine the upload URL by issuing the following command:

```
$AOM_HOME/scripts/emsctl status
```

Example:

```
$AOM_HOME/scripts/emsctl statusOracle Enterprise Manager Cloud Control 13c Release 5
Copyright (c) 2017, 2021, Oracle Corporation. All rights reserved.
------------------------------------------------------------------
Always-On Monitoring Version : 13.5.0.0.0
Always-On Monitoring Home : /mycomputer/ems
Started At : December 23, 2019 2:24:25 PM PST
Last Repository Sync : December 23, 2019 2:25:49 PM PST
Upload URL : https://myserver:8081/upload
Always-On Monitoring Process ID : 24924
Always-On Monitoring Repository : myserver.myco.com:15044:sarview
Enterprise Manager Repository : myserver.myco.com:15044:sarview
Notifications Enabled : true
Notification filters - Alert types
    Target Availability (targetAvail) : true
    Metric Alerts (metricAlert) : true
    Metric Errors (metricError) : false
Notification filters - Severity
    Fatal (fatal) : true
    Critical (critical) : true
    Warning (warning) : true
    Advisory (advisory) : false
    Informational (informational) : false
    Clear (clear) : false
Total Downtime Contacts Configured : 1
```

You can verify the setting in the OMS by using the `emctl get property` command.

You can verify the setting on a particular Agent by looking at *$AGENT_HOME/sysman/config/emd.properties* and searching for *EMS_URL*.

Controlling the Service

Starting Always-On Monitoring

```
% $AOM_HOME/scripts/emsctl start
Oracle Enterprise Manager Cloud Control 13c Release 5
Copyright (c) 2017, 2021 Oracle Corporation. All rights reserved.
------------------------------------------------------------------
Pinging Always-On Monitoring...
Always-On Monitoring is not running.
Starting Always-On Monitoring.
Waiting for process to start
Retrying...
Notifications Enabled : false
Total Downtime Contacts Configured : 1
Always-On Monitoring is up.
```

Getting Status & Logs

Use the `emsctl status` command to ascertain the operational status of Always-On Monitoring.

Example:

```
% $AOM_HOME/scripts/emsctl status

Oracle Enterprise Manager Cloud Control 13c Release 5
Copyright (c) 2017, 2021 Oracle Corporation. All rights reserved.
------------------------------------------------------------------
Pinging Always-On Monitoring...
Always-On Monitoring is not running.
Starting Always-On Monitoring.
Waiting for process to start
Retrying...
Notifications Enabled : false
Total Downtime Contacts Configured : 1
Always-On Monitoring is up.
```

Pinging Always-On Monitoring

In addition to running the emsctl status command, you can also issue the ping command if you just want to see that Always-On Monitoring service is up without listing all the operational details.

Example:

```
$AOM_HOME/scripts/emsctl ping

Oracle Enterprise Manager Cloud Control 13c Release 5
Copyright (c) 2017, 2021 Oracle Corporation. All rights reserved.
------------------------------------------------------------------
Always-On Monitoring is running
```

Log Files

Log files record Always-On Monitoring events that occur during operation and are generated as follows:

- **emsca logs**: emsca.err (only errors), emsca.log.0 (rotating log file that contains all output including errors).
- **ems logs**: ems.err (only errors), ems.log.0 (rotating log file that contains all output including errors).

These files are located in the `$AOM_HOME/logs` directory.

Log levels determine the type and operational severity of information recorded in the log files. Levels can be set to any one of the following:

- **INFO**: Always-On Monitoring operational events such and login, logout, or any other events tied to Always-On Monitoring lifecycle.
- **DEBUG**: Information considered useful when tracing the flow of Always-On Monitoring events to isolate specific problems.
- **WARN** (default setting): Unexpected operational Always-On Monitoring events. These are Always-On Monitoring events that should be tracked, but may not require immediate administrator intervention.
- **ERROR**: Always-On Monitoring operational malfunction.

To change the log level, add the `logLevel=... property to $AOM_HOME/conf/emsConfig.properties`. Note that this applies only to the current Always-On Monitoring instance.

> **Note:**
>
> You must bounce Always-On Monitoring in order for the log level changes to take affect.

Stopping the Service

Use the `emsctl stop` command to stop Always-On Monitoring.

```
$AOM_HOME/scripts/emsctl stop
Oracle Enterprise Manager Cloud Control 13c Release 5
Copyright (c) 2017, 2021 Oracle Corporation. All rights reserved.
------------------------------------------------------------------
Shutting down Always-On Monitoring.
Always-On Monitoring is stopped.
```

Always-On Monitoring Commands

The following EMSCTL commands are used to control Always-On Monitoring.

Command	What it does
emsctl start	Starts Always-On Monitoring.
emsctl stop	Shuts down Always-On Monitoring.

Command	What it does
emsctl enable_notification -nosync (optional)	Allow Always-On Monitoring to send email notifications. By default, an incremental SYNC is performed when enabling email notifications. To prevent synchronization with the Enterprise Manager repository, specify the -nosync option.
emsctl disable_notification	Prevent Always-On Monitoring from sending email notifications.
emsctl list_agents	Lists the URLs of Agents and a count of Agents that have communicated with the Always-On Monitoring Service in the past hour.
emsctl ping	Ping the Always-On Monitoring service to determine whether it is up or down.
emsctl sync	Synchronize Always-On Monitoring target and notification data with the Enterprise Manager repository.
emsctl secure [-wallet=<absolute path to the wallet location>] | [-reg_pwd=<Agent registration password>] [-host=<Fully qualified Always-On Monitoring hostname>] [-cert_validity=<validity period in days - defaults to 10 years] [-key_strength=<certificate key strength - defaults to 1024] [sign_alg=<md5|sha1|sha256|sha384|sha512 - defaults to sha512>	Generates a wallet file for Always-On Monitoring service to run on a secured port.
emsctl status	Display information related to Always-On Monitoring operational status.
emsctl getstats	Display the current Always-On Monitoring performance statistics
emsctl list_properties	Display Always-On Monitoring configuration properties.
emsctl set_property -name=<property name> -value=<property value>	Set Always-On Monitoring configuration properties. For a complete list of configuration properties, see Modifiable Always-On Monitoring Properties
emsctl unset_property -name=<property name>	Unset an existing Always-On Monitoring configuration property. For a complete list of configuration properties, see Modifiable Always-On Monitoring Properties.
emsctl update_task -task_name=<SYNC|PURGE> -run_interval=<time in minutes between task executions>	Update the intervals for performing Always-On Monitoring tasks such as SYNC or PURGE.
emsctl set_em_repos_conn -username=<em username> -password=<em password> -connect_string=<em connect descriptor>	Update the username, password and connect string used to connect to Enterprise Manager repository database.
emsctl set_ems_repos_conn -username=<ems username> -password=<ems password> -connect_string=<ems connect descriptor>	Update username, password and connect string used to connect to the Always-On Monitoring repository database.

Updating Always-On Monitoring

In order to carry out the notification function of Enterprise Manager, Always-On Monitoring requires target monitoring and administrator notification configuration data to be synchronized with the Enterprise Manager repository. Full synchronization is performed when you install Always-On Monitoring and run the following command:

```
$AOM_HOME/scripts/emsctl sync
```

Incremental synchronization must be performed thereafter in order to keep Always-On Monitoring monitoring/notification configuration current.

The following changes to Enterprise Manager will require an incremental Always-On Monitoring synchronization:

- New hosts (Agents) added
- New targets added
- Changes to downtime contacts
- Changes to administrator email addresses
- Changes to email server configuration

Always-On Monitoring stays in SYNC with Enterprise Manager automatically via an incremental synchronization job that runs every 24 hours. You can change the synchronization interval using the emsctl update_task command.

```
$AOM_HOME/scripts/emsctl update_task -task_name=<SYNC|PURGE> -run_interval=<time in
minutes between task executions>
```

Incremental synchronization can be run manually via the *emsctl sync* command.

Running an Incremental Synchronization Manually

You can, at any time, perform an incremental Always-On Monitoring synchronization by running the *sync* command.

```
% $AOM_HOME/scripts/emsctl sync
```

When performing a manual incremental synchronization, you do not need to shut down Always-On Monitoring: You can run the SYNC command regardless of whether Always-On Monitoring is up or down.

> **Note:**
>
> An incremental synchronization is automatically performed whenever you run the *enable_notification* command. You can bypass the synchronization by specifying the *-nosync* option.

Data Maintenance

The Always-On Monitoring repository contains historical availability and metric violations data. This data is maintained for 6 months (default interval). Data older than 6 months is deleted (purged) from the repository.

Data maintenance is controlled by following settings:

- **Purge Interval**—Determines how often the data is purged from Always-On Monitoring repository.
- **Purge Duration**—Determines how much data is kept during the purge process.

You can change the purge interval via the *update_task* command:

```
$AOM_HOME/scripts/emsctl update_task -task_name=PURGE -run_interval=<the number of
minutes between purge runs>
```

You can change the purge duration via the *set_property* command:

```
$AOM_HOME/scripts/emsctl set_property -name=purge.Duration -value=<the number of minutes
of data to keep>
```

Controlling Always-On Monitoring Configuration Settings

Always-On Monitoring configuration settings define how your Always-On Monitoring deployment operates within your environment.

Properties that are specific to a particular instance of Always-On Monitoring can be found at the following location:

```
$AOM_HOME/conf/emsConfig.properties
```

Properties are added/updated in the `emsConfig.properties` file using a text editor. Always-On Monitoring must be restarted in order for the changes to take effect. Global properties that are shared by all Always-On Monitoring instances are stored in the Always-On Monitoring repository and are manipulated using the `emsctl` commands described below.

The following commands are used to view and set the global configuration properties:

To list all properties and their values:

```
$AOM_HOME/scripts/emsctl list_properties
```

To set a property value:

```
$AOM_HOME/scripts/emsctl set_property -name=<propertyname> -value=<propertyvalue>
```

To remove a property value:

```
$AOM_HOME/scripts/emsctl unset_property -name=<propertyname>
```

Getting Performance Information

You can see how well your Always-On Monitoring installation is operating by viewing Always-On Monitoring performance information such as load and throughput. Use the getstats command to display Always-On Monitoring performance statistics:

```
$AOM_HOME/scripts/emsctl getstats
```

Modifiable Always-On Monitoring Properties

Always-On Monitoring allows you to change various properties that define how your Always-On Monitoring deployment operates. The following table lists the global properties that are available. These properties can be set as global properties or locally in the `emsConfig.properties` file..

Table 17-4 Always-On Monitoring Global Properties

Property Name	Description	Default Value	Range of Values
connPool.initialSize	The initial number of Always-On Monitoring repository connections to be created in the repository connection pool.	1	1-connPool.maxSize

Table 17-4 (Cont.) Always-On Monitoring Global Properties

Property Name	Description	Default Value	Range of Values
connPool.maxSize	The maximum number of Always-On Monitoring repository connections to be created in the repository connection pool.	100	connPool.minSize – 500
connPool.minSize	The minimum number of Always-On Monitoring repository connections to be kept in the repository connection pool after growth.	10	connPool.initialSize – connPool.maxSize
failover.heartbeatInterval	The interval in seconds between heartbeats sent from the running Always-On Monitoring to other Always-On Monitoring instances.	60	30 – 3600
http.idleTimeout	Time in milliseconds before idle web server connections in the Always-On Monitoring are released and terminated.	60000	10000 – 120000
http.maxThreads	The maximum number of Always-On Monitoring web server threads to be created in the web server connection pool. This value is of particular interest on large sites with 10000 hosts or more.	50	http.minThreads – 100
http.minThreads	The minimum number of Always-On Monitoring web server threads to be created in the web server connection pool.	10	1 – http.maxThreads
nls.useSystemLocale	Reserved for future use.	false	
notif.enabled	If email notifications are enabled or not; typically sent using `emsctl enable_notification`	true	true/false
notif.senderAddress	Override sender email address for emails sent from Always-On Monitoring. Otherwise the address is the same as the Enterprise Manager email address.	null	
notif.senderName	Override sender email username for e-mails sent from Always-On Monitoring. Otherwise name is the same as the Enterprise Manager email name.	null	
purge.Duration	Amount of data (in minutes) kept in all historical Always-On Monitoring tables.	259200	10080 – 0 (infinite)

Table 17-4 (Cont.) Always-On Monitoring Global Properties

Property Name	Description	Default Value	Range of Values
sync.allowIncrSync	Reserved for future use.	false	true/false
tasks.poolSize	Number of threads to use in the background task execution pool.	20	1-50
tasks.purgeRunInterval	Interval in minutes between instances of the data purge task.	10080	3600 – 10080
tasks.syncRunInterval	Interval in minutes between instances of the incremental sync task.	1440	60 - 3600

The following table lists the Always-On Monitoring local properties that can only be defined locally in the `emsConfig.properties` file.

Table 17-5 Always-On Monitoring Local Properties

Property Name	Description	Default Value	Range of Values
logLevel	The level of log messages to include in the Always-On Monitoring logs	INFO	DEBUG, INFO, WARN, ERROR
emsPort	The port number to use when listening for requests from the agent.	8081	1024-65535
oracle.sysman.core.notification.smtp.retry_enabled	Property that determines whether failed email notifications should be retried.	false	true/false
oracle.sysman.core.notification.max_email_delivery_time	Property in seconds that determines how long failed email notifications should be retried.	600 seconds	60-86400

Creating an SSO Wallet and JKS for CA Certificates

Creating an SSO Wallet and JKS for CA Certificates

You can create a single sign-on (SSO) wallet and Java Keystore (JKS) for certificates issued by an external Certificate Authority for use with your Always-On Monitoring environment. To do so, perform the following:

1. Copy the EM key from the Enterprise Manager repository by running:

   ```
   emctl config emkey -copy_to_repos
   ```

2. Generate the certificate and the wallet for the host by running:

   ```
   $AOM_HOME/bin/emsctl secure
   ```

If a certificate is issued by a third-party Certificate Authority, then you need to store the certificate in the wallet. Once the wallet has the certificate, you must run the following:

```
$AOM_HOME/bin/emsctl secure -wallet=<absolute location of the wallet file>
```

Diagnosing Problems

Though robust and easy to install and configure, you may encounter problems while using Always-On Monitoring. In general, Always-On Monitoring problems can be classified into three types:

- Problems starting Always-On Monitoring
- Problems with Always-On Monitoring-Enterprise Manager repository SYNC.
- Not receiving email

By performing diagnostic tasks listed below, you can resolve a majority of issues that may occur when using Always-On Monitoring.

1. Check the log files in the $AOM_HOME/log sub-directory
2. Verify that the URL is set in Enterprise Manager by running the following

   ```
   emctl get property -name "oracle.sysman.core.events.ems.emsURL" -sysman_pwd <sysman-password>
   ```

3. Verify that the URL is set on the Management Agent by checking the $AGENT_HOME/sysman/config/emd.properties file and searching for EMS_URL
4. Verify that Always-On Monitoring is running and notifications are enabled by running emsctl status or emsctl ping
5. Verify that downtime contacts have been specified in Enterprise Manager and Always-On Monitoring by running emsctl status.
6. Check the version of the Agent by running the following command:

   ```
   emctl status agent
   ```

   The `emctl` command can be run form the following directory:

   ```
   $AGENT_HOME/agent_inst/bin
   ```

High Availability and Disaster Recovery

The purpose of Always-On Monitoring is to make sure your Enterprise Manager environment continues to be monitored while Enterprise Manager OMS is down. However, to protect against a single point of failure within the Always-On Monitoring environment itself, the Always-On Monitoring service can be made highly available by adhering to the Enterprise Manager Maximum Availability Architecture (MAA) guidelines. By setting up Always-On Monitoring as a highly available service, you can ensure that targets will continue to be monitored in the event of catastrophic system failure.

Running Multiple Always-On Monitoring Instances

Always-On Monitoring provides you with the ability to continue alert monitoring while Enterprise Manager is down. By itself, Always-On Monitoring can be made highly available by implementing multiple instances for load sharing/high availability. It is also important to note that the Always-On Monitoring repository should be configured as a cluster database since you need to implement basic High Availability (HA) for this component as well.

Adding multiple Always-On Monitoring instances provides the following HA advantages:

- Shares the work of receiving and recording alerts from agents. If an instance goes down for some reason, a server load balancer (SLB) can redirect those requests to surviving instances
- Shares the work of sending of notifications. For ordering reasons, all notifications for a target are directed to one Always-On Monitoring instance. If that Always-On Monitoring instance goes down, the responsibility for its "queues" are transferred to another instance.

Once you have configured an Always-On Monitoring instance, you can add another instance by running the following emsca command:

```
emsca add_ems
```

The following diagram shows a typical Always-On Monitoring HA deployment.

Figure 17-1 Always-On Monitoring HA Deployment

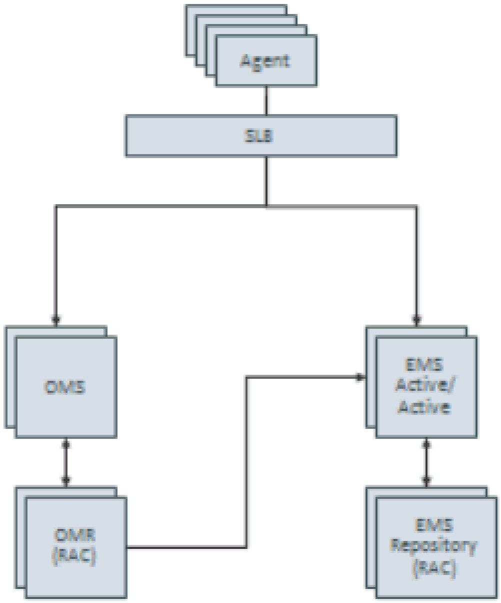

Shared Configuration Storage for the Multiple Instances

Instead of maintaining Always-On Monitoring configuration in the file system on each node, all instances can share a common configuration. This allows a single configuration at one location to be modified and read from all the instances.

You use the following EMSCTL commands to modify properties:

- `emsctl set_property`
- `emsctl unset_property`

To obtain a list of current property values:

- `emsctl list_properties`

When a property is modified on one instance, all other Always-On Monitoring instances receive a JMS AQ message that tells them to refresh the value from the repository.

Notification Queues for Tracking Incoming Alerts

Each Always-On Monitoring instance that receives an alert stores the alert and records an entry in one of 16 event notification queues. Ownership of these notifications queues is distributed among the running Always-On Monitoring instances.

A failover system notices when Always-On Monitoring instances are not up (and heartbeating) and switches ownership away from the instance that is down to other Always-On Monitoring instances that are up. The failover system then sends a JMS AQ message notifying all up instances of the change.

Task Scheduler System

A Task scheduler system remembers shared responsibilities among the Always-On Monitoring instances and is responsible for reminding the instances to do the work on a schedule. The task request is sent to an Always-On Monitoring instance that is known to be up. This reduces the possibility of multiple instances trying to do the same work and running into conflicts.

Configuring an SLB

When multiple Always-On Monitoring instances are present, we expect that traffic to these instances is directed through a SLB.

When HA is added to an Always-On Monitoring instance that was configured as a single instance, you need to repopulate the Always-On Monitoring instance's certificates with fresh copies that have the SLB host name in them. For Always-On Monitoring released with Enterprise Manager 13.3, you generate a certificate for the SLB and store that in a wallet by running the `emsctl secure` command.

1. Copy the EM key from the Enterprise Manager repository by running the following command:

   ```
   emctl config emkey -copy_to_repos
   ```

2. Run `emsctl secure`.
3. Restart the Always-On Monitoring instance.
4. Repeat for all Always-On Monitoring instances already configured.
5. Make sure to run emctl set property for the emsUrl property to use the newly configured Always-On Monitoring upload URL set up on the SLB. For example: `./emctl set property -name "oracle.sysman.core.events.ems.emsURL" -value "https://<slb hostname>:<slb aom port>/upload"`

The URL displayed from emsca and/or emsctl status is the local URL for that Always-On Monitoring instance. In an HA configuration, this local URL is NOT used as the Always-On Monitoring URL, but is used to set up the Always-On Monitoring pool in the SLB. The host and port for each Always-On Monitoring instance is used in the pool with the /upload PUT URL as the target for that pool.

When configuring the SLB, the monitoring URL is the /upload GET URL. This URL returns the string "Always On Monitoring is active."

Note:

Always-on monitoring is supported on F5 from version 11.4.x onwards.

Always-On Monitoring Disaster Recovery

Disaster Recovery (DR) allows Always-On Monitoring functionality to transition to a different location in the case of a site-wide failure.

There are existing procedures that allow an Enterprise Manager site to have a standby created for DR reasons. These procedures involve having the primary repository database's contents mirrored at the standby site using DataGuard, and the ability to link to the file system contents from the secondary site.

A very similar implementation can be used for Always-On Monitoring instances as well. An Always-On Monitoring DR setup is active-passive: If the primary site goes down, the virtual IP fails over to the standby node and, once the standby Always-On Monitoring instance is brought up, it runs exactly as it did on the primary node (because of the file system sharing and the virtual IP). In summary, the Always-On Monitoring repository needs to be available at the standby site, and the Always-On Monitoring instances' file systems need to be replicated. The following diagram shows the typical DR deployment for Always-On Monitoring.

Figure 17-2 Always-On Monitoring Disaster Recovery Deployment

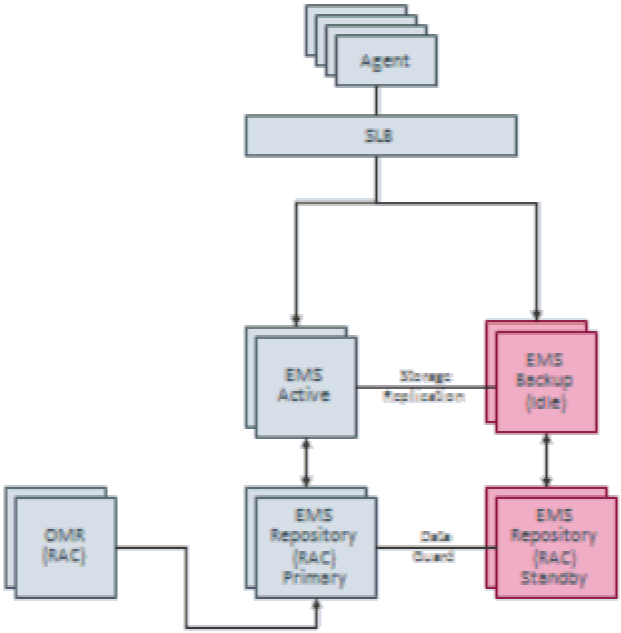

Setting Up Multiple Always-On Monitoring Instances

You can set up multiple Always-On Monitoring (AOM) instances to ensure Always-On Monitoring is available in the event of outages.

1. Uninstall any 13.1 Always-On Monitoring instances, if any.

   Except for the last running AOM instance (in this example, there are three AOM instances AOM1, AOM2 and AOM3) you must perform the following for AOM1 and AOM2):

   - Shut down the instances (AOM1 and AOM2) by running `$AOM_HOME/scripts/emsctl stop`
   - Remove all the files and sub-directories under $AOM_HOME

   And on the last AOM instance (AOM3):

   - Shut down the instance (AOM3) by running `$AOM_HOME/scripts/emsctl stop`
   - Uninstall AOM by running `$AOM_HOME/scripts/emsca uninstall`
   - Remove all files and sub-directories under $AOM_HOME
   - To remove any trace of the previous AOM installation, you can drop the AOM repository user in the repository by executing drop user ems cascade.
2. Obtain the latest Always-On Monitoring release. The ZIP file can be found on the OMS' `/sysman/ems` directory. Copy the *ems_13.5.x.x.x.zip* file to a location where you want AOM to be installed and unzip the file.

   ```
   unzip ems.zip
   ```

   > **Note:**
   >
   > For AOM 13.5 to run, both the Management Agent and Enterprise Manager installation MUST be upgraded to 13.5 as well.
3. Install AOM.

   For setting up the very first AOM instance, the user must run `$AOM_HOME/scripts/emsca`

   For installing subsequent AOM instances the user must run `$AOM_HOME/scripts/emsca add_ems`
4. Finally, you must run the following steps on all instances before starting AOM.

   - To generate the certificate and the wallet for the SLB host, run `$AOM_HOME/scripts/emsctl secure`
   - Configure the downtime contacts. For more information, see Configuring Downtime Contacts in Enterprise Manager.
   - To sync the AOM repository schema with the Enterprise Manager repository schema, execute `$AOM_HOME/scripts/emsctl sync`
   - Set the upload URL, by running `emctl config emkey -copy_to_repos` on the Enterprise Manager host.
   - Start the AOM instances by running, `$AOM_HOME/scripts/emsctl start`

Uninstalling Always-On Monitoring

If the purpose of uninstalling Always-On Monitoring is to upgrade to a new release, you should back up the following files.

- `$AOM_HOME/conf/emsConfig.properties`

- Custom certificates (if any) located at `$AOM_HOME/conf` (cwallet.sso, cwallet.sso.lck, ewallet.p12, and ewallet.p12.lck)

Use *emsca* to uninstall Always-On Monitoring as shown in the following example:

```
$ ./emsca uninstall
Oracle Enterprise Manager Cloud Control 13c Release 5
Copyright (c) 1996, 2021 Oracle Corporation. All rights reserved.
---------------------------------------------------------------
Loading configuration from: /homedirectory/ems/conf/emsConfig.properties
 Connecting to Always-On Monitoring repository.
 Execute the following:
 emctl delete property -name "oracle.sysman.core.events.ems.emsURL" -
sysman_pwd <pwd> to unset the value of the property.
 Uninstall completed.
```

Remove all files and sub-directories under $AOM_HOME.

To remove any trace of the previous AOM installation, you can drop the AOM repository user in the repository by running the following command:

```
drop user ems cascade
```

Configuring the Always-On Monitoring Application for Secure Communication Using the TLSv1.2 Protocol

Transport Layer Security (TLS) is a cryptographic protocol used to increase security over computer networks by providing communication privacy and data integrity between applications. In the case of Always-On Monitoring (AOM), these *secure* communication channels are between the following components:

- The AOM application and the AOM repository
- The AOM application and the Enterprise Manager repository

The following instructions cover how to enable TLSv1.2 communication for Always-On Monitoring.

Storing CA Certificates

The server CA certificates of the Always-On Monitoring repository and Enterprise Manager repository can be stored either in an external Trust Store or in the Oracle Management Service's JDK Trust Store `(JAVA_HOME/jre/lib/security/cacerts`).

Storing the Certificates in an External Trust Store

If you choose to store the certificates in an external Trust Store, then you need to set the following environment variables:

- **AOM_DB_WALLET_LOC -** Absolute path to the external Trust Store. For example: `/home/aom/externalTrustSTore.jks`
- **AOM_DB_WALLET_TYPE** - The type of Trust Store being used. JKS. PKCS12 (For Enterprise Manager 13.3, the SSO Trust Store is not supported)
- **AOM_DB_WALLET_PASSWORD -** The Trust Store password (For JKS and PKCS12)

> **Note:**
>
> In CSH, to set the environment variable, use the *setenv* command. For example, to set the AOM_DB_WALLET_LOC environment variable, run the following:
>
> ```
> % setenv AOM_DB_WALLET_LOC /home/aom/externalTrustStore.jks
> ```
>
> To set the same environment variable in a Bash environment, run the following:
>
> ```
> % export AOM_DB_WALLET_LOC=/home/aom/externalTrustStore.jks
> ```

Storing the Certificates in the Oracle Management Service's JDK Trust Store

If you choose to store the certificates in the Oracle Management Service's JDK Trust Store, then you can use the *emsca* or *emsctl* scripts without additional configuration.

Guidelines for Configuring Always-On Monitoring to use TLSv1.2

- During initial Always-On Monitoring setup with *emsca*, when prompted for the DB connection string, you must specify the connection string using the *long form* that has the protocol type (TCPS) being used and not the *short form* as shown in the following examples.

  Long Form

  ```
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)(HOST=myserver.myco.com)
  (PORT=15044)) (CONNECT_DATA =(SID=dbview))
  ```

  Short Form (will not work with TCPS):

  ```
  myserver.myco.com:15044:dbview
  ```

- If Always-On Monitoring was initially configured to use the TCP protocol as part of *emsca* configuration, and at a later point you want to switch over to TCPS, you can re-configure Always-On Monitoring by performing the following steps:

  1. Change the Always-On Monitoring database connection string in the `emsConfig.properties` file (`$AOM_HOME/conf/emsConfig.properties`).
  2. Change Enterprise Manager Repository connection string in Always-On Monitoring database which is stored in table - EMS_SYNC_CONNECT_PROPS.connect_string.
  3. Ensure that either the JDK wallet (under `JAVA_HOME/jre/lib/security/cacerts`) or the external Trust Store has the root CA certificates for both the Always-On Monitoring and Enterprise Manager repositories.

- If Always-On Monitoring was initially configured to use the TCPS protocol as part of *emsca* configuration, and at a later point you want to switch over to TCP, you can re-configure Always-On Monitoring by performing the following steps:

  1. In the `AOM_HOME/conf/emsConfig.properties` file, replace the correct AOM database connection string with the property *emsRepConnectString*. If TCP protocol is used to connect to the database, you can use either short or long form for the connection string. If TCPS is used, you must provide the database connection string in the long form as discussed previously.

2. To change the Enterprise Manager connection string, you must update the `EMS_SYNC_CONNECT_PROPS.connect_string` field in the Always-On Monitoring database.

18 Zero Downtime Monitoring

The Zero Downtime (ZDT) Monitoring Service is a new service in Enterprise Manager 24ai. It is responsible for processing and handling all event management, incident management and notification capabilities in Enterprise Manager. Because it is a ZDT Service, this means it will continue to be operational even during Enterprise Manager 24ai planned maintenance events such as updating to a new Enterprise Manager 24ai Release Update (RU).

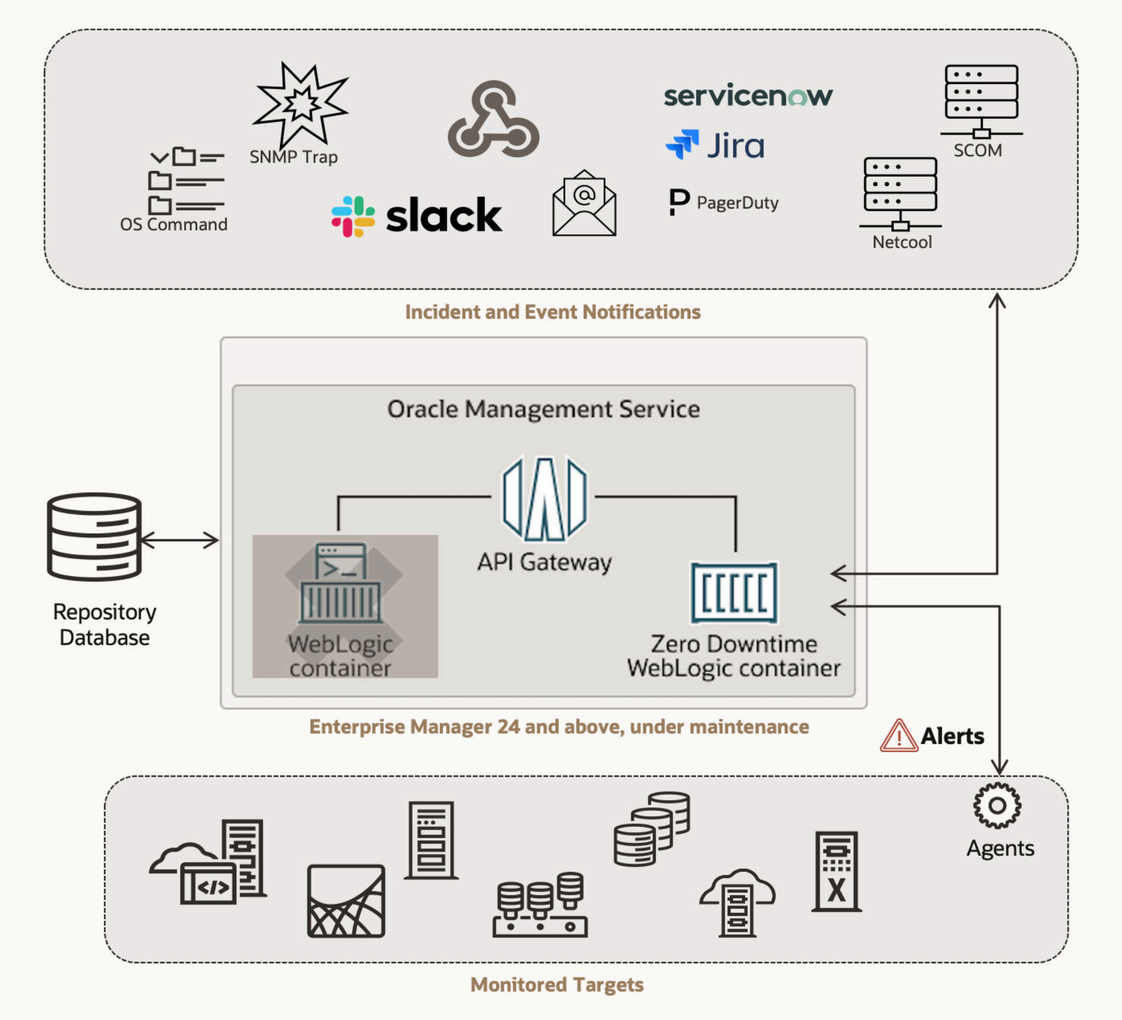
SNMP Trap
OS Command
slack
servicenow
Jira
PagerDuty
SCOM
Netcool
Incident and Event Notifications
Oracle Management Service
API Gateway
Repository Database
WebLogic container
Zero Downtime WebLogic container
Enterprise Manager 24 and above, under maintenance
Alerts
Agents
Monitored Targets

With the Zero Downtime Monitoring Service, you will have uninterrupted monitoring, alerting and notification capabilities even when Enterprise Manager 24ai is under planned maintenance for a new release update.

The features provided by Zero Downtime Monitoring Service include:

- Receiving of alerts from agents

- All actions specified in Incident Rule Sets for the targets and groups specified in these Incident Rule Sets. Actions include creation of incidents, event compression, setting of Incident properties, sending of notifications.
- Event compression using Event Compression Policies or event compression rules specified in Incident Rule Sets
- Sending of notifications: Email, SNMP traps, OS Command, Webhooks, Slack
- Opening of tickets using Ticketing connectors
- Sending of events to external systems using Event Connectors
- Priority processing of events based on the targets' Lifecycle Status property

Targets that are in blackout will continue to have monitoring suspended until the blackout ends. After the blackout has ended, monitoring for the target resumes.

> **Note:**
>
> When upgrading from Enterprise Manager 13.5 to Enterprise Manager 24ai, the Zero Downtime Monitoring Service will not be available. It is only available once you are using Enterprise Manager 24ai. To provide monitoring visibility during the upgrade from Enterprise Manager 13.5 to Enterprise Manager 24ai, use Always-On Monitoring to send email notifications for critical alerts that may occur during the upgrade.

Part IV

Extensibility: SNMP Support

This section contains the following chapters:

Introducing Enterprise Manager Support for SNMP

This chapter provides a brief overview of Enterprise Manager support for SNMP. It includes the following sections:

The Simple Network Management Protocol (SNMP) is a protocol used for managing or monitoring devices, where many of these devices are network-type devices such as routers, switches, and so on. SNMP enables a single application to first retrieve information, then push new information between a wide range of systems independent of the underlying hardware.

Designed primarily for database, network, and system administrators, SNMP support integrates Enterprise Manager into a number of existing, widely-used management systems. Also, Enterprise Manager can extend its monitoring scope to devices that can be monitored using SNMP.

Benefits of SNMP Support

The primary benefits of SNMP support include the following:

- The monitoring of key Oracle products is quickly integrated into any management framework based upon SNMP.
- These Oracle products are located, identified, and monitored in real time across enterprise networks of any size.
- Administrators see standard Oracle icons that represent Oracle products in a network map. You can dynamically customize this map.

- Administrators see the current status of Oracle products, as shown by several status variables that are defined for each product in a management information base (MIB), or they can select which elements to view by their status.
- Administrators can anticipate exceptional conditions by defining thresholds and alerts, to respond to special situations as soon as they occur or to enable automatic responses.
- Administrators can store and analyze historical data that has been obtained through SNMP.
- Providers of management applications can easily build customized solutions for Oracle customers because SNMP is an open standard.

Strictly speaking, SNMP support is intended more for monitoring Oracle products than for managing them. SNMP support is invaluable for tracking the status of an entire network of Oracle applications — first, to verify normal operations, and second, to spot and react to potential problems as soon as they are detected. However, for purposes of investigating and solving some problems, other Oracle tools such as Oracle SQL *Plus Worksheet may be more appropriate. This is because SNMP support is designed to query status, but not to change system parameters, whereas other tools are designed to set or tune system parameters.

> **Note:**
>
> Oracle SNMP is not supported on HP OpenVMS platform.

About the SNMP Management Station

The SNMP management station refers to a node from which managed elements are monitored using the SNMP protocol. Typically, it is a standalone workstation that is on the same network as the managed elements. While this book will consistently use the term SNMP management station, other terms used for it include management console, management system, or managing node.

Because most frameworks use SNMP as a basis for communication, Oracle products that support SNMP can be integrated into virtually every management framework. Third-party products such as CA Unicenter, HP OpenView, Tivoli NetView, Aprisma Spectrum, Sun Solstice, and Castle Rock SNMPc Network Manager provide SNMP Management Station functionality.

How Enterprise Manager Supports SNMP

Enterprise Manager supports SNMP by integrating with third-party management systems, sharing event or incident information through SNMP traps and extending the monitoring scope of Enterprise Manager by monitoring new devices and targets using SNMP.

There are number of ways that Enterprise Manager uses SNMP as illustrated in Figure 1:

1. Sharing event or incident information with third-party management systems by generating SNMP trap notifications from Enterprise Manager to an SNMP management station. For example, you can use SNMP to notify a third-party application that a selected metric has exceeded its threshold.

   This method supports SNMP version 1 (SNMPv1) and SNMP version 3 (SNMPv3).

   For more information, see Sending SNMP Trap Notifications and Using Notifications .

2. Extending the monitoring scope of Enterprise Manager to new entities by receiving SMNP traps or fetching SNMP data from or to the managed entity. By developing a metadata plug-in to receive SNMP traps, you can enable Enterprise Manager to monitor a product capable of throwing SNMP traps.

   This method supports SNMP version 1 (SNMPv1), SNMP version 2 (SNMPv2c), and SNMP version 3 (SNMPv3).

   For more information, see Monitoring External Devices Using SNMP and the *Oracle Enterprise Manager Extensibility Programmer's Reference*. .

3. Creating new metrics to query SNMP agents by using SNMP adapters to allow Management Agents to query native SNMP agents on host targets for Management Information Base (MIB) variable information to be used as metric data.

   This method supports SNMP version 1 (SNMPv1), SNMP version 2 (SNMPv2c), and SNMP version 3 (SNMPv3).

   For more information, see About Metric Extensions and Using Metric Extensions .

Figure 1 How Enterprise Manager Supports SNMP

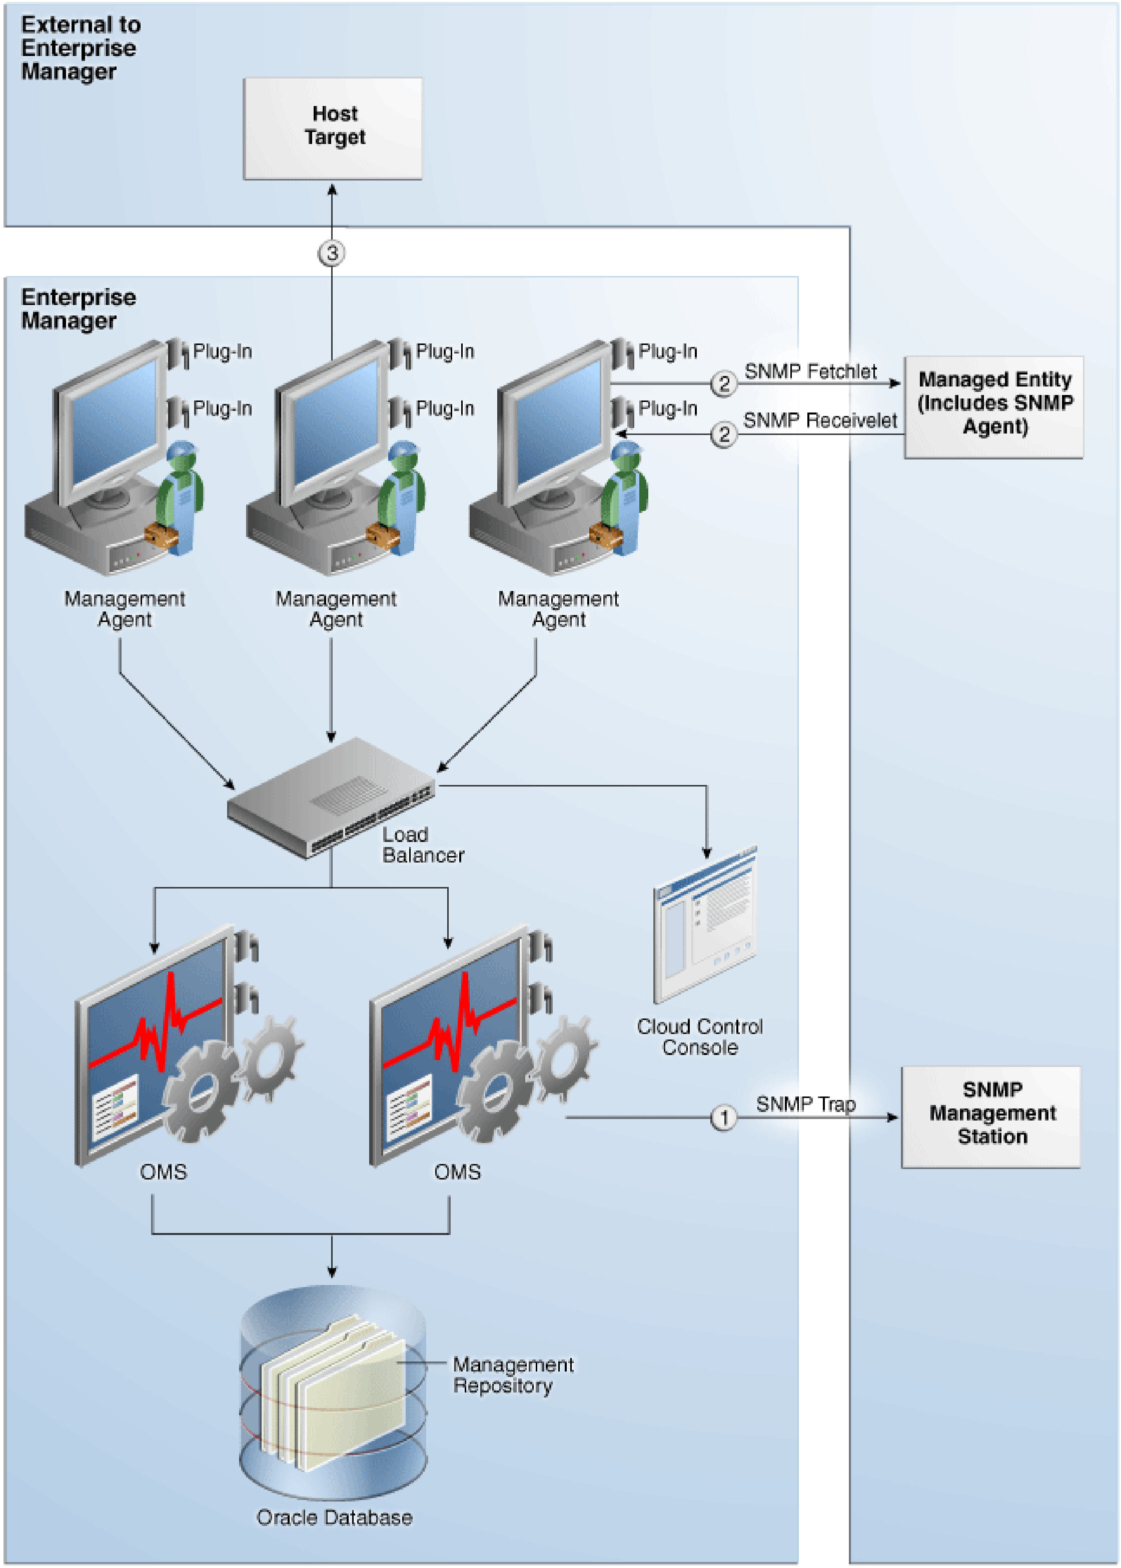

Sending SNMP Trap Notifications

Using the Enterprise Manager notification system, you can share Enterprise Manager event or incident information with other SNMP-enabled third-party applications through SNMP traps. Enterprise Manager supports the SNMPv1 protocol for sending traps. For example, you might want to send event information as traps to a third-party applications from Enterprise Manager when one of the following events takes place:

- a certain metric has exceeded a threshold (metric alert event)
- a target is down (target availability event)
- a job fails (job status change event)

> **Note:**
>
> For a full list and description of Enterprise Manager event types, see Event Management.

Using SNMP traps with the notification system is a matter of:

1. Defining a notification method that uses an SNMP trap. For more information, see Sending SNMP Traps to Third Party Systems.
2. Assigning the notification method to a rule. You can edit an existing rule or create a new incident rule. For more information, see Creating a Rule to Send SNMP Traps for Events to Third Party Systems or Creating a Rule to Send SNMP Traps for Incidents to Third Party Systems.

About the Management Information Base (MIB)

While SNMP allows Enterprise Manager to send information to third-party SNMP-enabled applications, there might be situations where you want SNMP-enabled applications to obtain information from Enterprise Manager. This is accomplished with the help of MIB variables, and by signing up for SNMP traps. Details of the trap contents can be obtained from the MIB variables.

For more information about the MIB, see Management Information Base (MIB) and Interpreting Variables of the Enterprise Manager MIB.

> **Note:**
>
> A valid Diagnostic Pack license is required to use the Enterprise Manager MIB variables.

Monitoring External Devices Using SNMP

It is often critical for an administrator to receive alerts from applications that are not managed by Enterprise Manager. Many of these applications can be configured to trigger SNMP traps when an alert condition takes place. You can receive these traps within Enterprise Manager and start monitoring those applications from Enterprise Manager. Having the capability to receive and analyze SNMP traps raised by such applications allows you extend Enterprise Manager's monitoring and alerting capabilities to these applications and reduce monitoring complexity in your IT environments.

You can configure Enterprise Manager to receive the SNMP traps raised by an application (not managed by Enterprise Manager) and display the traps as alerts in Enterprise Manager.

To receive these traps, you must develop a metadata plug-in to represent the managed entity. Then use an SNMP receivelet or SNMP fetchlet to receive or get monitoring data about that entity.

For more information about developing metadata plug-ins, see the *Oracle Enterprise Manager Cloud Control Extensibility Programmer's Reference*.

About SNMP Receivelets

While monitoring third-party entities in your managed environment, if the status of a third-party network element turns unavailable or if its metric severity conditions (metric thresholds) are met or exceeded, the SNMP Agent of that third-party network element sends a notification to the Management Agent. These notifications are in the form of SNMP traps that get triggered asynchronously upon reaching the performance thresholds, and without any requests from the Management Agent.

Since these traps are based on SNMP, the Management Agent uses SNMP Receivelets to receive and translate these SNMP traps into a form compatible with Oracle Management Service.

For more information about the SNMP receivelet, see the *Oracle Enterprise Manager Extensibility Programmer's Reference*.

About SNMP Fetchlets

Fetchlets are parameterized data access mechanisms available to map relevant data from a managed element into Enterprise Manager's metric format. In the standards area, Enterprise Manager currently uses SNMP Fetchlets to fetch information from SNMP-enabled entities within your managed environment.

The SNMP fetchlet queries the SNMP Agent for data about the managed entity as defined in the target type's Management Information Base (MIB). For more information about the MIB, see About the Management Information Base (MIB) and Management Information Base (MIB).

For more information about the SNMP fetchlet, see the *Oracle Enterprise Manager Extensibility Programmer's Reference*.

About Metric Extensions

Metric extensions provide you with the ability to extend Oracle's monitoring capabilities to monitor conditions specific to your IT environment. For example, you can create a new metric for a host target type. Use an SNMP Adapter to allow Enterprise Manager Management Agents to query native SNMP agents on target hosts for Management Information Base (MIB) variable information to be used as the metric data.

For more information about Metric Extensions and the SNMP Adapter, see Using Metric Extensions .

Interpreting Variables of the Enterprise Manager 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. This appendix provides information to help you interpret the variables of the private Oracle Enterprise Manager MIB.

For information about the format of MIB variable descriptions, see Reading the MIB Variable Descriptions.

This appendix contains the following sections:

- oraEMNGEvent
- oraEM4AlertTable
- oraEM4JobAlertTable

Note:

SNMP trap notification methods created from Cloud Control 12*c* and later send the oraEMNGEvent trap.

SNMP trap notification methods created before Enterprise Manager Release 12*c* send oraEM4Alert and oraEM4JobAlert traps. After upgrading to Enterprise Manager Release 12*c* or later, existing methods from earlier releases continue to deliver the old traps for backward compatibility.

oraEMNGEvent

The following sections describe the SNMP traps sent from SNMP trap notification methods created from Cloud Control 12*c* and later.

Table 1 oraEMNGEvent Variables and Corresponding Object IDs

Variable Name	Object ID
oraEMNGEventIndex	1.3.6.1.4.1.111.15.3.1.1.1.1
oraEMNGEventNotifType	1.3.6.1.4.1.111.15.3.1.1.2.1
oraEMNGEventMessage	1.3.6.1.4.1.111.15.3.1.1.3.1
oraEMNGEventMessageURL	1.3.6.1.4.1.111.15.3.1.1.4.1
oraEMNGEventSeverity	1.3.6.1.4.1.111.15.3.1.1.5.1
oraEMNGEventSeverityCode	1.3.6.1.4.1.111.15.3.1.1.6.1
oraEMNGEventRepeatCount	1.3.6.1.4.1.111.15.3.1.1.7.1
oraEMNGEventActionMsg	1.3.6.1.4.1.111.15.3.1.1.8.1
oraEMNGEventOccurrenceTime	1.3.6.1.4.1.111.15.3.1.1.9.1
oraEMNGEventReportedTime	1.3.6.1.4.1.111.15.3.1.1.10.1

Table 1   (Cont.) oraEMNGEvent Variables and Corresponding Object IDs

Variable Name	Object ID
oraEMNGEventCategories	1.3.6.1.4.1.111.15.3.1.1.11.1
oraEMNGEventCategoryCodes	1.3.6.1.4.1.111.15.3.1.1.12.1
oraEMNGEventType	1.3.6.1.4.1.111.15.3.1.1.13.1
oraEMNGEventName	1.3.6.1.4.1.111.15.3.1.1.14.1
oraEMNGAssocIncidentId	1.3.6.1.4.1.111.15.3.1.1.15.1
oraEMNGAssocIncidentOwner	1.3.6.1.4.1.111.15.3.1.1.16.1
oraEMNGAssocIncidentAcked	1.3.6.1.4.1.111.15.3.1.1.17.1
oraEMNGAssocIncidentStatus	1.3.6.1.4.1.111.15.3.1.1.18.1
oraEMNGAssocIncidentPriority	1.3.6.1.4.1.111.15.3.1.1.19.1
oraEMNGAssocIncidentEscLevel	1.3.6.1.4.1.111.15.3.1.1.20.1
oraEMNGEventTargetName	1.3.6.1.4.1.111.15.3.1.1.21.1
oraEMNGEventTargetNameURL	1.3.6.1.4.1.111.15.3.1.1.22.1
oraEMNGEventTargetType	1.3.6.1.4.1.111.15.3.1.1.23.1
oraEMNGEventHostName	1.3.6.1.4.1.111.15.3.1.1.24.1
oraEMNGEventTargetOwner	1.3.6.1.4.1.111.15.3.1.1.25.1
oraEMNGEventTgtLifeCycleStatus	1.3.6.1.4.1.111.15.3.1.1.26.1
oraEMNGEventTargetVersion	1.3.6.1.4.1.111.15.3.1.1.27.1
oraEMNGEventUserDefinedTgtProp	1.3.6.1.4.1.111.15.3.1.1.28.1
oraEMNGEventSourceObjName	1.3.6.1.4.1.111.15.3.1.1.29.1
oraEMNGEventSourceObjNameURL	1.3.6.1.4.1.111.15.3.1.1.30.1
oraEMNGEventSourceObjType	1.3.6.1.4.1.111.15.3.1.1.31.1
oraEMNGEventSourceObjSubType	1.3.6.1.4.1.111.15.3.1.1.32.1
oraEMNGEventSourceObjOwner	1.3.6.1.4.1.111.15.3.1.1.33.1
oraEMNGEventCAJobName	1.3.6.1.4.1.111.15.3.1.1.34.1
oraEMNGEventCAJobStatus	1.3.6.1.4.1.111.15.3.1.1.35.1
oraEMNGEventCAJobOwner	1.3.6.1.4.1.111.15.3.1.1.36.1
oraEMNGEventCAJobStepOutput	1.3.6.1.4.1.111.15.3.1.1.37.1
oraEMNGEventCAJobType	1.3.6.1.4.1.111.15.3.1.1.38.1
oraEMNGEventRuleSetName	1.3.6.1.4.1.111.15.3.1.1.39.1
oraEMNGEventRuleName	1.3.6.1.4.1.111.15.3.1.1.40.1
oraEMNGEventRuleOwner	1.3.6.1.4.1.111.15.3.1.1.41.1
oraEMNGEventSequenceId	1.3.6.1.4.1.111.15.3.1.1.42.1
oraEMNGEventRCADetails	1.3.6.1.4.1.111.15.3.1.1.43.1
oraEMNGEventContextAttrs	1.3.6.1.4.1.111.15.3.1.1.44.1
oraEMNGEventUserComments	1.3.6.1.4.1.111.15.3.1.1.45.1
oraEMNGEventUpdates	1.3.6.1.4.1.111.15.3.1.1.46.1
oraEMNGEventTotalOccurrenceCount	1.3.6.1.4.1.111.15.3.1.1.47.1

Table 1 (Cont.) oraEMNGEvent Variables and Corresponding Object IDs

Variable Name	Object ID
oraEMNGEventCurrOccurrenceCount	1.3.6.1.4.1.111.15.3.1.1.48.1
oraEMNGEventCurrFirstOccurDate	1.3.6.1.4.1.111.15.3.1.1.49.1
oraEMNGEventCurrLastOccurDate	1.3.6.1.4.1.111.15.3.1.1.50.1
oraEMNGEventRCAStatus	1.3.6.1.4.1.111.15.3.1.1.51.1
oraEMNGEventReportedState	1.3.6.1.4.1.111.15.3.1.1.52.1
oraEMNGEventTypeAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1
oraEMNGEventTypeAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1
oraEMNGEventTypeAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1
oraEMNGEventTypeAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1
oraEMNGEventTypeAttr5	1.3.6.1.4.1.111.15.3.1.1.65.1
oraEMNGEventTypeAttr6	1.3.6.1.4.1.111.15.3.1.1.66.1
oraEMNGEventTypeAttr7	1.3.6.1.4.1.111.15.3.1.1.67.1
oraEMNGEventTypeAttr8	1.3.6.1.4.1.111.15.3.1.1.68.1
oraEMNGEventTypeAttr9	1.3.6.1.4.1.111.15.3.1.1.69.1
oraEMNGEventTypeAttr10	1.3.6.1.4.1.111.15.3.1.1.70.1
oraEMNGEventTypeAttr11	1.3.6.1.4.1.111.15.3.1.1.71.1
oraEMNGEventTypeAttr12	1.3.6.1.4.1.111.15.3.1.1.72.1
oraEMNGEventTypeAttr13	1.3.6.1.4.1.111.15.3.1.1.73.1
oraEMNGEventTypeAttr14	1.3.6.1.4.1.111.15.3.1.1.74.1
oraEMNGEventTypeAttr15	1.3.6.1.4.1.111.15.3.1.1.75.1
oraEMNGEventTypeAttr16	1.3.6.1.4.1.111.15.3.1.1.76.1
oraEMNGEventTypeAttr17	1.3.6.1.4.1.111.15.3.1.1.77.1
oraEMNGEventTypeAttr18	1.3.6.1.4.1.111.15.3.1.1.78.1
oraEMNGEventTypeAttr19	1.3.6.1.4.1.111.15.3.1.1.79.1
oraEMNGEventTypeAttr20	1.3.6.1.4.1.111.15.3.1.1.80.1
oraEMNGEventTypeAttr21	1.3.6.1.4.1.111.15.3.1.1.81.1
oraEMNGEventTypeAttr22	1.3.6.1.4.1.111.15.3.1.1.82.1
oraEMNGEventTypeAttr23	1.3.6.1.4.1.111.15.3.1.1.83.1
oraEMNGEventTypeAttr24	1.3.6.1.4.1.111.15.3.1.1.84.1
oraEMNGEventTypeAttr25	1.3.6.1.4.1.111.15.3.1.1.85.1

oraEMNGEventIndex

Syntax

Integer

Access

Read-only

Status

Mandatory

Description

The index of a particular event, unique only at the moment an event is generated.

oraEMNGEventNotifType

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The notification type. Possible values are:

- NOTIF_NORMAL
- NOTIF_RETRY
- NOTIF_DURATION
- NOTIF_REPEAT
- NOTIF_CA
- NOTIF_RCA

oraEMNGEventMessage

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The message associated with this event.

oraEMNGEventMessageURL

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The Enterprise Manager console URL for the event message. It is populated for events with severities other than INFORMATIONAL. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventSeverity

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The severity of the event, such as Fatal, Critical, Warning, Advisory, Information, or Clear.

oraEMNGEventSeverityCode

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The internal code of the severity, such as Fatal, Critical, Warning, Advisory, Informational, or Clear.

oraEMNGEventRepeatCount

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The repeat notification counter for the event.

oraEMNGEventActionMsg

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The action message for this event.

oraEMNGEventOccurrenceTime

Syntax

DisplayString

Access

read-only

Status

Mandatory

Explanation

The time when this event occurred (optional). This is only populated for events that have an occurrence time.

oraEMNGEventReportedTime

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The time when this event was reported.

oraEMNGEventCategories

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The list of categories to which this event belongs. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventCategoryCodes

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The list of internal category codes to which this event belongs. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventType

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of the event type to which this event belongs.

Available Event Types

- metric_alert
- target_availability
- job_status_change
- metric_error
- user_reported
- cs_core
- sla_alert
- mext_update
- selfupdate
- cs_rule_violation

oraEMNGEventName

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of this event.

oraEMNGAssocIncidentId

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The ID of the associated incident with the event (optional).

oraEMNGAssocIncidentOwner

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

Owner of the associated incident with the event (optional).

oraEMNGAssocIncidentAcked

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

Acknowledged status of the associated incident with the event. 1 indicates acknowledged. 0 indicates unacknowledged.

oraEMNGAssocIncidentStatus

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The status of the associated incident with the event.

oraEMNGAssocIncidentPriority

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The priority of the associated incident with the event.

oraEMNGAssocIncidentEscLevel

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The escalation level of the associated incident with the event.

oraEMNGEventTargetName

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of the target to which this event applies. Populated for events that are about a target only.

oraEMNGEventTargetNameURL

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The Enterprise Manager Console URL of the target to which this event applies. Populated for events that are about a target only. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventTargetType

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The type of the target to which this event applies. Populated for events that are about a target only.

oraEMNGEventHostName

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of the host on which this event originated. Populated for events that are about a target only.

oraEMNGEventTargetOwner

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The primary administrator of the target on which this event originated. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventTgtLifeCycleStatus

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The life cycle status of the target on which this event originated.

oraEMNGEventTargetVersion

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The version of the target on which this event originated.

oraEMNGEventUserDefinedTgtProp

Syntax

DisplayString

Access

read-only

Status

Mandatory

Explanation

The user defined target properties [name,value pair list] of the associated target with this event. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventSourceObjName

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of the source object to which this event belongs to. Populated for events that are about a non-target object only, such as Jobs.

oraEMNGEventSourceObjNameURL

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

Enterprise Manager Console URL for the source object to which this event belongs. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventSourceObjType

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The type of the source object to which this event belongs.

oraEMNGEventSourceObjSubType

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The subtype of the source object to which this event belongs. (Optional). This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventSourceObjOwner

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The primary administrator of the source object to which this event belongs. (Optional). This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventCAJobName

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of the corrective action job associated with this event.

oraEMNGEventCAJobStatus

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The status of the corrective action job associated with this event.

oraEMNGEventCAJobOwner

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The owner of the corrective action job associated with this event.

oraEMNGEventCAJobStepOutput

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The job step output from the corrective action job associated with this event.

oraEMNGEventCAJobType

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The job type from the corrective action job associated with this event.

oraEMNGEventRuleSetName

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of the ruleset that caused this notification. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventRuleName

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of the rule within the ruleset that caused this notification.

oraEMNGEventRuleOwner

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The owner of the ruleset that caused this notification.

oraEMNGEventSequenceId

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

An Enterprise Manager-generated identifier that uniquely identifies the current issue until it is cleared.

oraEMNGEventRCADetails

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

Root Cause Analysis information associated with this event if it exists.

oraEMNGEventContextAttrs

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The context attributes associated with this event. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventUserComments

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The user comments associated with this event. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventUpdates

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The updates associated with this event. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventTotalOccurrenceCount

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

Total number of occurrences of the same event on a target across all open deduplicated events. This attribute applies only to deduplicated events.

oraEMNGEventCurrOccurrenceCount

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

Total occurrences of the event in this collection period. This attribute applies only to deduplicated events.

oraEMNGEventCurrFirstOccurDate

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

Timestamp when the event first occurred in this collection period. This attribute applies only to deduplicated events.

oraEMNGEventCurrLastOccurDate

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

Timestamp when the event last occurred in this collection period. This attribute applies only to deduplicated events.

oraEMNGRCAStatus

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

Summary of Root Cause Analysis, if applicable.

oraEMNGEventReportedState

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

This is an optional value reporting the latest state of an entity and is only applicable for events that are representing a state transition for a specified entity. For example, for Target availability and Job state change events, this value would be the latest state of the target or job, respectively.

oraEMNGEventTypeAttr(1-71)

The following tables list oraEMNGEventType MIB variables 1 through 71. Each table categorizes the MIB variables by specific event type.

Table 2 Metric Alert Event Type

Variable Name	OID Number	Event Type Attribute	Description
oraEMNGEventTypeAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Metric GUID	A unique ID for the metric.
oraEMNGEventTypeAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Severity GUID	A unique ID for the alert record.
oraEMNGEventTypeAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1	Cycle GUID	A unique ID for the alert cycle.
oraEMNGEventTypeAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1	Collection Name	The name of the collection collecting the metric.
oraEMNGEventTypeAttr5	1.3.6.1.4.1.111.15.3.1.1.65.1	Metric Group	The name of the metric.
oraEMNGEventTypeAttr6	1.3.6.1.4.1.111.15.3.1.1.66.1	Metric	The name of the metric column.
oraEMNGEventTypeAttr7	1.3.6.1.4.1.111.15.3.1.1.67.1	Metric Description	A brief description of the metric.
oraEMNGEventTypeAttr8	1.3.6.1.4.1.111.15.3.1.1.68.1	Metric value	The value of the metric when the event triggered.
oraEMNGEventTypeAttr9	1.3.6.1.4.1.111.15.3.1.1.69.1	Key Value	The monitored object for the metric corresponding to the Metric Alert event.
oraEMNGEventTypeAttr10	1.3.6.1.4.1.111.15.3.1.1.70.1	Key Column 1	Key Column 1
oraEMNGEventTypeAttr11	1.3.6.1.4.1.111.15.3.1.1.71.1	Key Column 1 Value	The value of Key Column 1.
oraEMNGEventTypeAttr12	1.3.6.1.4.1.111.15.3.1.1.72.1	Key Column 2	Key Column 2
oraEMNGEventTypeAttr13	1.3.6.1.4.1.111.15.3.1.1.73.1	Key Column 2 Value	The value of Key Column 2.
oraEMNGEventTypeAttr14	1.3.6.1.4.1.111.15.3.1.1.74.1	Key Column 3	Key Column 3
oraEMNGEventTypeAttr15	1.3.6.1.4.1.111.15.3.1.1.75.1	Key Column 3 Value	The value of Key Column 3.
oraEMNGEventTypeAttr16	1.3.6.1.4.1.111.15.3.1.1.76.1	Key Column 4	Key Column 4
oraEMNGEventTypeAttr17	1.3.6.1.4.1.111.15.3.1.1.77.1	Key Column 4 Value	The value of Key Column 4.
oraEMNGEventTypeAttr18	1.3.6.1.4.1.111.15.3.1.1.78.1	Key Column 5	Key Column 5
oraEMNGEventTypeAttr19	1.3.6.1.4.1.111.15.3.1.1.79.1	Key Column 5 Value	The value of Key Column 5.
oraEMNGEventTypeAttr20	1.3.6.1.4.1.111.15.3.1.1.80.1	Key Column 6	Key Column 6
oraEMNGEventTypeAttr21	1.3.6.1.4.1.111.15.3.1.1.81.1	Key Column 6 Value	The value of Key Column 6.
oraEMNGEventTypeAttr22	1.3.6.1.4.1.111.15.3.1.1.82.1	Key Column 7	Key Column 7

Table 2 (Cont.) Metric Alert Event Type

Variable Name	OID Number	Event Type Attribute	Description
oraEMNGEventTypeAttr23	1.3.6.1.4.1.111.15.3.1.1.83.1	Key Column 7 Value	The value of Key Column 7.
oraEMNGEventTypeAttr24	1.3.6.1.4.1.111.15.3.1.1.84.1	Number of keys	The number of key metric columns in the metric.

Table 3 Target Availability Event Type

Variable Name	OID Number	Event Type Attribute	Description
oraEMNGEventTypeAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Availability status	The current availability status of the target.
oraEMNGEventTypeAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Severity GUID	The GUID of the severity record associated with this availability status.
oraEMNGEventTypeAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1	Availability Sub-status	The sub-status of a target for the current status.
oraEMNGEventTypeAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1	Transition Severity	The severity that resulted in the target's status change to the current availability status.
oraEMNGEventTypeAttr5	1.3.6.1.4.1.111.15.3.1.1.65.1	Response metric GUID	The Metric GUID of response metric.
oraEMNGEventTypeAttr6	1.3.6.1.4.1.111.15.3.1.1.66.1	Severity GUID of the first severity in the availability cycle	The GUID of the first severity record in this availability cycle.

Table 4 Job Status Change Event Type

Variable Name	OID Number	Event Type Attribute	Description
oraEMNGEventTypeAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Execution ID	The unique ID of the job execution.
oraEMNGEventTypeAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Job Status	The status of the job execution.
oraEMNGEventTypeAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1	Execution Log	The job output of the last step executed.
oraEMNGEventTypeAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1	Job Status Code	The execution status code of job execution.
oraEMNGEventTypeAttr5	1.3.6.1.4.1.111.15.3.1.1.65.1	State Change ID	The unique ID of the last status change.

Table 5   Compliance Standard Rule Violation Event Type

Variable Name	OID Number	Event Type Attribute	Description
oraEMNGEventTypeAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Root Compliance Standard	The root compliance standard node display name.
oraEMNGEventTypeAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Root Compliance Standard Version	The root compliance standard version.
oraEMNGEventTypeAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1	Root Compliance Standard Author	The author of the root compliance standard.
oraEMNGEventTypeAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1	Parent Compliance Standard	The parent compliance standard node display name.
oraEMNGEventTypeAttr5	1.3.6.1.4.1.111.15.3.1.1.65.1	Compliance Standard Version	The compliance standard version.
oraEMNGEventTypeAttr6	1.3.6.1.4.1.111.15.3.1.1.66.1	Parent Compliance Standard Author	The author of a parent compliance standard.
oraEMNGEventTypeAttr7	1.3.6.1.4.1.111.15.3.1.1.67.1	Root Target Name	The root target name.
oraEMNGEventTypeAttr8	1.3.6.1.4.1.111.15.3.1.1.68.1	Root Target Type	The root target type.
oraEMNGEventTypeAttr9	1.3.6.1.4.1.111.15.3.1.1.69.1	Rule Name	The rule display name

Table 6   Compliance Standard Score Event Type

Variable Name	OID Number	Event Type Attribute	Description
oraEMNGEventTypeAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Root Compliance Standard	The root compliance standard node display name.
oraEMNGEventTypeAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Root Compliance Standard Author	The author of the root compliance standard.
oraEMNGEventTypeAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1	Root Compliance Standard Version	The version of the root compliance standard.
oraEMNGEventTypeAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1	Compliance Standard	The compliance standard node display name.
oraEMNGEventTypeAttr5	1.3.6.1.4.1.111.15.3.1.1.65.1	Compliance Standard Version	The version of a compliance standard.
oraEMNGEventTypeAttr6	1.3.6.1.4.1.111.15.3.1.1.66.1	Compliance Standard Author	The author of a compliance standard.
oraEMNGEventTypeAttr7	1.3.6.1.4.1.111.15.3.1.1.67.1	Root Target Name	The root target name.
oraEMNGEventTypeAttr8	1.3.6.1.4.1.111.15.3.1.1.68.1	Root Target Type	The root target type.
oraEMNGEventTypeAttr10	1.3.6.1.4.1.111.15.3.1.1.70.1	Warning Threshold	The warning threshold of a compliance score.
oraEMNGEventTypeAttr9	1.3.6.1.4.1.111.15.3.1.1.69.1	Compliance Score	The compliance score.
oraEMNGEventTypeAttr11	1.3.6.1.4.1.111.15.3.1.1.71.1	Critical Threshold	The critical threshold of a compliance score.

Table 7 Metric Error Event Type

Variable Name	OID Number	Event Type Attribute	Description
oraEMNGEventTypeAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Metric Group	The name of the metric.
oraEMNGEventTypeAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Collection Name	The name of the collection collecting the metric.

Table 8 Metric Extension Event Type

Variable Name	OID Number	Event Type Attribute	Description
oraEMNGEventTypeAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Metric Extension Version attribute	The version of the metric extension.

Table 9 Self-update Event Type

Variable Name	OID Number	Event Type Attribute	Description
oraEMNGEventTypeAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Type	Type
oraEMNGEventTypeAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Description	Description
oraEMNGEventTypeAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1	Version	Version
oraEMNGEventTypeAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1	Status	Status

Table 10 Service Level Agreement Alert Event Type

Variable Name	OID Number	Event Type Attribute	Description
oraEMNGEventTypeAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Service Level Agreement Name	Service Level Agreement Name
oraEMNGEventTypeAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Service Level Objective Name	Service Level Objective Name
oraEMNGEventTypeAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1	Service Level Objective Type	The type of the Service Level Objective which will be either Performance or Availability
oraEMNGEventTypeAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1	Value at Event Triggered	The value at Event Triggered
oraEMNGEventTypeAttr5	1.3.6.1.4.1.111.15.3.1.1.65.1	Customer Name	Customer Name

Table 11 User-reported Event Type

Variable Name	OID Number	Event Type Attribute	Description
oraEMNGEventTypeAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Name	The name describing the nature of the issue.
oraEMNGEventTypeAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	key	The optional key describing a sub-component within the target that this event is about.

oraEM4AlertTable

The oraEM4AlertTable describes the SNMP traps sent from Oracle Enterprise Manager for both metric severity alerts and policy violations.

Table 12 lists the variables of the oraEM4AlertTable and their corresponding Object IDs.

Table 12   oraEM4AlertTable Variables and Corresponding Object IDs

Variable Name	Object ID
oraEM4AlertTargetName	1.3.6.1.4.1.111.15.1.1.1.2.1
oraEM4AlertTargetType	1.3.6.1.4.1.111.15.1.1.1.3.1
oraEM4AlertHostName	1.3.6.1.4.1.111.15.1.1.1.4.1
oraEM4AlertMetricName	1.3.6.1.4.1.111.15.1.1.1.5.1
oraEM4AlertKeyName	1.3.6.1.4.1.111.15.1.1.1.6.1
oraEM4AlertKeyValue	1.3.6.1.4.1.111.15.1.1.1.7.1
oraEM4AlertTimeStamp	1.3.6.1.4.1.111.15.1.1.1.8.1
oraEM4AlertSeverity	1.3.6.1.4.1.111.15.1.1.1.9.1
oraEM4AlertMessage	1.3.6.1.4.1.111.15.1.1.1.10.1
oraEM4AlertRuleName	1.3.6.1.4.1.111.15.1.1.1.11.1
oraEM4AlertRuleOwner	1.3.6.1.4.1.111.15.1.1.1.12.1
oraEM4AlertMetricValue	1.3.6.1.4.1.111.15.1.1.1.13.1
oraEM4AlertContext	1.3.6.1.4.1.111.15.1.1.1.14.1
oraEM4AlertCycleGuid	1.3.6.1.4.1.111.15.1.1.1.15.1
oraEM4AlertRepeatCount	1.3.6.1.4.1.111.15.1.1.1.16.1
oraEM4AlertUDTargetProperties	1.3.6.1.4.1.111.15.1.1.1.17.1
oraEM4AlertAck	1.3.6.1.4.1.111.15.1.1.1.18.1
oraEM4AlertAckBy	1.3.6.1.4.1.111.15.1.1.1.19.1
oraEM4AlertNotifType	1.3.6.1.4.1.111.15.1.1.1.20.1
oraEM4AlertViolationGuid	1.3.6.1.4.1.111.15.1.1.1.21.1

A description of each of these variables follows.

oraEM4AlertTargetName

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the target to which this alert applies.

Typical Range

Not applicable

Significance

Very important

oraEM4AlertTargetType

Syntax

DisplayString

Max-Access

read-only

Status

Mandatory

Explanation

The type of the target to which this alert applies.

Typical Range

Not applicable

Significance

Very important

oraEM4AlertHostName

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the host on which this alert originated.

Typical Range

Not applicable

Significance

Very important

oraEM4AlertMetricName

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the metric or policy which generated this alert.

Typical Range

Not applicable

Significance

Very important

oraEM4AlertKeyName

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the key-column, if present, for the metric which generated this alert.

Typical Range

Not applicable

Significance

Very important

oraEM4AlertKeyValue

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The value of the key-column, if present, for the metric which generated this alert.

Typical Range

Not applicable

Significance

Very important

oraEM4AlertTimeStamp

Syntax

DisplayString

Max-Access

read-only

Status

Mandatory

Explanation

The time at which this alert was generated.

Typical Range

Not applicable

Significance

Important

oraEM4AlertSeverity

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The severity of the alert (for example, Clear, Informational, Warning, Critical, Unreachable Start, Blackout End, Blackout Start, Metric Error Clear, Metric Error Start, Status Pending).

Typical Range

Critical, warning, clear

Significance

Very important

oraEM4AlertMessage

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The message associated with the alert.

Typical Range

Not applicable

Significance

Very important

oraEM4AlertRuleName

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the notification rule that caused this notification.

Typical Range

Not applicable

Significance

Important

oraEM4AlertRuleOwner

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The owner of the notification rule that caused this notification.

Typical Range

Not applicable

Significance

Important

oraEM4AlertMetricValue

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The value of the metric which caused this alert to be generated.

Typical Range

Not applicable

Significance

Important

oraEM4AlertContext

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

A comma separated list of metric column names and values associated with the metric that caused this alert to be generated.

Typical Range

Not applicable

Significance

Important

oraEM4AlertCycleGuid

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

An Enterprise Manager-generated identifier that is unique for the lifecycle of an alert.

Typical Range

Not applicable

Significance

Important

oraEM4AlertRepeatCount

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The repeat notification counter for the alert.

Typical Range

Not applicable

Significance

Important

oraEM4AlertUDTargetProperties

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

User-defined target properties associated with the target.

Typical Range

Not applicable

Significance

Important

oraEM4AlertAck

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

Acknowledged status flag associated with the alert. 1 indicates acknowledged. 0 indicates unacknowledged.

Typical Range

Not applicable

Significance

Important

oraEM4AlertAckBy

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

Acknowledged By value associated with the alert.

Typical Range

Not applicable

Significance

Important

oraEM4AlertNotifType

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

Notification type.

Possible values:

- 1 - Normal
- 4 - Repeat
- 9 - Duration

Typical Range

Not applicable

Significance

Important

oraEM4AlertViolationGuid

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

An Enterprise Manager-generated identifier that identifies a particular alert.

Typical Range

Not applicable

Significance

Important

oraEM4JobAlertTable

The oraEM4JobAlertTable describes changes in the status of either a Job or a Corrective Action that is running as part of the Oracle Enterprise Manager Job system.

Table 13 lists the variables of the oraEM4JobAlertTable and their corresponding Object IDs.

Table 13   oraEM4JobAlertTable Variables and Corresponding Object IDs

Variable Name	Object ID
oraEM4JobAlertJobName	1.3.6.1.4.1.111.15.1.2.1.2.1
oraEM4JobAlertJobOwner	1.3.6.1.4.1.111.15.1.2.1.3.1
oraEM4JobAlertJobType	1.3.6.1.4.1.111.15.1.2.1.4.1
oraEM4JobAlertJobStatus	1.3.6.1.4.1.111.15.1.2.1.5.1
oraEM4JobAlertTargets	1.3.6.1.4.1.111.15.1.2.1.6.1
oraEM4JobAlertTimeStamp	1.3.6.1.4.1.111.15.1.2.1.7.1
oraEM4JobAlertRuleName	1.3.6.1.4.1.111.15.1.2.1.8.1
oraEM4JobAlertRuleOwner	1.3.6.1.4.1.111.15.1.2.1.9.1
oraEM4JobAlertMetricName	1.3.6.1.4.1.111.15.1.2.1.10.1
oraEM4JobAlertMetricValue	1.3.6.1.4.1.111.15.1.2.1.11.1
oraEM4JobAlertContext	1.3.6.1.4.1.111.15.1.2.1.12.1
oraEM4JobAlertKeyName	1.3.6.1.4.1.111.15.1.2.1.13.1
oraEM4JobAlertKeyValue	1.3.6.1.4.1.111.15.1.2.1.14.1
oraEM4JobAlertSeverity	1.3.6.1.4.1.111.15.1.2.1.15.1
oraEM4JobAlertJobId	1.3.6.1.4.1.111.15.1.2.1.16.1
oraEM4JobAlertJobExecId	1.3.6.1.4.1.111.15.1.2.1.17.1

A description of each of these variables follows.

oraEM4JobAlertJobName

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the job to which this alert applies.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertJobOwner

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The owner of the job to which this alert applies.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertJobType

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The type of the job to which this alert applies.

Typical Range

Not applicable

Significance

Important

oraEM4JobAlertJobStatus

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The status of the job to which this alert applies.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertTargets

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

A comma separated list of target to which this alert applies.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertTimeStamp

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The time at which this job status changed causing this alert.

Typical Range

Not applicable

Significance

Important

oraEM4JobAlertRuleName

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the notification rule that caused this notification.

Typical Range

Not applicable

Significance

Important

oraEM4JobAlertRuleOwner

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The owner of the notification rule that caused this notification.

Typical Range

Not applicable

Significance

Important

oraEM4JobAlertMetricName

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the metric or policy which caused the Corrective Action to run that caused this alert.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertMetricValue

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The value of the metric which caused the Corrective Action to run that caused this alert.

Typical Range

Not applicable

Significance

Important

oraEM4JobAlertContext

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

A comma separated list of metric column names and values associated with the metric which caused the Corrective Action to run that caused this alert.

Typical Range

Not applicable

Significance

Important

oraEM4JobAlertKeyName

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the key-column, if present, for the metric which caused the Corrective Action to run that generated this alert.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertKeyValue

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The value of the key-column, if present, for the metric which caused the Corrective Action to run that generated this alert.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertSeverity

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The severity of the metric which caused the Corrective Action to run that generated this alert (for example, Critical).

Typical Range

Critical, warning, clear

Significance

Very important

oraEM4JobAlertJobId

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The job ID of the Enterprise Manager job that triggered this notification.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertJobExecId

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The job execution ID of the Enterprise Manager job that triggered this notification.

Typical Range

Not applicable

Significance

Very important

Enterprise Manager MIB Definition

The following MIB definition is the latest version at the time of publication. For the most recent version of the Enterprise Manager 13*c* MIB definition, view your installation MIB definition file at:

OMS_HOME/network/doc/omstrap.v1

MIB Definition

```
ORACLE-ENTERPRISE-MANAGER-4-MIB DEFINITIONS ::= BEGIN

IMPORTS
    TRAP-TYPE
        FROM RFC-1215
    DisplayString
        FROM RFC1213-MIB
    OBJECT-TYPE
        FROM RFC-1212
    enterprises
        FROM RFC1155-SMI;

oracle OBJECT IDENTIFIER ::= { enterprises  111 }

oraEM4 OBJECT IDENTIFIER ::= { oracle  15 }

oraEM4Objects OBJECT IDENTIFIER ::= { oraEM4  1 }

oraEM4AlertTable OBJECT-TYPE
    SYNTAX  SEQUENCE OF OraEM4AlertEntry
    ACCESS  not-accessible
    STATUS  mandatory
    DESCRIPTION
```

```
     "Information on alerts generated by Oracle Enterprise Manager. This
table is not queryable; it exists only to document the variables included in
the oraEM4Alert trap.  Each trap contains a single instance of each variable
in the table."
    ::= { oraEM4Objects  1 }

oraEM4AlertEntry OBJECT-TYPE
    SYNTAX  OraEM4AlertEntry
    ACCESS  not-accessible
    STATUS  mandatory
    DESCRIPTION
     "Information about a particular Oracle Enterprise Manager alert."
    INDEX   { oraEM4AlertIndex }
    ::= { oraEM4AlertTable  1 }

OraEM4AlertEntry ::=
    SEQUENCE {
        oraEM4AlertIndex
            INTEGER,

        oraEM4AlertTargetName
       DisplayString,

        oraEM4AlertTargetType
       DisplayString,

        oraEM4AlertHostName
       DisplayString,

        oraEM4AlertMetricName
       DisplayString,

        oraEM4AlertKeyName
       DisplayString,

        oraEM4AlertKeyValue
       DisplayString,

        oraEM4AlertTimeStamp
       DisplayString,

        oraEM4AlertSeverity
       DisplayString,

        oraEM4AlertMessage
       DisplayString,

        oraEM4AlertRuleName
       DisplayString,

        oraEM4AlertRuleOwner
       DisplayString,

    oraEM4AlertMetricValue
           DisplayString,
        oraEM4AlertContext
           DisplayString,

    oraEM4AlertCycleGuid
           DisplayString,

    oraEM4AlertRepeatCount
           DisplayString,

    oraEM4AlertUDTargetProperties
           DisplayString,

    oraEM4AlertAck
           DisplayString,

    oraEM4AlertAckBy
           DisplayString,

    oraEM4AlertNotifType
           DisplayString,

    oraEM4AlertViolationGuid
           DisplayString
 }

oraEM4AlertIndex OBJECT-TYPE
    SYNTAX  INTEGER (0..2147483647)
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Index of a particular alert, unique only at the moment an alert is
generated."
    ::= { oraEM4AlertEntry  1 }

oraEM4AlertTargetName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the target to which this alert applies."
    ::= { oraEM4AlertEntry  2 }

oraEM4AlertTargetType OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The type of the target to which this alert applies."
    ::= { oraEM4AlertEntry  3 }

oraEM4AlertHostName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the host on which this alert originated."
    ::= { oraEM4AlertEntry  4 }

oraEM4AlertMetricName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the metric or policy which generated this alert."
    ::= { oraEM4AlertEntry  5 }

oraEM4AlertKeyName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the key-column, if present, for the metric which generated
this alert."
    ::= { oraEM4AlertEntry  6 }

oraEM4AlertKeyValue OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The value of the key-column, if present, for the metric which generated
this alert."
    ::= { oraEM4AlertEntry  7 }

oraEM4AlertTimeStamp OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The time at which this alert was generated."
    ::= { oraEM4AlertEntry  8 }

oraEM4AlertSeverity OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The severity of the alert e.g. Clear, Informational, Warning, Critical,
Unreachable Clear, Unreachable Start, Blackout End, Blackout Start, Metric
Error Clear, Metric Error Start, Status Pending."
    ::= { oraEM4AlertEntry  9 }

oraEM4AlertMessage OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The message associated with the alert."
    ::= { oraEM4AlertEntry  10 }

oraEM4AlertRuleName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the notification rule that caused this notification."
    ::= { oraEM4AlertEntry  11 }

oraEM4AlertRuleOwner OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The owner of the notification rule that caused this notification."
    ::= { oraEM4AlertEntry  12 }

oraEM4AlertMetricValue OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The value of the metric which caused this alert to be generated."
    ::= { oraEM4AlertEntry  13 }

oraEM4AlertContext OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "A comma separated list of metric column names and values associated
with the metric that caused this alert to be generated."
    ::= { oraEM4AlertEntry  14 }

oraEM4AlertCycleGuid OBJECT-TYPE
     SYNTAX  DisplayString
     ACCESS  read-only
     STATUS  mandatory
     DESCRIPTION
      "An EM generated identifier that is unique for the lifecyle of an
alert."
     ::= { oraEM4AlertEntry  15 }

oraEM4AlertRepeatCount OBJECT-TYPE
     SYNTAX  DisplayString
     ACCESS  read-only
     STATUS  mandatory
     DESCRIPTION
      "The repeat notification counter for the alert."
     ::= { oraEM4AlertEntry  16 }

oraEM4AlertUDTargetProperties OBJECT-TYPE
     SYNTAX  DisplayString
     ACCESS  read-only
     STATUS  mandatory
     DESCRIPTION
      "User-defined target properties associated with the target."
     ::= { oraEM4AlertEntry  17 }
```

```
oraEM4AlertAck OBJECT-TYPE
     SYNTAX  DisplayString
     ACCESS  read-only
     STATUS  mandatory
     DESCRIPTION
      "Acknowledged status flag associated with the alert. 1 indicates
acknowledged, 0 indicates unacknowledged."
     ::= { oraEM4AlertEntry  18 }

oraEM4AlertAckBy OBJECT-TYPE
     SYNTAX  DisplayString
     ACCESS  read-only
     STATUS  mandatory
     DESCRIPTION
      "Acknowledged By value  associated with the alert."
     ::= { oraEM4AlertEntry  19 }

oraEM4AlertNotifType OBJECT-TYPE
     SYNTAX  DisplayString
     ACCESS  read-only
     STATUS  mandatory
     DESCRIPTION
      "Notification Type. 1 - Normal, 4 - Repeat, 9 - Duration"
     ::= { oraEM4AlertEntry  20 }

oraEM4AlertViolationGuid OBJECT-TYPE
     SYNTAX  DisplayString
     ACCESS  read-only
     STATUS  mandatory
     DESCRIPTION
      "An EM generated identifier that identifies a particular alert."
     ::= { oraEM4AlertEntry  21 }

oraEM4Traps OBJECT IDENTIFIER ::= { oraEM4  2 }

oraEM4Alert TRAP-TYPE
    ENTERPRISE  oraEM4Traps
    VARIABLES   { oraEM4AlertTargetName, oraEM4AlertTargetType,
                  oraEM4AlertHostName, oraEM4AlertMetricName,
                  oraEM4AlertKeyName, oraEM4AlertKeyValue,
oraEM4AlertTimeStamp,
                  oraEM4AlertSeverity, oraEM4AlertMessage,
                  oraEM4AlertRuleName, oraEM4AlertRuleOwner,
                  oraEM4AlertMetricValue, oraEM4AlertContext,
oraEM4AlertCycleGuid,
                  oraEM4AlertRepeatCount,
                  oraEM4AlertUDTargetProperties, oraEM4AlertAck,
oraEM4AlertAckBy,
                  oraEM4AlertNotifType, oraEM4AlertViolationGuid }
    DESCRIPTION
     "The variables included in the oraEM4Alert trap."
    ::= 1


oraEM4JobAlertTable OBJECT-TYPE
    SYNTAX  SEQUENCE OF OraEM4JobAlertEntry
    ACCESS  not-accessible
    STATUS  mandatory
    DESCRIPTION
     "Information on alerts generated by Oracle Enterprise Manager. This
table is not queryable; it exists only to document the variables included in
the oraEM4JobAlert trap.  Each trap contains a single instance of each
variable in the table."
    ::= { oraEM4Objects  2 }

oraEM4JobAlertEntry OBJECT-TYPE
    SYNTAX  OraEM4JobAlertEntry
    ACCESS  not-accessible
    STATUS  mandatory
    DESCRIPTION
     "Information about a particular Oracle Enterprise Manager alert."
    INDEX   { oraEM4JobAlertIndex }
    ::= { oraEM4JobAlertTable  1 }

OraEM4JobAlertEntry ::=
    SEQUENCE {
        oraEM4JobAlertIndex
            INTEGER,

        oraEM4JobAlertJobName
       DisplayString,

        oraEM4JobAlertJobOwner
       DisplayString,

        oraEM4JobAlertJobType
       DisplayString,

        oraEM4JobAlertJobStatus
       DisplayString,

        oraEM4JobAlertTargets
       DisplayString,

        oraEM4JobAlertTimeStamp
       DisplayString,

        oraEM4JobAlertRuleName
       DisplayString,

        oraEM4JobAlertRuleOwner
       DisplayString,

    oraEM4JobAlertMetricName
           DisplayString,

    oraEM4JobAlertMetricValue
           DisplayString,

        oraEM4JobAlertContext
           DisplayString,
        oraEM4JobAlertKeyName
       DisplayString,

        oraEM4JobAlertKeyValue
       DisplayString,

        oraEM4JobAlertSeverity
       DisplayString,

        oraEM4JobAlertJobId
       DisplayString,

        oraEM4JobAlertJobExecId
       DisplayString
    }


oraEM4JobAlertIndex OBJECT-TYPE
    SYNTAX  INTEGER (0..2147483647)
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Index of a particular alert, unique only at the moment an alert is
generated."
    ::= { oraEM4JobAlertEntry  1 }

oraEM4JobAlertJobName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the job to which this alert applies."
    ::= { oraEM4JobAlertEntry  2 }

oraEM4JobAlertJobOwner OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The owner of the job to which this alert applies."
    ::= { oraEM4JobAlertEntry  3 }

oraEM4JobAlertJobType OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The type of the job to which this alert applies."
    ::= { oraEM4JobAlertEntry  4 }

oraEM4JobAlertJobStatus OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The status of the job to which this alert applies."
    ::= { oraEM4JobAlertEntry  5 }

oraEM4JobAlertTargets OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "A comma separated list of target to which this alert applies."
    ::= { oraEM4JobAlertEntry  6 }

oraEM4JobAlertTimeStamp OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The time at which this job status changed causing this alert."
    ::= { oraEM4JobAlertEntry  7 }

oraEM4JobAlertRuleName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the notification rule that caused this notification."
    ::= { oraEM4JobAlertEntry  8 }

oraEM4JobAlertRuleOwner OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The owner of the notification rule that caused this notification."
    ::= { oraEM4JobAlertEntry  9 }

oraEM4JobAlertMetricName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the metric or policy which caused the Corrective Action to
run that caused this alert."
    ::= { oraEM4JobAlertEntry  10 }

oraEM4JobAlertMetricValue OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The value of the metric which caused the Corrective Action to run that
caused this alert."
    ::= { oraEM4JobAlertEntry  11 }

oraEM4JobAlertContext OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "A comma separated list of metric column names and values associated
with the metric which caused the Corrective Action to run that caused this
alert."
    ::= { oraEM4JobAlertEntry  12 }

oraEM4JobAlertKeyName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the key-column, if present, for the metric which caused the
Corrective Action to run that generated this alert."
    ::= { oraEM4JobAlertEntry  13 }

oraEM4JobAlertKeyValue OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The value of the key-column, if present, for the metric which caused
the Corrective Action to run that generated this alert."
    ::= { oraEM4JobAlertEntry  14 }

oraEM4JobAlertSeverity OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The severity of the metric which caused the Corrective Action to run
that generated this alert e.g. Critical."
    ::= { oraEM4JobAlertEntry  15 }

oraEM4JobAlertJobId OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The Job Id of the EM Job that triggered this notification."
    ::= { oraEM4JobAlertEntry  16 }

oraEM4JobAlertJobExecId OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The Job Execution Id of the EM Job that triggered this notification."
    ::= { oraEM4JobAlertEntry  17 }

oraEM4JobAlert TRAP-TYPE
    ENTERPRISE  oraEM4Traps
    VARIABLES   { oraEM4JobAlertJobName, oraEM4JobAlertJobOwner,
                  oraEM4JobAlertJobType, oraEM4JobAlertJobStatus,
                  oraEM4JobAlertTargets, oraEM4JobAlertTimeStamp,
                  oraEM4JobAlertRuleName, oraEM4JobAlertRuleOwner,
                  oraEM4JobAlertMetricName, oraEM4JobAlertMetricValue,
                  oraEM4JobAlertContext, oraEM4JobAlertKeyName,
                  oraEM4JobAlertKeyValue, oraEM4JobAlertSeverity,
                  oraEM4JobAlertJobId, oraEM4JobAlertJobExecId }
    DESCRIPTION
     "The variables included in the oraEM4JobAlert trap."
    ::= 2

oraEMNGObjects OBJECT IDENTIFIER ::= { oraEM4  3 }

oraEMNGEventTable OBJECT-TYPE
    SYNTAX  SEQUENCE OF OraEMNGEventEntry
    ACCESS  not-accessible
    STATUS  mandatory
    DESCRIPTION
     "Information on events published to Oracle Enterprise Manager. This
table is not queryable; it exists only to document the variables included in
the oraEMNGEventTrap trap.  Each trap can contain a single instance of each
variable in the table."
    ::= { oraEMNGObjects  1 }

oraEMNGEventEntry OBJECT-TYPE
    SYNTAX  OraEMNGEventEntry
    ACCESS  not-accessible
    STATUS  mandatory
    DESCRIPTION
     "Information about a particular Oracle Enterprise Manager event."
    INDEX   { oraEMNGEventIndex }
    ::= { oraEMNGEventTable  1 }

OraEMNGEventEntry ::=
    SEQUENCE {
        oraEMNGEventIndex
            INTEGER,

        oraEMNGEventNotifType
            DisplayString,

        oraEMNGEventMessage
            DisplayString,

        oraEMNGEventMessageURL
            DisplayString,

        oraEMNGEventSeverity
            DisplayString,

        oraEMNGEventSeverityCode
            DisplayString,

        oraEMNGEventRepeatCount
            DisplayString,

        oraEMNGEventActionMsg
            DisplayString,

        oraEMNGEventOccurrenceTime
            DisplayString,

        oraEMNGEventReportedTime
            DisplayString,

        oraEMNGEventCategories
            DisplayString,

        oraEMNGEventCategoryCodes
            DisplayString,

        oraEMNGEventType
            DisplayString,

        oraEMNGEventName
            DisplayString,

        oraEMNGAssocIncidentId
            DisplayString,

        oraEMNGAssocIncidentOwner
            DisplayString,

        oraEMNGAssocIncidentAcked
            DisplayString,

        oraEMNGAssocIncidentStatus
            DisplayString,

        oraEMNGAssocIncidentPriority
            DisplayString,

        oraEMNGAssocIncidentEscLevel
            DisplayString,

        oraEMNGEventTargetName
            DisplayString,

        oraEMNGEventTargetNameURL
            DisplayString,

        oraEMNGEventTargetType
            DisplayString,

        oraEMNGEventHostName
            DisplayString,

        oraEMNGEventTargetOwner
            DisplayString,

        oraEMNGEventTgtLifeCycleStatus
            DisplayString,

        oraEMNGEventTargetVersion
            DisplayString,
        oraEMNGEventUserDefinedTgtProp
            DisplayString,

        oraEMNGEventSourceObjName
            DisplayString,

        oraEMNGEventSourceObjNameURL
            DisplayString,

        oraEMNGEventSourceObjType
            DisplayString,

        oraEMNGEventSourceObjSubType
            DisplayString,

        oraEMNGEventSourceObjOwner
            DisplayString,

        oraEMNGEventCAJobName
            DisplayString,

        oraEMNGEventCAJobStatus
            DisplayString,

        oraEMNGEventCAJobOwner
            DisplayString,

        oraEMNGEventCAJobStepOutput
            DisplayString,

        oraEMNGEventCAJobType
            DisplayString,

        oraEMNGEventRuleSetName
            DisplayString,

        oraEMNGEventRuleName
            DisplayString,

        oraEMNGEventRuleOwner
            DisplayString,

        oraEMNGEventSequenceId
            DisplayString,

        oraEMNGEventRCADetails
            DisplayString,

        oraEMNGEventContextAttrs
            DisplayString,

        oraEMNGEventUserComments
            DisplayString,

        oraEMNGEventUpdates
            DisplayString,
        oraEMNGEventTotalOccurrenceCount
            DisplayString,

        oraEMNGEventCurrOccurrenceCount
            DisplayString,

        oraEMNGEventCurrFirstOccurDate
            DisplayString,

        oraEMNGEventCurrLastOccurDate
            DisplayString,

        oraEMNGEventRCAStatus
            DisplayString,

        oraEMNGEventReportedState
            DisplayString,

        oraEMNGEventTypeAttr1
            DisplayString,

        oraEMNGEventTypeAttr2
            DisplayString,

        oraEMNGEventTypeAttr3
            DisplayString,

        oraEMNGEventTypeAttr4
            DisplayString,

        oraEMNGEventTypeAttr5
            DisplayString,

        oraEMNGEventTypeAttr6
            DisplayString,

        oraEMNGEventTypeAttr7
            DisplayString,

        oraEMNGEventTypeAttr8
            DisplayString,

        oraEMNGEventTypeAttr9
            DisplayString,

        oraEMNGEventTypeAttr10
            DisplayString,

        oraEMNGEventTypeAttr11
            DisplayString,

        oraEMNGEventTypeAttr12
            DisplayString,

        oraEMNGEventTypeAttr13
```

```
            DisplayString,

        oraEMNGEventTypeAttr14
            DisplayString,

        oraEMNGEventTypeAttr15
            DisplayString,

        oraEMNGEventTypeAttr16
            DisplayString,

        oraEMNGEventTypeAttr17
            DisplayString,

        oraEMNGEventTypeAttr18
            DisplayString,

        oraEMNGEventTypeAttr19
            DisplayString,

        oraEMNGEventTypeAttr20
            DisplayString,

        oraEMNGEventTypeAttr21
            DisplayString,

        oraEMNGEventTypeAttr22
            DisplayString,

        oraEMNGEventTypeAttr23
            DisplayString,

        oraEMNGEventTypeAttr24
            DisplayString,

        oraEMNGEventTypeAttr25
            DisplayString
  }

oraEMNGEventIndex OBJECT-TYPE
    SYNTAX INTEGER (0..2147483647)
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Index of a particular event, unique only at the moment an event is
generated."
    ::= { oraEMNGEventEntry  1 }

oraEMNGEventNotifType  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
      "Notification Type. NOTIF_NORMAL, NOTIF_RETRY, NOTIF_DURATION,
NOTIF_REPEAT, NOTIF_CA, NOTIF_RCA"
    ::= { oraEMNGEventEntry  2 }
```

```
oraEMNGEventMessage  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The message associated with this event."
    ::= { oraEMNGEventEntry  3 }

oraEMNGEventMessageURL  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "EM Console URL for the event message. Populated for events with
severity other than INFORMATIONAL. Empty if trap size exceeds configured snmp
packet size."
    ::= { oraEMNGEventEntry  4 }

oraEMNGEventSeverity  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The severity of the event e.g. Fatal, Critical, Warning, Advisory,
Information, Clear."
    ::= { oraEMNGEventEntry  5 }

oraEMNGEventSeverityCode  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Internal code of the severity: FATAL, CRITICAL, WARNING, ADVISORY,
INFORMATIONAL, CLEAR."
    ::= { oraEMNGEventEntry  6 }

oraEMNGEventRepeatCount  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The repeat notification counter for the event."
    ::= { oraEMNGEventEntry  7 }

oraEMNGEventActionMsg  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The action message for this event."
    ::= { oraEMNGEventEntry  8 }

oraEMNGEventOccurrenceTime  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The time when this event occurred (optional), this is only populated
for events that have occurrence time."
    ::= { oraEMNGEventEntry  9 }

oraEMNGEventReportedTime  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The time when this event was reported."
    ::= { oraEMNGEventEntry  10 }

oraEMNGEventCategories  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The list of categories to which this event belongs to. Empty if trap
size exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry  11 }

oraEMNGEventCategoryCodes  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The list of internal category codes to which this event belongs to.
Empty if trap size exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry  12 }

oraEMNGEventType  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
      "The name of the event type to which this event belongs to."
    ::= { oraEMNGEventEntry  13 }

oraEMNGEventName  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
      "The name of this event."
    ::= { oraEMNGEventEntry  14 }

oraEMNGAssocIncidentId  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "ID of the associated incident with the event (optional)."
    ::= { oraEMNGEventEntry  15 }
```

```
oraEMNGAssocIncidentOwner  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Owner of the associated incident with the event (optional)."
    ::= { oraEMNGEventEntry  16 }

oraEMNGAssocIncidentAcked  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Acknowledged status of the associated incident with the event. 1
indicates acknowledged, 0 indicates unacknowledged."
    ::= { oraEMNGEventEntry  17 }

oraEMNGAssocIncidentStatus  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The status of the associated incident with the event."
    ::= { oraEMNGEventEntry  18 }

oraEMNGAssocIncidentPriority  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The proirity of the associated incident with the event."
    ::= { oraEMNGEventEntry  19 }

oraEMNGAssocIncidentEscLevel  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The Escalation Level of the associated incident with the event."
    ::= { oraEMNGEventEntry  20 }

oraEMNGEventTargetName  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the target to which this event applies. Populated for
events that are about a target only."
    ::= { oraEMNGEventEntry  21 }

oraEMNGEventTargetNameURL  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "EM Console URL of the target to which this event applies. Populated for
```

```
events that are about a target only. Empty if trap size exceeds configured
snmp packet size."
    ::= { oraEMNGEventEntry  22 }

oraEMNGEventTargetType  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The type of the target to which this event applies. Populated for
events that are about a target only."
    ::= { oraEMNGEventEntry  23 }

oraEMNGEventHostName  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the host on which this event originated. Populated for
events that are about a target only."
    ::= { oraEMNGEventEntry  24 }

oraEMNGEventTargetOwner  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The primary administrator of the target on which this event originated.
Empty if trap size exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry  25 }

oraEMNGEventTgtLifeCycleStatus  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The life cycle status of the target on which this event originated."
    ::= { oraEMNGEventEntry  26 }

oraEMNGEventTargetVersion  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The version of the target on which this event originated."
    ::= { oraEMNGEventEntry  27 }

oraEMNGEventUserDefinedTgtProp  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The user defined target properties [name,value pair list] of the
associated target with this event. Empty if trap size exceeds configured snmp
packet size."
    ::= { oraEMNGEventEntry  28 }
```

```
oraEMNGEventSourceObjName  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the source object to which this event belongs to. Populated
for events that are about a non-target object only, such as Job."
    ::= { oraEMNGEventEntry  29 }

oraEMNGEventSourceObjNameURL  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "EM Console URL for the source object to which this event belongs to.
Empty if trap size exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry  30 }

oraEMNGEventSourceObjType  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The type of the source object to which this event belongs to."
    ::= { oraEMNGEventEntry  31 }

oraEMNGEventSourceObjSubType  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The sub type of the source object to which this event belongs to
(Optional property). Empty if trap size exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry  32 }

oraEMNGEventSourceObjOwner  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The primary adminstrator of the source object to which this event
belongs to. (Optional property). Empty if trap size exceeds configured snmp
packet size."
    ::= { oraEMNGEventEntry  33 }

oraEMNGEventCAJobName  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The Name of the Corrective Action Job associated with this event."
    ::= { oraEMNGEventEntry  34 }

oraEMNGEventCAJobStatus  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The Status of the Corrective Action Job associated with this event."
    ::= { oraEMNGEventEntry  35 }

oraEMNGEventCAJobOwner  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The Owner of the Corrective Action Job associated with this event."
    ::= { oraEMNGEventEntry  36 }

oraEMNGEventCAJobStepOutput  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The job step output from the Corrective Action Job associated with this
event."
    ::= { oraEMNGEventEntry  37 }

oraEMNGEventCAJobType  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The job type from the Corrective Action Job associated with this event."
    ::= { oraEMNGEventEntry  38 }

oraEMNGEventRuleSetName  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the ruleset that caused this notification. Empty if trap
size exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry  39 }

oraEMNGEventRuleName  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the rule within the ruleset that caused this notification."
    ::= { oraEMNGEventEntry  40 }

oraEMNGEventRuleOwner  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The owner of the ruleset that caused this notification."
    ::= { oraEMNGEventEntry  41 }
```

```
oraEMNGEventSequenceId  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "An EM generated identifier that uniquely identifies current issue until 
it is cleared."
    ::= { oraEMNGEventEntry  42 }

oraEMNGEventRCADetails  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Root Cause Analysis details associated with this event if it exists."
    ::= { oraEMNGEventEntry  43 }

oraEMNGEventContextAttrs  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The context attributes associated with this event. Empty if trap size 
exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry  44 }

oraEMNGEventUserComments  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The user comments associated with this event. Empty if trap size 
exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry  45 }

oraEMNGEventUpdates  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The updates associated with this event. Empty if trap size exceeds 
configured snmp packet size."
    ::= { oraEMNGEventEntry  46 }

oraEMNGEventTotalOccurrenceCount  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Total number of occurrences of the same event on a target across all 
open deduplicated events. This attribute applies only to deduplicated events."
    ::= { oraEMNGEventEntry  47 }

oraEMNGEventCurrOccurrenceCount  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Total occurrences of the event in this collection period. This
attribute applies only to deduplicated events."
    ::= { oraEMNGEventEntry  48 }

oraEMNGEventCurrFirstOccurDate  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Timestamp when the event first occurred in this collection period. This
attribute applies only to deduplicated events."
    ::= { oraEMNGEventEntry  49 }

oraEMNGEventCurrLastOccurDate  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Timestamp when the event last occurred in this collection period. This
attribute applies only to deduplicated events."
    ::= { oraEMNGEventEntry  50 }

oraEMNGEventRCAStatus  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Summary of Root Cause Analysis, if applicable."
    ::= { oraEMNGEventEntry  51 }

oraEMNGEventReportedState  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "This is an optional value reporting the latest state of an entity and
is only applicable for events that are representing a state transition for a
given entity. For example, for Target availability and Job state change
events, this value would be the latest state of the target or job,
respectively."
    ::= { oraEMNGEventEntry  52 }

oraEMNGEventTypeAttr1  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#1."
    ::= { oraEMNGEventEntry  61 }

oraEMNGEventTypeAttr2  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#2."
    ::= { oraEMNGEventEntry  62 }

oraEMNGEventTypeAttr3  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#3."
    ::= { oraEMNGEventEntry  63 }

oraEMNGEventTypeAttr4  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#4."
    ::= { oraEMNGEventEntry  64 }

oraEMNGEventTypeAttr5  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#5."
    ::= { oraEMNGEventEntry  65 }

oraEMNGEventTypeAttr6  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#6."
    ::= { oraEMNGEventEntry  66 }

oraEMNGEventTypeAttr7  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#7."
    ::= { oraEMNGEventEntry  67 }

oraEMNGEventTypeAttr8  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#8."
    ::= { oraEMNGEventEntry  68 }

oraEMNGEventTypeAttr9  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#9."
    ::= { oraEMNGEventEntry  69 }

oraEMNGEventTypeAttr10  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#10."
    ::= { oraEMNGEventEntry  70 }

oraEMNGEventTypeAttr11  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#11."
    ::= { oraEMNGEventEntry  71 }

oraEMNGEventTypeAttr12  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#12."
    ::= { oraEMNGEventEntry  72 }

oraEMNGEventTypeAttr13  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#13."
    ::= { oraEMNGEventEntry  73 }

oraEMNGEventTypeAttr14  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#14."
    ::= { oraEMNGEventEntry  74 }

oraEMNGEventTypeAttr15  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#15."
    ::= { oraEMNGEventEntry  75 }

oraEMNGEventTypeAttr16  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#16."
    ::= { oraEMNGEventEntry  76 }

oraEMNGEventTypeAttr17  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#17."
    ::= { oraEMNGEventEntry  77 }

oraEMNGEventTypeAttr18  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#18."
    ::= { oraEMNGEventEntry  78 }

oraEMNGEventTypeAttr19  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#19."
    ::= { oraEMNGEventEntry  79 }

oraEMNGEventTypeAttr20  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#20."
    ::= { oraEMNGEventEntry  80 }

oraEMNGEventTypeAttr21  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#21."
    ::= { oraEMNGEventEntry  81 }

oraEMNGEventTypeAttr22  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#22."
    ::= { oraEMNGEventEntry  82 }

oraEMNGEventTypeAttr23  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#23."
    ::= { oraEMNGEventEntry  83 }

oraEMNGEventTypeAttr24  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#24."
    ::= { oraEMNGEventEntry  84 }

oraEMNGEventTypeAttr25  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Name and value pair as name=value for event type specific attribute#25."
    ::= { oraEMNGEventEntry  85 }

oraEMNGEvent TRAP-TYPE
    ENTERPRISE  oraEM4Traps
    VARIABLES   {
          oraEMNGEventNotifType,
          oraEMNGEventMessage, oraEMNGEventMessageURL,
          oraEMNGEventSeverity, oraEMNGEventSeverityCode,
          oraEMNGEventRepeatCount, oraEMNGEventActionMsg,
          oraEMNGEventOccurrenceTime, oraEMNGEventReportedTime,
          oraEMNGEventCategories, oraEMNGEventCategoryCodes,
          oraEMNGEventType, oraEMNGEventName,
          oraEMNGAssocIncidentId, oraEMNGAssocIncidentOwner,
          oraEMNGAssocIncidentAcked, oraEMNGAssocIncidentStatus,
          oraEMNGAssocIncidentPriority, oraEMNGAssocIncidentEscLevel,
          oraEMNGEventTargetName, oraEMNGEventTargetNameURL,
          oraEMNGEventTargetType, oraEMNGEventHostName,
          oraEMNGEventTargetOwner, oraEMNGEventTgtLifeCycleStatus,
          oraEMNGEventTargetVersion, oraEMNGEventUserDefinedTgtProp,
          oraEMNGEventSourceObjName, oraEMNGEventSourceObjNameURL,
          oraEMNGEventSourceObjType, oraEMNGEventSourceObjSubType,
          oraEMNGEventSourceObjOwner, oraEMNGEventCAJobName,
          oraEMNGEventCAJobStatus, oraEMNGEventCAJobOwner,
          oraEMNGEventCAJobStepOutput, oraEMNGEventCAJobType,
          oraEMNGEventRuleSetName, oraEMNGEventRuleName,
          oraEMNGEventRuleOwner, oraEMNGEventSequenceId,
          oraEMNGEventRCADetails, oraEMNGEventContextAttrs,
          oraEMNGEventUserComments, oraEMNGEventUpdates,
          oraEMNGEventTotalOccurrenceCount, oraEMNGEventCurrOccurrenceCount,
          oraEMNGEventCurrFirstOccurDate, oraEMNGEventCurrLastOccurDate,
          oraEMNGEventRCAStatus,
          oraEMNGEventTypeAttr1, oraEMNGEventTypeAttr2,
          oraEMNGEventTypeAttr3, oraEMNGEventTypeAttr4,
          oraEMNGEventTypeAttr5, oraEMNGEventTypeAttr6,
          oraEMNGEventTypeAttr7, oraEMNGEventTypeAttr8,
          oraEMNGEventTypeAttr9, oraEMNGEventTypeAttr10,
          oraEMNGEventTypeAttr11, oraEMNGEventTypeAttr12,
          oraEMNGEventTypeAttr13, oraEMNGEventTypeAttr14,
          oraEMNGEventTypeAttr15, oraEMNGEventTypeAttr16,
          oraEMNGEventTypeAttr17, oraEMNGEventTypeAttr18,
          oraEMNGEventTypeAttr19, oraEMNGEventTypeAttr20,
          oraEMNGEventTypeAttr21, oraEMNGEventTypeAttr22,
          oraEMNGEventTypeAttr23, oraEMNGEventTypeAttr24,
          oraEMNGEventTypeAttr25,
          oraEMNGEventReportedState
}
    DESCRIPTION
     "The variables included in the oraEMNGAlert trap."
    ::= 3

oraEMIncObjects OBJECT IDENTIFIER ::= { oraEM4  4 }

oraEMNGIncidentTable OBJECT-TYPE
    SYNTAX  SEQUENCE OF OraEMNGIncidentEntry
    ACCESS  not-accessible
    STATUS  mandatory
    DESCRIPTION
     "Information on Incidents published to Oracle Enterprise Manager. This
table is not queryable; it exists only to document the variables included in
the oraEMNGIncidentTrap trap.  Each trap can contain a single instance of
each variable in the table."
    ::= { oraEMIncObjects  1 }

oraEMNGIncidentEntry OBJECT-TYPE
    SYNTAX  OraEMNGIncidentEntry
    ACCESS  not-accessible
    STATUS  mandatory
    DESCRIPTION
     "Information about a particular Oracle Enterprise Manager Incident."
    INDEX   { oraEMNGIncidentIndex }
    ::= { oraEMNGIncidentTable  1 }

OraEMNGIncidentEntry ::=
    SEQUENCE {
        oraEMNGIncidentIndex
            INTEGER,

        oraEMNGIncidentNotifType
            DisplayString,

        oraEMNGIncidentMessage
            DisplayString,

        oraEMNGIncidentMessageURL
            DisplayString,

        oraEMNGIncidentSeverity
            DisplayString,

        oraEMNGIncidentSeverityCode
            DisplayString,

        oraEMNGIncidentRepeatCount
            DisplayString,
        oraEMNGIncidentCreationTime
            DisplayString,

        oraEMNGIncidentLastUpdatedTime
            DisplayString,

        oraEMNGIncidentCategories
            DisplayString,

        oraEMNGIncidentCategoryCodes
            DisplayString,

        oraEMNGIncidentId
            DisplayString,

        oraEMNGIncidentOwner
            DisplayString,

        oraEMNGIncidentAcked
            DisplayString,

        oraEMNGIncidentStatus
            DisplayString,

        oraEMNGIncidentPriority
            DisplayString,

        oraEMNGIncidentEscLevel
            DisplayString,

        oraEMNGIncidentTargetName
            DisplayString,

        oraEMNGIncidentTargetNameURL
            DisplayString,

        oraEMNGIncidentTargetType
            DisplayString,

        oraEMNGIncidentHostName
            DisplayString,

        oraEMNGIncidentTargetOwner
            DisplayString,

        oraEMNGIncidentTgtLifeCycleStatus
            DisplayString,

        oraEMNGIncidentTargetVersion
            DisplayString,

        oraEMNGIncidentUserDefinedTgtProp
            DisplayString,

        oraEMNGIncidentSourceObjName
```

```
    DisplayString,

oraEMNGIncidentSourceObjNameURL
    DisplayString,

oraEMNGIncidentSourceObjType
    DisplayString,

oraEMNGIncidentSourceObjSubType
    DisplayString,

oraEMNGIncidentSourceObjOwner
    DisplayString,

oraEMNGIncidentSRId
    DisplayString,

oraEMNGIncidentTicketId
    DisplayString,

oraEMNGIncidentTicketStatus
    DisplayString,

oraEMNGIncidentTicketType
    DisplayString,

oraEMNGIncidentRuleSetName
    DisplayString,

oraEMNGIncidentRuleName
    DisplayString,

oraEMNGIncidentRuleOwner
    DisplayString,

oraEMNGIncidentRCADetails
    DisplayString,

oraEMNGIncidentUserComments
    DisplayString,

oraEMNGIncidentUpdates
    DisplayString,

oraEMNGIncidentRCAStatus
    DisplayString,

oraEMNGIncidentAssocEventCount
    DisplayString,

oraEMNGIncidentCompressedEventMsg1
    DisplayString,

oraEMNGIncidentCompressedEventType1
    DisplayString,
oraEMNGIncidentCompressedEventSeverity1
    DisplayString,

oraEMNGIncidentCompressedEventCreationDate1
    DisplayString,

oraEMNGIncidentCompressedEventTargetName1
    DisplayString,

oraEMNGIncidentCompressedEventTargetType1
    DisplayString,

oraEMNGIncidentCompressedEventMsg2
    DisplayString,

oraEMNGIncidentCompressedEventType2
    DisplayString,

oraEMNGIncidentCompressedEventSeverity2
    DisplayString,

oraEMNGIncidentCompressedEventCreationDate2
    DisplayString,

oraEMNGIncidentCompressedEventTargetName2
    DisplayString,

oraEMNGIncidentCompressedEventTargetType2
    DisplayString,

oraEMNGIncidentCompressedEventMsg3
    DisplayString,

oraEMNGIncidentCompressedEventType3
    DisplayString,

oraEMNGIncidentCompressedEventSeverity3
    DisplayString,

oraEMNGIncidentCompressedEventCreationDate3
    DisplayString,

oraEMNGIncidentCompressedEventTargetName3
    DisplayString,

oraEMNGIncidentCompressedEventTargetType3
    DisplayString,

oraEMNGIncidentCompressedEventMsg4
    DisplayString,

oraEMNGIncidentCompressedEventType4
    DisplayString,

oraEMNGIncidentCompressedEventSeverity4
    DisplayString,
        oraEMNGIncidentCompressedEventCreationDate4
            DisplayString,

        oraEMNGIncidentCompressedEventTargetName4
            DisplayString,

        oraEMNGIncidentCompressedEventTargetType4
            DisplayString,

        oraEMNGIncidentCompressedEventMsg5
            DisplayString,

        oraEMNGIncidentCompressedEventType5
            DisplayString,

        oraEMNGIncidentCompressedEventSeverity5
            DisplayString,

        oraEMNGIncidentCompressedEventCreationDate5
            DisplayString,

        oraEMNGIncidentCompressedEventTargetName5
            DisplayString,

        oraEMNGIncidentCompressedEventTargetType5
            DisplayString
  }

oraEMNGIncidentIndex OBJECT-TYPE
    SYNTAX INTEGER (0..2147483647)
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Index of a particular Incident, unique only at the moment an Incident 
is generated."
    ::= { oraEMNGIncidentEntry  1 }

oraEMNGIncidentNotifType  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
      "Notification Type. NOTIF_NORMAL, NOTIF_RETRY, NOTIF_DURATION, 
NOTIF_REPEAT, NOTIF_CA, NOTIF_RCA"
    ::= { oraEMNGIncidentEntry  2 }

oraEMNGIncidentMessage  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The message of this Incident."
    ::= { oraEMNGIncidentEntry  3 }

oraEMNGIncidentMessageURL  OBJECT-TYPE
```

```
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "EM Console URL for the Incident message. Populated for Incidents with
severity other than INFORMATIONAL. Empty if trap size exceeds configured snmp
packet size."
    ::= { oraEMNGIncidentEntry  4 }

oraEMNGIncidentSeverity  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The severity of the Incident e.g. Fatal, Critical, Warning, Advisory,
Information, Clear."
    ::= { oraEMNGIncidentEntry  5 }

oraEMNGIncidentSeverityCode  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Internal code of the severity: FATAL, CRITICAL, WARNING, ADVISORY,
INFORMATIONAL, CLEAR."
    ::= { oraEMNGIncidentEntry  6 }

oraEMNGIncidentRepeatCount  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The repeat notification counter for the Incident."
    ::= { oraEMNGIncidentEntry  7 }

oraEMNGIncidentCreationTime  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The time when this Incident was created in EM."
    ::= { oraEMNGIncidentEntry  8 }

oraEMNGIncidentLastUpdatedTime  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The time when this Incident was last updated."
    ::= { oraEMNGIncidentEntry  9 }

oraEMNGIncidentCategories  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
```

```
     "The list of categories to which this Incident belongs to. Empty if trap
size exceeds configured snmp packet size."
    ::= { oraEMNGIncidentEntry  10 }

oraEMNGIncidentCategoryCodes  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The list of internal category codes to which this Incident belongs to.
Empty if trap size exceeds configured snmp packet size."
    ::= { oraEMNGIncidentEntry  11 }

oraEMNGIncidentId  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "ID of the Incident."
    ::= { oraEMNGIncidentEntry  12 }

oraEMNGIncidentOwner  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Owner of the Incident."
    ::= { oraEMNGIncidentEntry  13 }

oraEMNGIncidentAcked  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Acknowledged status of the Incident. 1 indicates acknowledged, 0
indicates unacknowledged."
    ::= { oraEMNGIncidentEntry  14 }

oraEMNGIncidentStatus  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The status of the Incident."
    ::= { oraEMNGIncidentEntry  15 }

oraEMNGIncidentPriority  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The priority of the Incident."
    ::= { oraEMNGIncidentEntry  16 }

oraEMNGIncidentEscLevel  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The Escalation Level of the Incident."
    ::= { oraEMNGIncidentEntry  17 }

oraEMNGIncidentTargetName  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the target to which this Incident applies. Populated for
Incidents that are about a target only."
    ::= { oraEMNGIncidentEntry  18 }

oraEMNGIncidentTargetNameURL  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "EM Console URL of the target to which this Incident applies. Populated
for Incidents that are about a target only. Empty if trap size exceeds
configured snmp packet size."
    ::= { oraEMNGIncidentEntry  19 }

oraEMNGIncidentTargetType  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The type of the target to which this Incident applies. Populated for
Incidents that are about a target only."
    ::= { oraEMNGIncidentEntry  20 }

oraEMNGIncidentHostName  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the host on which this Incident originated. Populated for
Incidents that are about a target only."
    ::= { oraEMNGIncidentEntry  21 }

oraEMNGIncidentTargetOwner  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The primary administrator of the target on which this Incident
originated. Empty if trap size exceeds configured snmp packet size."
    ::= { oraEMNGIncidentEntry  22 }

oraEMNGIncidentTgtLifeCycleStatus  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The life cycle status of the target on which this Incident originated."
    ::= { oraEMNGIncidentEntry  23 }

oraEMNGIncidentTargetVersion  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The version of the target on which this Incident originated."
    ::= { oraEMNGIncidentEntry  24 }

oraEMNGIncidentUserDefinedTgtProp  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The user defined target properties [name,value pair list] of the
associated target with this Incident. Empty if trap size exceeds configured
snmp packet size."
    ::= { oraEMNGIncidentEntry  25 }

oraEMNGIncidentSourceObjName  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the source object to which this Incident belongs to.
Populated for Incidents that are about a non-target object only, such as Job."
    ::= { oraEMNGIncidentEntry  26 }

oraEMNGIncidentSourceObjNameURL  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "EM Console URL for the source object to which this Incident belongs to.
Empty if trap size exceeds configured snmp packet size."
    ::= { oraEMNGIncidentEntry  27 }

oraEMNGIncidentSourceObjType  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The type of the source object to which this Incident belongs to."
    ::= { oraEMNGIncidentEntry  28 }

oraEMNGIncidentSourceObjSubType  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The sub type of the source object to which this Incident belongs to
(Optional property). Empty if trap size exceeds configured snmp packet size."
    ::= { oraEMNGIncidentEntry  29 }
```

```
oraEMNGIncidentSourceObjOwner  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The primary adminstrator of the source object to which this Incident
belongs to. (Optional property). Empty if trap size exceeds configured snmp
packet size."
    ::= { oraEMNGIncidentEntry  30 }

oraEMNGIncidentSRId  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The Service Request ID associated with this Incident, if applicable."
    ::= { oraEMNGIncidentEntry  31 }

oraEMNGIncidentTicketId  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The Ticket ID associated with this Incident, if applicable."
    ::= { oraEMNGIncidentEntry  32 }

oraEMNGIncidentTicketStatus  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The status of the Ticket associated with this Incident, if applicable."
    ::= { oraEMNGIncidentEntry  33 }

oraEMNGIncidentTicketType  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The type of ticket associated with this Incident, if applicable"
    ::= { oraEMNGIncidentEntry  34 }

oraEMNGIncidentRuleSetName  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the ruleset that triggered this notification. Empty if trap
size exceeds configured snmp packet size."
    ::= { oraEMNGIncidentEntry  35 }

oraEMNGIncidentRuleName  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The name of the rule within the ruleset that triggered this
notification."
    ::= { oraEMNGIncidentEntry  36 }

oraEMNGIncidentRuleOwner  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The owner of the ruleset that triggered this notification."
    ::= { oraEMNGIncidentEntry  37 }

oraEMNGIncidentRCADetails  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Root Cause Analysis details associated with this Incident if it exists."
    ::= { oraEMNGIncidentEntry  38 }

oraEMNGIncidentUserComments  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The user comments associated with this Incident. Empty if trap size
exceeds configured snmp packet size."
    ::= { oraEMNGIncidentEntry  39 }

oraEMNGIncidentUpdates  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The updates associated with this Incident. Empty if trap size exceeds
configured snmp packet size."
    ::= { oraEMNGIncidentEntry  40 }

oraEMNGIncidentRCAStatus  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "Summary of Root Cause Analysis, if applicable."
    ::= { oraEMNGIncidentEntry  41 }

oraEMNGIncidentAssocEventCount  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The number of underlying events associated to this incident. In most
cases, this is set to 1 which means it is a single-event incident. When there
are multiple events associated to this incident, the most recent 5 events
details (message, event type, creation date, target name, target type) are
```

```
set as part of the subsequent oraEMNGIncidentCompressedEvent% fields."
    ::= { oraEMNGIncidentEntry  42 }

oraEMNGIncidentCompressedEventMsg1  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The message of the most recent event that was associated to this
incident, applicable only if the value of oraEMNGIncidentAssocEventCount is
greater than 1"
    ::= { oraEMNGIncidentEntry  43 }

oraEMNGIncidentCompressedEventType1  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The event type of the most recent event that was associated to this
incident, applicable only if the value of oraEMNGIncidentAssocEventCount is
greater than 1"
    ::= { oraEMNGIncidentEntry  44 }

oraEMNGIncidentCompressedEventSeverity1  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The severity of the most recent event that was associated to this
incident, applicable only if the value of oraEMNGIncidentAssocEventCount is
greater than 1"
    ::= { oraEMNGIncidentEntry  45 }

oraEMNGIncidentCompressedEventCreationDate1  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The creation date of the most recent event that was associated to this
incident, applicable only if the value of oraEMNGIncidentAssocEventCount is
greater than 1"
    ::= { oraEMNGIncidentEntry  46 }

oraEMNGIncidentCompressedEventTargetName1  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The target name of the most recent event that was associated to this
incident, applicable only if the value of oraEMNGIncidentAssocEventCount is
greater than 1"
    ::= { oraEMNGIncidentEntry  47 }

oraEMNGIncidentCompressedEventTargetType1  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The target type of the most recent event that was associated to this
incident, applicable only if the value of oraEMNGIncidentAssocEventCount is
greater than 1"
    ::= { oraEMNGIncidentEntry  48 }

oraEMNGIncidentCompressedEventMsg2  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The message of the second most recent event that was associated to this
incident, applicable only if the value of oraEMNGIncidentAssocEventCount is
greater than or equals to 2"
    ::= { oraEMNGIncidentEntry  49 }

oraEMNGIncidentCompressedEventType2  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The event type of the second most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 2"
    ::= { oraEMNGIncidentEntry  50 }

oraEMNGIncidentCompressedEventSeverity2  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The severity of the second most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 2"
    ::= { oraEMNGIncidentEntry  51 }

oraEMNGIncidentCompressedEventCreationDate2  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The creation date of the second most recent event that was associated
to this incident, applicable only if the value of
oraEMNGIncidentAssocEventCount is greater than or equals to 2"
    ::= { oraEMNGIncidentEntry  52 }

oraEMNGIncidentCompressedEventTargetName2  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The target name of the second most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 2"
    ::= { oraEMNGIncidentEntry  53 }
```

```
oraEMNGIncidentCompressedEventTargetType2  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The target type of the second most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 2"
    ::= { oraEMNGIncidentEntry  54 }

oraEMNGIncidentCompressedEventMsg3  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The message of the third most recent event that was associated to this
incident, applicable only if the value of oraEMNGIncidentAssocEventCount is
greater than or equals to 3"
    ::= { oraEMNGIncidentEntry  55 }

oraEMNGIncidentCompressedEventType3  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The event type of the third most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 3"
    ::= { oraEMNGIncidentEntry  56 }

oraEMNGIncidentCompressedEventSeverity3  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The severity of the third most recent event that was associated to this
incident, applicable only if the value of oraEMNGIncidentAssocEventCount is
greater than or equals to 3"
    ::= { oraEMNGIncidentEntry  57 }

oraEMNGIncidentCompressedEventCreationDate3  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The creation date of the third most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 3"
    ::= { oraEMNGIncidentEntry  58 }

oraEMNGIncidentCompressedEventTargetName3  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
```

```
     "The target name of the third most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 3"
    ::= { oraEMNGIncidentEntry  59 }

oraEMNGIncidentCompressedEventTargetType3  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The target type of the third most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 3"
    ::= { oraEMNGIncidentEntry  60 }

oraEMNGIncidentCompressedEventMsg4  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The message of the fourth most recent event that was associated to this
incident, applicable only if the value of oraEMNGIncidentAssocEventCount is
greater than or equals to 4"
    ::= { oraEMNGIncidentEntry  61 }

oraEMNGIncidentCompressedEventType4  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The event type of the fourth most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 4"
    ::= { oraEMNGIncidentEntry  62 }

oraEMNGIncidentCompressedEventSeverity4  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The severity of the fourth most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 4"
    ::= { oraEMNGIncidentEntry  63 }

oraEMNGIncidentCompressedEventCreationDate4  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The creation date of the fourth most recent event that was associated
to this incident, applicable only if the value of
oraEMNGIncidentAssocEventCount is greater than or equals to 4"
    ::= { oraEMNGIncidentEntry  64 }

oraEMNGIncidentCompressedEventTargetName4  OBJECT-TYPE
```

```
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The target name of the fourth most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 4"
    ::= { oraEMNGIncidentEntry  65 }

oraEMNGIncidentCompressedEventTargetType4  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The target type of the fourth most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 4"
    ::= { oraEMNGIncidentEntry  66 }

oraEMNGIncidentCompressedEventMsg5  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The message of the fifth most recent event that was associated to this
incident, applicable only if the value of oraEMNGIncidentAssocEventCount is
greater than or equals to 5"
    ::= { oraEMNGIncidentEntry  67 }

oraEMNGIncidentCompressedEventType5  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The event type of the fifth most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 5"
    ::= { oraEMNGIncidentEntry  68 }

oraEMNGIncidentCompressedEventSeverity5  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The severity of the fifth most recent event that was associated to this
incident, applicable only if the value of oraEMNGIncidentAssocEventCount is
greater than or equals to 5"
    ::= { oraEMNGIncidentEntry  69 }

oraEMNGIncidentCompressedEventCreationDate5  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The creation date of the fifth most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
```

```
is greater than or equals to 5"
    ::= { oraEMNGIncidentEntry  70 }

oraEMNGIncidentCompressedEventTargetName5  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The target name of the fifth most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 5"
    ::= { oraEMNGIncidentEntry  71 }

oraEMNGIncidentCompressedEventTargetType5  OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
     "The target type of the fifth most recent event that was associated to
this incident, applicable only if the value of oraEMNGIncidentAssocEventCount
is greater than or equals to 5"
    ::= { oraEMNGIncidentEntry  72 }

oraEMNGIncident TRAP-TYPE
    ENTERPRISE  oraEM4Traps
    VARIABLES   {
          oraEMNGIncidentNotifType, oraEMNGIncidentMessage,
          oraEMNGIncidentMessageURL, oraEMNGIncidentSeverity,
          oraEMNGIncidentSeverityCode, oraEMNGIncidentRepeatCount,
          oraEMNGIncidentCreationTime, oraEMNGIncidentLastUpdatedTime,
          oraEMNGIncidentCategories, oraEMNGIncidentCategoryCodes,
          oraEMNGIncidentId, oraEMNGIncidentOwner,
          oraEMNGIncidentAcked, oraEMNGIncidentStatus,
          oraEMNGIncidentPriority, oraEMNGIncidentEscLevel,
          oraEMNGIncidentTargetName, oraEMNGIncidentTargetNameURL,
          oraEMNGIncidentTargetType, oraEMNGIncidentHostName,
          oraEMNGIncidentTargetOwner, oraEMNGIncidentTgtLifeCycleStatus,
          oraEMNGIncidentTargetVersion, oraEMNGIncidentUserDefinedTgtProp,
          oraEMNGIncidentSourceObjName, oraEMNGIncidentSourceObjNameURL,
          oraEMNGIncidentSourceObjType, oraEMNGIncidentSourceObjSubType,
          oraEMNGIncidentSourceObjOwner, oraEMNGIncidentSRId,
          oraEMNGIncidentTicketId, oraEMNGIncidentTicketStatus,
          oraEMNGIncidentTicketType, oraEMNGIncidentRuleSetName,
          oraEMNGIncidentRuleName, oraEMNGIncidentRuleOwner,
          oraEMNGIncidentRCADetails, oraEMNGIncidentUserComments,
          oraEMNGIncidentUpdates, oraEMNGIncidentRCAStatus,
          oraEMNGIncidentAssocEventCount, oraEMNGIncidentCompressedEventMsg1,
          oraEMNGIncidentCompressedEventType1,
oraEMNGIncidentCompressedEventSeverity1,
          oraEMNGIncidentCompressedEventCreationDate1,
oraEMNGIncidentCompressedEventTargetName1,
          oraEMNGIncidentCompressedEventTargetType1,
oraEMNGIncidentCompressedEventMsg2,
          oraEMNGIncidentCompressedEventType2,
oraEMNGIncidentCompressedEventSeverity2,
          oraEMNGIncidentCompressedEventCreationDate2,
oraEMNGIncidentCompressedEventTargetName2,
          oraEMNGIncidentCompressedEventTargetType2,
oraEMNGIncidentCompressedEventMsg3,
          oraEMNGIncidentCompressedEventType3,
oraEMNGIncidentCompressedEventSeverity3,
          oraEMNGIncidentCompressedEventCreationDate3,
oraEMNGIncidentCompressedEventTargetName3,
          oraEMNGIncidentCompressedEventTargetType3,
oraEMNGIncidentCompressedEventMsg4,
          oraEMNGIncidentCompressedEventType4,
oraEMNGIncidentCompressedEventSeverity4,
          oraEMNGIncidentCompressedEventCreationDate4,
oraEMNGIncidentCompressedEventTargetName4,
          oraEMNGIncidentCompressedEventTargetType4,
oraEMNGIncidentCompressedEventMsg5,
          oraEMNGIncidentCompressedEventType5,
oraEMNGIncidentCompressedEventSeverity5,
          oraEMNGIncidentCompressedEventCreationDate5,
oraEMNGIncidentCompressedEventTargetName5,
          oraEMNGIncidentCompressedEventTargetType5
}
    DESCRIPTION
     "The variables included in the oraEMNGAlert trap."
    ::= 4


END
```

SNMP Trap Mappings

The following tables list SNMP trap mappings between Enterprise Manager 12c and later and previous releases.

Pre-12c Enterprise Manager Metric Alerts

Before Enterprise Manager 12*c*, metric alerts were sent using the oraEM4Alert trap type. From Enterprise Manager 12*c* onwards, the event type corresponding to these alerts is metric alert. The value for oraEMNGEventType in an Enterprise Manager 12*c* or 13*c* SNMP trap would be set to 'Metric Alert'.

Table 14 Metric Alert Mappings

Pre-12C OID Number	Pre-12C OID Name	12C and later OID Number	12C and later OID Name
1.3.6.1.4.1.111.15.1.1.1.2.1	oraEM4AlertTargetName	1.3.6.1.4.1.111.15.3.1.1.21.1	oraEMNGEventTargetName
1.3.6.1.4.1.111.15.1.1.1.3.1	oraEM4AlertTargetType	1.3.6.1.4.1.111.15.3.1.1.23.1	oraEMNGEventTargetType
1.3.6.1.4.1.111.15.1.1.1.4.1	oraEM4AlertHostName	1.3.6.1.4.1.111.15.3.1.1.24.1	oraEMNGEventHostName
1.3.6.1.4.1.111.15.1.1.1.5.1	oraEM4AlertMetricName	1.3.6.1.4.1.111.15.3.1.1.65.1	oraEMNGEventTypeAttr5
1.3.6.1.4.1.111.15.1.1.1.6.1	oraEM4AlertKeyName	1.3.6.1.4.1.111.15.3.1.1.66.1	oraEMNGEventTypeAttr6
1.3.6.1.4.1.111.15.1.1.1.7.1	oraEM4AlertKeyValue	* See the note below for details	* See the note below for details

Table 14 (Cont.) Metric Alert Mappings

Pre-12C OID Number	Pre-12C OID Name	12C and later OID Number	12C and later OID Name
1.3.6.1.4.1.111.15.1.1.1.8.1	oraEM4AlertTimeStamp	1.3.6.1.4.1.111.15.3.1.1.10.1	oraEMNGEventReportedTime
1.3.6.1.4.1.111.15.1.1.1.9.1	oraEM4AlertSeverity	1.3.6.1.4.1.111.15.3.1.1.5.1	oraEMNGEventSeverity
1.3.6.1.4.1.111.15.1.1.1.10.1	oraEM4AlertMessage	1.3.6.1.4.1.111.15.3.1.1.3.1	oraEMNGEventMessage
1.3.6.1.4.1.111.15.1.1.1.11.1	oraEM4AlertRuleName	1.3.6.1.4.1.111.15.3.1.1.39.1	oraEMNGEventRuleSetName
1.3.6.1.4.1.111.15.1.1.1.12.1	oraEM4AlertRuleOwner	1.3.6.1.4.1.111.15.3.1.1.41.1	oraEMNGEventRuleOwner
1.3.6.1.4.1.111.15.1.1.1.13.1	oraEM4AlertMetricValue	1.3.6.1.4.1.111.15.3.1.1.68.1	oraEMNGEventTypeAttr8
1.3.6.1.4.1.111.15.1.1.1.14.1	oraEM4AlertContext	1.3.6.1.4.1.111.15.3.1.1.44.1	oraEMNGEventContextAttrs
1.3.6.1.4.1.111.15.1.1.1.15.1	oraEM4AlertCycleGuid	1.3.6.1.4.1.111.15.3.1.1.70.1	oraEMNGEventTypeAttr3
1.3.6.1.4.1.111.15.1.1.1.16.1	oraEM4AlertRepeatCount	1.3.6.1.4.1.111.15.3.1.1.7.1	oraEMNGEventRepeatCount
1.3.6.1.4.1.111.15.1.1.1.17.1	oraEM4AlertUDTargetProperties	1.3.6.1.4.1.111.15.3.1.1.28.1	oraEMNGEventUserDefinedTgtProp
1.3.6.1.4.1.111.15.1.1.1.18.1	oraEM4AlertAck	1.3.6.1.4.1.111.15.3.1.1.17.1	oraEMNGAssocIncidentAcked
1.3.6.1.4.1.111.15.1.1.1.19.1	oraEM4AlertAckBy	1.3.6.1.4.1.111.15.3.1.1.16.1	oraEMNGAssocIncidentOwner
1.3.6.1.4.1.111.15.1.1.1.20.1	oraEM4AlertNotifType	1.3.6.1.4.1.111.15.3.1.1.2.1	oraEMNGEventNotifType
1.3.6.1.4.1.111.15.1.1.1.21.1	oraEM4AlertViolationGuid	1.3.6.1.4.1.111.15.3.1.1.42.1	oraEMNGEventSequenceId

Pre-12C Target Availability Alerts

Before Enterprise Manager 12*c*, target availability alerts were sent using oraEM4Alert SNMP trap type. From Enterprise Manager 12c onwards, the event type corresponding to these alerts is target_availability. Value for oraEMNGEventType in an Enterprise Manager 12*c* or 13*c* trap would be set to 'Target Availability'.

Table 15 Target Availability Alert Mappings

Pre-12C OID Number	Pre-12C OID Name	12C and later OID Number	12C and later OID Name
1.3.6.1.4.1.111.15.1.1.1.2.1	oraEM4AlertTargetName	1.3.6.1.4.1.111.15.3.1.1.21.1	oraEMNGEventTargetName
1.3.6.1.4.1.111.15.1.1.1.3.1	oraEM4AlertTargetType	1.3.6.1.4.1.111.15.3.1.1.23.1	oraEMNGEventTargetType
1.3.6.1.4.1.111.15.1.1.1.4.1	oraEM4AlertHostName	1.3.6.1.4.1.111.15.3.1.1.24.1	oraEMNGEventHostName
1.3.6.1.4.1.111.15.1.1.1.5.1	oraEM4AlertMetricName	N/A	N/A
1.3.6.1.4.1.111.15.1.1.1.6.1	oraEM4AlertKeyName	// deprecated in 12C, was always null in 11GC	// deprecated in 12C, was always null in 11GC
1.3.6.1.4.1.111.15.1.1.1.7.1	oraEM4AlertKeyValue	N/A	N/A

Table 15 (Cont.) Target Availability Alert Mappings

Pre-12C OID Number	Pre-12C OID Name	12C and later OID Number	12C and later OID Name
1.3.6.1.4.1.111.15.1.1.1.8.1	oraEM4AlertTimeStamp	1.3.6.1.4.1.111.15.3.1.1.10.1	oraEMNGEventReportedTime
1.3.6.1.4.1.111.15.1.1.1.9.1	oraEM4AlertSeverity	1.3.6.1.4.1.111.15.3.1.1.61.1	oraEMNGEventTypeAttr1 // target_status
1.3.6.1.4.1.111.15.1.1.1.10.1	oraEM4AlertMessage	1.3.6.1.4.1.111.15.3.1.1.3.1	oraEMNGEventMessage
1.3.6.1.4.1.111.15.1.1.1.11.1	oraEM4AlertRuleName	1.3.6.1.4.1.111.15.3.1.1.39.1	oraEMNGEventRuleSetName
1.3.6.1.4.1.111.15.1.1.1.12.1	oraEM4AlertRuleOwner	1.3.6.1.4.1.111.15.3.1.1.41.1	oraEMNGEventRuleOwner
1.3.6.1.4.1.111.15.1.1.1.13.1	oraEM4AlertMetricValue	N/A	N/A
1.3.6.1.4.1.111.15.1.1.1.14.1	oraEM4AlertContext	1.3.6.1.4.1.111.15.3.1.1.44.1	oraEMNGEventContextAttrs
1.3.6.1.4.1.111.15.1.1.1.15.1	oraEM4AlertCycleGuid	1.3.6.1.4.1.111.15.3.1.1.66.1	oraEMNGEventTypeAttr6
1.3.6.1.4.1.111.15.1.1.1.16.1	oraEM4AlertRepeatCount	1.3.6.1.4.1.111.15.3.1.1.7.1	oraEMNGEventRepeatCount
1.3.6.1.4.1.111.15.1.1.1.17.1	oraEM4AlertUDTargetProperties	1.3.6.1.4.1.111.15.3.1.1.28.1	oraEMNGEventUserDefinedTgtProp
1.3.6.1.4.1.111.15.1.1.1.18.1	oraEM4AlertAck	1.3.6.1.4.1.111.15.3.1.1.17.1	oraEMNGAssocIncidentAcked
1.3.6.1.4.1.111.15.1.1.1.19.1	oraEM4AlertAckBy	1.3.6.1.4.1.111.15.3.1.1.16.1	oraEMNGAssocIncidentOwner
1.3.6.1.4.1.111.15.1.1.1.20.1	oraEM4AlertNotifType	1.3.6.1.4.1.111.15.3.1.1.2.1	oraEMNGEventNotifType
1.3.6.1.4.1.111.15.1.1.1.21.1	oraEM4AlertViolationGuid	1.3.6.1.4.1.111.15.3.1.1.42.1	oraEMNGEventSequenceId

Pre-12C Corrective Action Results for Metric Alerts

Before Enterprise Manager 12*c*, corrective action results for metric alerts were sent using the oraEM4JobAlert trap type. From Enterprise Manager 12*c* onwards, the event type corresponding to these alerts is metric alert. The value for oraEMNGEventType in an Enterprise Manager 12*c* or 13*c* trap would be set to 'Metric Alert'.

Table 16 Corrective Action Results for Metric Alert Mappings

Pre-12c OID Number	Pre-12c OID Name	12c and later OID Number	12c and later OID Name
1.3.6.1.4.1.111.15.1.2.1.2.1	oraEM4JobAlertJobName	1.3.6.1.4.1.111.15.3.1.1.34.1	oraEMNGEventCAJobName
1.3.6.1.4.1.111.15.1.2.1.3.1	oraEM4JobAlertJobOwner	1.3.6.1.4.1.111.15.3.1.1.36.1	oraEMNGEventCAJobOwner
1.3.6.1.4.1.111.15.1.2.1.4.1	oraEM4JobAlertJobType	1.3.6.1.4.1.111.15.3.1.1.38.1	oraEMNGEventCAJobType
1.3.6.1.4.1.111.15.1.2.1.5.1	oraEM4JobAlertJobStatus	1.3.6.1.4.1.111.15.3.1.1.35.1	oraEMNGEventCAJobStatus
1.3.6.1.4.1.111.15.1.2.1.6.1	oraEM4JobAlertTargets	1.3.6.1.4.1.111.15.3.1.1.23.1	oraEMNGEventTargetType
NA	NA	1.3.6.1.4.1.111.15.3.1.1.21.1	oraEMNGEventTargetName

Table 16 (Cont.) Corrective Action Results for Metric Alert Mappings

Pre-12c OID Number	Pre-12c OID Name	12c and later OID Number	12c and later OID Name
1.3.6.1.4.1.111.15.1.2.1.7.1	oraEM4JobAlertTimeStamp	1.3.6.1.4.1.111.15.3.1.1.10.1	oraEMNGEventReportedTime
1.3.6.1.4.1.111.15.1.2.1.8.1	oraEM4JobAlertRuleName	1.3.6.1.4.1.111.15.3.1.1.39.1	oraEMNGEventRuleSetName
1.3.6.1.4.1.111.15.1.2.1.9.1	oraEM4JobAlertRuleOwner	1.3.6.1.4.1.111.15.3.1.1.41.1	oraEMNGEventRuleOwner
1.3.6.1.4.1.111.15.1.2.1.10.1	oraEM4JobAlertMetricName	1.3.6.1.4.1.111.15.3.1.1.65.1	oraEMNGEventTypeAttr5
1.3.6.1.4.1.111.15.1.2.1.11.1	oraEM4JobAlertMetricValue	1.3.6.1.4.1.111.15.3.1.1.68.1	oraEMNGEventTypeAttr8
1.3.6.1.4.1.111.15.1.2.1.12.1	oraEM4JobAlertContext	1.3.6.1.4.1.111.15.3.1.1.44.1	oraEMNGEventContextAttrs
1.3.6.1.4.1.111.15.1.2.1.13.1	oraEM4JobAlertKeyName	1.3.6.1.4.1.111.15.3.1.1.66.1	oraEMNGEventTypeAttr6
1.3.6.1.4.1.111.15.1.2.1.14.1	oraEM4JobAlertKeyValue	1.3.6.1.4.1.111.15.3.1.1.69.1	oraEMNGEventTypeAttr9
1.3.6.1.4.1.111.15.1.2.1.15.1	oraEM4JobAlertSeverity	1.3.6.1.4.1.111.15.3.1.1.5.1	oraEMNGEventSeverity
1.3.6.1.4.1.111.15.1.2.1.16.1	oraEM4JobAlertJobId	NA	NA
1.3.6.1.4.1.111.15.1.2.1.17.1	oraEM4JobAlertJobExecId	NA	NA

Corrective Action Results for Target Availability

Before Enterprise Manager 12c, corrective action results for target availability alerts were sent using the oraEM4JobAlert trap type. From Enterprise Manager 12c onwards, the event type corresponding to these alerts is target_availability alert. The value for oraEMNGEventType in an Enterprise Manager 12*c* or 13*c* trap would be set to 'Metric Alert'.

Table 17 Target Availability Mappings

Pre-12c OID Number	Pre-12c OID Name	12c and later OID Number	12c and later OID Name
1.3.6.1.4.1.111.15.1.2.1.2.1	oraEM4JobAlertJobName	1.3.6.1.4.1.111.15.3.1.1.34.1	oraEMNGEventCAJobName
1.3.6.1.4.1.111.15.1.2.1.3.1	oraEM4JobAlertJobOwner	1.3.6.1.4.1.111.15.3.1.1.36.1	oraEMNGEventCAJobOwner
1.3.6.1.4.1.111.15.1.2.1.4.1	oraEM4JobAlertJobType	1.3.6.1.4.1.111.15.3.1.1.38.1	oraEMNGEventCAJobType
1.3.6.1.4.1.111.15.1.2.1.5.1	oraEM4JobAlertJobStatus	1.3.6.1.4.1.111.15.3.1.1.35.1	oraEMNGEventCAJobStatus
1.3.6.1.4.1.111.15.1.2.1.6.1	oraEM4JobAlertTargets	1.3.6.1.4.1.111.15.3.1.1.23.1	oraEMNGEventTargetType and
N/A	N/A	1.3.6.1.4.1.111.15.3.1.1.21.1	oraEMNGEventTargetName
1.3.6.1.4.1.111.15.1.2.1.7.1	oraEM4JobAlertTimeStamp	1.3.6.1.4.1.111.15.3.1.1.10.1	oraEMNGEventReportedTime
1.3.6.1.4.1.111.15.1.2.1.8.1	oraEM4JobAlertRuleName	1.3.6.1.4.1.111.15.3.1.1.39.1	oraEMNGEventRuleSetName
1.3.6.1.4.1.111.15.1.2.1.9.1	oraEM4JobAlertRuleOwner	1.3.6.1.4.1.111.15.3.1.1.41.1	oraEMNGEventRuleOwner
1.3.6.1.4.1.111.15.1.2.1.10.1	oraEM4JobAlertMetricName	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.11.1	oraEM4JobAlertMetricValue	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.12.1	oraEM4JobAlertContext	1.3.6.1.4.1.111.15.3.1.1.44.1	oraEMNGEventContextAttrs
1.3.6.1.4.1.111.15.1.2.1.13.1	oraEM4JobAlertKeyName	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.14.1	oraEM4JobAlertKeyValue	N/A	N/A

Table 17   (Cont.) Target Availability Mappings

Pre-12c OID Number	Pre-12c OID Name	12c and later OID Number	12c and later OID Name
1.3.6.1.4.1.111.15.1.2.1.15.1	oraEM4JobAlertSeverity	1.3.6.1.4.1.111.15.3.1.1.61.1	oraEMNGEventTypeAttr5 // target_status
1.3.6.1.4.1.111.15.1.2.1.16.1	oraEM4JobAlertJobId	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.17.1	oraEM4JobAlertJobExecId	N/A	N/A

Job Status Change

Before Enterprise Manager 12*c* , job status change was sent using oraEM4JobAlert trap type. From Enterprise Manager 12*c* onwards, the event type corresponding to these alerts is the job_status_change alert. The value for the oraEMNGEventType in an Enterprise Manager 12*c* or 13*c* trap would be set to 'Job Status Change'.

Table 18   Job Status Change Mappings

Pre-12c OID Number	Pre-12c OID Name	12c and later OID Number	12c and later OID Name
1.3.6.1.4.1.111.15.1.2.1.2.1	oraEM4JobAlertJobName	1.3.6.1.4.1.111.15.3.1.1.29.1	oraEMNGEventSourceObjName
1.3.6.1.4.1.111.15.1.2.1.3.1	oraEM4JobAlertJobOwner	1.3.6.1.4.1.111.15.3.1.1.33.1	oraEMNGEventSourceObjOwner
1.3.6.1.4.1.111.15.1.2.1.4.1	oraEM4JobAlertJobType	1.3.6.1.4.1.111.15.3.1.1.32.1	oraEMNGEventSourceObjSubType
1.3.6.1.4.1.111.15.1.2.1.5.1	oraEM4JobAlertJobStatus	1.3.6.1.4.1.111.15.3.1.1.62.1	oraEMNGEventTypeAttr2
1.3.6.1.4.1.111.15.1.2.1.6.1	oraEM4JobAlertTargets	1.3.6.1.4.1.111.15.3.1.1.23.1	oraEMNGEventTargetType and
NA	NA	1.3.6.1.4.1.111.15.3.1.1.21.1	oraEMNGEventTargetName
1.3.6.1.4.1.111.15.1.2.1.7.1	oraEM4JobAlertTimeStamp	1.3.6.1.4.1.111.15.3.1.1.10.1	oraEMNGEventReportedTime
1.3.6.1.4.1.111.15.1.2.1.8.1	oraEM4JobAlertRuleName	1.3.6.1.4.1.111.15.3.1.1.39.1	oraEMNGEventRuleSetName
1.3.6.1.4.1.111.15.1.2.1.9.1	oraEM4JobAlertRuleOwner	1.3.6.1.4.1.111.15.3.1.1.41.1	oraEMNGEventRuleOwner
1.3.6.1.4.1.111.15.1.2.1.10.1	oraEM4JobAlertMetricName	NA	NA
1.3.6.1.4.1.111.15.1.2.1.11.1	oraEM4JobAlertMetricValue	NA	NA
1.3.6.1.4.1.111.15.1.2.1.12.1	oraEM4JobAlertContext	1.3.6.1.4.1.111.15.3.1.1.44.1	oraEMNGEventContextAttrs
1.3.6.1.4.1.111.15.1.2.1.13.1	oraEM4JobAlertKeyName	NA	NA
1.3.6.1.4.1.111.15.1.2.1.14.1	oraEM4JobAlertKeyValue	NA	NA
1.3.6.1.4.1.111.15.1.2.1.15.1	oraEM4JobAlertSeverity	NA	NA
1.3.6.1.4.1.111.15.1.2.1.16.1	oraEM4JobAlertJobId	NA	NA
1.3.6.1.4.1.111.15.1.2.1.17.1	oraEM4JobAlertJobExecId	1.3.6.1.4.1.111.15.3.1.1.61.1	oraEMNGEventTypeAttr1

* **Note**: When mapping 1.3.6.1.4.1.111.15.1.1.1.7.1 oraEM4AlertKeyValue to a12*c* or 13*c* metric_alert event to 1.3.6.1.4.1.111.15.1.1.1.7.1 oraEM4AlertKeyValue, you must look at 1.3.6.1.4.1.111.15.3.1.1.84.1 oraEMNGEventTypeAttr24.

```
if oraEMNGEventTypeAttr24 is null
       then
         oraEM4AlertKeyValue is null

       if oraEMNGEventTypeAttr24 value =  "Number of keys=1"
         oraEM4AlertKeyValue --> oraEMNGEventTypeAttr8

       if oraEMNGEventTypeAttr24 value =  "Number of keys=x" where x is greater than 1
         => check the values for the following pairs of attributes.
            <oraEMNGEventTypeAttr10, oraEMNGEventTypeAttr11>
            <oraEMNGEventTypeAttr12, oraEMNGEventTypeAttr13>
            <oraEMNGEventTypeAttr14, oraEMNGEventTypeAttr15>
            <oraEMNGEventTypeAttr16, oraEMNGEventTypeAttr17>
            <oraEMNGEventTypeAttr18, oraEMNGEventTypeAttr19>
            <oraEMNGEventTypeAttr20, oraEMNGEventTypeAttr21>
            <oraEMNGEventTypeAttr22, oraEMNGEventTypeAttr23>
            ...
            ...
```

As many pairs as the number of parts present in the key would be populated, the rest of it will be set to null.

For each non-null pair of attributes, the first attribute provides the name for that part of the key and second attribute provides the value for that part of the key.

> **Note:**
>
> OID 1.3.6.1.4.1.111.15.3.1.1.13.1 specifies the event type. Examples:
>
> For a metric_alert event type
>
> oraEMNGEventType=Metric Alert
>
> For a target_availability event type,
>
> oraEMNGEventType=Target Availability
>
> For a job_status_change event type
>
> oraEMNGEventType=Job Status Change

Incidents

The following tables list incident SNMP trap mappings for Enterprise Manager 13.5 RU16.

OID Number	OID Name
1.3.6.1.4.1.111.15.4.1.1.1.1	oraEMNGincidentsIndex
1.3.6.1.4.1.111.15.4.1.1.10.1	oraEMNGIncidentCategories
1.3.6.1.4.1.111.15.4.1.1.11.1	oraEMNGIncidentCategoryCodes
1.3.6.1.4.1.111.15.4.1.1.12.1	oraEMNGIncidentId
1.3.6.1.4.1.111.15.4.1.1.13.1	oraEMNGIncidentOwner

OID Number	OID Name
1.3.6.1.4.1.111.15.4.1.1.14.1	oraEMNGIncidentAcked
1.3.6.1.4.1.111.15.4.1.1.15.1	oraEMNGIncidentStatus
1.3.6.1.4.1.111.15.4.1.1.16.1	oraEMNGIncidentPriority
1.3.6.1.4.1.111.15.4.1.1.17.1	oraEMNGIncidentEscLevel
1.3.6.1.4.1.111.15.4.1.1.18.1	oraEMNGIncidentTargetName
1.3.6.1.4.1.111.15.4.1.1.19.1	oraEMNGIncidentTargetNameURL
1.3.6.1.4.1.111.15.4.1.1.2.1	oraEMNGIncidentNotifType
1.3.6.1.4.1.111.15.4.1.1.20.1	oraEMNGIncidentTargetType
1.3.6.1.4.1.111.15.4.1.1.21.1	oraEMNGIncidentHostName
1.3.6.1.4.1.111.15.4.1.1.22.1	oraEMNGIncidentTargetOwner
1.3.6.1.4.1.111.15.4.1.1.23.1	oraEMNGIncidentTgtLifeCycleStatus
1.3.6.1.4.1.111.15.4.1.1.24.1	oraEMNGIncidentTargetVersion
1.3.6.1.4.1.111.15.4.1.1.25.1	oraEMNGIncidentUserDefinedTgtProp
1.3.6.1.4.1.111.15.4.1.1.26.1	oraEMNGIncidentSourceObjName
1.3.6.1.4.1.111.15.4.1.1.27.1	oraEMNGIncidentSourceObjNameURL
1.3.6.1.4.1.111.15.4.1.1.28.1	oraEMNGIncidentSourceObjType
1.3.6.1.4.1.111.15.4.1.1.29.1	oraEMNGIncidentSourceObjSubType
1.3.6.1.4.1.111.15.4.1.1.3.1	oraEMNGIncidentMessage
1.3.6.1.4.1.111.15.4.1.1.30.1	oraEMNGIncidentSourceObjOwner
1.3.6.1.4.1.111.15.4.1.1.31.1	oraEMNGIncidentSRId
1.3.6.1.4.1.111.15.4.1.1.32.1	oraEMNGIncidentTicketId
1.3.6.1.4.1.111.15.4.1.1.33.1	oraEMNGIncidentTicketStatus
1.3.6.1.4.1.111.15.4.1.1.34.1	oraEMNGIncidentTicketType
1.3.6.1.4.1.111.15.4.1.1.35.1	oraEMNGIncidentRuleSetName
1.3.6.1.4.1.111.15.4.1.1.36.1	oraEMNGIncidentRuleName
1.3.6.1.4.1.111.15.4.1.1.37.1	oraEMNGIncidentRuleOwner
1.3.6.1.4.1.111.15.4.1.1.38.1	oraEMNGIncidentRCADetails
1.3.6.1.4.1.111.15.4.1.1.39.1	oraEMNGIncidentUserComments
1.3.6.1.4.1.111.15.4.1.1.4.1	oraEMNGIncidentMessageURL
1.3.6.1.4.1.111.15.4.1.1.40.1	oraEMNGIncidentUpdates
1.3.6.1.4.1.111.15.4.1.1.41.1	oraEMNGIncidentRCAStatus
1.3.6.1.4.1.111.15.4.1.1.42.1	oraEMNGIncidentAssocEventCount*
1.3.6.1.4.1.111.15.4.1.1.43.1	oraEMNGIncidentCompressedEventMsg1*
1.3.6.1.4.1.111.15.4.1.1.44.1	oraEMNGIncidentCompressedEventType1*
1.3.6.1.4.1.111.15.4.1.1.45.1	oraEMNGIncidentCompressedEventSeverity1*
1.3.6.1.4.1.111.15.4.1.1.46.1	oraEMNGIncidentCompressedEventCreationDate1*
1.3.6.1.4.1.111.15.4.1.1.47.1	oraEMNGIncidentCompressedEventTargetName1*
1.3.6.1.4.1.111.15.4.1.1.48.1	oraEMNGIncidentCompressedEventTargetType1*
1.3.6.1.4.1.111.15.4.1.1.49.1	oraEMNGIncidentCompressedEventMsg2*
1.3.6.1.4.1.111.15.4.1.1.5.1	oraEMNGIncidentSeverity
1.3.6.1.4.1.111.15.4.1.1.50.1	oraEMNGIncidentCompressedEventType2*
1.3.6.1.4.1.111.15.4.1.1.51.1	oraEMNGIncidentCompressedEventSeverity2*
1.3.6.1.4.1.111.15.4.1.1.52.1	oraEMNGIncidentCompressedEventCreationDate2*
1.3.6.1.4.1.111.15.4.1.1.53.1	oraEMNGIncidentCompressedEventTargetName2*
1.3.6.1.4.1.111.15.4.1.1.54.1	oraEMNGIncidentCompressedEventTargetType2*

OID Number	OID Name
1.3.6.1.4.1.111.15.4.1.1.55.1	oraEMNGIncidentCompressedEventMsg3*
1.3.6.1.4.1.111.15.4.1.1.56.1	oraEMNGIncidentCompressedEventType3*
1.3.6.1.4.1.111.15.4.1.1.57.1	oraEMNGIncidentCompressedEventSeverity3*
1.3.6.1.4.1.111.15.4.1.1.58.1	oraEMNGIncidentCompressedEventCreationDate3*
1.3.6.1.4.1.111.15.4.1.1.59.1	oraEMNGIncidentCompressedEventTargetName3*
1.3.6.1.4.1.111.15.4.1.1.6.1	oraEMNGIncidentSeverityCode
1.3.6.1.4.1.111.15.4.1.1.60.1	oraEMNGIncidentCompressedEventTargetType3*
1.3.6.1.4.1.111.15.4.1.1.61.1	oraEMNGIncidentCompressedEventMsg4*
1.3.6.1.4.1.111.15.4.1.1.62.1	oraEMNGIncidentCompressedEventType4*
1.3.6.1.4.1.111.15.4.1.1.63.1	oraEMNGIncidentCompressedEventSeverity4*
1.3.6.1.4.1.111.15.4.1.1.64.1	oraEMNGIncidentCompressedEventCreationDate4*
1.3.6.1.4.1.111.15.4.1.1.65.1	oraEMNGIncidentCompressedEventTargetName4*
1.3.6.1.4.1.111.15.4.1.1.66.1	oraEMNGIncidentCompressedEventTargetType4*
1.3.6.1.4.1.111.15.4.1.1.67.1	oraEMNGIncidentCompressedEventMsg5*
1.3.6.1.4.1.111.15.4.1.1.68.1	oraEMNGIncidentCompressedEventType5*
1.3.6.1.4.1.111.15.4.1.1.69.1	oraEMNGIncidentCompressedEventSeverity5*
1.3.6.1.4.1.111.15.4.1.1.7.1	oraEMNGIncidentRepeatCount
1.3.6.1.4.1.111.15.4.1.1.70.1	oraEMNGIncidentCompressedEventCreationDate5*
1.3.6.1.4.1.111.15.4.1.1.71.1	oraEMNGIncidentCompressedEventTargetName5*
1.3.6.1.4.1.111.15.4.1.1.72.1	oraEMNGIncidentCompressedEventTargetType5*
1.3.6.1.4.1.111.15.4.1.1.8.1	oraEMNGIncidentCreationTime
1.3.6.1.4.1.111.15.4.1.1.9.1	oraEMNGIncidentLastUpdatedTime

* = = These rows will only be populated for compressed incidents. A compressed incident will have 'oraEMNGIncidentAssocEventCount' populated with count greater than 1. For more information on incidents with compression, see

Sample SNMP V3 Trap for Incidents

```
**************V3 TRAP***[24]*****************

ContextEngineId : 80:00:00:2a:01:64:4b:0e:72:00:00:04:93

ContexName :

SecurityLevel : 1

MsgMaxSize : 65429

MsgSecurityModel : 3

1.3.6.1.2.1.1.3.0: 0:21:15.17

1.3.6.1.4.1.111.15.4.1.1.2.1: NOTIF_NORMAL

1.3.6.1.4.1.111.15.4.1.1.3.1: There are 21 agent unreachable events on
targets monitored by agent: sample:1838
1.3.6.1.4.1.111.15.4.1.1.4.1: https://sampleserver.oracle.com:5416/em/
```

```
redirect?pageType=sdk-core-event-console-
detailIncident&issueID=FC538FAFD181116BE053720E4B64E6C6

1.3.6.1.4.1.111.15.4.1.1.5.1: Warning

1.3.6.1.4.1.111.15.4.1.1.6.1: WARNING

1.3.6.1.4.1.111.15.4.1.1.7.1: 0

1.3.6.1.4.1.111.15.4.1.1.8.1: May 22, 2023 6:46:52 PM PDT

1.3.6.1.4.1.111.15.4.1.1.9.1: May 22, 2023 6:46:58 PM PDT

1.3.6.1.4.1.111.15.4.1.1.10.1: Availability

1.3.6.1.4.1.111.15.4.1.1.11.1: Availability

1.3.6.1.4.1.111.15.4.1.1.12.1: 12

1.3.6.1.4.1.111.15.4.1.1.13.1:

1.3.6.1.4.1.111.15.4.1.1.14.1: No

1.3.6.1.4.1.111.15.4.1.1.15.1: New

1.3.6.1.4.1.111.15.4.1.1.16.1: None

1.3.6.1.4.1.111.15.4.1.1.17.1: 0

1.3.6.1.4.1.111.15.4.1.1.18.1:

1.3.6.1.4.1.111.15.4.1.1.19.1:

1.3.6.1.4.1.111.15.4.1.1.20.1:

1.3.6.1.4.1.111.15.4.1.1.21.1:

1.3.6.1.4.1.111.15.4.1.1.22.1:

1.3.6.1.4.1.111.15.4.1.1.23.1:

1.3.6.1.4.1.111.15.4.1.1.24.1:

1.3.6.1.4.1.111.15.4.1.1.25.1:

1.3.6.1.4.1.111.15.4.1.1.26.1:

1.3.6.1.4.1.111.15.4.1.1.27.1:

1.3.6.1.4.1.111.15.4.1.1.28.1:

1.3.6.1.4.1.111.15.4.1.1.29.1:

1.3.6.1.4.1.111.15.4.1.1.30.1:

1.3.6.1.4.1.111.15.4.1.1.31.1:
```

```
1.3.6.1.4.1.111.15.4.1.1.32.1:

1.3.6.1.4.1.111.15.4.1.1.33.1:

1.3.6.1.4.1.111.15.4.1.1.34.1:

1.3.6.1.4.1.111.15.4.1.1.35.1: Agent unreachable ruleset

1.3.6.1.4.1.111.15.4.1.1.36.1: Agent unreachable ruleset,agent unreachable
incident snmp v3 rule

1.3.6.1.4.1.111.15.4.1.1.37.1: SYSMAN

1.3.6.1.4.1.111.15.4.1.1.38.1: This incident has 0 cause events, 0 events
that are neither causes nor symptoms, 21 events for which causal analysis
does not apply and 0 symptom events.

1.3.6.1.4.1.111.15.4.1.1.39.1:
1.3.6.1.4.1.111.15.4.1.1.40.1: Event associated to incident Status by
compression policy (Name= Agent unreachable events for targets monitored by
the same agent; Owner=System Generated) and rule (Name= Agent unreachable
ruleset, AUR incident creation rule; Owner=SYSMAN).

1.3.6.1.4.1.111.15.4.1.1.41.1:

1.3.6.1.4.1.111.15.4.1.1.42.1: 21

1.3.6.1.4.1.111.15.4.1.1.43.1: Agent is Unreachable (REASON = Unable to
connect to the agent at https://sampleserver.oracle.com ...

1.3.6.1.4.1.111.15.4.1.1.44.1: target_availability

1.3.6.1.4.1.111.15.4.1.1.45.1: Warning

1.3.6.1.4.1.111.15.4.1.1.46.1: May 22, 2023 6:46:48 PM PDT

1.3.6.1.4.1.111.15.4.1.1.47.1: Management Services and Repository

1.3.6.1.4.1.111.15.4.1.1.48.1: oracle_emrep

1.3.6.1.4.1.111.15.4.1.1.49.1: Agent is Unreachable (REASON = Unable to
connect to the agent at https://sampleserver.oracle.com ...

1.3.6.1.4.1.111.15.4.1.1.50.1: target_availability

1.3.6.1.4.1.111.15.4.1.1.51.1: Warning

1.3.6.1.4.1.111.15.4.1.1.52.1: May 22, 2023 6:46:48 PM PDT

1.3.6.1.4.1.111.15.4.1.1.53.1: sampleserver1.oracle.com

1.3.6.1.4.1.111.15.4.1.1.54.1: j2ee_application

1.3.6.1.4.1.111.15.4.1.1.55.1: Agent is Unreachable (REASON = Unable to
connect to the agent at https://sampleserver.oracle.com ...
```

```
1.3.6.1.4.1.111.15.4.1.1.56.1: target_availability

1.3.6.1.4.1.111.15.4.1.1.57.1: Warning

1.3.6.1.4.1.111.15.4.1.1.58.1: May 22, 2023 6:46:48 PM PDT

1.3.6.1.4.1.111.15.4.1.1.59.1: sampleserver2.oracle.com

1.3.6.1.4.1.111.15.4.1.1.60.1: j2ee_application

1.3.6.1.4.1.111.15.4.1.1.61.1: Agent is Unreachable (REASON = Unable to
connect to the agent at https://sampleserver.oracle.com ...

1.3.6.1.4.1.111.15.4.1.1.62.1: target_availability

1.3.6.1.4.1.111.15.4.1.1.63.1: Warning

1.3.6.1.4.1.111.15.4.1.1.64.1: May 22, 2023 6:46:48 PM PDT

1.3.6.1.4.1.111.15.4.1.1.65.1: sampleserver3.oracle.com

1.3.6.1.4.1.111.15.4.1.1.66.1: oracle_coherence_cache

1.3.6.1.4.1.111.15.4.1.1.67.1: Agent is Unreachable (REASON = Unable to
connect to the agent at https://sampleserver.oracle.com ...

1.3.6.1.4.1.111.15.4.1.1.68.1: target_availability

1.3.6.1.4.1.111.15.4.1.1.69.1: Warning

1.3.6.1.4.1.111.15.4.1.1.70.1: May 22, 2023 6:46:48 PM PDT

1.3.6.1.4.1.111.15.4.1.1.71.1: Oemrep_Database

1.3.6.1.4.1.111.15.4.1.1.72.1: oracle_database

1.3.6.1.6.3.1.1.4.1.0: 1.3.6.1.4.1.111.15.2

1.3.6.1.6.3.18.1.3.0: 127.0.0.1

**************END V3 NOTIFICATION*****************
```

Part V

Systems Infrastructure

This section contains the following chapters:

- Working with Systems Infrastructure Targets
- Managing Networks
- Managing Storage
- Monitoring Servers
- Managing the PDU
- Managing the Rack
- Managing Oracle SuperCluster
- Monitoring Oracle Operating Systems
- Monitoring Oracle Solaris Zones
- Monitoring Oracle VM Server for SPARC

19 Working with Systems Infrastructure Targets

This chapter describes the Oracle Enterprise Manager Systems Infrastructure plug-in. This chapter covers the following:

Overview of Enterprise Manager Systems Infrastructure

A host is a computer where managed databases and other services reside. A host is one of many components or targets that Oracle Enterprise Manager Cloud Control monitors. Oracle Enterprise Manager Cloud Control is an enterprise-level data center management solution for the Oracle product stack, from applications to storage disks, as shown in Figure 19-1.

Figure 19-1 Oracle Product Stack

The Enterprise Manager Systems Infrastructure (EMSI) plug-in is a fully integrated software plug-in that provides monitoring and an enterprise-wide view the bottom half of the stack, including Oracle Solaris and Linux operating systems, virtualized operating systems (zones) and virtual machines (logical domains), servers, storage appliances, storage for a host, and network resources.

The Oracle Enterprise Manager Cloud Control console interface provides detailed information about managed components and shows the relationship between the components. You can drill down to view greater details about specific components or metrics.

For server targets, you can view server details, including power usage, network information, service processor configuration, and fan and temperature information. Other hardware targets include chassis, racks, power distribution units, and network equipment. You can view the storage resources and their components for a host or a storage appliance, including storage pools, storage hardware, filesystems, logical volumes, and Logical Units (LUNs).

For Oracle Solaris and Linux operating system targets, you can view resource and process information, top consumers, performance, and open incidents details. In addition, you can view details about Oracle Solaris alternate boot environments and zones.

When you use Oracle VM Server for SPARC to virtualize hardware or Oracle Solaris to virtualize operating systems, the details appear on virtualization platform pages. Virtualization platform pages provide high-level details such as overall guest, CPU and memory, incident, and configuration information. You can drill down to individual virtual server pages to view metrics specific to a selected logical domain or zone.

The Enterprise Manager Systems Infrastructure also provides a view of Oracle SuperCluster engineered systems. Oracle SuperCluster integrates SPARC compute nodes, an Oracle ZFS Storage Appliance, InfiniBand switches, PDUs, and Exadata Storage Servers into a multi-rack system. You can view all of the components of an Oracle SuperCluster including Compute Nodes, Exadata Storage Servers, InfiniBand Switches, Power Distribution Units, and ZFS Storage Servers. You can expand the Compute Nodes to view the logical domains and zones.

About Monitoring for the Systems Infrastructure Targets

A series of monitoring rules and parameters monitor your managed targets. Alerts and incidents are raised for resources that are not performing as expected.

Each target has a dashboard that displays details based on the target type. All dashboards show the number of warning, critical, and fatal open incidents with links that enable you to drill down to the Incident Manager for the specific incident.

A Management Agent deployed on the host gathers information and keeps track of activity, status, performance and the health of the targets. You can view metrics for the following discovered Oracle target types:

- Servers
- Storage servers
- Networks
- Racks and power distribution units (PDUs)
- Oracle Solaris and Linux operating systems
- Oracle Solaris Zones
- Oracle VM Server for SPARC and logical domains
- Oracle SuperCluster engineered systems

About Dynamic Views for the Systems Infrastructure Targets

Depending on the type of target, a target home page might include graphs, charts and tables to provide greater detail at a glance. The home pages of more complex targets include dynamic photorealistic views and relationship charts. In some cases, you can interact with the images to better understand how hardware is deployed and how resources are utilized.

The following charts and views enable you to assess a target quickly and determine the relationships at a glance:

- Relationship Chart: Displays how resources are allocated among guests. For example, the Systems Infrastructure Virtualization Platform page of an Oracle VM Server for SPARC includes a Core Distribution tab that displays the vCPU and core allocation. The chart contains concentric circles with segments that display which CPUs and cores are allocated to which guests, and which CPUs and cores are not allocated. You can click a guest in the outer ring to view detailed information about that guest's resource consumption.

- Photorealistic View: Displays the components and ports of a hardware target, and if there are open incidents. For example, the Oracle SuperCluster engineered system monitoring pages provide a photorealistic view, which enables you to see how the system is laid out in the rack. All active targets in the system appear in the image. You can view greater detail by hovering your mouse over a target in the image. When a component of the engineered system has an open incident, the component appears in the image with a red border.
- Schematic View: Displays a symbolic view that displays the labels of an engineered system's components. At a glance, you can see the LED status (up, down, or blackout) and temperature of the server, ZFS Storage Appliance Server, InfiniBand Switch, and PDU in the engineered system.

Overview of the Systems Infrastructure User Interface

Information gathered by the Systems Infrastructure plug-in appears in an updated user interface. Each target home page is slightly different, depending on the type of target and whether the target utilizes Oracle VM Server for SPARC or Oracle Solaris Zones virtualization technology.

Open incidents, resource utilization and metrics for a target appear in a dashboard, helping you to maintain high availability and optimized performance. Tabs in the user interface contain more detailed metric information. Information appears in graphs, tables, charts, and schematic, and photorealistic views to help you to quickly understand the status and relationships between components.

About the Target Home Page

The home page enables you to quickly view the status, identify potential resource issues, and view the service request and configuration history of a specific target. From this page, you can drill down to specifics for an open incident, view detailed metrics, and guest details.

You can access the home pages from the All Targets menu. Each home page includes a dashboard across the top. Greater details and historic graphs are available in the tabs on the right side of the page.

The dashboard is designed to display an overview of important information and information that you might want to monitor closely. The information appears in a series of sections, called dashlets. At least four dashlets appear in the dashboard. The target type determines the number of dashlets and the content. Click the small button below the row of dashlets to toggle to the next series of dashlets. Below the navigation buttons is a single button that gives you the option to minimize the dashboard.

About the Target Home Page is an example of the first three dashlets on an Oracle ZFS Storage Server target home page. This example shows target type and date the page was last refreshed. The first dashlet contains target details, the second dashlet shows the number of open incidents and the severity level. The third dashlet shows the amount of space used and available for this target. One or more buttons appears beneath the Open Incidents dashlet, as indicated by the red box. Click a button to navigate to the next series of dashlets.

Figure 19-2 Dashboard

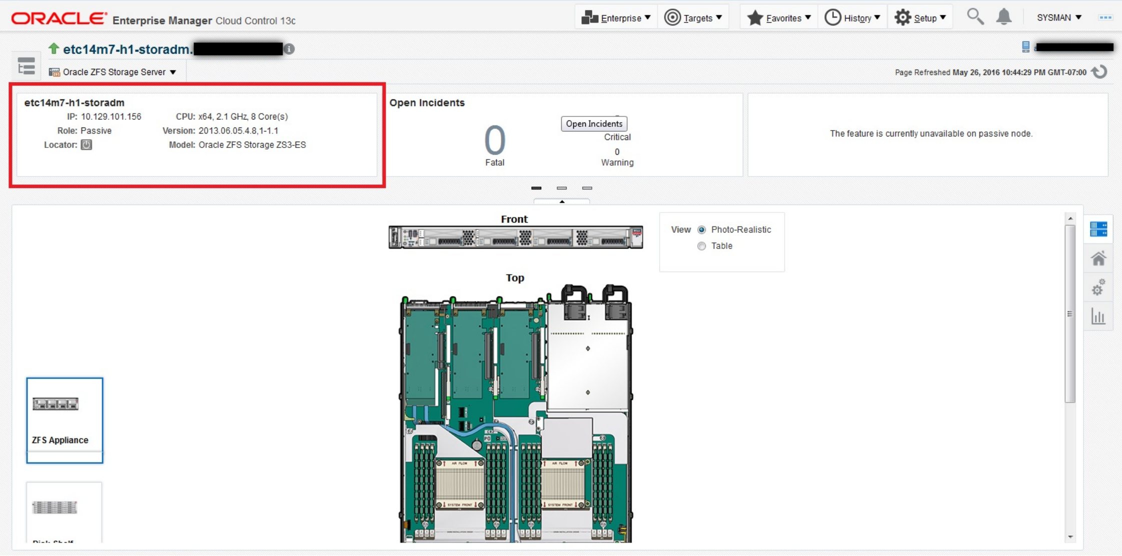

etc14m7-h1-storadm.[redacted]

[redacted]

Oracle ZFS Storage Server ▾

Page Refreshed May 26, 2016 10:44:29 PM GMT-07:00

etc14m7-h1-storadm

IP: 10.129.101.156
Role: Passive
Locator:

CPU: x64, 2.1 GHz, 8 Core(s)
Version: 2013.06.05.4.8,1-1.1
Model: Oracle ZFS Storage ZS3-ES

Open Incidents

0 Fatal

Open Incidents
Critical
0 Warning

The feature is currently unavailable on passive node.

View ● Photo-Realistic ○ Table

While the content will differ, each home page provides a consolidated view of the status, resource utilization, and metrics for the target. The following are some of the details that appear in dashlets:

- Target details, such as the name and status of the selected target, appears in the first dashlet.
- Incident information appears in the second dashlet. The number of Critical, Warning and Informational incidents that are currently open appear in this dashlet. Each number is a link to greater detail.
- Resource usage, top statistics, and metrics information specific to the target type appears in one or more dashlets.
- The number of service requests for the selected target that were filed with Oracle in the past 30 days and the past 5 days appears in a dashlet.
- Last configuration change and last reported incident time appear in the last dashlet.

In some cases, you can click content in a dashlet to drill down to get more information. For example, the Open Incidents dashlet links to more detailed information. When you click a number next to the type of incident, the dashboard flips and shows a table of incidents with the target name, a synopsis, and additional information depending on the type of incident. You can drill down further to navigate to the Incident Manager for the highest level of detail.

See About the Virtualization Home Page and About the Oracle Engineered Systems Home Page for some differences in the home pages for these types of targets.

About the Virtualization Home Page

When you are using the Oracle VM Server for SPARC or Oracle Solaris Zone virtualization technology, the target home pages are different.

The following are the target home pages for virtualization:

- Virtualization Platform page
- Virtual Server page

The Virtualization Platform page is the home page for the Control Domain or for the Global Zone. The Virtualization Platform page has a Target Navigation icon in the upper left corner that contains links to the Virtual Server home page for each associated logical domain or zone.

For zones and logical domains, the number of incidents that appear on the Virtualization Platform page includes incidents for all associated zones and domains. For example, the control domain and all associated logical domains.

About the Oracle Engineered Systems Home Page

For Oracle SuperCluster engineered systems, you can monitor the components of the engineered system from the target's home page. The monitoring pages provide a photorealistic view and a schematic view of the engineered system. The photorealistic view is helpful in seeing how the system is physically laid out in the rack. You can hover over the image to view details about the components. The schematic view displays the component labels and LED status.

Creating Roles for Systems Infrastructure Administration

The Systems Infrastructure plug-in does not define its own roles for access controls. To manage the plug-in, you must create roles and administrators, and then assign roles to administrators. The roles restrict a user's privileges.

> **Note:**
>
> For security reasons, Oracle recommends that the SYSMAN account be used only as a template to create other accounts, and not used directly.

To create roles to provide management rights to users:

1. Log in to the Enterprise Manager Cloud Control as the super administrator user.
2. Click **Setup**, then **Security**.
3. Select **Roles**.

   On the Security page, a list of predefined roles is provided. These roles can serve as basis to define custom roles to suite specific site level requirements.

   > **Note:**
   >
   > The predefined roles provided cannot be edited or deleted.

4. Select a role that closely matches the role you wish to create. Click **Create Like**.
5. On the Properties page, enter a name for the new role. You can optionally add a description. Click **Next**.
6. On the Roles page, select the roles from the list of Available Roles. Click **Move** to add the role to Selected Roles. Click **Next**.
7. On the Target Privileges page, select the privilege you want to grant to the new role. Click **Next**.

8. On the Resource Privileges page, you can edit specific privileges to be explicitly granted. Click the Manage Privilege Grant edit icon to make the changes. Click **Next**.
9. On the Administrators page, select the administrators from the list of Available Administrators that you want to grant the new role to. Click **Move** to add the administrator to Selected Administrators. Click **Next**.
10. On the Review page, a complete summary of the new role you have created is displayed. Click **Back** to go to previous screens to make changes. Click **Finish** to complete the role creation.

When the newly created administrator logs in, unlike SYSMAN, the administrator is restricted by the privileges set.

Related Resources for Systems Infrastructure Targets

See the following for more information:

- Discovering and Adding Host and Non-Host Targets
- Discovering, Promoting, and Adding System Infrastructure Targets
- Using Incident Management

20 Managing Networks

The following information is included:

Get Started with Managing Networks

Enterprise Manager discovers and manages targets that populate Layers 1, 2, and 3 of the Open Systems Interconnection (OSI) model, shown in Figure 20-1. In the OSI model, a layer supports the layer above it and is provisioned by the layer below it. Any faults in a layer affect the layer above it.

Figure 20-1 OSI Model

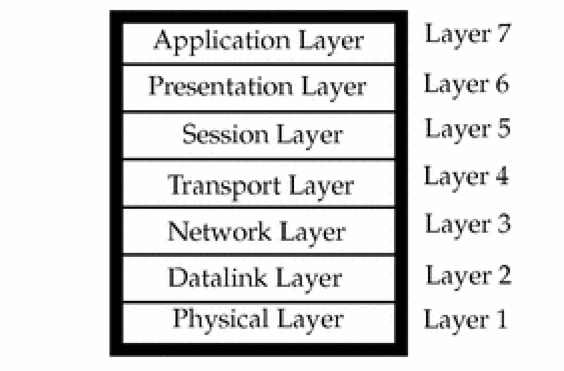

Application Layer	Layer 7
Presentation Layer	Layer 6
Session Layer	Layer 5
Transport Layer	Layer 4
Network Layer	Layer 3
Datalink Layer	Layer 2
Physical Layer	Layer 1

- Layer 1: Fabric targets are in the physical layer of the OSI model, such as network switches.
- Layer 2: Datalink targets are in the link layer of the model. For Ethernet fabrics, datalinks are VLAN IDs. For InfiniBand fabrics, datalinks are partition keys. This layer also includes combinations of links, such as link aggregations.
- Layer 3: Network targets are in the network layer of the model, which are network interfaces plumbed as IP addresses and IPMP groups.

Network management includes all the activities for making network resources available to targets:

- Discover fabrics and their datalinks and networks.
- Collect metrics for layer-related configuration and performance.
- Detect faults at each level.
- View the topology, a graphic view of the relationships in the fabrics and networks.

Location of Network Information in the User Interface

Table 20-1 shows where to find information.

Table 20-1 Location of Network Information in the BUI

Resource	Location
To see fabrics	View All Targets and select **Ethernet/InfiniBand Fabric**.
To see networks	View All Targets and select **Systems Infrastructure Network**. Select a network.
To see datalinks	View All Targets, then select either **Systems Infrastructure Network** or **Host**. Select a network, click **Network Connectivity** tab, then click **Data links** option.
To see network switches	View All Targets and select **System Infrastructure Oracle InfiniBand Switch** or **System Infrastructure Cisco Switch**.
To see the virtualization host that is using a network	Display a network and view Connected Nodes section or click **Network Members** tab.
To see incidents for a network	View Networks Incident dashlet.
To see status of the network switch	View All Targets and select **Systems Infrastructure Switch** then view the Temperature, the Fan performance, and Throughput dashlets.
To view the performance of a network switch	View All Targets and select **Systems Infrastructure Switch** then select a network switch. Click the **Performance** tab. You can refresh the display at any time.
To view the metrics that are being monitored for a network switch	View All Targets and select **Systems Infrastructure Switch**. Right-click on Systems Infrastructure Switch and select **Monitoring,** then **All Metrics**.
To see ports of the network switch	View All Targets and select **Systems Infrastructure Switch**. Select a network switch. A grid shows the used and available ports.
To discover a fabric	Use one of the **Add Target** procedures.
To delete a fabric, datalink, or network	Use the **Remove Target** action.

Actions for Network Management

You can perform the following actions, depending on the requirements.

- Discover switches, and network-connected targets. When you discover a network-connected target such as a fabric, its datalinks and networks also become discovered assets.
- View the configuration of fabrics, datalinks, and networks.
- Diagnose problems using incidents and performance metrics.
- Modify how the targets are monitored and how performance is measured.

View Topology

The topology shows the relationships among assets. The network topology consists of ports, datalinks, and network interfaces.

1. Select **Targets** and then **All Targets**.
2. For a fabric, select **Ethernet/InfiniBand Fabric**.

   For a network switch, select **Systems Infrastructure Switch**.

   For a network, select **Systems Infrastructure Network**.
3. Select one of the targets.
4. On the target's landing page, click the down arrow next the type of target to display a menu.
5. For a fabric, select **Fabric Topology** as shown in Figure 20-2.

Figure 20-2 Menu Selections for Topology

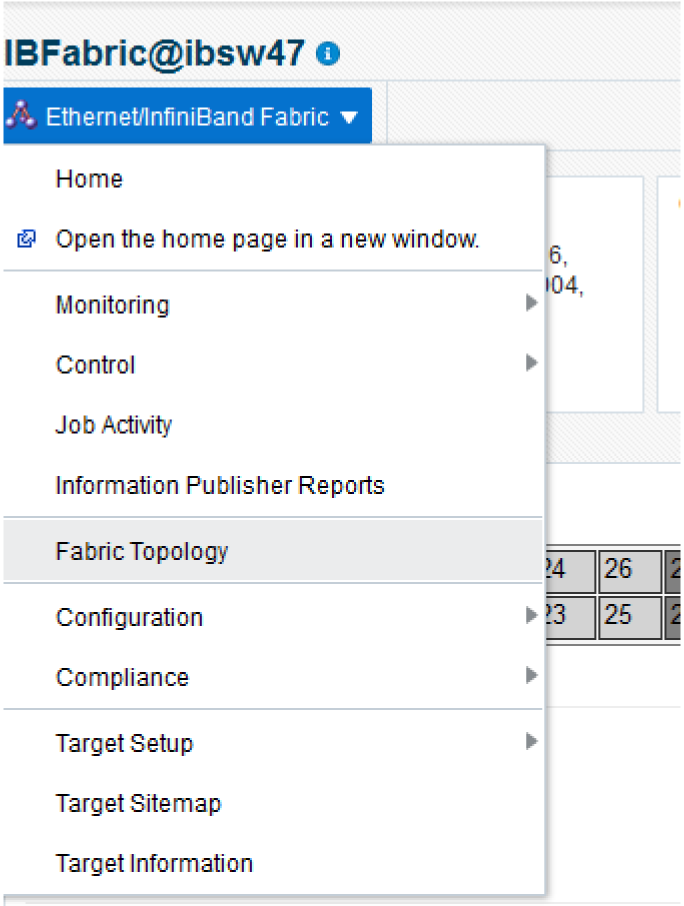

6. The **Topology** page shows the relationships of this target. You can adjust the display of the topology to learn more about the target. Figure 20-3 shows the assets that use the selected fabric because the **Filter** options include them.

Figure 20-3 Fabric Topology

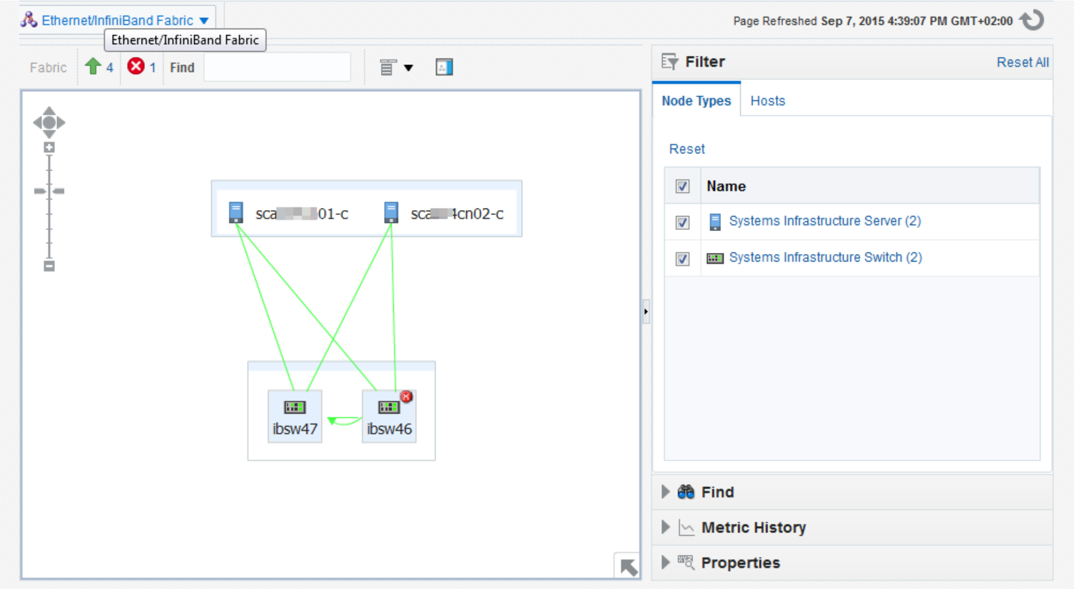
Ethernet/InfiniBand Fabric
Ethernet/InfiniBand Fabric
Page Refreshed Sep 7, 2015 4:39:07 PM GMT+02:00
Fabric
4
1
Find
Filter
Reset All
Node Types
Hosts
Reset
Name
Systems Infrastructure Server (2)
Systems Infrastructure Switch (2)
sca[redacted]01-c
sca[redacted]4cn02-c
ibsw47
ibsw46
Find
Metric History
Properties

7. Hover on an object in the topology to display the asset's information. Figure 20-4 shows the result of hovering over a switch of a fabric.

Figure 20-4 Target Topology

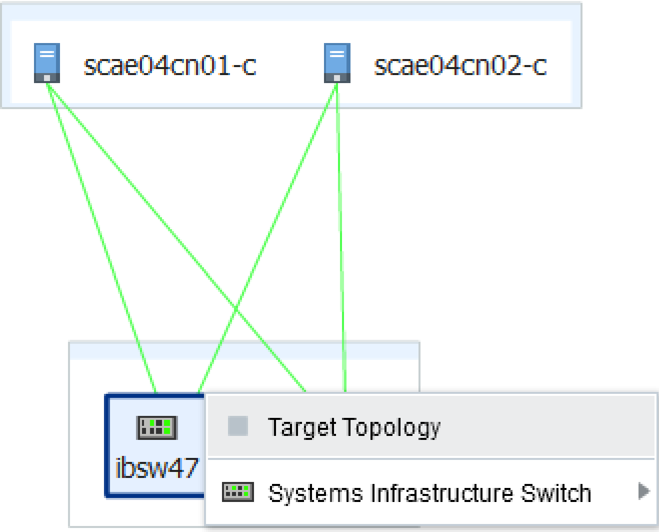

8. To view a different target, click **All Targets** again and then click the **Remove** icon next to the target type to remove the filter.

Fabric

- About Fabrics
- View Information About Fabrics
- About Fabric Information

- About Performance of Fabrics
- Delete a Fabric

About Fabrics

A fabric represents the physical network targets including the connection between targets, that is, a port. For a switch, all ports support the fabric. For a network-enabled device such as a server or a storage appliance, its port is its connection to the fabric. As the physical layer of the OSI model, any fault in the fabric affects the datalinks and networks that rely on the fabric.

A fabric's name is based on the name of the target. For example, when you discover an InfiniBand switch named `abcp01sw-ib02` , the fabric is named `IBFabric@abcp01sw-ib02.us.example.com`. When the switch is a component of a system, the fabric is named for the first target in the system to be discovered. A fabric with the name `ETHFabric@abcp01cn02.us.example.com` indicates that this Ethernet fabric followed the discovery of a compute node.

Enterprise Manager can discover and manage Ethernet fabrics and InfiniBand fabrics. For information about discovering fabrics, see Discovering, Promoting, and Adding System Infrastructure Targets. After discovery, you can view the attributes that these fabrics have in common and also their unique attributes.

View Information About Fabrics

To see information about a fabric:

1. Select **Targets** and then **All Targets**.
2. Select **Ethernet/InfiniBand Fabric**.
3. Select one of the fabrics. The landing page for the fabric shows the incidents for this fabric and the nodes that use this fabric.

Figure 20-5 Fabric Landing Page

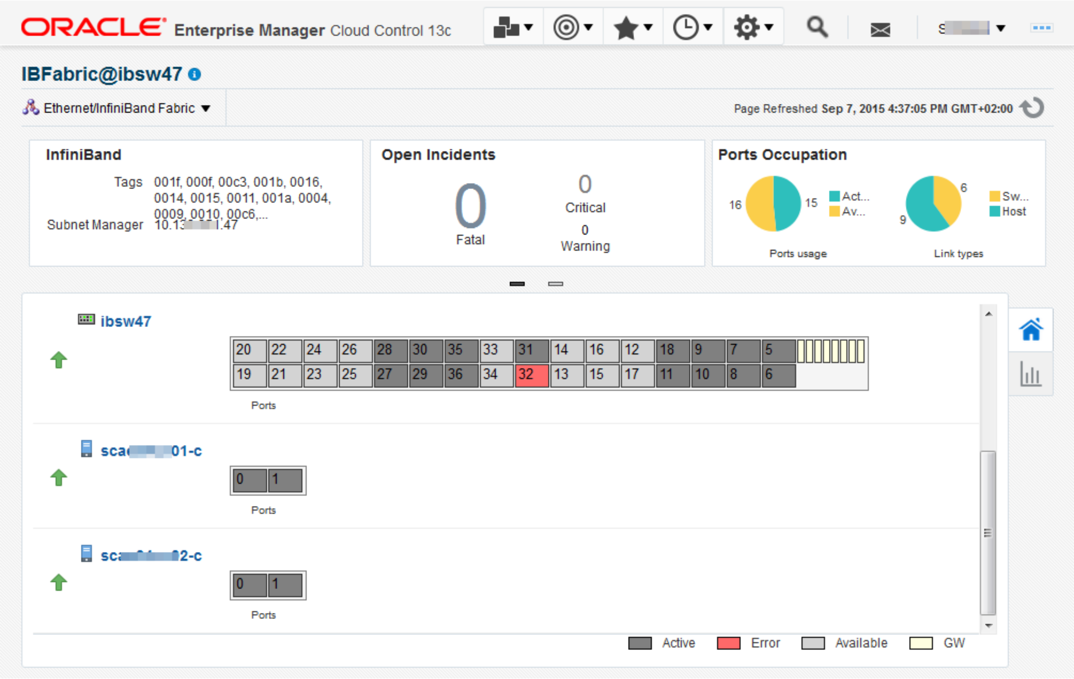

IBFabric@ibsw47

Ethernet/InfiniBand Fabric ▼

Page Refreshed Sep 7, 2015 4:37:05 PM GMT+02:00

InfiniBand

Tags 001f, 000f, 00c3, 001b, 0016, 0014, 0015, 0011, 001a, 0004, 0009, 0010, 00c6,...

Subnet Manager 10.13[redacted].47

Open Incidents

0 Fatal

0 Critical

0 Warning

ibsw47

20	22	24	26	28	30	35	33	31	14	16	12	18	9	7	5
19	21	23	25	27	29	36	34	32	13	15	17	11	10	8	6

Ports

sca[redacted]01-c

0	1

Ports

sc[redacted]2-c

0	1

Ports

4. Click the **Fabric Details** icon or the **Fabrics Performance** icon to view more information about this fabric.

Figure 20-6 Fabric Performance Icon

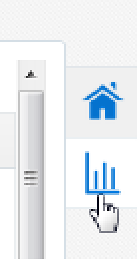

About Fabric Information

The Fabric dashboard's dashlets include the following information:

- VLAN IDs or partition keys for this fabric:
  - For an Ethernet fabric with tagged VLANs, this section lists each VLAN ID. If there are no tagged VLANs, this section displays `Untagged`.
  - For an InfiniBand fabric, this section lists each partition key.
- Open Incidents
- Ports Occupation
- Last Configuration time

The Connected Nodes section lists every node, or network switch, in the fabric. You can also display specific information by hovering on the node, as show in Figure 20-7.

Figure 20-7 Fabric's Nodes

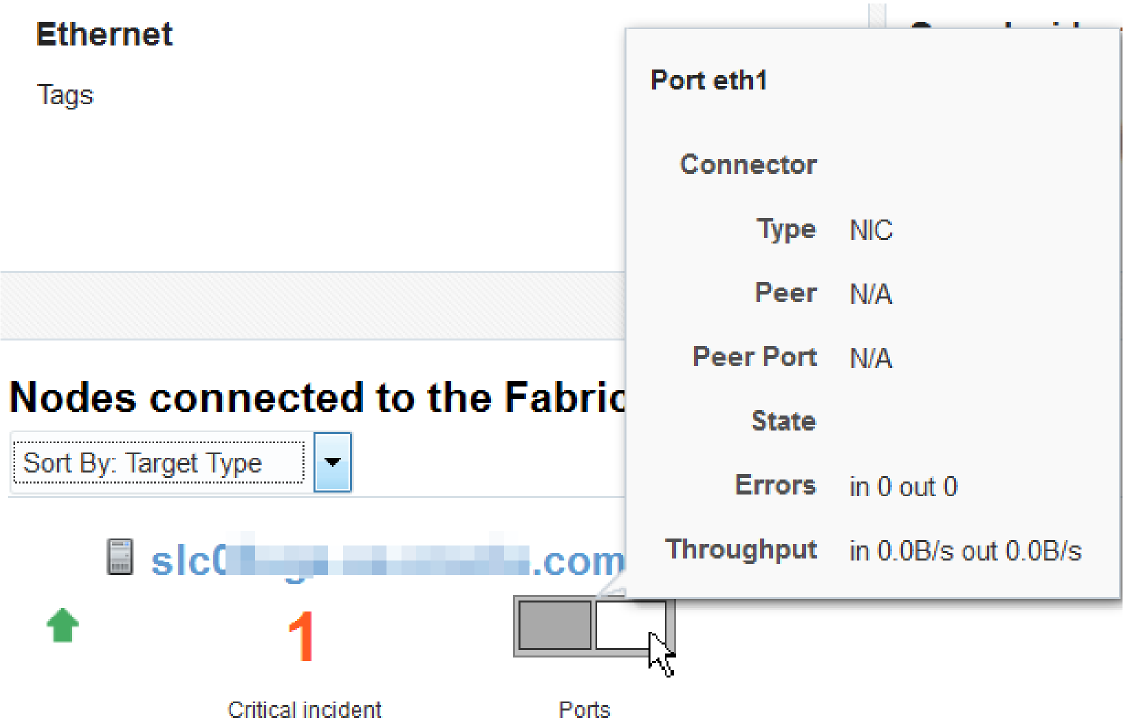

Ethernet

Tags

Nodes connected to the Fabric

Sort By: Target Type

slc[redacted].com

1

Critical incident

Ports

Port eth1

Connector

Type	NIC
Peer	N/A
Peer Port	N/A
State	
Errors	in 0 out 0
Throughput	in 0.0B/s out 0.0B/s

About Performance of Fabrics

Metrics are collected and displayed in real-time when possible. For the ILOM service processors of networks switches, the collection of metrics is not in real time. In these cases, the Last Collected metrics are displayed with a timestamp.

You can view incidents for a fabric or a network switch in the fabric and view a record of performance.

One factor that effects both performance and metric collection is the version of SNMP that you specify when you discover a network switch using the **Add Using Guided Process** wizard. Although all three versions of SNMP are supported, Version 3 is recommended because it is the most secure and the most efficient. Version 2c, which uses public community strings instead of credentials to get access to the network switch, is the default version and provides efficient performance. Version 1 of SNMP, which is not recommended, is not secure and collects metrics using multiple inquiries. If you require Version 1, you must increase SNMP's timeout value to at least 180 seconds. Table 20-2 summarizes the options for SNMP specification.

Table 20-2 Supported Versions of SNMP

Version of SNMP	Level of Security	Level of Performance	Recommendation
Version 3	High: requires username and password	High: Metrics are collected in bulk by one inquiry.	Recommended
Version 2	Not secure: Public community string	Acceptable if Version 2c is used, which can collect metrics in bulk. Other sub-versions require extended expiration intervals.	Performance can be acceptable, but security is poor.
Version 1	Not secure: Public community string	Requires expiration interval of at least 180 seconds.	Not recommended

Delete a Fabric

A fabric is deleted automatically when all targets connected to the fabric are removed. To remove a fabric, use the **Remove Target** action to remove each target. After the last target is removed, the fabric is deleted.

Datalinks

- About Datalinks
- View Information About Datalinks

About Datalinks

When you discover a network switch or network-enabled device, its datalinks are discovered including:

- Target Components

- LAG Configuration
- LAG Membership
- Bonding Configuration
- IPMP Configuration
- IPMP Membership

The datalink layer manages networks that share resources by means VLANs (tagged and untagged VLAN IDs), InfiniBand partitions (default and assigned partition keys), and Fibre Channel zones. Datalinks can also depend on other datalinks to increase throughput or reliability.

View Information About Datalinks

To see information about datalinks:

1. Select **Targets** and then **All Targets**.
2. Select **Host**.
3. Select a host.
4. In the summary, click the **Network Connectivity** tab.
5. Click the **Data links** option.

Figure 20-8 Data links

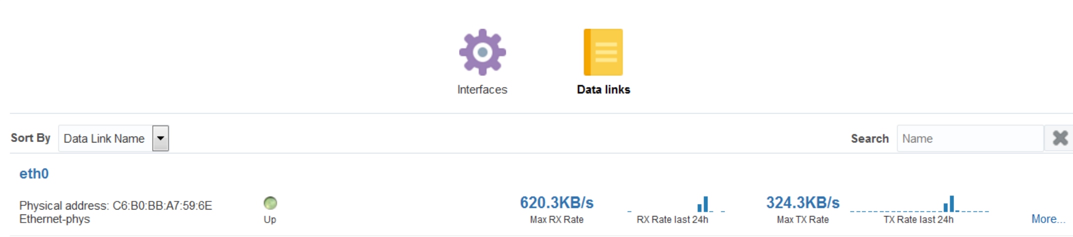

eth0

Physical address: C6:B0:BB:A7:59:6E
Ethernet-phys

Up

620.3KB/s
Max RX Rate

324.3KB/s
Max TX Rate

6. Click the **More...** option to see the path and the transmission rate.

Networks

- About Networks
- View Information About Networks
- Delete Networks

About Networks

When you discover a network switch or network-enabled device, its networks are discovered. You can view the following information about a selected network:

- Metric and Collection Settings
- Metric Collection Errors
- Status History
- Incident Manager

- Alert History
- Blackouts and Brownouts
- Create Blackout...
- End Blackout...
- All Metrics
- Create Brownout...
- End Brownout...
- Job Activity
- Information Publisher Reports
- Last Collected
- Comparison & Drift Management
- Compare...
- Search...
- History
- Save...
- Saved
- Topology
- Compliance Results
- Standard Associations
- Real-time Observations
- Monitoring Configuration
- Administrator Access
- Remove Target...
- Add to Group...
- Properties
- Target Sitemap
- Target Information

When you discover an operating system, you can view the following information about an OS target's network connectivity:

- Network interfaces: MAC address, IP address, netmask and broadcast address
- IPV4 and Ipv6
- Routing tables
- Open network ports

View Information About Networks

To see information about a network switch:

1. Select **Targets** and then **All Targets**.
2. Select **Ethernet/InfiniBand Fabric**.

3. Select one of the fabrics.
4. Select one of the network switches.
5. Use the tabs, such as the **Performance** tab, to view information about this target.

To see home page of **Systems Infrastructure Switch** :

1. Select **Targets** and then **All Targets**.
2. Select **Systems Infrastructure Network**.
3. Select a subnet.

Delete Networks

A network is deleted automatically when all targets connected to the network are removed. To remove a network, use the **Remove Target** action to remove each target. After the last target is removed, the network is deleted.

View Network Details of a Host Target

You can select the host and then its operating system or you can select the operating system and then choose the server from the list, as described in this procedure.

1. Select **Targets** and then **All Targets**.
2. Scroll down to the Operating System section and select the type of operating system. All servers with that operating system are listed.
3. Select a host target.
4. In the OS dashboard, scroll down to the Network Usage chart for an overview of network activity. You can change the values of this graph and its time interval.

Figure 20-9 Graph of Network Usage

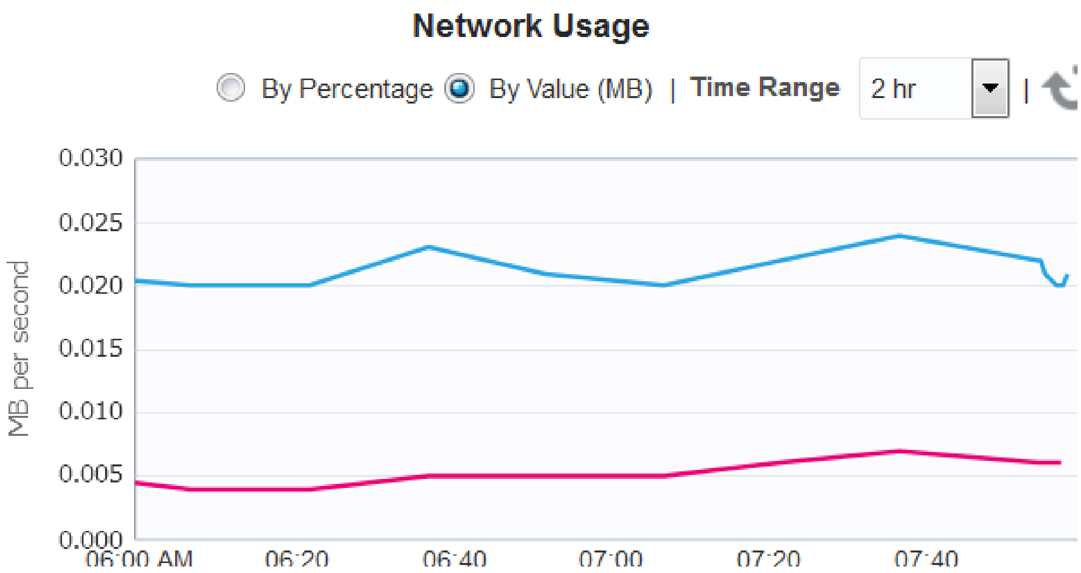

5. Click the **Network Connectivity** tab to view details. You can then click the **Interfaces** or the **Data links** options.

Related Resources for Network Management

See the following for more information:

- Discovering and Adding Host and Non-Host Targets
- Discovering Fabrics
- Using Incident Management
- Monitoring and Managing Targets

21 Managing Storage

The following information is included in this chapter:

Get Started with Managing Storage

Oracle Enterprise Manager discovers and manages storage resources and their components, including storage pools, storage hardware, filesystems, and logical volumes.

A storage server exposes its resources as Logical Units (LUNs) or Shares. A storage client has access to these LUNs and shares, and represents them as its own filesystems or devices and also exposes these filesystems to its own clients, such as applications, databases, or virtual machines.

The System Infrastructure plug-in manages the storage resources for a Storage Appliance and for a Host.

Storage management includes all the activities for making storage resources available to operating systems and virtual assets. The following are some of the activities:

- Discover storage targets and derive their relationship.
- Collect configuration and performance metrics for all storage target components.
- Detect incidents.
- Monitor the storage target components such as, LUNs, filesystems, projects, and pools.

Location of Storage Information in the User Interface

Table 21-1 shows where to find information.

Table 21-1   Location of Storage Information in the User Interface

Object	Location
To view Oracle ZFS Storage Appliance target	Select **Oracle ZFS Storage Server** in the All Targets selector. Click the Oracle ZFS Storage Server to display the Summary page.
To view host's storage information	Select **Hosts** in the All Targets selector. Click the specific host to display the Summary page. Click the Storage icon to the right. To view more details, click **Disks**, **Filesystems**, **SAN Configuration**, **Linux LVM Volume Group(s)**, or **ZFS Storage Pool(s)** icons which appear to the left.
To view storage topology	Right-click on a storage appliance in the All Targets selector. Select **Configuration** and then click **Topology**.
To view storage performance metrics	Right-click on a storage appliance in the All Targets selector. Select **Monitoring** and then click **All Metrics**.
To view storage configuration metrics	Right-click on a storage appliance in the All Targets selector. Select **Configuration** and then click **Last Collected**.
To view storage cluster membership	Select a storage appliance cluster in the All Targets selector. Click on **Target Navigation** icon.

Actions for Storage Management

You can perform the following actions, depending on the requirements:

- Discover storage targets.
- View the configuration of storage targets.
- Diagnose incidents and performance metrics.
- Modify how the targets are monitored and how performance is measured.
- Blackout a storage appliance target which in turn, performs blackout on all of its associated automatically promoted target members.

About Storage Appliance Dashboard

The storage appliance dashboard contains basic information about the storage appliance's status, including incidents, space usage, disk I/O activity graphical representations, number of LUNs, number of shares, and last configuration details in the form of eight dashlets. See Viewing the Storage Appliance Dashboard for the steps to view the dashboard of a selected storage appliance.

The dashlets are displayed as shown in Figure 21-1. Click the icon beneath the dashboard to switch to another set of dashlets.

Figure 21-1 Dashlets View

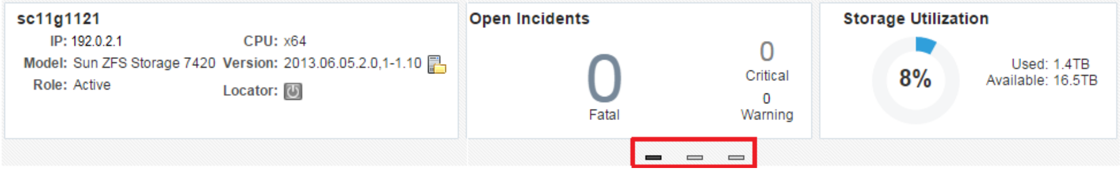

sc11g1121

IP: 192.0.2.1 CPU: x64

Model: Sun ZFS Storage 7420 Version: 2013.06.05.2.0,1-1.10

Role: Active Locator:

Open Incidents

0 Fatal

0 Critical

0 Warning

Storage Utilization

8%

Used: 1.4TB
Available: 16.5TB

You can view the following dashlets in the storage appliance dashboard:

- Basic Hardware Information: Displays the model name of the appliance or storage server, IP Address, CPU, Version, Role, and Locator details. It also provides information if a new appliance version is available for download.
- Open Incidents: Displays the number of open incidents in the fatal, critical, and warning categories. Click the number displayed as the open incident in this dashlet to view detailed information about the incidents in that category.
- Storage Utilization: Displays the percentage of used space and available space of the appliance or storage server. It displays the storage utilization at the appliance level.
- Disk I/O Bytes/Sec (Last 24 Hour): Displays the graphical representation of overall disk I/O activity in bytes per second on the storage appliance.
- Distribution of Logical Units Utilization: Displays the total number of LUNs used in the form of graphical representation. It also provides distribution of LUNs based on storage utilization percentage.
- Distribution of Share Utilization: Displays the total number of shares used in the form of graphical representation. It also provides distribution of shares based on storage utilization percentage.
- Disk I/O Ops/Sec (Last 24 Hour): Displays the graphical representation of the overall disk I/O operations per second on the storage appliance.
- Last Configuration details: Displays the date and time information of the last configuration change and last reported incident.

Viewing the Storage Appliance Dashboard

To view a storage appliance's dashboard, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Servers, Storage and Network** target type.

   You can view a list of appliance targets under Servers, Storage and Network.
3. Click a specific appliance target under Servers, Storage and Network.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. Click on a specific target name from the list in the right pane to view the storage appliance's dashboard.

Viewing Storage Appliance Cluster Dashboard

To view a storage appliance cluster's dashboard, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Groups, Systems and Services** target type.

   You can view a list of cluster targets under Groups, Systems and Services.
3. Click a specific cluster target under Groups, Systems and Services.

   The target name, target type, and the target status of the selected cluster target is displayed in the right pane.
4. Click on a specific cluster target name from the list in the right pane to view the storage appliance cluster's dashboard.

About Photorealistic Image

The Hardware tab in the main window of the storage appliance user interface displays a photorealistic view of the selected appliance, including the front, top, and rear as shown in Figure 21-2. See Viewing the Photorealistic Image for the steps to view the photorealistic image of a selected storage appliance.

You can hover over any component displayed in the photorealistic image to view additional information about that component.

Figure 21-2   Hardware Tab View

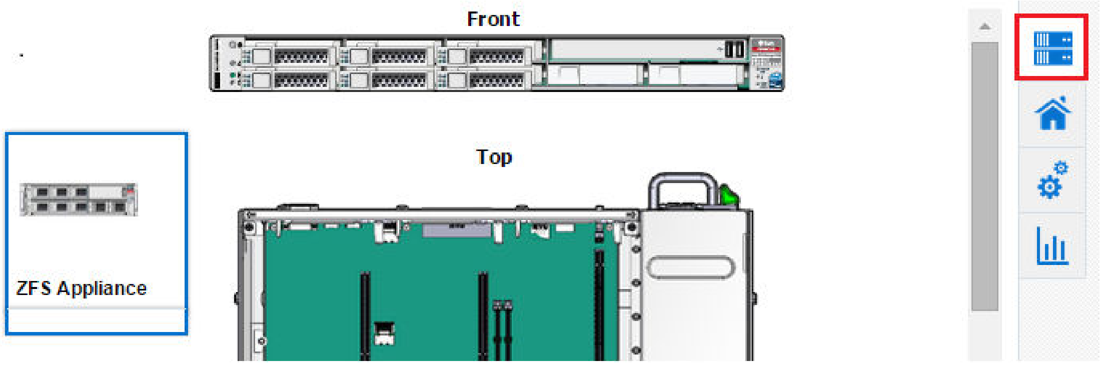

You can view the following hardware information of an appliance, if it is available and relevant to the component:

- Component Name
- Model Name
- Architecture
- Manufacturer
- Serial Number
- Critical incidents information
- Part Number
- Size of memory components
- Total Cores of the CPU
- Enabled Cores of the CPU

Viewing the Photorealistic Image

To view the Hardware tab details of an appliance, perform the following steps:

> **Note:**
>
> Photorealistic view of a ZFS appliance is not supported in ZS5.2.

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Servers, Storage and Network** target type.

   You can view a list of appliance targets under Servers, Storage and Network.

   > **Note:**
   >
   > To view the cluster targets, expand **Groups, Systems and Services** in the left navigation pane. Select a specific cluster target in the right pane.

   > **Note:**
   >
   > Photorealistic view of the ZFS appliance not supported in ZS5.2.

3. Click a specific appliance target under Servers, Storage and Network.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. Click on a specific target name from the list in the right pane.
5. Click the **Hardware** tab which appears to the right of the storage appliance user interface to view the photorealistic image of the target selected.

About Summary

The Summary tab in the main window of the storage appliance user interface displays detailed information about the hardware components and the active Storage Services. See Viewing the Summary for the steps to view the summary of a selected storage appliance.

When you select an appliance or a storage server to view detailed information, the user interface opens with the Summary tab view, as shown is Figure 21-3.

Figure 21-3 Summary Tab View

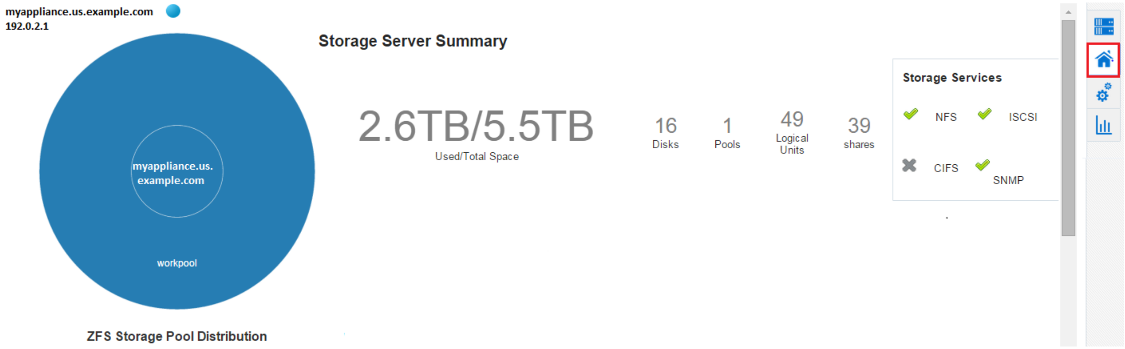

You can view the following storage summary information of an appliance in the Summary tab:

- Appliance Name with an IP address and a locator light indicating the status of the storage resource.
- A pie chart showing the information about storage pool distribution. You can click the center node of the pie chart to view detailed summary of storage appliance. The following details of the storage appliance are displayed:
  - Used/Total Space
  - Number of disks attached to the storage appliance
  - Number of pools in the storage appliance
  - Number of LUNs in the storage appliance
  - Number of shares in the storage appliance
  - Storage Services which are active for the storage appliance
  - Storage Usage graph indicating the storage utilization for last 5 hours, last 24 hours, and last 7 days.
  - Disks I/O ops per second graph for last 5 hours, last 24 hours, and last 7 days.
  - Disks I/O bytes per second graph for last 5 hours, last 24 hours, and last 7 days.
- You can click the pool node of the appliance pie chart to view detailed summary of the storage pool. The following details of the storage pool are displayed:
  - Used/Total Space information of the pool
  - Number of projects in the pool
  - Number of LUNs in the pool
  - Number of shares in the pool
  - Status of the pool
  - ZFS Storage Pool Storage Usage graph indicating the size usage of the physical and logical group storage pools
  - ZFS Storage Pool Utilization graph for last 5 Hours, last 24 Hours, and last 7 days

Viewing the Summary

To view the Summary tab details of an appliance or a cluster, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Servers, Storage and Network** target type.

   You can view a list of appliance targets under Servers, Storage and Network.

   > **Note:**
   >
   > To view the cluster targets, expand **Groups, Systems and Services** in the left navigation pane. Select a specific cluster target in the right pane.

3. Click a specific appliance target under Servers, Storage and Network.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. Click on a specific target name from the list in the right pane.
5. Click the **Summary** tab to view the summary details of the selected target.

About Projects

The Projects tab in the main window of the storage appliance user interface displays detailed information about the storage space used by the filesystems and LUNs as shown in Figure 21-4. It also indicates the Storage Services which are active for the particular filesystem or LUN, such as NFS, SMB, HTTP, FTP, TFTP, and SFTP. See Viewing the Projects for the steps to view the projects of a selected storage appliance.

Figure 21-4 Projects Tab View

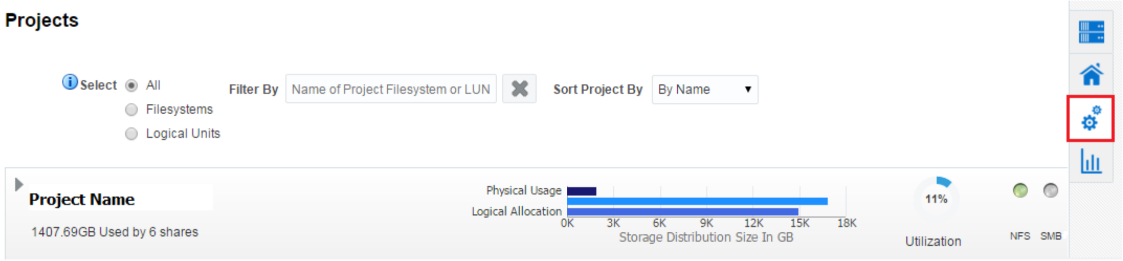

Projects

Select ● All ○ Filesystems ○ Logical Units

Filter By: Name of Project Filesystem or LUN

Sort Project By: By Name

Project Name

1407.69GB Used by 6 shares

You can view the following storage project information of an appliance in the Projects tab:

- List of filesystems and LUNs of a specific storage appliance and the storage space used by them.
- Storage Space graph which indicates Logical Allocation, Physical Size, and Physical Usage in GB.
  - Logical Allocation bar in the graph indicates the size allocated for storage.
  - Physical Size bar above the Logical Allocation bar in the graph indicates the actual storage size.
  - Physical Usage bar indicates the space that is used by storage.
- Storage Utilization in percentage.
- Storage Services that are active are indicated in Green.

Viewing the Projects

To view the Projects tab details of an appliance, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Servers, Storage and Network** target type.

   You can view a list of appliance targets under Servers, Storage and Network.

   > **Note:**
   >
   > To view the cluster targets, expand **Groups, Systems and Services** in the left navigation pane. Select a specific cluster target in the right pane.

3. Click a specific appliance target under Servers, Storage and Network.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. Click on a specific target name from the list in the right pane.
5. Click the **Projects** tab to view the storage details of the selected target.

About Charts

The Charts tab of a storage appliance user interface displays the Resources, Devices, SAN Usage, NAS Usage, and ZFS Storage Pool(s) tabs. You can select one of these tabs to view the respective graphical representation of a specific appliance.

The Charts tab appears to right of the storage appliance user interface. Click the Charts tab to view the different graphical representation tabs of the storage appliance to the left as shown in Figure 21-5.

Figure 21-5 Charts Tab View

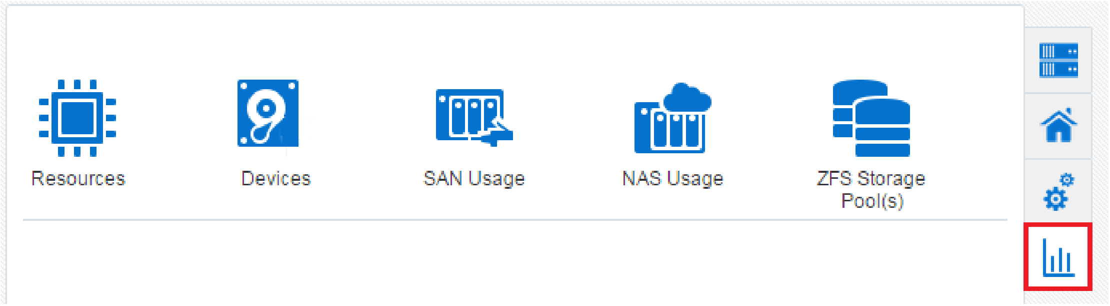

You can view the following items that displays respective graphical representation of an appliance in the Charts tab:

- Resources: The Resources tab displays the graphical representation of CPU Utilization (%) and Memory Usage Details with respect to time.

- Devices: The Devices tab displays the graphical representation of total disk I/O operations across all disks per second and total disk I/O bytes across all disks in bytes per second with respect to time
- SAN Usage: The SAN Usage tab displays the graphical representation of total operations per second and total bytes per second with respect to time.
- NAS Usage: The NAS Usage tab displays two graphical representations of total operations per second with respect to time for NAS Utilization and SMB Utilization.
- ZFS Storage Pools: The ZFS Storage Pool(s) tab displays the graphical representation of used space and allocated space with respect to time.

Viewing Resources Chart

To view the Resources chart details of an appliance, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Servers, Storage and Network** target type.

   You can view a list of appliance targets under Servers, Storage and Network.

   > **Note:**
   >
   > To view the cluster targets, expand **Groups, Systems and Services** in the left navigation pane. Select a specific cluster target in the right pane.

3. Click a specific appliance target under Servers, Storage and Network.

   The target name, target type, and the target status of the selected target is displayed in the right pane.

4. Click on a specific target name from the list in the right pane.
5. Click the **Charts** tab for the different options available for viewing graphical representation of the selected target.
6. Click the **Resources** tab to view the respective graphical representation of a selected target.
7. (Optional): You can select a specific duration for which you need to view the graph. In the **Duration** menu, you can select to view the graphical representation for last 1 hour, last 24 hours, or last 5 days. By default, it displays the graph for last 24 hours duration.
8. (Optional): You can click **Table View** to the lower right corner of the individual graph to view the details in the form of a table.

Viewing Devices Chart

To view the Devices chart details of an appliance, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Servers, Storage and Network** target type.

   You can view a list of appliance targets under Servers, Storage and Network.

> **Note:**
>
> To view the cluster targets, expand **Groups, Systems and Services** in the left navigation pane. Select a specific cluster target in the right pane.

3. Click a specific appliance target under Servers, Storage and Network.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. Click on a specific target name from the list in the right pane.
5. Click the **Charts** tab for the different options available for viewing graphical representation of the selected target.
6. Click the **Devices** tab to view the respective graphical representation of a selected target.
7. (Optional): You can select a specific duration for which you need to view the graph. In the **Duration** menu, you can select to view the graphical representation for last 1 hour, last 24 hours, or last 5 days. By default, it displays the graph for last 24 hour duration.
8. (Optional): You can click **Table View** to the lower right corner of the individual graph to view the details in the form of a table.

Viewing SAN Usage Chart

To view the SAN Usage chart details of an appliance, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Servers, Storage and Network** target type.

   You can view a list of appliance targets under Servers, Storage and Network.

> **Note:**
>
> To view the cluster targets, expand **Groups, Systems and Services** in the left navigation pane. Select a specific cluster target in the right pane.

3. Click a specific appliance target under Servers, Storage and Network.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. Click on a specific target name from the list in the right pane.
5. Click the **Charts** tab for the different options available for viewing graphical representation of the selected target.
6. Click the **SAN Usage** tab to view the respective graphical representation of a selected target.
7. (Optional): You can select a specific duration for which you need to view the graph. In the **Duration** menu, you can select to view the graphical representation for last 1 hour, last 24 hours, or last 5 days. By default, it displays the graph for last 24 hours duration.
8. (Optional): You can click **Table View** to the lower right corner of the individual graph to view the details in the form of a table.

Viewing NAS Usage Chart

To view the NAS Usage chart details of an appliance, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Servers, Storage and Network** target type.

   You can view a list of appliance targets under Servers, Storage and Network.

   > **Note:**
   >
   > To view the cluster targets, expand **Groups, Systems and Services** in the left navigation pane. Select a specific cluster target in the right pane.

3. Click a specific appliance target under Servers, Storage and Network.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. Click on a specific target name from the list in the right pane.
5. Click the **Charts** tab for the different options available for viewing graphical representation of the selected target.
6. Click the **NAS Usage** tab to view the NAS Utilization and SMB Utilization graphical representation of a selected target.
7. (Optional): You can select a specific duration for which you need to view the graph. In the **Duration** menu, you can select to view the graphical representation for last 1 hour, last 24 hours, or last 5 days. By default, it displays the graph for last 24 hours duration.
8. (Optional): You can click **Table View** to the lower right corner of the individual graph to view the details in the form of a table.

Viewing ZFS Storage Pools Chart

To view the ZFS Storage Pool(s) chart details of an appliance, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Servers, Storage and Network** target type.

   You can view a list of appliance targets under Servers, Storage and Network.

   > **Note:**
   >
   > To view the cluster targets, expand **Groups, Systems and Services** in the left navigation pane. Select a specific cluster target in the right pane.

3. Click a specific appliance target under Servers, Storage and Network.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. Click on a specific target name from the list in the right pane.

5. Click the **Charts** tab for the different options available for viewing graphical representation of the selected target.
6. Click the **ZFS Storage Pool(s)** tab to view the respective graphical representation of a selected target.
7. In the **ZFS Storage Pool(s)** drop-down, select the specific storage pool for which you want to view the graphical representation of a selected target.
8. (Optional): You can select a specific duration for which you need to view the graph. In the **Duration** menu, you can select to view the graphical representation for last 1 hour, last 24 hours, or last 5 days. By default, it displays the graph for last 24 hours duration.
9. (Optional): You can click **Table View** to the lower right corner of the individual graph to view the details in the form of a table.

> **Note:**
>
> The unit mentioned in y-axis should be read as follows:
>
> - M stands for Million
> - B stands for Billion
> - T stands for Trillion

About Host Storage Information

You can view the storage information of a discovered host target. The Storage tab in the main window of the host's user interface displays the Disks, Filesystems, Linux Volume Group(s), ZFS Storage Pool(s), and SAN Configuration tabs. You can select one of these tabs to view more detailed information.

The Storage tab appears to right of the host's user interface. Select the Storage tab to view the different host storage tabs to the left as shown in Figure 21-6.

Figure 21-6 Host Storage Tab View

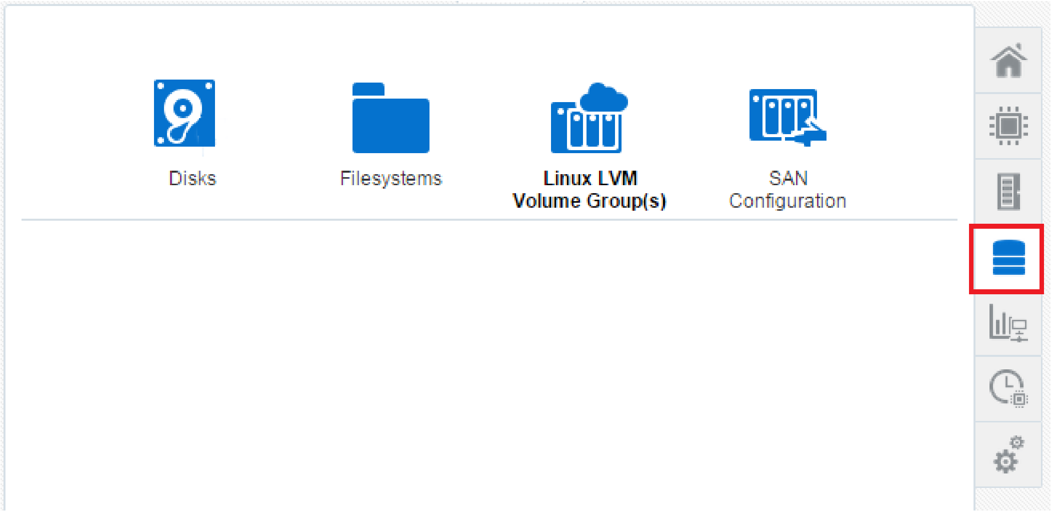

You can view the following host's storage information in the Storage tab:

- Disks: Displays the Logical disks information and Physical disks information.
- Filesystems: Displays the Filesystem name, their mounted location, size, incidents, and Storage Utilization.
- Linux LVM Volume Group(s): Displays different volumes names. You can click any of the volume boxes to view the Total Space in GB, number of Incidents, number of Volumes, Storage Utilization in percentage, and the status of the volume. This tab appears for a Linux host only.
- ZFS Storage Pool(s): Displays different storage pool volumes. You can click any of the volume boxes to view the Total Space in GB, number of Incidents, number of Volumes, Storage Utilization in percentage, and the status of the volume. This tab appears for a Oracle Solaris host only.
- SAN Configuration: Displays the Network hardware name, number of targets connected, number of Devices connected, Session Response Time, Maximum Connection Retry Time, FC Port ID, FC Port State, and so on.

Disks of a Host

The Storage tab in the main window of the host lists the Disks tab. You can view the disks information of the Linux or Oracle Solaris host.

You can view the following host's disks information in the Disks tab:

- Type of disk, such as local or remote disk
- Name of each disk
- Size of each disk
- Device location of each disk
- Number of incidents for the respective disk if incident count is greater than 0
- Read Operations per second for last 24 hours
- Max Read Operations per second
- Write Operations per second for last 24 hours
- Max Write Operations per second
- Storage Utilization
- SnapshotCount and SnapshotUtilization for Oracle Solaris host if the snapshot count is greater than 0
- Vendor name

Viewing Disks of a Host

To view the host's disks information, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand the **Servers, Storage, and Network** target type.

   You can view the host target listed.
3. Click **Host**.

The target name, target type, and the target status of the selected target is displayed in the right pane.

4. Click on the specific host target from the list in the right pane.
5. Click the **Storage** tab which appears to the right side of the user interface.
6. Click the **Disks** tab.
7. (Optional): You can select to view all the disks, only the local disks, or only the remote disks. Also, you can filter the list of disks by entering the Volume Name in the **Filtered By** field.

By default, the number of disk volumes displayed are restricted to 25 by utilization and by the volumes that have alerts. You can change this default value. After selecting Host in the left navigation pane, right click on the specific host listed in the right pane. Select **Monitoring** and click on **Metric and Collection Settings**. In the Metric and Collection Settings window, locate the Volume Statistics metric and click **Edit** in the Space Utilization (%) of a Volume row. Under Metric Collection Properties, the Property Names and Property Values are displayed. The value of NumberOfVolumes property is displayed as 25. You can edit this Property Value to the desired value.

Filesystems of a Host

The Storage tab in the main window of the host lists the Filesystems tab. You can view the filesystems information of the Linux or Oracle Solaris host.

You can view the following host's filesystem information in the Filesystems tab:

- Type of filesystem, such as local or remote filesystem
- Name of each filesystem
- Size of each filesystem
- Mounted location of each filesystem
- Number of incidents for respective filesystem
- Read Operations per second for last 24 hours
- Max Read Operations per second
- Write Operations per second for last 24 hours
- Max Write Operations per second
- Storage Utilization in percentage
- Compression Ratio if enabled for Oracle Solaris host

Viewing Filesystems of a Host

To view the host's filesystems information, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand the **Servers, Storage, and Network** target type.

   You can view the host target listed.
3. Click **Host**.

   The target name, target type, and the target status of the selected target is displayed in the right pane.

4. Click on the specific host target from the list in the right pane.
5. Click the **Storage** tab which appears to the right side of the user interface.
6. Click the **Filesystems** tab.
7. (Optional): You can select to view all the filesystems, only the local filesystems, or only the remote filesystems. Also, you can filter the list of disks by entering the Filesystem Name in the **Filtered By** field.

By default, the number of filesystems displayed are restricted to 25 by utilization. You can change this default value. After selecting Host in the left navigation pane, right click on the specific host listed in the right pane. Select **Monitoring** and click on **Metric and Collection Settings**. In the Metric and Collection Settings window, locate the Filesystem Statistics metric and click **Edit** in the Space Utilization (%) of a filesystem row. Under Metric Collection Properties, the Property Names and Property Values are displayed. The value of NumberOfFilesystems property is displayed as 25. You can edit this Property Value to the desired value. You can also set the filesystems property to specify the filesystems that must be collected irrespective of utilization.

SAN Configuration of a Host

The Storage tab in the main window of the host lists the SAN Configuration tab. You can view the SAN Configuration information of the Linux or Oracle Solaris host.

You can view the following host's SAN Configuration information in the SAN Configuration tab:

- SAN initiator type, such as, iSCSI or FC
- IQN name for an iSCSI SAN initiator type
- Total number of targets connected
- Total number of devices attached

Viewing SAN Configuration of a Host

To view the host's SAN Configuration information, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand the **Servers, Storage, and Network** target type.

   You can view the host target listed.
3. Click **Host**.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. Click on the specific host target from the list in the right pane.
5. Click the **Storage** tab which appears to the right side of the user interface.
6. Click the **SAN Configuration** tab.

Linux Volume Groups of a Host

The Storage tab in the main window of the host lists the Linux LVM Volume Group(s) tab. You can view the Linux volume group information of the Linux host.

Note:

The Linux LVM Volume Group(s) tab is only available for a Linux environment.

You can view the following information of the selected volume in the Linux LVM Volume Group(s) tab:

- Total Space in GB
- Target Status of LVM to indicate if up or down
- Incidents
- Volumes
- Storage Utilization in percentage
- Linux LVM Volume Group Usage graph
- Linux LVM Volume Group Hierarchy

Viewing Linux Volume Groups of a Host

To view the host's linux volume group information, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand the **Servers, Storage, and Network** target type.

   You can view the host target listed.
3. Click **Host**.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. Click on the specific host target from the list in the right pane.
5. Click the **Storage** tab which appears to the right side of the user interface.
6. Click the **Linux LVM Volume Group(s)** tab. The different volumes are listed in the form of blocks to the left.
7. Click on a specific volume to view detailed information.

ZFS Storage Pools of a Host

The Storage tab in the main window of the host lists the ZFS Storage Pool(s) tab. You can view the detailed information of different zpools.

Note:

The ZFS Storage Pool(s) tab is only available for an Oracle Solaris environment.

You can view the following information of the selected zpool in the ZFS Storage Pool(s) tab:

- Total Space in GB
- Target Status of the pool to indicate if up or down

- Incidents
- Filesystems
- Volumes
- Dedup Ratio
- Storage Utilization in percentage
- Zpool status indicating if it is online, offline, or faulted
- ZFS Storage Pool Usage graph
- ZFS Storage Pool Hierarchy

Viewing ZFS Storage Pools of a Host

To view the host's ZFS storage pool information, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand the **Servers, Storage, and Network** target type.

   You can view the host target listed.
3. Click **Host**.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. Click on the specific host target from the list in the right pane.
5. Click the **Storage** tab which appears to the right side of the user interface.
6. Click the **ZFS Storage Pool(s)** tab. The different zpools used are listed in the form of blocks to the left.
7. Click on a specific zpool to view detailed information.

About Storage Configuration Topology

Enterprise Manager allows you to view the relationship between storage servers and storage clients. The Topology option displays the relationships between each exposed element such as LUNs or shares of the storage server and it's presence on a given host as devices and filesystems. It also shows the relationship between various applications such as database, or virtualization entities like VMs, and so on.

You can view the storage configuration topology of the storage appliance or the host. You can also view the storage configuration topology of a specific storage server element or a specific host element. Select the respective element and choose to view the topology details. See Viewing Storage Configuration Topology for the steps to view the topology of a selected storage appliance or a host.

Viewing Storage Configuration Topology

To view topology information of a selected target, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Servers, Storage and Network** target type.

   You can view a list of appliance targets under Servers, Storage and Network.

> **Note:**
>
> To view the cluster targets, expand **Groups, Systems and Services** in the left navigation pane. Select a specific cluster target in the right pane.

3. Click a specific appliance target under Servers, Storage and Network.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. You can view the storage configuration topology information by following one of these methods:
   - Right-click on a target name in the right pane to view the related options.
   - Click a specific target in the right pane, and then click the target drop-down list, which appears to the top-left of the user interface.
5. Select **Configuration** and click **Topology**.

About Storage Metrics

This section describes the steps to view the Storage Configuration metrics and Storage Performance metrics. It also describes the steps to change the Storage Performance metrics. Following are the items listed in this section:

- Viewing Storage Performance Metrics
- Viewing Storage Configuration Metrics
- Changing Metric Collection

Viewing Storage Performance Metrics

To view the storage performance metrics, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Servers, Storage and Network** target type.

   You can view a list of appliance targets under Servers, Storage and Network.

> **Note:**
>
> To view the cluster targets, expand **Groups, Systems and Services** in the left navigation pane. Select a specific cluster target in the right pane.

3. Click a specific appliance target under Servers, Storage and Network.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. You can view the performance metrics of the selected target by following one of these methods:
   - Right-click on a target name in the right pane to view the related options.
   - Click a specific target in the right pane, and then click the target drop-down list which appears to the top-left of the user interface.

5. Select **Monitoring** and click **All Metrics**.

   You can view the list of metric collection overview, different metric events such as open alerts and top five altering metrics in the last 7 days, and deployed metric extension information.

Viewing Storage Configuration Metrics

To view a storage configuration metrics, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Servers, Storage and Network** target type.

   You can view a list of appliance targets under Servers, Storage and Network.

   > **Note:**
   >
   > To view the cluster targets, expand **Groups, Systems and Services** in the left navigation pane. Select a specific cluster target in the center pane.

3. Click a specific appliance target under Servers, Storage and Network.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. You can proceed to view the storage configuration metrics by following one of these methods:
   - Right-click on a target name in the right pane to view the related options.
   - Click a specific target name in the right pane, and then click the target drop-down list which appears to the top-left of the user interface.
5. Select **Configuration** and click **Last Collected**.
6. Select a configuration metric from the left navigation pane to view the configuration properties of the selected metric.

Changing Metric Collection

Oracle Enterprise Manager monitors storage resources and collects information about them to track performance. You can change what metrics are collected and, in some cases, change how the metrics are collected, for example, the frequency.

To change the performance metrics of a selected target, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Servers, Storage and Network** target type.

   You can view a list of appliance targets under Servers, Storage and Network.

   > **Note:**
   >
   > To view the cluster targets, expand **Groups, Systems and Services** in the left navigation pane. Select a specific cluster target in the right pane.

3. Click a specific appliance target under Servers, Storage and Network.

The target name, target type, and the target status of the selected target is displayed in the right pane.

4. You can proceed to view the performance metrics by following one of these methods:
   - Right-click on any of the target name in the right pane to view the related options.
   - Click a specific target name in the right pane, and then click the target drop-down list which appears to the top-left of the user interface.
5. Select **Monitoring** and click **All Metrics**.

   You can view details of the different metrics.
6. In the left navigation pane, select a specific metric. If this metric can be changed, a **Modify** button appears in the right pane for Collection Schedule.
7. Click **Modify**.
8. In the Modify Collection Schedule window, change the value to adjust the collection of the metric, usually the frequency and how the metric is used. Make a note of the metrics that are dependent on this change, so that you can change their collection, if necessary.
9. Click **OK** to save the changes.

About Storage Cluster Membership

You can have a cluster configuration setup with two Oracle ZFS Storage Appliances. When such configuration is in use, you can view the cluster membership information of the two appliances.

Viewing Storage Cluster Membership

To view cluster membership information of the appliances which are in a cluster configuration setup, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Groups, Systems and Services** target type.

   You can view a list of cluster targets under Groups, Systems and Services.
3. Click a specific cluster target under Groups, Systems and Services.

   The target name, target type, and the target status of the selected target is displayed in the right pane.
4. Click on a specific cluster target name from the list in the right pane.
5. Click the **Target Navigation** tab which appears to the left of the cluster drop-down list, as shown in Figure 21-7.

Figure 21-7 Target Navigation Tab View

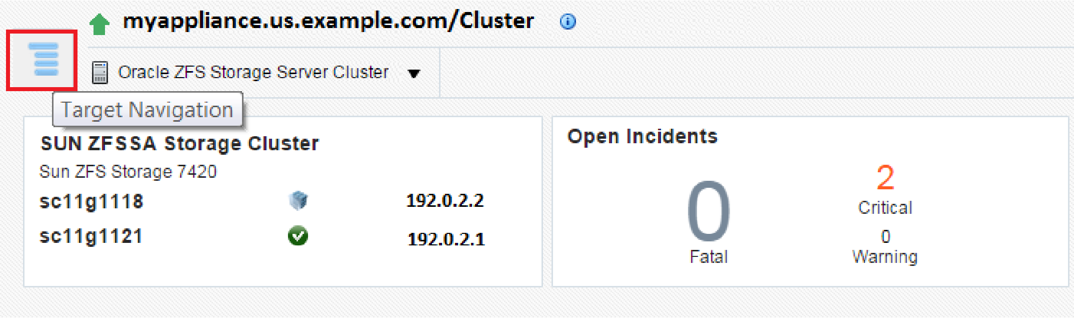
myappliance.us.example.com/Cluster
Oracle ZFS Storage Server Cluster
Target Navigation
Open Incidents
0
Fatal
2
Critical
0
Warning

SUN ZFSSA Storage Cluster
Sun ZFS Storage 7420

sc11g1118	192.0.2.2
sc11g1121	192.0.2.1

6. Select a specific cluster to view the storage summary details of the cluster. You can view the cluster information in the dashboard, images of the storage appliance which have a cluster configuration setup, different networks and pools information.

About Storage Resource Deletion

When an appliance or a host is discovered as a target, it automatically promotes some of the target members. See Target Members of an Oracle ZFS Storage Appliance to view the list of automatically promoted target members. You can choose to either remove the target or remove its members. If you choose to remove a target, such as an appliance or a host, all the associated members are also removed. See Removing a Storage Resource for the steps to remove a storage resource.

You must start the deletion process from the target level to delete all the members.

> **Note:**
>
> If the target members have a proper association with an appliance or a host, the deletion of these members is successful. For example, a diskshelf target member is associated to a ZFS Storage Server. If the association of these target members are not correct, diskshelf target members are not deleted.

Removing a Storage Resource

To remove a storage resource, perform the following steps:

1. From the Targets menu, select **All Targets**.
2. In the left navigation pane, expand **Servers, Storage and Network** target type.

   You can view a list of appliance targets under Servers, Storage and Network.

   > **Note:**
   >
   > To view the cluster targets, expand **Groups, Systems and Services** in the left navigation pane. Select a specific cluster target in the right pane.

3. Click a specific appliance target under Servers, Storage and Network.

The target name, target type, and the target status of the selected target is displayed in the right pane.

4. You can proceed to delete a storage resource by following one of these methods:
   - Right-click on the target name which you want to remove in the right pane.
   - Click a specific target name in the right pane, and then click the target drop-down list which appears to the top-left of the user interface.
5. Select **Target Setup** and click **Remove Target**.
6. In the confirmation dialog box, click **Yes** to remove the specific target.

Removing an Oracle ZFS Storage Appliance Cluster

You can delete an Oracle ZFS Storage Appliance which is part of a cluster. See Removing a Storage Resource for steps to delete an Oracle ZFS Storage Appliance. When two Oracle ZFS Storage Servers have a cluster configuration setup, the deletion of cluster target initiates deletion of both the storage servers in the cluster and the associated target members, such as, diskshelves. The members of an Oracle ZFS Storage Appliance include the following:

- ZFS Storage Server
- Diskshelf
- ZFS Storage Appliance Cluster

When you initiate deletion of one of the storage server that is a part of the cluster configuration setup, the associated cluster target is deleted. The other storage server is not deleted. The diskshelves are not deleted if there is a second storage server.

When you initiate deletion of a storage server that is not a part of the cluster configuration setup, the associated diskshelves are deleted.

Using Oracle ZFS Storage Appliance in Engineered Systems

This section describes the prerequisite for an Oracle ZFS Storage Appliance on engineered systems. A non-root user should have required privileges to view the analytics dataset. If the specific privileges are not assigned, then the analytic metrics displays an error.

If you need data to be collected for the following metrics, the corresponding dataset must be enabled in Oracle ZFS Storage Appliance.

Metric name	DatasetName
Replication statistics by Direction	`repl.bytes[direction],repl.ops[direction]`
Replication Ops by Latency	`repl.ops[latency]`
IOopsByLatency	`io.ops[latency]`
L2ARC statistics	`arc.l2_accesses[hit/miss],arc.l2_bytes,arc.l2_size`
NAS IO statistics	`nfs2.bytes, nfs3.bytes,nfs4.bytes`

You can add an extra privilege to a non-root user to access analytics dataset.

To add a privilege in the appliance's user interface, perform the following steps:

1. Login to Oracle ZFS Storage Appliance user interface as the `root` user.

2. Click **Configuration**, and click **Users**.
3. Select the non-root user and click the edit icon.
4. Select the **Exceptions** tab.
5. From the **Scope** drop-down, select **Analytics**.
6. From the **List of drilldowns** menu, select *****, and check the **read** check box.
7. Click **Add** and log out of the appliance.
8. Login to Oracle ZFS Storage Appliance user interface as the `non-root` user.
9. Click **Configuration**, and click **Preferences**.
10. Check the check box **Make available advanced analytics statistics** and log out of the appliance.

You can also add an exception to a non-root user using Command Line Interface for reading analytics dataset. For example:

```
zfssa:> configuration users
zfssa:configuration users> select oemuser
zfssa:configuration users oemuser> exceptions
zfssa:configuration users oemuser exceptions> create
zfssa:configuration users oemuser auth (uncommitted)> set scope=stat
                         scope = stat
zfssa:configuration users oemuser auth (uncommitted)> set allow_read=true
                    allow_read = true (uncommitted)
zfssa:configuration users oemuser auth (uncommitted)> commit
zfssa:configuration users oemuser exceptions>
```

Related Resources for Storage

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources.

- To view the *Sun ZFS Storage 7000 System Administration Guide*, log in to the Unified Storage System software interface and click Help in the top right corner of any screen. You can also access this guide at the host name or IP address of the storage system:
  - `https://`*`hostname`*`:215/wiki`
  - `https://`*`ipaddress`*`:215/wiki`
- See Oracle ZFS Storage Appliance Software for more information.
- See Discovering and Promoting Oracle ZFS Storage for more information on Storage Discovery.
- See the Discovering, Promoting, and Adding System Infrastructure Targets chapter for more information on discovering and promoting system infrastructure targets.

22
Monitoring Servers

The following features and topics are covered in this chapter:

Get Started With Server Management

You can discover hardware assets in Oracle Enterprise Manager by using the existing discovery method, then deploying the Enterprise Manager agent onto the system. Once they are discovered, you can view monitoring information about the hardware, including incidents, power usage, network information, service processor configuration, fan and temperature information. The relationship between managed hardware and the operating systems, virtualization platforms, and other software installed on it is also represented in the user interface.

> **Note:**
>
> - The remote agent connects to ILOM of the server through **SSH port (22)**.
> - The agent connects to Server target using the **HTTPS port (443)** when ReST `AccessPointResponse` is used for monitoring the ILOM.

Systems Infrastructure Server target can be monitored using non-root user credentials. Create a non-root user in the ILOM shell and use it for discovery or for changing the monitoring credentials. If you use the non-root user credentials, then the SNMP user creation and SNMP subscription will not happen automatically during discovery. You have to make these changes manually for the SNMP alerts to be processed.

Location of Server Information in the UI

You can select any target that is a child of a server (virtual platform, guest, or host) and the server will appear in the Navigation pane of the target.

From the All Targets page, you can also click **Systems Infrastructure Server** to see the list of servers. You can click any server in this list to open the server's home page.

Actions for Server Management

You can perform the following actions:

- Discover a server
- View the server's hardware components
- View the server's configuration
- View the server's utilization of resources

About the Hardware Dashboard

The dashboard is located near the top of the main window. It contains basic information about the server's status, including incidents, power usage, temperature information, core information, and recent events.

The information in the dashboard is automatically displayed. Three dashlets are visible. Click the icon beneath the dashboard to switch to another set of dashlets.

The following dashlets are displayed:

- About Basic Hardware Information
- About Open Incidents
- About Fan and Temperature Information
- About Power Usage
- About Core Information
- About the Last Configuration Change and Incident

About Basic Hardware Information

The first dashlet contains basic information about the hardware. The heading includes the full hardware name and power status.

The following fields are displayed:

- IP Address
- Model
- Serial Number
- Health

- CPU
- Memory
- Firmware
- Locator

About Open Incidents

The second dashlet contains information about open incidents for the hardware.

The following fields are displayed:

- Fatal
- Critical
- Warning

Click on a category to view a detailed view of incidents within this category, including their target, summary, date of last update, whether they have been acknowledged, and status. Click the X icon in the upper right to close this detailed view.

About Fan and Temperature Information

The fourth dashlet displays fan and temperature information.

The following fields are displayed:

- Fan Usage: This displays the fan usage as a percentage of its maximum.
- Temperature: This displays the temperature of the hardware in degrees Celsius.

About Power Usage

The third dashlet displays power usage information. A chart displays the power usage as a percentage of the maximum.

The following fields are displayed:

- Available Power
- Peak Permitted
- Used Power
- Power Policy (SPARC servers only)

About Core Information

The fifth dashlet displays a pie chart showing the number of active and inactive cores.

About the Last Configuration Change and Incident

The sixth dashlet displays the date and time of the last configuration change and the last reported incident.

Viewing the Hardware Dashboard

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Server**.

   A list of the target servers is displayed.
3. Click the target name to open the Summary page for the server. The dashlets appear on the top of the page and provide the summary information.

About Server Metrics

You can view a complete list of the metrics for a selected server.

Viewing Server Metrics

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Server**.

   A list of the target servers is displayed.
3. Click the target name to open the Summary page for the server.
4. Click **Systems Infrastructure Server** in the upper left corner of the page. Click **Monitoring**, then click **All Metrics**.
5. Click a metric to view details, collection schedule, upload interval and other details.

About the Photorealistic Image of the Hardware

The Hardware View tab in the main window displays a photorealistic view of compatible hardware, including the front, top, and rear, and a table view of the hardware's components.

Select Photorealistic View to display the photorealistic view of the hardware. Components with incidents are outlined in red.

You can click any component displayed in the photorealistic view to view additional information about that component. The following information is displayed if it is available and relevant to the component:

- Component Name
- Manufacturer
- Serial Number
- Part Number
- Total Cores: The number of cores for a CPU
- Enabled Cores: The number of enabled cores for a CPU
- Size: The size of memory components in GB

Select Table to display a table of the hardware components. The following information is displayed for each component:

- Slot Number

- Component Name
- Component Type

Viewing the Photorealistic Image of the Hardware

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Server**.

   A list of the target servers is displayed.
3. Click the target name to open the Summary page for the server.
4. Click the **Hardware View** tab.

About the Logical View

The Logical View tab in the main window displays detailed information about the hardware components and capabilities. You can select one of the tabs to view detailed information about it.

About CPU Information

The CPU tab shows CPU and CPU usage information.

The top section shows summary information. A pie chart displays the number of installed and available CPUs.

The following fields are displayed:

- Architecture
- Clock Speed
- Model
- CPU Power Consumption (Watts)
- Overall Status

The bottom section displays a table showing the available processors, including name, active cores, serial numbers, part number, overall cache in KB, component location, and operational status.

About Memory Information

The Memory tab shows overall memory and DIMM-specific information.

The top section shows summary information. A pie chart displays the number of installed and available DIMMs.

The following fields are displayed:

- Memory (GB)
- Memory Power Consumption (Watts)
- Overall Status

The bottom section displays a table showing the memory modules, including the memory component name, size in GB, manufacturer, part number, serial number, location, and operational status.

About Power Information

The Power tab shows power and power supply information.

The top section shows summary information. A pie chart displays the number of installed and available power supplies. The overall status of the power supply is displayed.

The bottom section displays a table showing the available power supplies, including name, manufacturer, part number, serial number, output power in watts, location, and operational status.

About Fan Information

The Fan tab shows cooling and fan information.

The top section shows summary information. A pie chart displays the number of total and available power supply unit fans. The overall status of the cooling is displayed.

The bottom section displays a table showing the available fans, including name, RPM as a percentage of maximum, location, and operational status.

About Storage Information

The Storage tab shows information about the available storage.

The top section shows summary information. The following fields are displayed:

- Total Installed Storage (GB)
- Installed Disk

The bottom section displays a table showing the available storage disks, including the name, size in gigabytes, manufacturer, serial number, part number, and operational status. This information is displayed only if the ILOM is discovered.

About Disk Controller Information

The Disk Controller tab displays a table of the available disk controllers, including name, model, manufacturer, serial number, and operational status.

This information is displayed only if the ILOM is discovered.

About Disk Expander Information

The Disk Expander tab displays a table of the available disk expanders, including name, manufacturer, version, model, firmware version, and chassis ID.

This information is displayed only if the ILOM is discovered.

About Network Ports Information

The Network Ports tab displays information about network interface controllers and network adapters.

The top section shows NIC and status information. The following fields are displayed:

- Installed Ethernet NICs
- Overall Status

The bottom section shows a table of network ports, including name, MAC address, description, and operational status.

About PCI Devices Information

The PCI Devices tab displays a table of the PCI devices, including name, description, device class, PCI device ID, PCI vendor ID, PCI end point, PCI sub device ID, and PCI sub device vendor ID.

This information is displayed only if an Agent is deployed on an operating system on the server.

About PDOMs Information

The PDOMs tab displays a table of the physical domains for M-series hardware, including name, configuration status, assigned DCUs, and operational status.

About DCUs Information

The DCUs tab displays a table of the DCUs for M-series hardware, including name, number of CPUs, memory in GB, number of fans, PDOM ID, power status, and operational status.

Viewing the Logical View

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Server**.

   A list of the target servers is displayed.
3. Click the target name to open the Summary page for the server.
4. Click the **Logical View** tab.

About Energy Consumption

The Energy tab in the main window displays information about the hardware's energy consumption.

The summary section displays three graphs showing basic temperature, fan speed, and power information. You can use the Time Range dropdown to select a different time interval to display.

The first graph shows the inlet and exhaust temperatures in degrees Celsius.

The second graph shows the fan speed as a percentage of the maximum.

The third graph shows the power consumption and utilization in watts.

You can click the Table View link to view a table of the data points used to create the graph.

Viewing the Energy Consumption

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Server**.

   A list of the target servers is displayed.
3. Click the target name to open the Summary page for the server.
4. Click the **Energy** tab.

About Network Connectivity

The Network Connectivity tab in the main window displays information about the hardware's network interfaces, data links, and ports. You can select one of these three options to view detailed information about it.

About Network Interfaces

The Network Interfaces page shows a table of the hardware's network interfaces, including IP address, netmask, and an icon indicating the current state.

You can sort the list by interface name or interface state.

Click the more link for additional information.

About Network Data Links

The Network Data Links page shows a table of the hardware's data links, including name, physical address, media, and VLAN ID.

You can sort the list by data link name or data link state.

Click the more link for additional information, including device and device path.

About Network Ports

The Network Ports page shows a table of the hardware's ports and their types.

You can sort the list by state, connector, or number and name.

Click the more link for additional information, including errors and throughput.

Viewing the Network Connectivity

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Server**.

   A list of the target servers is displayed.
3. Click the target name to open the Summary page for the server.
4. Click the **Network Connectivity** tab.

About the Service Processor Configuration

The Service Processor Configuration tab in the main window displays information about the firmware, host policy configuration, power on self test configuration, SP alert configuration, and DNS and NTP settings. You can select one of these tabs to view detailed information about it.

About Firmware Information

The Firmware Information tab shows a table with the component identifier for all installed firmware, the type, the version, and the release date.

About the Host Policy Configuration

The Host Policy Configuration tab shows a table with a list of the host policy names and their current values.

About the Power On Self Test Configuration

The Power On Self Test Configuration tab shows a table with a list of the power on self test setting names and their current values.

About the SP Alert Configuration

The SP Alert Configuration tab shows a table with the service processor alert names. For each alert, the table provides the alert type, alert level, destination address, destination port, SNMP version, and community.

About the DNS & NTP Information

The DNS & NTP tab shows information about the DNS and NTP settings. The following fields are displayed:

- Auto DNS/DHCP: Indicates whether DNS or DHCP is being used.
- DNS Servers: Lists the DNS servers in use.
- Search Path: Lists the search path for the DNS servers.
- Time: Lists the current time and time zone for the hardware.
- Use NTP Server: Indicates whether an NTP server is being used.
- NTP Server 1: The IP address of the first NTP server.
- NTP Server 2: The IP address of the second NTP server.

Viewing the Service Processor Configuration

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Server**.

   A list of the target servers is displayed.
3. Click the target name to open the Summary page for the server.
4. Click the **Service Processor Configuration** tab.

Managing Metrics and Incident Notifications

You can perform the following tasks to manage monitoring and incident notification:

- Viewing Metric Collection Errors
- Editing Metric and Collection Settings
- Editing a Monitoring Configuration
- Suspending Monitoring Notifications
- Suspending Monitoring for Maintenance
- Ending a Monitoring Brownout or Blackout

Viewing Metric Collection Errors

Metric collection errors are usually caused by installation or configuration issues. You can view errors for a server.

1. Click **Systems Infrastructure Server** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Systems Infrastructure Server** in the upper left corner of the page. Click **Monitoring**, then click **Metric Collection Errors**.

Editing Metric and Collection Settings

The Metrics tab contains displays all of the monitored attributes. The default view is metrics with thresholds. For these types of monitored attributes, you can modify the comparison operator, the threshold limits, the corrective action, and the collection schedule.

1. Click **Systems Infrastructure Server** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Systems Infrastructure Server** in the upper left corner of the page. Click **Monitoring**, then click **Metric and Collection Settings**.
4. Modify threshold limits or collection schedule. When a threshold field is empty, the alert is disabled for that metric.
5. Click the **Edit** icon for advanced settings.

   Click the **Other Collected Items** tab to view non-threshold monitored attributes. You can modify the collection period for these attributes, or disable monitoring.
6. Click **OK** to save your changes.

Editing a Monitoring Configuration

1. Click **Systems Infrastructure Server** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Systems Infrastructure Server** in the upper left corner of the page. Click **Target Setup**.
4. Click **Monitoring Configuration**.

Suspending Monitoring Notifications

Brownouts enable you to temporarily suppress notifications on a target. The Agent continues to monitor the target under brownout. You can view the actual target status along with an indication that the target is currently under brownout.

You can create a brownout for a server.

1. Click **Systems Infrastructure Server** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Systems Infrastructure Server** in the upper left corner of the page. Click **Control**.
4. Click **Create Brownout**.
5. Enter a name for the brownout event.
6. Select a reason from the menu and add comments, as needed.
7. Click the options to define how jobs will run and the maintenance window.
8. Click **Submit**.

Suspending Monitoring for Maintenance

Blackouts enable you to suspend monitoring on one or more targets in order to perform maintenance operations. To place a target under blackout, you must have at least the Blackout Target privilege on the target. If you select a host, then by default all the targets on that host are included in the blackout. Similarly, if you select a target that has members, then by default all the members are included in the blackout.

1. Click **Systems Infrastructure Server** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Systems Infrastructure Server** in the upper left corner of the page. Click **Control**.
4. Click **Create Blackout**.
5. Select a reason from the menu.
6. Add comments, as needed.
7. Click **Submit**.

Ending a Monitoring Brownout or Blackout

You can end a blackout or brownout for a server.

1. Click **Systems Infrastructure Server** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Systems Infrastructure Server** in the upper left corner of the page. Click **Control**.
4. Click **End Blackout** or **End Brownout**.

Administering Servers

You can perform the following tasks to manage and administer servers:

- Viewing Compliance

- Identifying Changes in a Server Configuration
- Editing Server Administrator Access
- Adding a Server to a Group
- Editing Server Properties

Viewing Compliance

The Compliance pages enable you to view the compliance framework, standards, and the server's compliance.

1. Click **Systems Infrastructure Server** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Systems Infrastructure Server** in the upper left corner of the page. Click **Compliance**.
4. Click the option to view **Results**, **Standard Associations**, or **Real-time Observations**.

Identifying Changes in a Server Configuration

When an administrator changes a system's configuration, it can be helpful to know the when the configuration was last changed. This information appears in the configuration dashlet on the Summary page.

To view more detailed information for a server:

1. Click **Systems Infrastructure Server** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Systems Infrastructure Server** in the upper left corner of the page. Click **Configuration**.
4. Click the option to view **Last Collected**, **Comparison and Drift Management**, **Compare**, **Search**, **History**, **Save**, **Saved**, or **Topology**.

Editing Server Administrator Access

1. Click **Systems Infrastructure Server** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Systems Infrastructure Server** in the upper left corner of the page. Click **Target Setup**.
4. Click **Administrator Access**.

Adding a Server to a Group

1. Click **Systems Infrastructure Server** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Systems Infrastructure Server** in the upper left corner of the page. Click **Target Setup**.
4. Click **Add to Group**.

Editing Server Properties

1. Click **Systems Infrastructure Server** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Systems Infrastructure Server** in the upper left corner of the page. Click **Target Setup**.
4. Click **Properties**.

Related Resources for Server Management

See the following chapters for more information:

- Discovering and Adding Host and Non-Host Targets
- Discovering, Promoting, and Adding System Infrastructure Targets
- Using Incident Management

23 Managing the PDU

The following information is included in this chapter:

Getting Started with PDU Management

Enterprise Manager Cloud Control 13.2 enables management and monitoring of Oracle hardware targets, including servers, switches, ZFS Storage appliance, Exadata storage cells, PDUs, racks, and Engineered Systems. PDU target provides monitoring information about the PDU powering the hardware in the rack. Discovering and managing your targets is a prerequisite for almost every action in the software. The discovery is made quick and easy with the Guided Discovery Wizard that guides you through the whole process. The discovery process requires only necessary information as input and helps you solve possible issues in order to successfully complete the discovery.

Location of PDU Information in the User Interface

Table 23-1 shows where to find information.

Table 23-1 Location of PDU Information in the BUI

Object	Location
Power Distribution Unit	In the Enterprise Manager user interface, under Targets, click **All Targets**. In the Refine Search section, under Target type, click **Servers, Storage, and Network**. Click **Systems Infrastructure PDU**, then select a PDU from the displayed list.

Actions for PDU

You can perform the following actions, depending on the requirements.

- Discover a PDU

- View the PDU

PDU Version Identification

PDU hardware is shipped in two versions, namely PDU v1 (Original PDU) and PDU v2 (Enhanced PDU).

For more information on enhanced PDU, see Monitoring Enhanced PDUs

For more information on original PDU, see Monitoring Original PDUs

You can distinguish the PDU version by accessing the PDU Management Interface. This interface is accessible using the web browser on IP address or DNS name that you have assigned to the PDU. Knowledge of PDU version is required to solve some PDU monitoring or discovery issues.

> **Note:**
>
> PDU may be on an isolated management network not reachable by your web browser. In that case, make sure you reach the PDU management interface from within the management network.

1. Open the web browser.
2. Enter the address of the PDU Management Interface in the web browser.

   For example, http://<IP or DNS name of your PDU> for PDU v1 or https://<IP or DNS name of your PDU> for the PDU v2.

   > **Note:**
   >
   > Whether to use **http://** or **https://** for the PDU v2 depends on how your PDU is configured. Try to use both if you are not sure. If neither of http:// or https:// work, the PDU is probably offline or the PDU address is incorrect. If the PDU Management Interface is turned off, you can turn it on using the PDU SNMP interface.

3. In the PDU Management Interface, you can distinguish the PDU version by checking the PDU Power Consumption section. The PDU Consumption details are displayed only for PDU v2. Figure 23-1 is an example for PDU v1 management interface and Figure 23-2 is an example for PDU v2 management interface.

Figure 23-1   PDU v1 Management Interface

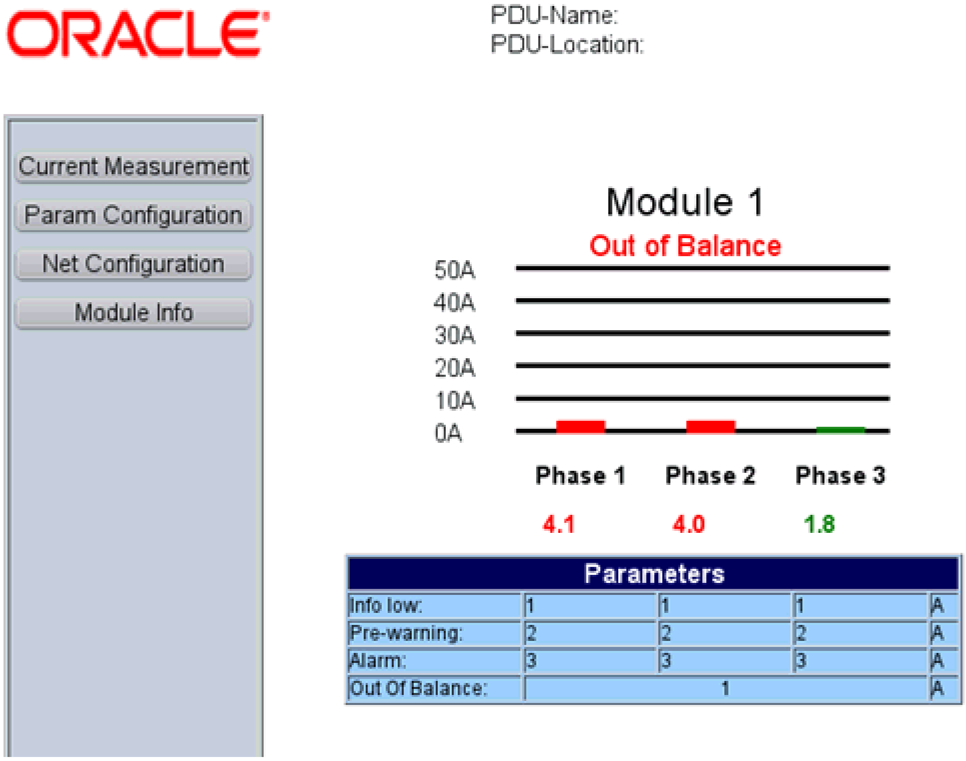

PDU-Name:
PDU-Location:

Parameters				
Info low:	1	1	1	A
Pre-warning:	2	2	2	A
Alarm:	3	3	3	A
Out Of Balance:	1			A

Figure 23-2   PDU v2 Management Interface

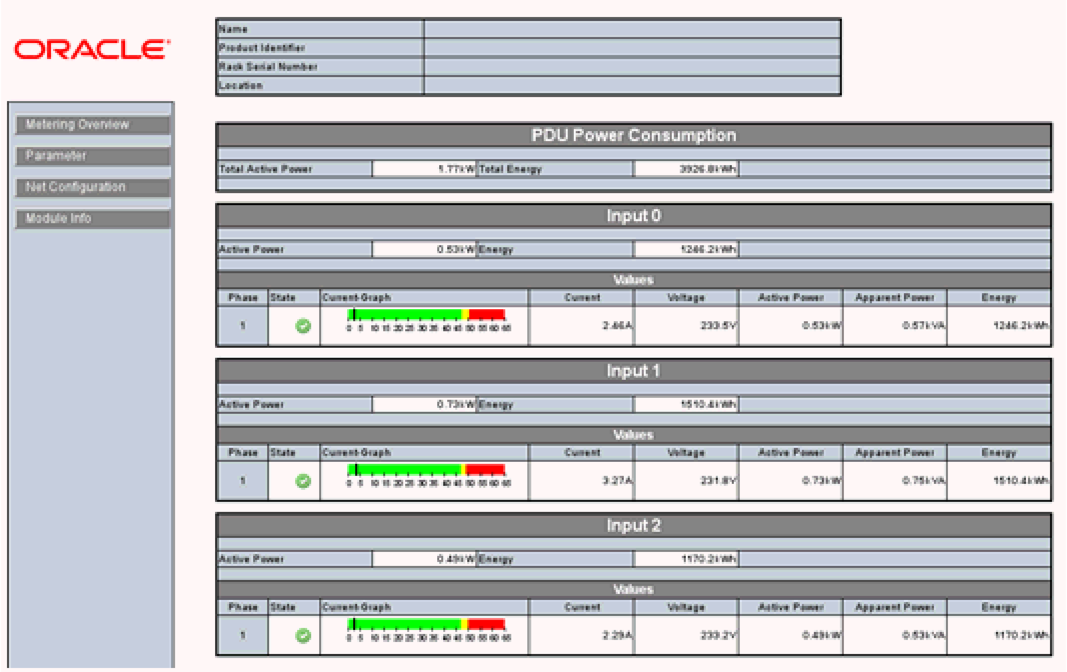

Name	
Product Identifier	
Rack Serial Number	
Location	

PDU Power Consumption

Total Active Power	1.77kW	Total Energy	3926.8kWh

Input 0

Active Power	0.53kW	Energy	1246.2kWh

Values

Phase	State	Current Graph	Current	Voltage	Active Power	Apparent Power	Energy
1	✓		2.46A	233.5V	0.53kW	0.57kVA	1246.2kWh

Input 1

Active Power	0.73kW	Energy	1510.4kWh

Values

Phase	State	Current Graph	Current	Voltage	Active Power	Apparent Power	Energy
1	✓		3.27A	231.8V	0.73kW	0.75kVA	1510.4kWh

Input 2

Active Power	0.49kW	Energy	1170.2kWh

Values

Phase	State	Current Graph	Current	Voltage	Active Power	Apparent Power	Energy
1	✓		2.29A	233.2V	0.49kW	0.53kVA	1170.2kWh

Viewing the PDU Information

The PDU information screen is divided into two parts. The top region consists of dashlets that provide you a general overview of the system. The main region displays more detailed information of the PDU.

The first dashlet in the first series displays the summary of the PDU. It displays PDU model, part number, serial number, IP address, firmware version, count of modules and status of SNMP communication with PDU. If EM Agent can communicate with PDU using SNMP, there is a green tick, if it cannot, there is a red cross.

The second dashlet in the first series displays all the open incidents relayed on the PDU. Click on one of the displayed numbers to view the list of incidents of a given severity.

The third dashlet in the first series displays the power usage of the PDU. It displays a summary of the power (current) in amperes per phase of all the modules of the PDU.

The first dashlet in the second series displays the last configuration changes made on the PDU. It also displays the time when the last incident was raised.

Physical View of the PDU

Click the first tab in the main section of the PDU landing page for a physical view of the PDU. Physical view of the PDU displays the front and side view of the PDU. The PDU consists of one or more modules. You can click on any of the modules or the PDU itself to view more detailed information.

You can switch between photorealistic and tabular representation of the view.

- Photorealistic view provides a realistic picture of the PDU, providing detailed graphics of all the components.
- Table view is a tabular representation of the physical view, providing a list of displayed components with the most important information.

Figure 23-3 is a representation of the physical view of the PDU.

Figure 23-3 PDU Physical View

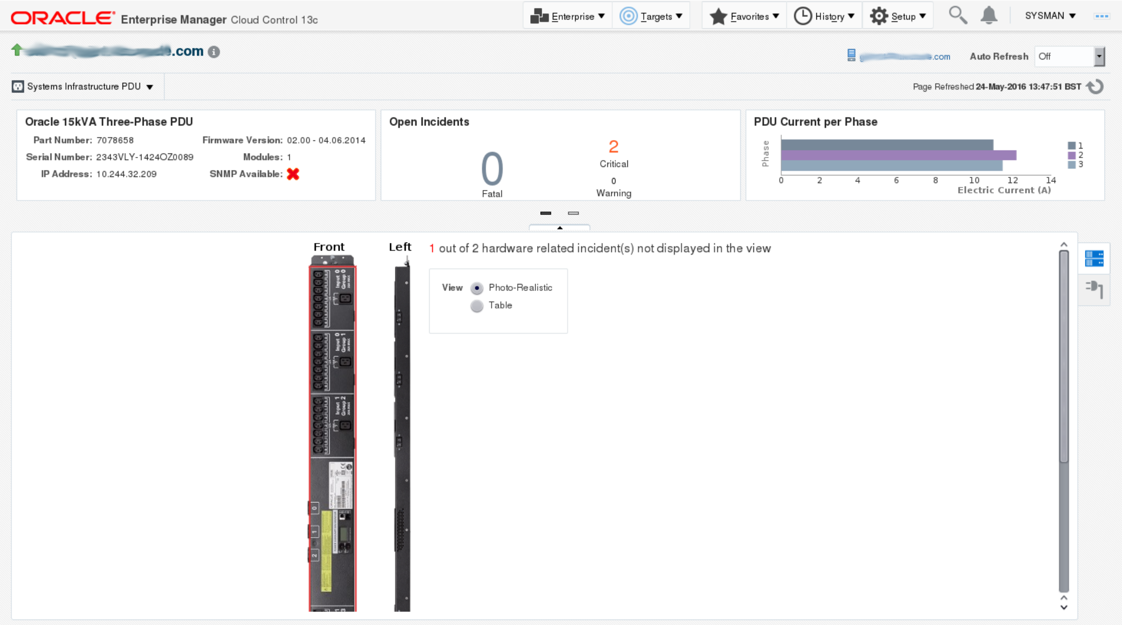

[redacted].com

[redacted].com Auto Refresh Off

Systems Infrastructure PDU

Page Refreshed 24-May-2016 13:47:51 BST

Oracle 15kVA Three-Phase PDU

Part Number: 7078658
Firmware Version: 02.00 - 04.06.2014
Serial Number: 2343VLY-1424OZ0089
Modules: 1
IP Address: 10.244.32.209
SNMP Available: ✗

Open Incidents

0 Fatal
2 Critical
0 Warning

1 out of 2 hardware related incident(s) not displayed in the view

View ● Photo-Realistic
○ Table

PDU Load View

Click the second tab of the main section of the PDU landing page to see the PDU Load view. The PDU load view displays historical data of the phase load (current, in Ampere) per module.

Changing PDU Monitoring Credentials

In case the HTTP and SNMP PDU credentials are changed on the PDU, you can change the credentials in the Enterprise Manager using the PDU's Management Interface.

Change the HTTP Credentials

Perform the following steps to change the http credentials:

1. Under Setup, click **Security**, then click **Monitoring Credentials**.
2. In the Target Type column, select Systems Infrastructure PDU, then click **Manage Monitoring Credentials**.

Figure 23-4 Select Target Type

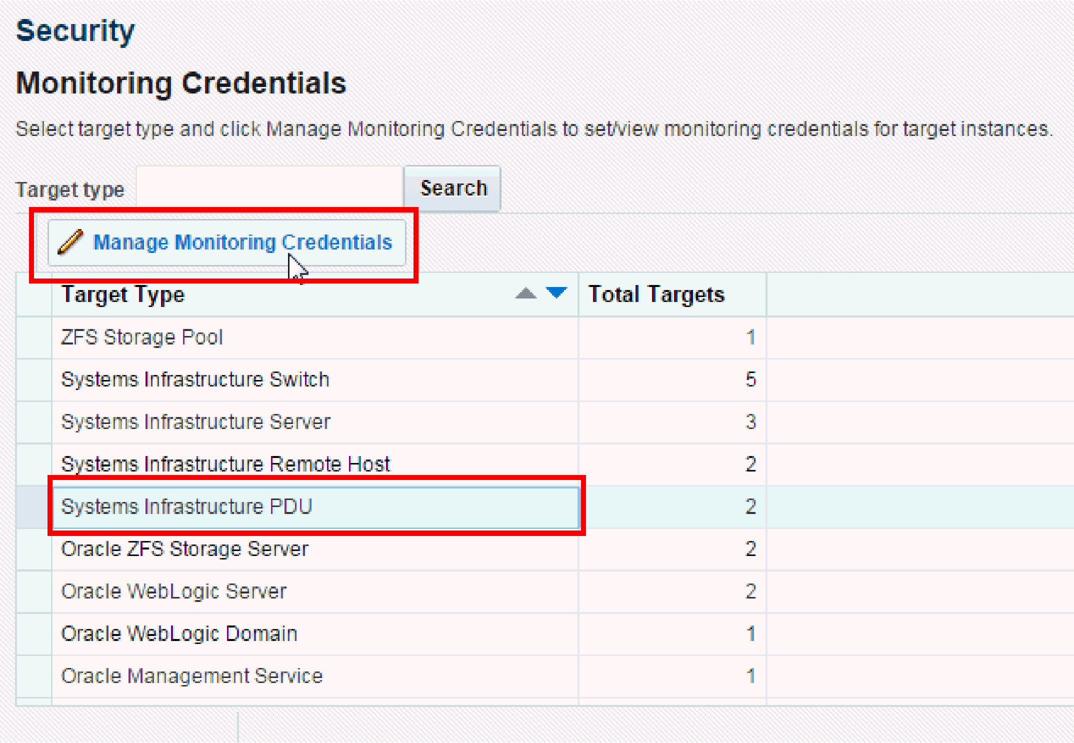

Security

Monitoring Credentials

Select target type and click Manage Monitoring Credentials to set/view monitoring credentials for target instances.

Target type [ ] Search

Manage Monitoring Credentials

Target Type	Total Targets
ZFS Storage Pool	1
Systems Infrastructure Switch	5
Systems Infrastructure Server	3
Systems Infrastructure Remote Host	2
Systems Infrastructure PDU	2
Oracle ZFS Storage Server	2
Oracle WebLogic Server	2
Oracle WebLogic Domain	1
Oracle Management Service	1

3. In the Systems Infrastructure PDU Monitoring Credentials screen, select the HTTP Monitoring Credentials (in the Credential Set column), then click **Set Credentials**.

Figure 23-5 Set HTTP Credentials

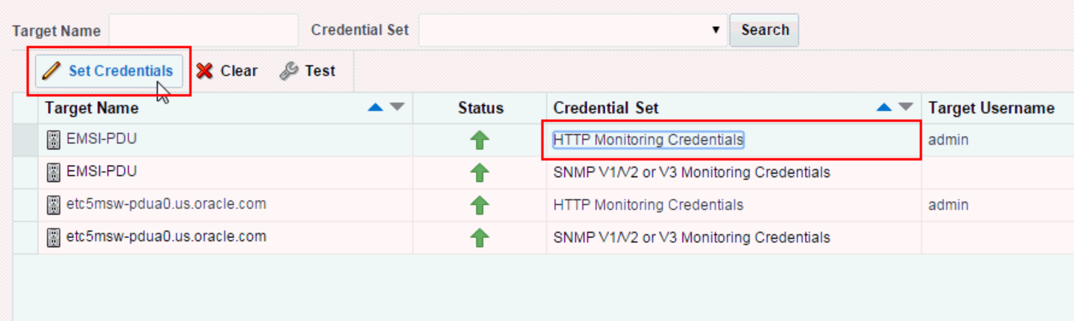

Target Name [ ] Credential Set [ ] Search

Set Credentials Clear Test

Target Name	Status	Credential Set	Target Username
EMSI-PDU	↑	HTTP Monitoring Credentials	admin
EMSI-PDU	↑	SNMP V1/V2 or V3 Monitoring Credentials	
etc5msw-pdua0.us.oracle.com	↑	HTTP Monitoring Credentials	admin
etc5msw-pdua0.us.oracle.com	↑	SNMP V1/V2 or V3 Monitoring Credentials	

4. In the Enter Monitoring Credentials screen, enter new credentials for the PDU HTTP, then click **Test and Save**.

Changing the SNMP Credentials

Perform the following steps to change the SNMP credentials:

1. Under Setup, click **Security**, then click **Monitoring Credentials**.
2. In the Target Type column, select Systems Infrastructure PDU, then click **Manage Monitoring Credentials**.

Figure 23-6 Select Target Type

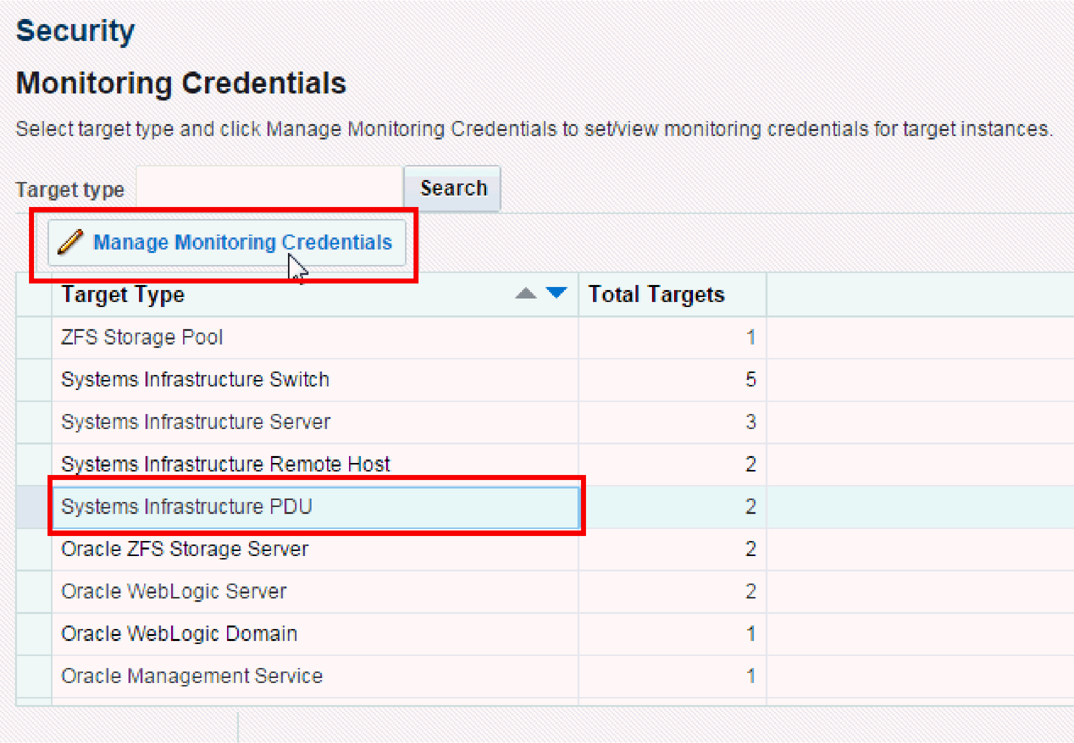

Select target type and click Manage Monitoring Credentials to set/view monitoring credentials for target instances.

Target type

Target Type	Total Targets
ZFS Storage Pool	1
Systems Infrastructure Switch	5
Systems Infrastructure Server	3
Systems Infrastructure Remote Host	2
Systems Infrastructure PDU	2
Oracle ZFS Storage Server	2
Oracle WebLogic Server	2
Oracle WebLogic Domain	1
Oracle Management Service	1

3. In the Systems Infrastructure PDU Monitoring Credentials screen, select the SNMP Monitoring Credentials (in the Credential Set column), then click **Set Credentials**.

Figure 23-7 Set SNMP Credentials

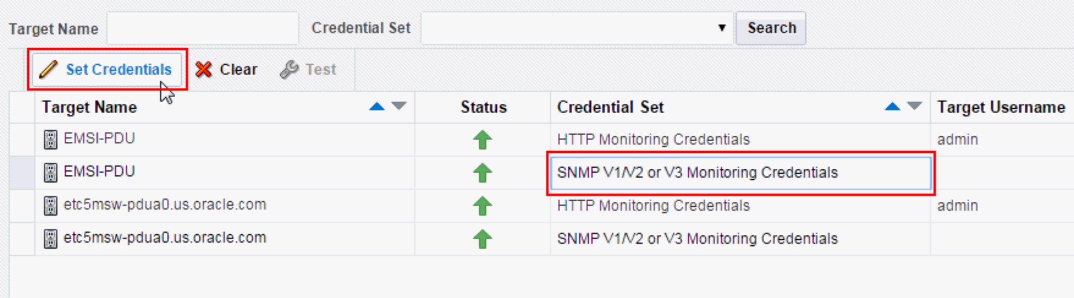

Target Name

Credential Set

Target Name	Status	Credential Set	Target Username
EMSI-PDU	↑	HTTP Monitoring Credentials	admin
EMSI-PDU	↑	SNMP V1/V2 or V3 Monitoring Credentials	
etc5msw-pdua0.us.oracle.com	↑	HTTP Monitoring Credentials	admin
etc5msw-pdua0.us.oracle.com	↑	SNMP V1/V2 or V3 Monitoring Credentials	

4. In the Enter Monitoring Credentials screen, enter new credentials for the PDU SNMP, then click **Save**.

PDU Test Connection and Metric Collection Error Troubleshooting

Collection of some PDU metric or test connection might fail. In that case, PDU test connection fails with an error message. This section explains how to troubleshoot such conditions.

Test Connection Error Identification

If a test connection fails, a message with a problem description is displayed. See the **PDU Error States and Resolution table** for possible PDU error states and their resolution and try to fix the problem source as described in the table and repeat the test connection.

Figure 23-8   Test Connection Error

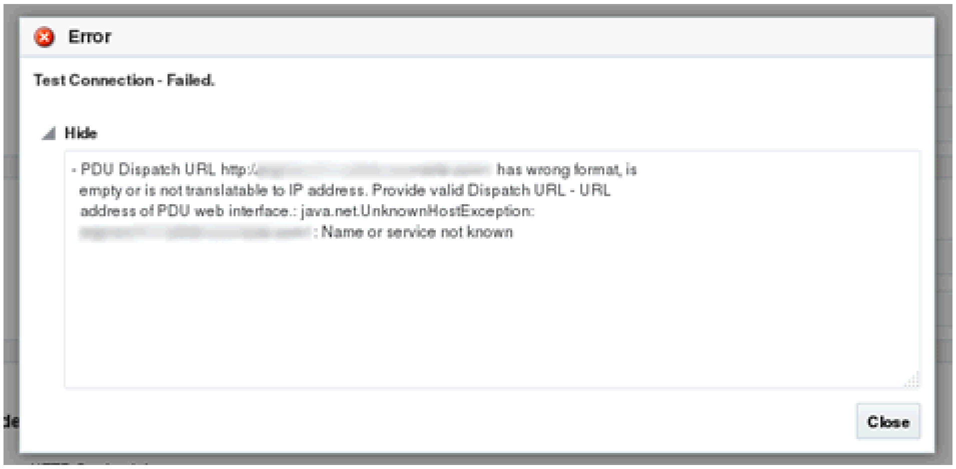

Test Connection - Failed.

- PDU Dispatch URL http:/[redacted] has wrong format, is empty or is not translatable to IP address. Provide valid Dispatch URL - URL address of PDU web interface.: java.net.UnknownHostException: [redacted]: Name or service not known

Metric Collection Error Identification

If collection of some metric fails, a metric collection error event is raised for the PDU.

> **Note:**
>
> Incidents are not generated from metric collection errors by default. You can turn on incidents generation under Setup > Incidents > Incident Rules.
>
> Search for the rule **Group metric collection error events for a target** and enable it.

Perform the following steps to identify and view and view all metric collection errors of PDU:

1. Go to Target Menu at the left top corner of the PDU landing page, click **Systems Infrastructure PDU**.
2. Click **Monitoring**, then click **All Metrics**.
3. In the Overview section, click on the number in Metric Collection Errors.
4. Select and click a metric collection error to view more details. A detailed error description is displayed.

See the **PDU Error States and Resolution table** for possible PDU error states and their resolution, try to fix the problem source as described in the Table 23-2. Repeat evaluation for every failed metric.

Metric Recollection

If you have evaluated all failed metrics, repeat the metric collection.

1. Go to Target Menu at the left top corner of the PDU landing page, click **Configuration**, then click **Last Collected**.
2. In the Latest Configuration screen, click the **Refresh** button.

Figure 23-9 PDU Target Menu

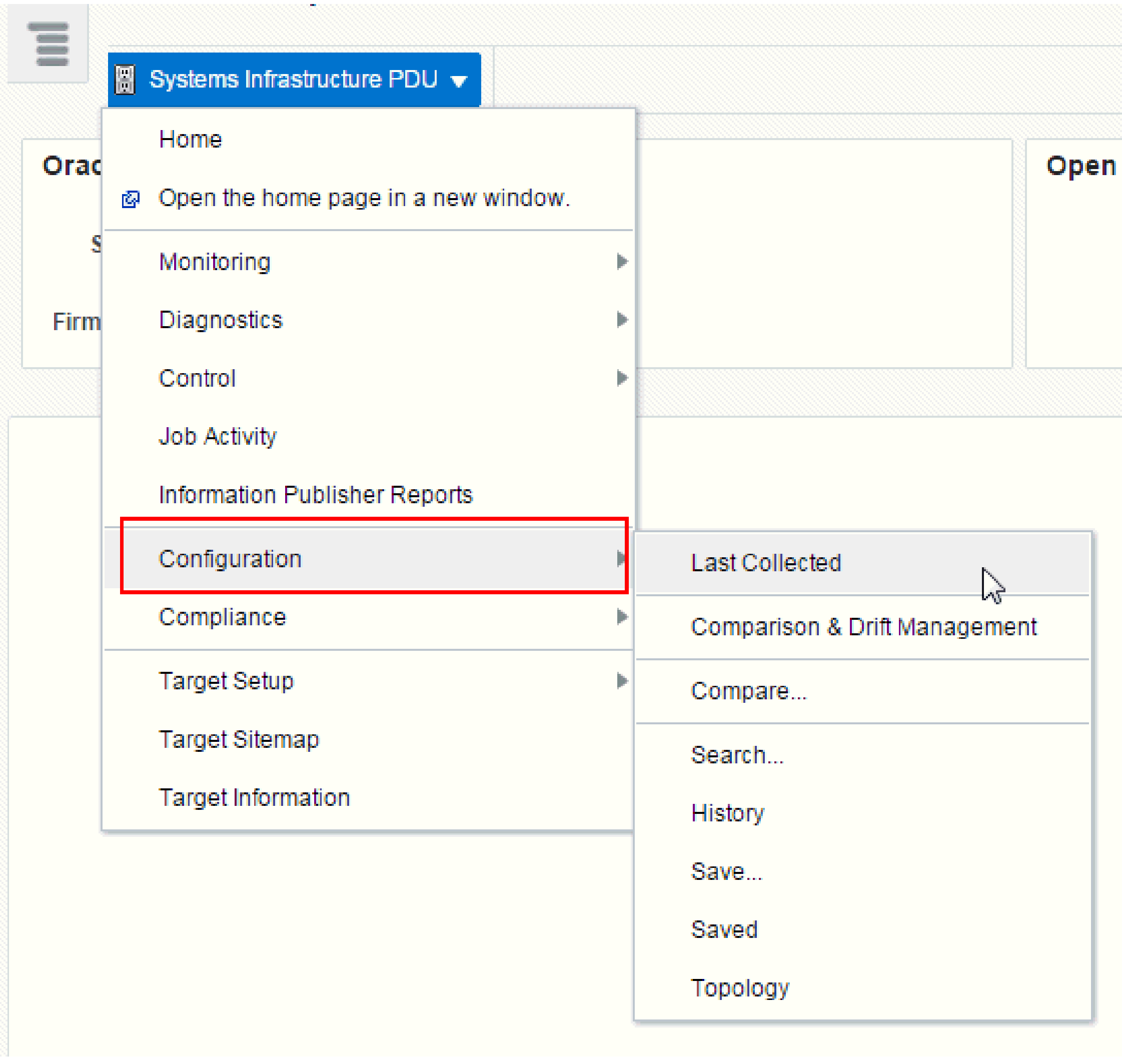

When metrics are collected again and all reasons of collection errors are resolved, the incident disappears from Incidents Manager, dashlet, and Photorealistic view.

If the failed metric is displayed with the word Status, it is a performance metric which cannot be refreshed. You have to wait for the next scheduled metric recollection. In a default configuration, PDU performance metrics are recollected every 15 minutes.

Note:

Metric recollection interval can be changed by user in Enterprise Manager in Target Menu, submenu Metrics, item Metric and Collection Settings.

PDU Error States

Most of the error states that can happen during test connection or metric collection are caused by PDU or network misconfiguration or unresponsive PDU. These errors can be fixed by user.

If you have completed problem resolution, repeat metric collection as described in Metric Recollection or test connection as described in Discovering and Promoting PDUs.

List of errors in the PDU Error States and Resolution table is not a complete list. New messages may be added or changed in the next releases.

Table 23-2 PDU Error States and Resolutions

Error Message	Resolution
PDU Dispatch URL http://pdu.example.com has incorrect format, is empty or is not translatable to IP address. Provide valid Dispatch URL - URL address of PDU web interface.: detailed exception	Check that you provided valid PDU DNS name in PDU discovery. If you are sure you provided valid DNS name, make sure you have got DNS correctly configured on hosts with monitoring and backup agents.
Cannot communicate with #PDU: 127.0.0.1. PDU is unreachable. PDU Web interface cannot be reached using HTTP or HTTPS. Check if PDU is online (Open address http://127.0.0.1 in web browser, try both http:// and https://) and try to repeat action (metric collection, connection test).: detailed exception	Check that PDU is up and running as described in PDU Version Identification. Check your network configuration if PDU is reachable from monitoring agent and backup agent if set.
Cannot communicate with #PDU: pdu.example.com using SNMP. PDU SNMP interface is down or unreachable or SNMP community string is incorrectly configured. Check if PDU is online (Open address https://pdu.example.com in web browser), check if community string is set correctly in Enterprise Manager and in PDU (Open address https://pdu.example.com in web browser, go to Net Configuration section, login, see NMS IPs-Communities table and Trap Hosts Setup IPs-Communities table), check if SNMP is enabled (Open address https://pdu.example.com in web browser, go to Net Configuration section, login, see if SNMP is enabled) and try to repeat action (metric collection, connection test).	Check that PDU is up as described in PDU Version Identification. If yes, check if NMS and Trap Hosts Setup tables are correct and SNMP Community string was not changed or SNMPv3 Access table and Trap Hosts Setup table are correct and SNMPv3 credentials were not changed as described in Verify PDU v1 NMS Table and Trap Hosts Setup Table and Verify PDU v2 NMS Table, SNMPv3 Access Table, and Trap Hosts Setup Table Then recollect SNMP Configuration metric as described in Metric Recollection. If SNMP Community or SNMPv3 credentials were changed in the PDU Management Interface, change it in the Enterprise Manager as well. SNMP credentials changes are described in Metric Collection Error Identification.

Table 23-2 (Cont.) PDU Error States and Resolutions

Error Message	Resolution
Cannot communicate with #PDU: pdu.example.com using SNMP v3	Check that PDU is up as described in PDU Version Identification. If yes, check if SNMPv3 Access table and Trap Hosts Setup tables are correct and SNMPv3 credentials were not changed as described in Verify PDU v2 NMS Table, SNMPv3 Access Table, and Trap Hosts Setup Table Then recollect SNMP Configuration metric as described in Metric Recollection. If SNMPv3 credentials were changed in the PDU Management Interface, change it in the Enterprise Manager as well. SNMP credentials change is described in Changing the SNMP Credentials
Cannot identify PDU model of #PDU: pdu.example.com. PDU model is not supported or it not possible to identify PDU model because the PDU is not reachable. Check that PDU is online (Open address https://pdu.example.com in web browser, try both http:// and https://) and try to repeat action (metric collection, connection test) or try a different PDU.	Check that PDU is up and running as described in PDU Version Identification. Check your network configuration if PDU is reachable. Make sure that IP address od DNS name entered during discovery point to PDU not some different hardware and to supported PDU model as described in PDU Version Identification.
Cannot login to #PDU: pdu.example.com. Wrong credentials. Please provide correct PDU user name and password in Enterprise Manager and try to repeat action (metric collection, connection test)	Check that you provided correct HTTP credentials during PDU discovery. If incident with this message was raised, PDU HTTP credentials were probably changed in the PDU Management Interface. You have to change them in the Enterprise Manager as well. HTTP credentials change is described in Change the HTTP Credentials.
Cannot login to #PDU: pdu.example.com. Another user is already logged in. Cannot proceed till another user is logged out. Ask another user to logout or wait until user is logged out automatically (approximately 30 minutes) and try to repeat action (metric collection, connection test).	Another user is logged in to the PDU Management Interface. If the user who is logged in is you, go to the PDU Management Interface (see Discovering and Promoting PDUs) and click Logout button. If the user logged in is someone else, you have to wait for automatic logout for approximately 30 minutes.
Unable to find PDU SNMP Community string for #PDU: pdu.example.com. Provide correct community string in Enterprise Manager and try to repeat action (metric collection, connection test).	Provide correct SNMP Community string in PDU discovery and repeat discovery. If incident with this message was raised, you have to change SNMP monitoring credentials in the Enterprise Manager. SNMP credentials change is described in Changing the SNMP Credentials.
SNMP version 3 is not supported for #PDU: pdu.example.com original (v1). Provide SNMP version 1 credentials in Enterprise Manager and try to repeat action (metric collection, connection test).	SNMP V3 credentials are not supported for Original PDU (PDU v1), you have to provide SNMP V1 credentials Provide correct SNMP Community string in PDU discovery and repeat discovery. If incident with this message was raised, you have to change SNMP monitoring credentials in the Enterprise Manager. SNMP credentials change is described in Changing the SNMP Credentials.

Table 23-2 (Cont.) PDU Error States and Resolutions

Error Message	Resolution
Cannot write monitoring EM Agent host IP address to the NMS IPs-Communities table of #PDU: pdu.example.com using PDU web interface. There is already entry with desired agent IP address 192.0.2.100 but with different community string. Remove this entry from table manually using the PDU web interface or change it to have correct community string (Open address https://pdu.example.com in web browser, go to Net Configuration section, login, see NMS IPs-Communities table) or provide correct community string in Enterprise Manager and repeat action (SNMP config metric collection, connection test).	Make sure you entered correct SNMP community string during PDU discovery. Check if NMS table is correct and SNMP Community string was not changed as described in Verify PDU v1 NMS Table and Trap Hosts Setup Table and Verify PDU v2 NMS Table, SNMPv3 Access Table, and Trap Hosts Setup Table. If incident with this message was raised, you have to change SNMP monitoring credentials in the Enterprise Manager. SNMP credentials change is described in Changing the SNMP Credentials.
Cannot write monitoring EM Agent host IP address to the Trap Hosts Setup IPs-Communities table of #PDU: pdu.example.com using PDU web interface. There is already entry with desired agent IP address 192.0.2.100 but with different community string. Remove this entry from table manually using the PDU web interface or change it to have correct community string (Open address https://pdu.example.com in web browser, go to Net Configuration section, login, see Trap Hosts Setup IPs-Communities table) or provide correct community string in Enterprise Manager and repeat action (SNMP config metric collection, connection test).	Make sure you entered correct SNMP community string during PDU discovery. Check if Trap Hosts Setup table is correct and SNMP Community string was not changed as described in Verify PDU v1 NMS Table and Trap Hosts Setup Table and Verify PDU v2 NMS Table, SNMPv3 Access Table, and Trap Hosts Setup Table. If incident with this message was raised, you have to change SNMP monitoring credentials in the Enterprise Manager. SNMP credentials change is described in Changing the SNMP Credentials.
Cannot write monitoring EM Agent host IP address to the NMS IPs-Communities table of #PDU: pdu.example.com using PDU web interface. Table is full. Remove some entries from table manually using the PDU web interface (Open address https://pdu.example.com in web browser, go to Net Configuration section, login, see NMS IPs-Communities table) and repeat action (SNMP config metric collection, connection test).	Check that there is an empty slot in the NMS table as described in chapters Verify PDU v1 NMS Table and Trap Hosts Setup Table and Verify PDU v2 NMS Table, SNMPv3 Access Table, and Trap Hosts Setup Table.

Table 23-2 (Cont.) PDU Error States and Resolutions

Error Message	Resolution
Cannot write monitoring EM Agent host IP address to the Trap Hosts Setup IPs-Communities table of #PDU: pdu.example.com using PDU web interface. Table is full. Remove some entries from table manually using the PDU web interface (Open address https://pdu.example.com in web browser, go to Net Configuration section, login, see Trap Hosts Setup IPs-Communities table) and repeat action (SNMP config metric collection, connection test).	Check that there is an empty slot in the Trap Hosts Setup table as described in Verify PDU v1 NMS Table and Trap Hosts Setup Table and Verify PDU v2 NMS Table, SNMPv3 Access Table, and Trap Hosts Setup Table.
Cannot write SNMP v3 monitoring credentials to the SNMP v3 Access table of #PDU: pdu.example.com using PDU web interface. There is already entry with desired user name user but with different password or security level. Remove this entry from table manually using the PDU web interface or change it to have correct password and security level (Open address https://pdu.example.com in web browser, go to Net Configuration section, login, see SNMP v3 Access table) or provide correct SNMP v3 credentials in Enterprise Manager and repeat action (SNMP config metric collection, connection test). Do not provide privacy password. Privacy is not suported for communication between PDU and EM agent.	Make sure you entered correct SNMPv3 credentials during PDU discovery. Check if SNMPv3 Access table and Trap Hosts Setup tables are correct and SNMPv3 credentials were not changed as described in Verify PDU v2 NMS Table, SNMPv3 Access Table, and Trap Hosts Setup Table. If incident with this message was raised, you have to change SNMP monitoring credentials in the Enterprise Manager. SNMP credentials change is described in Changing the SNMP Credentials.
Cannot write SNMP v3 monitoring credentials to the SNMP v3 Access table of #PDU: pdu.example.com using PDU web interface. Table is full. Remove some entries from table manually using the PDU web interface (Open address https://pdu.example.com in web browser, go to Net Configuration section, login, see SNMP v3 Access) and repeat action (SNMP config metric collection, connection test).	Check if there is an empty row in SNMPv3 Access table as described in Verify PDU v2 NMS Table, SNMPv3 Access Table, and Trap Hosts Setup Table.

Table 23-2 (Cont.) PDU Error States and Resolutions

Error Message	Resolution
Cannot write monitoring EM Agent host IP address to the Trap Hosts Setup IPs-Communities table of #PDU: pdu.example.com using PDU web interface. There is already entry with desired agent IP address 127.0.0.1 but with different SNMP v3 user. Remove this entry from table manually using the PDU web interface or change it to have correct SNMP v3 user (Open address {2} in web browser, go to Net Configuration section, login, see Trap Hosts Setup IPs-Communities table) or provide correct SNMP v3 credentials in Enterprise Manager and repeat action (SNMP config metric collection, connection test). Do not provide privacy password. Privacy is not suported for communication between PDU and EM agent.	Make sure you entered correct SNMPv3 credentials during PDU discovery. Check if SNMPv3 Access table and Trap Hosts Setup table are correct and SNMPv3 credentials were not changed as described in Verify PDU v2 NMS Table, SNMPv3 Access Table, and Trap Hosts Setup Table. If incident with this message was raised, you have to change SNMP monitoring credentials in the Enterprise Manager. SNMP credentials change is described in Changing the SNMP Credentials.
Too short SNMP Authentication password for #PDU: pdu.example.com. Password must be at least 8 characters long. Provide correct SNMP credentials in Enterprise Manager and try to repeat action (metric collection, connection test).	If incident with this message was raised, you have to change SNMP monitoring credentials in the Enterprise Manager. SNMP credentials change is described in Changing the SNMP Credentials. Provide SNMPv3 password at least 8 characters long.
SNMP Credentials for #PDU: pdu.example.com were not provided. Provide correct SNMP credentials in Enterprise Manager and try to repeat action (metric collection, connection test).	If incident with this message was raised, you have to change SNMP monitoring credentials in the Enterprise Manager. SNMP credentials change is described in Changing the SNMP Credentials.
You have to provide SNMPv3 authentication password for #PDU: pdu.example.com. Provide SNMP version 3 credentials in Enterprise Manager and try to repeat action (metric collection, connection test).	If incident with this message was raised, you have to change SNMP monitoring credentials in the Enterprise Manager. SNMP credentials change is described in Changing the SNMP Credentials.
Do not provide privacy password in SNMP v3 credentials for #PDU: pdu.example.com. Privacy is not suported for communication between PDU and EM agent. Remove privacy password from SNMP version 3 credentials in Enterprise Manager and try to repeat action (metric collection, connection test).	If incident with this message was raised, you have to change SNMP monitoring credentials in the Enterprise Manager. SNMP credentials change is described in Changing the SNMP Credentials. Remove privacy password from SNMP version 3 credentials.

Table 23-2 (Cont.) PDU Error States and Resolutions

Error Message	Resolution
#PDU: pdu.example.com reports it has zero modules or phases. PDU web interface does not work correctly. Please restart PDU web interface. Refer to Sun Rack II Power Distribution Units User''s Guide for how to (depending on PDU version, PDU can be restarted using dedicated hardware button or from PDU web interface https://pdu.example.com). Try to repeat action (metric collection, connection test) after restart.	Follow error message text.
You have specified wrong SNMP MIB version to be used to monitor #PDU: pdu.example.com. Please enter 'Enhanced' or 'Original'. You entered something.	Follow error message text.
Exception during lookup for IP address of network interface on EM Agent host through which is #PDU: pdu.example.com with address #PDU: pdu.example.com reachable. PDU is not reachable from Agent host. Exception during lookup for IP addresses of network interfaces on EM Agent host through which #PDU: pdu.example.com could be reachable. #PDU: pdu.example.com with address 203.0.113.200 is not reachable through any network interface on EM Agent host. Select another agent or try to resolve issues causing that PDU is not reachable.	Check that PDU is up and running as described in PDU Version Identification. Check your network configuration if PDU is reachable.
#PDU: pdu.example.com: Problem description. You probably tried to access unsupported PDU hardware not a supported PDU.	Check that PDU is up and running as described in PDU Version Identification. Check your network configuration if PDU is reachable. Make sure that IP address od DNS name entered during discovery point to PDU not some different hardware and to supported PDU model as described in PDU Version Identification.
Unsupported model of #PDU: pdu.example.com. Only PDU with maximum count of 4 modules is supported. Detected count of modules: 5	Check that PDU is up and running as described in PDU Version Identification. Check your network configuration if PDU is reachable. Make sure that IP address od DNS name entered during discovery point to PDU not some different hardware and to supported PDU model as described in PDU Version Identification.

Table 23-2 (Cont.) PDU Error States and Resolutions

Error Message	Resolution
PDU does not respond. It is overloaded or offline. Cannot get/find something. Check if PDU is online (Open address https://pdu.example.com in web browser) and try to repeat action (metric collection, connection test).	Check that PDU is up and running as described in chapter PDU Version Identification. Check your network configuration if PDU is reachable. Make sure that IP address od DNS name entered during discovery point to PDU not some different hardware and to supported PDU model as described in PDU Version Identification.

PDU Alerts and Configuration

PDU and Enterprise Manager can be configured to report two kinds of incidents to the user:

- If current level of a phase of some module in amperes crossed some set warning or alarm threshold.
- If difference between current level in ampere of phases of PDU module is bigger than the set threshold.

You can set warning and alarm thresholds in both PDU and Enterprise Manager. These settings are independent on each other. Incidents are generated independently from warning and alarm thresholds set in PDU and in Enterprise Manager.

Configuring Alerts in a PDU

Perform the following steps to configure alarm and warning thresholds in a legacy PDU Management Interface:

1. Open the PDU Management Interface in the web browser (See PDU Version Identification).
2. In the PDU User Interface, click **Param Configuration**.
3. Login with your user name and password.
4. Set alarm and warning current thresholds in amperes, then click **Submit**.

> **Note:**
>
> Info low is not used for PDU monitoring in Enterprise Manager.

5. Repeat the above steps for all modules.

Configuring Alerts in Enterprise Manager

Perform the following steps to configure alarm and warning thresholds in Enterprise Manager.

1. Open PDU landing page as described in Physical View of the PDU.
2. Click the Target Menu Systems Infrastructure PDU.
3. Under Systems Infrastructure PDU, select **Monitoring**, then click **Metric and Collection Settings**.

Figure 23-10 Metric and Collection Settings

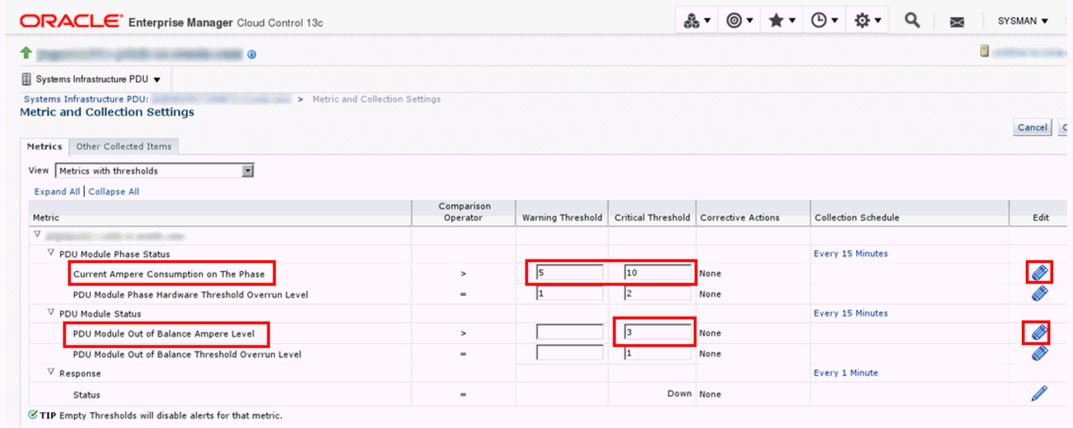

ORACLE Enterprise Manager Cloud Control 13c

SYSMAN

[redacted]

Systems Infrastructure PDU

Systems Infrastructure PDU: [redacted] > Metric and Collection Settings

Metric and Collection Settings

Cancel

Metrics | Other Collected Items

View: Metrics with thresholds

Expand All | Collapse All

Metric	Comparison Operator	Warning Threshold	Critical Threshold	Corrective Actions	Collection Schedule	Edit
[redacted]						
PDU Module Phase Status					Every 15 Minutes	
Current Ampere Consumption on The Phase	>	5	10	None		✎
PDU Module Phase Hardware Threshold Overrun Level	=	1	2	None		✎
PDU Module Status					Every 15 Minutes	
PDU Module Out of Balance Ampere Level	>		3	None		✎
PDU Module Out of Balance Threshold Overrun Level	=		1	None		✎
Response					Every 1 Minute	
Status	=		Down	None		✎

✓ TIP Empty Thresholds will disable alerts for that metric.

4. In the table, click the Edit icon against Current Ampere Consumption on the Phase and edit the values of Warning Threshold and Critical Threshold.
5. Click the Edit icon against PDU Module out of Balance Ampere Level and edit the value of Critical Threshold.

> **Note:**
>
> Do not change Warning Threshold and Critical Threshold for PDU Module Phase Hardware Threshold Overrun Level and PDU Module out of Balance Threshold Overrun Level. It those values are changed, incidents will not to be generated based on alarm and warning levels set directly in the PDU or they will be generated incorrectly.
>
> If values were already changed, set PDU Module Phase Hardware Threshold Overrun Level Warning Threshold to 1, Critical Threshold to 2, and PDU Module out of Balance Threshold Overrun Level Critical Threshold to 1.

6. Click **OK**.

Viewing Alert Incidents

To view incidents generated from alarm and warning threshold and identify the incident, click Open Incidents dashlet to view a summary of all the incidents on the PDU. Click a specific incident to view incident details.

- For example, any incident generated from alarm and warning thresholds set directly in PDU contains the following text:

  PDU Module 0, Phase 1 crossed the Module Phase Ampere Level alarm or warning threshold set in PDU Web interface on Parameter page. Module Phase Ampere Level was 2.4 A.

- For example, any incident generated from alarm and warning thresholds set in the Enterprise Manager contains the following text:

  PDU Module 1, Phase 1 crossed the Module Phase Ampere Level alarm 2 A or warning 1 A threshold set by user in monitoring template. Module Phase Ampere Level was 3.2 A.

Incidents are automatically cleared when current level goes under the set alarm and warning levels.

SNMP Traps Forwarding

PDU is by default configured by Enterprise Manager to generate and send traps to IP address where Enterprise Manager monitoring agent and backup reside.

Enterprise Manager can receive these traps and generate Alert Incidents immediately after alert condition is met in PDU.

PDU can send traps only to default SNMP traps port UDP 162 as is defined in SNMP standard. Enterprise Manager agent however listens for SNMP traps on a different port.

Because of this port mismatch, SNMP traps generated by PDU are not delivered to Enterprise Manager by default.

Alert Incidents are still generated from alert and warning thresholds set in PDU but they are not generated immediately. They are generated when PDU performance metrics are recollected. By default performance metrics are recollected every 15 minutes.

> **Note:**
>
> *Metric recollection interval can be changed by user in Enterprise Manager in Target Menu, submenu Metrics, item Metric and Collection Settings.*

To deliver PDU SNMP traps from PDU to Enterprise Manager, you have to set port forwarding on hosts where monitoring and backup agents are deployed. Forward UDP port 162 to UDP port where Enterprise Manager agent is listening for SNMP traps.

How to set port forwarding depends on operation system you use on host where agent is deployed. Consult your operating system documentation on how to set UDP port forwarding.

> **Note:**
>
> *Do not use Linux tool snmptrapd to forward traps. The traps from snmptrapd do not contain IP of the originating PDU, but of the snmptrapd forwarder. Enterprise Manager uses PDU IP address to couple received SNMP trap with monitored PDU. If received SNMP trap has wrong originator IP address, it is thrown away by Enterprise Manager and not considered any more.*

To find the port where PDU monitoring and backup agent listen for SNMP traps you can:

- Get the port from EMD_URL property in the *AGENT_INST/sysman/config/emd.properties* file on hosts where primary and monitoring agents are deployed. (This is the port at which agent will listen over UDP for traps).
- Open All Targets page in the Enterprise Manager UI and find monitoring and backup agent in the targets list. Number in the agent name after colon (:) is the UDP port which the agent will listen for SNMP traps.

Related Resources for PDU Management

See the following for more information:

- Discovering and Adding Host and Non-Host Targets
- Discovering, Promoting, and Adding System Infrastructure Targets
- Using Incident Management
- Enterprise Monitoring

24

Managing the Rack

The following information is included in this chapter:

Getting Started with Rack Management

Target management is the process through which Enterprise Manager begins to manage and monitor your targets including server hardware, chassis, racks, power distribution unit, network equipment, operating systems, virtualization software, and clustering software. Discovering and managing your targets is a prerequisite for almost every action in the software. The discovery feature makes adding targets quick and easy. You can discover power distribution units and racks using the guided process.

Location of Rack Information in the User Interface

Table 24-1 shows where to find information.

Table 24-1 Location of Rack Information in the BUI

Object	Location
Rack	In the Enterprise Manager user interface, under Targets, click **All Targets**. In the Refine Search section, under Target type, click **Servers, Storage, and Network**. Click **Systems Infrastructure Rack**, then select a rack from the displayed list.
Physical View of Rack	In the Enterprise Manager user interface, under Targets, click **All Targets**. In the Refine Search section, under Target type, click **Servers, Storage, and Network**. Click **Systems Infrastructure Rack**, then select a rack from the displayed list. Click the Dashboard tab. Select Photorealistic View radio button on the right side viewing option.

Actions for Rack

You can perform the following actions, depending on the requirements.

- Create a rack
- View the rack
- Place a target to the rack
- Edit a target in the rack
- Remove a target from the rack
- Delete a rack

Target Navigation for Rack Management

The target navigation tree displays the rack and all targets in the rack in a tree structure. Click on any of the targets to navigate to the landing page of the selected target.

> **Note:**
>
> The target navigation tree is active only if the rack is populated with targets.

Figure 24-1 Target Navigation for Rack Management

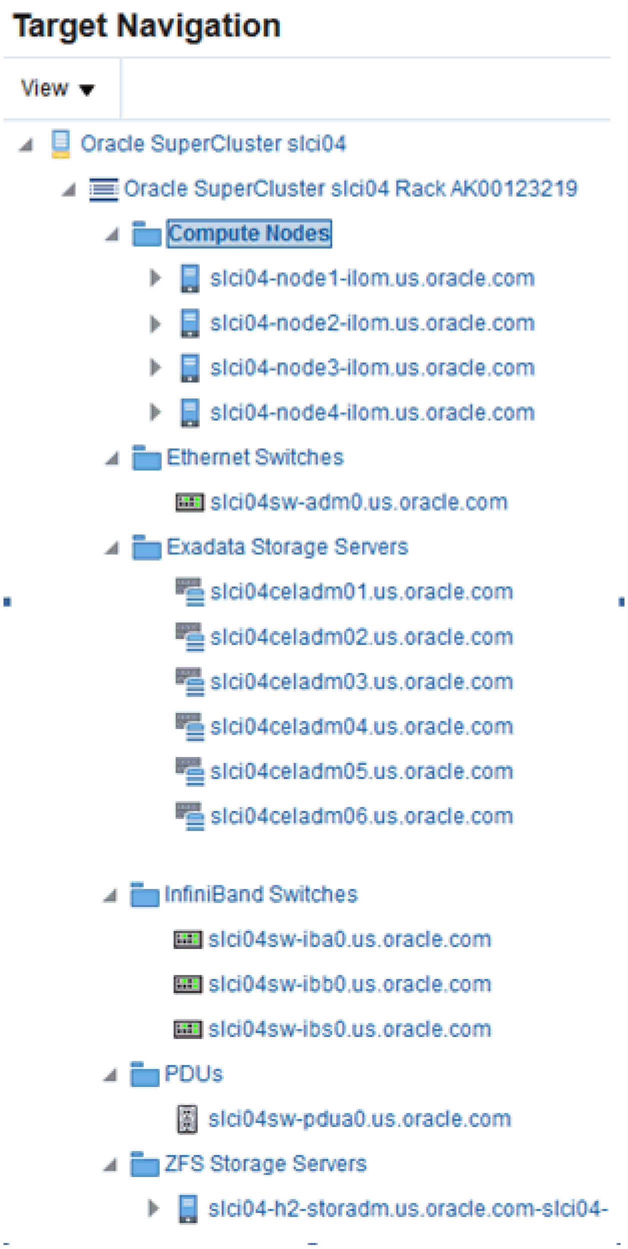

Target Navigation

View ▾

- Oracle SuperCluster slci04
  - Oracle SuperCluster slci04 Rack AK00123219
    - Compute Nodes
      - slci04-node1-ilom.us.oracle.com
      - slci04-node2-ilom.us.oracle.com
      - slci04-node3-ilom.us.oracle.com
      - slci04-node4-ilom.us.oracle.com
    - Ethernet Switches
      - slci04sw-adm0.us.oracle.com
    - Exadata Storage Servers
      - slci04celadm01.us.oracle.com
      - slci04celadm02.us.oracle.com
      - slci04celadm03.us.oracle.com
      - slci04celadm04.us.oracle.com
      - slci04celadm05.us.oracle.com
      - slci04celadm06.us.oracle.com
    - InfiniBand Switches
      - slci04sw-iba0.us.oracle.com
      - slci04sw-ibb0.us.oracle.com
      - slci04sw-ibs0.us.oracle.com
    - PDUs
      - slci04sw-pdua0.us.oracle.com
    - ZFS Storage Servers
      - slci04-h2-storadm.us.oracle.com-slci04-

Creating a Rack

A rack target serves as a container for other hardware targets that are managed by Enterprise Controller. To create a rack, perform the following steps:

1. Log in to Enterprise Manager.
2. Under Setup, click **Add Target**, then click **Add Targets Manually**.
3. In the Overview section, click **Add Targets using Guided Process**.
4. In the Add Using Guided Process screen, scroll down to Systems Infrastructure Rack, then click **Add**.

5. In the Systems Infrastructure Rack Discovery screen, enter the required information.
   a. In the Target Name field, enter a name for the rack.
   b. In the Type field, select the type of rack and additional information if required.
   c. (Optional) You can also set additional information, such as Location, in the Global Properties section.
6. Click **Add** in the top right corner of the screen. Once the job is successfully run, an empty rack is created. You can now navigate to the rack landing screen and add hardware targets to the rack. The following figure is an image of an empty rack.

Figure 24-2 Empty Rack

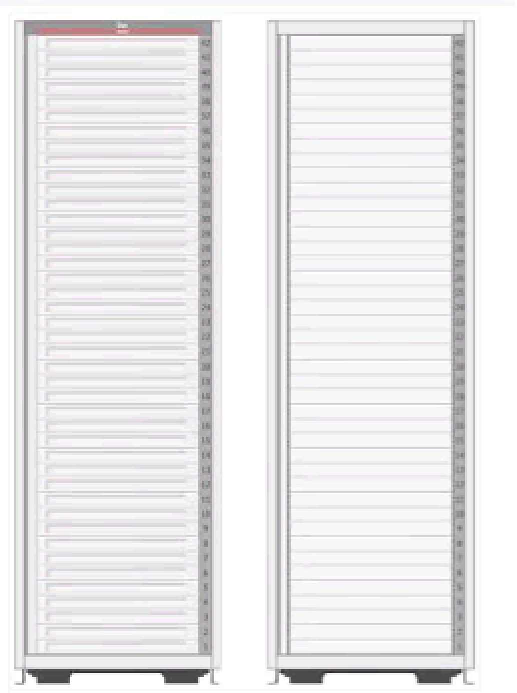

Creating a Rack Using Command Line Interface

You can create a rack using the command line interface.

Perform the following steps to create a rack using CLI:

1. Open the command line interface on the host where OMS is running.
2. Log in to emcli using the following command: *emcli login –username=<your user name>*
3. Type the password when prompted.
4. Execute `emcli sync`.
5. Add a new rack using the following command:

```
emcli add_target \

-name="Name of your Rack" \

-type=oracle_si_rack \

-subseparator=properties='=' \
-separator=properties=';' \
-
properties='EngineeredSystemId=SomeID;RackType=SomeType;RackSubtype=SomeSubtyp
e;TotalSlots=42'
```

6. Set the following in the *emcli add_target* command:
   - Replace **Name of your Rack** with **name of your Rack**
   - Set values for properties

Properties of Rack

Rack has four properties, namely, EngineeredSystemId, RackType, RackSubtype, and TotalSlots. See Table 24-2 for description of the properties.

Table 24-2 Properties Description

Property	Description	Allowed Values	Mandatory	Note
EngineeredSystemId	Unique identifier of an Engineered System	Arbitrary string	No	Does not have to be provided if Rack is standalone not belonging to some Engineered System
RackType	Rack Type (Generic 42U Cabinet or well-known type e.g. Oracle Exalogic, SPARC SuperCluster, or Oracle Database Appliance)	• GENERIC: Generic 42U rack cabinet • EXALOGIC: Exalogic Sun Rack II 42U rack cabinet • EXADATA: Exadata Sun Rack II 42U rack cabinet • SUPERCLUSTER: SuperCluster Sun Rack II 42U rack cabinet • BIGDATA: Oracle Big Data Appliance Sun Rack II 42U rack cabinet • OPCA: Oracle Private Cloud Appliance Sun Rack II 42U rack cabinet • ZDLRA: Oracle Zero Data Loss Recovery Appliance Sun Rack II 42U rack cabinet • EXADATA_STORAGE_EXPANSION: Oracle Exadata Storage Expansion Sun Rack II 42U rack cabinet	Yes	Provide GENERIC if Rack is standalone not belonging to some Engineered System

Table 24-2 (Cont.) Properties Description

Property	Description	Allowed Values	Mandatory	Note
RackSubtype	Optional specification of the RackType size according to Rack Type. For example, Full, Quarter etc. This property is applicable only for Engineered Systems racks.	• UNDEFINED • FULL • HALF • QUARTER • EIGHTH	No	NA
TotalSlots	Total count of slots in the rack.	42	Yes	Rack with 42 slots only is supported.

Following is a sample command to create a generic rack:

```
emcli add_target \

-name="Name of your Rack" \

-type=oracle_si_rack \

-subseparator=properties='=' \

-separator=properties=';' \

-properties=RackType=GENERIC;TotalSlots=42'
```

Viewing the Rack Information

The rack information screen is divided into two parts. The top region consists of dashlets that provide you a general overview of the system. The main region displays more detailed information of the rack.

The dashlets are grouped by three into one or more series that can be switched using navigation buttons.

The first dashlet in the first series displays the summary of the rack, including the serial number of the rack, total number, and number of occupied slots. It also displays the location of the rack if the information is available.

The second dashlet in the first series displays number of open incidents relayed on the rack and all the targets in the rack. Click on one of the displayed numbers to view the list of the incidents of a given severity.

The third dashlet in the first series displays the current (in Ampere) usage per phase of the rack.

The first dashlet in the second series displays the number of targets that are online, offline, and aggregated number of targets in any other status.

The second dashlet in the second series displays the occupancy of the rack. It displays the slot occupancy and PDU occupancy of the rack in a graphical representation.

The third dashlet in the second series displays the last configuration changes made to the rack.

> **Note:**
>
> Click the Second dashlet series icon (displayed below the dashlets) to view more dashlets. Click the First dashlet series icon to view the first set of dashlets.

Physical View of the Rack

Click the first tab in the main section of the rack landing page for a physical view of the rack. Photorealistic view of the rack displays the rack and it's content providing an overview of how the targets are placed into the rack. You can click on any of the targets in the rack to view more detailed information about the selected targets. Click on the rack to view the information about the rack itself. You can switch between several representations of the physical view.

- Photorealistic view provides a realistic picture of the system with detailed graphics of all the components.
- Schematic view is data oriented and displays the most important information such as locator light, status, temperature, and host name. Each component in the view has its color based on the type for easy identification of the component. Click Show Temperature option on the right side to view the temperature of targets that provide that information. To place targets into the empty slots, check the Empty Slot option below the Temperature toggle. If unchecked, the empty slots are not active.
- Table view is a tabular representation of the physical view, providing a list of displayed components with the most important information.

Figure 24-3 Photorealistic View of the Rack

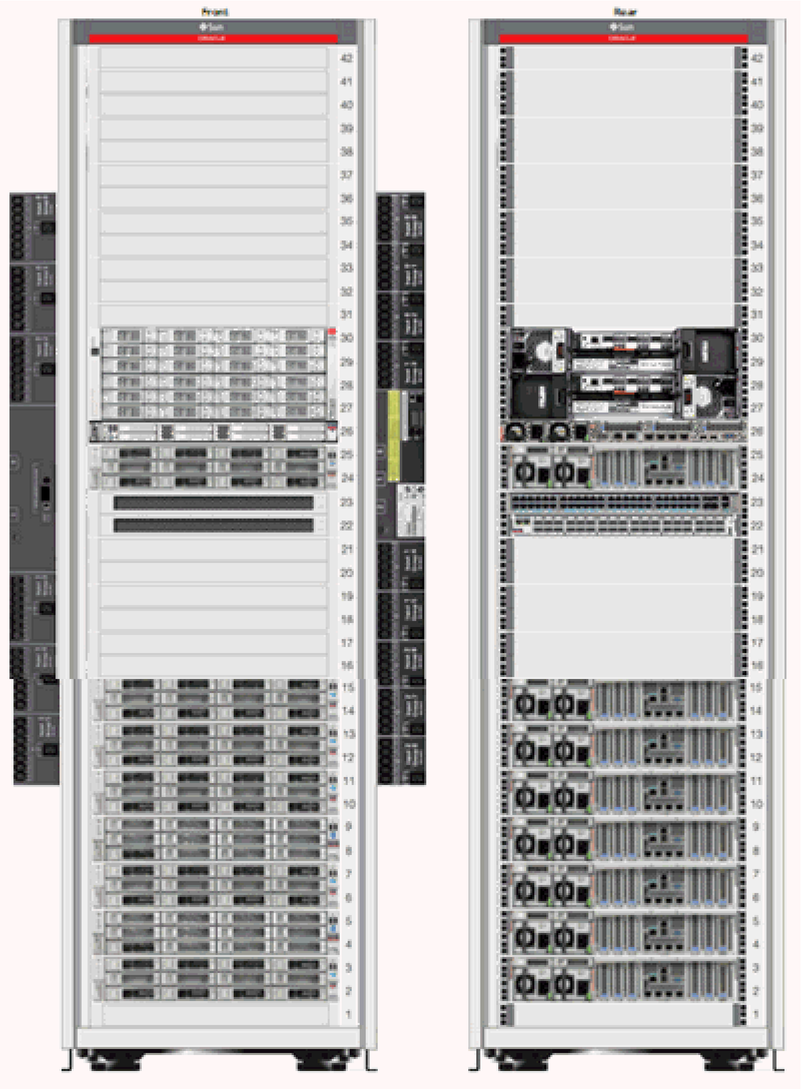

Firmware View

Click the second tab in the main section of the rack landing page to display the Firmware view. Firmware View displays a tabular view of all the firmware of targets placed in the rack. It displays the Target Name, Target Type, Firmware Type, and Firmware Version in a tabular format. The search option is available on the top of the table header for all columns such as Target Name, Target Type, Firmware Type, and Firmware Version. In the Target Type search field, select the type of target from the drop-down list.You can enter a firmware type in the firmware search field to view targets with the selected firmware type only.

Load View

Load View displays detailed information of the load of the PDUs in the rack. It displays a graphical view of the PDU phase load per module where the time history of the PDU phase load is displayed.

Temperature View

Temperature View displays the temperature of all the targets in the rack. Historical data such as average and maximum temperatures for last seven days is also displayed.

Placing Targets in the Rack

After the rack is created, you can place (add) hardware targets in the rack. A photorealistic image of the empty rack is displayed with front and rear views. There are two types of slots in the rack, rack slot and PDU slot. In the rack slot, you can place switches, servers, ZFS appliances, storage cells, and other targets that can be placed into a rack. In the PDU slots, you can place the PDU.

You can invoke the action menu by right clicking on the empty or occupied rack or the PDU slot.

Place a Target in the Rack

To place a target in the rack, perform the following steps:

1. Right click on an empty slot in the rack and click **Place Target**. The Place Target into Rack wizard opens.
2. In the Selected Target field, click the search icon to find the target that you want to add.
   a. In the Target Type field, select the check box against the type of target that you want to add. The discovered targets are listed.
   b. Select a target from the list, then click **Select**.
3. In the Position field, the value is automatically filled from the slot position you clicked. If you want to place the target in a different slot, enter the position of the empty slot where you want the target to be placed.
4. In the Height field, specify the height of the target. The suggested height displayed is based on the type of the target that you selected.
5. In the Facing field, specify the orientation of that target such as Full, Front, or Rear. (Some targets might occupy only half of the slot, for example, switches).
6. In the Horizontal position field, specify whether Full, Left, or Right. This is based on the occupancy of that target. Some targets might occupy the full slot, some might occupy either the left or right side of the slot.
7. Click **OK**. The target is placed into the rack.

Edit Target Placement in the Rack

You can edit the target position in the rack and place it to a different slot if needed.

To edit a target position, perform the following steps:

1. Right click on the target, then select **Edit Target**.
2. Edit the corresponding values and click **OK**.

Remove a Target from the Rack

You can delete targets placed in the rack.

Perform the following steps to remove a target from the rack.

1. Navigate to **Physical View of the Rack**.
2. Right click on the target that you want to delete, then click **Remove Target**.
3. Click **OK** to confirm. The target is removed from the rack.

> **Note:**
>
> Though the target is removed from the rack, Enterprise Manager continues to monitor the target.

Delete a Rack

The rack can be removed in Enterprise Manager without targets placed in it. If the rack target is deleted and the rack contains targets placed in it, you will be asked if you want to remove the targets from Enterprise Manager.

Perform the following steps to delete a rack:

1. In the rack landing page, click the Target menu drop-down list.

Figure 24-4   Delete Rack

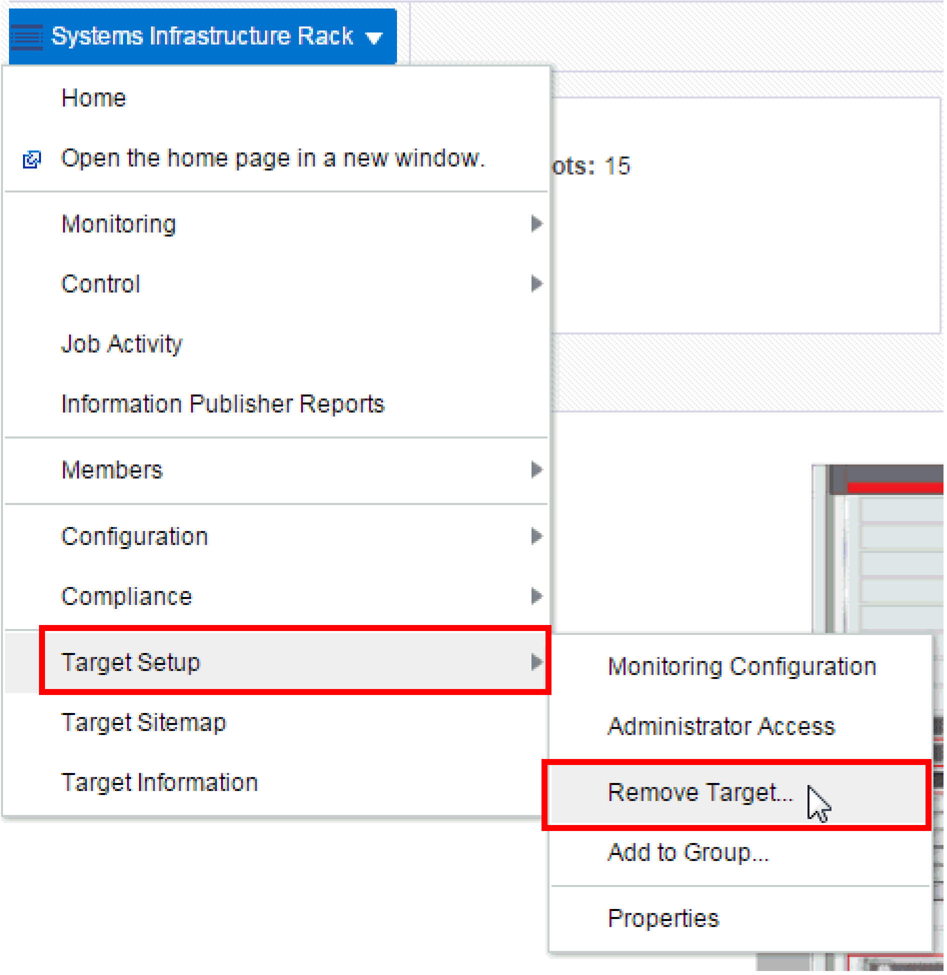

2. Click **Target Setup**, then click **Remove Target**.
3. If you want remove targets placed in rack along with the removed rack, select “Remove all targets associated to target Rack”. List of targets placed in rack will be shown.
4. Some targets cannot be removed. Such targets have a red cross in the column Removal Details including a reason as to why target cannot be removed. Unremovable targets will not be removed when rack target is removed.
5. Some targets provide a custom removal flow (like SNMP monitoring unsubscription). Such targets have a yellow exclamation mark in the column Removal Details. If you want to remove a target with a custom removal flow, go to the target’s landing page and remove the target from there. If you don't remove targets with a custom removal flow, they will be removed along with rack without custom removal options available in their removal flows.
6. Click **Yes** to confirm. The rack is deleted with or without the assigned targets depending on your selection.

Related Resources for Rack Management

See the following for more information:

- Discovering and Adding Host and Non-Host Targets
- Discovering, Promoting, and Adding System Infrastructure Targets

- Using Incident Management
- Enterprise Monitoring

25
Managing Oracle MiniCluster

The following information is included in this chapter:

Getting Started with Oracle MiniCluster

Oracle MiniCluster is an Oracle Engineered System that integrates two SPARC S-7 servers and Oracle Storage Drive Enclosure DE3-24C.

Oracle MiniCluster is supported in the following configurations:

- Oracle MiniCluster S7-2

To be able to start target monitoring, it is necessary to discover it. The discovery is made quick and easy with the Guided Discovery Wizard that guides you through the whole process, requests only for necessary information, and helps you solve possible issues in order to successfully complete the discovery.

The following MiniCluster Software releases are supported:

- 1.1.0
- 1.2.0
- 1.3.0
- 1.4.0

Actions for Oracle MiniCluster

You can perform the following actions, depending on the requirements.

- Discover Oracle MiniCluster
- View the Oracle MiniCluster system
- Monitor the Oracle MiniCluster system

Target Navigation for Oracle MiniCluster

The target navigation tree helps you to navigate between targets in the Oracle MiniCluster system. The Disk shelves are not listed in the navigation tree, information about Disk Shelf devices is available in the Storage tab of the MiniCluster target landing page.

Figure 25-1 Target Navigation of Oracle MiniCluster

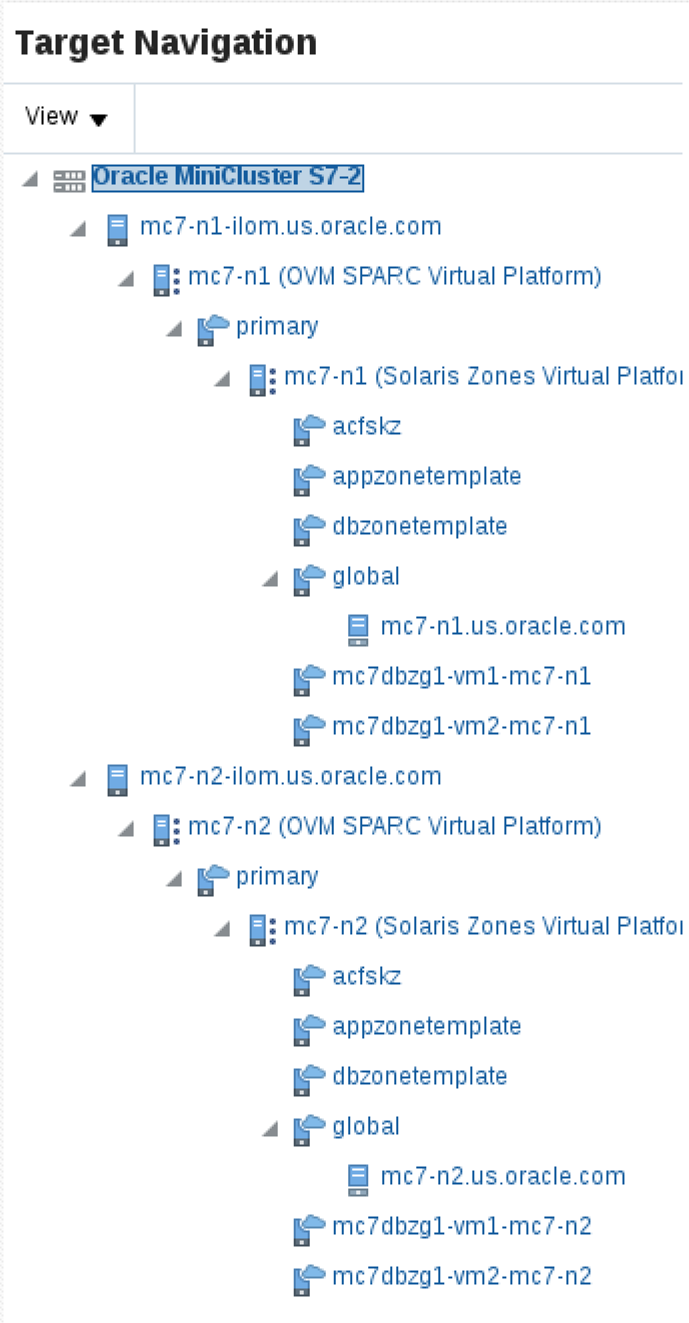

Target Navigation

View ▾

- Oracle MiniCluster S7-2
  - mc7-n1-ilom.us.oracle.com
    - mc7-n1 (OVM SPARC Virtual Platform)
      - primary
        - mc7-n1 (Solaris Zones Virtual Platfo
          - acfskz
          - appzonetemplate
          - dbzonetemplate
          - global
            - mc7-n1.us.oracle.com
          - mc7dbzg1-vm1-mc7-n1
          - mc7dbzg1-vm2-mc7-n1
  - mc7-n2-ilom.us.oracle.com
    - mc7-n2 (OVM SPARC Virtual Platform)
      - primary
        - mc7-n2 (Solaris Zones Virtual Platfo
          - acfskz
          - appzonetemplate
          - dbzonetemplate
          - global
            - mc7-n2.us.oracle.com
          - mc7dbzg1-vm1-mc7-n2
          - mc7dbzg1-vm2-mc7-n2

After the Enterprise Manager Agents are deployed to the Oracle MiniCluster Oracle Solaris global zones, you can view the Virtualization stack for each compute node. See Monitoring Oracle Solaris Zones and Monitoring Oracle VM Server for SPARC for details about the Virtualization stack.

Viewing the Oracle MiniCluster System

The Oracle MiniCluster landing page provides information about the system and aggregated information about the targets it contains. The page is divided into two parts. The top region consists of dashlets which provides you the general overview of the system.

The first dashlet is the Summary dashlet. It displays the summary of the system, that is, total number of most important target types (servers, disk shelves).

The second dashlet displays the number of open incidents grouped by severity. Oracle MiniCluster collects incidents from all associated targets. You can get more details about incidents by clicking on the numbers in the dashlet. The details are displayed in a table with information such as Target, Summary, Last Updated, Acknowledged, and Status. You can click on a target to view more information of the particular target or click on the summary to view the details of the incident.

The third dashlet displays the time of the last configuration change and last reported incident.

The main page consists of two tabs

The first tab displays the physical view of Oracle MiniCluster servers and disk shelves. Switch between servers and disk shelves using the carousel on the left side of the screen. Click on any target to view important information and to open the target's landing page.

The second tab displays information about disk shelves attached to the servers. Select one or more disk shelves to see details about configuration and status of the disks mounted in the disk shelf.

Physical View of Oracle MiniCluster

Physical view of the Oracle MiniCluster system provides an overview of the Oracle MiniCluster system components. You can switch between the MiniCluster system servers and disk shelves using the carousel on the left side. You can click any target or component to get more information about it. If there are any incidents on a target, that particular target is highlighted by a red border to indicate it needs attention. Click on the target to view incidents grouped by severity. You can then click the severity to open the incident manager where you can further interact with the incident.

You can switch between several representations of the physical view.

- Photorealistic view provides a realistic picture of the system, providing detailed graphics of all the components.
- Schematic view is data oriented and displays the most important information such as locator light, status, temperature, host name. Each component in the view has its color based on the type for easy identification of the component.
- Table view is a tabular representation of the physical view, providing a list of displayed components with the most important information.

Figure 25-2 Oracle MiniCluster Rack View

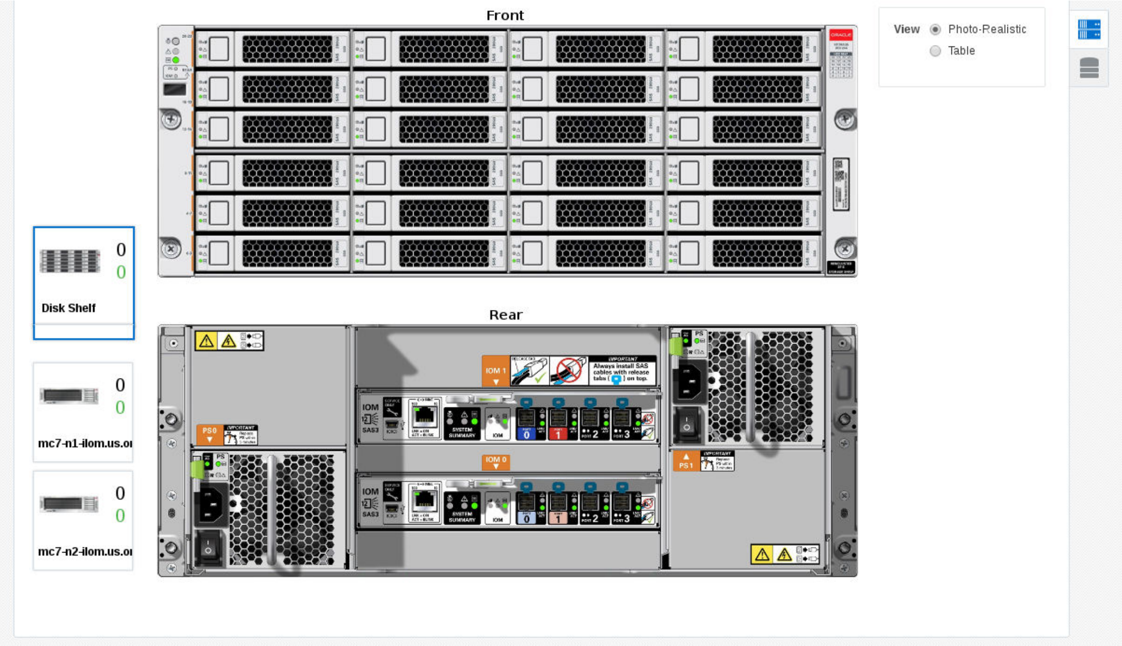

Storage View of Oracle MiniCluster

Storage view of the Oracle MiniCluster system displays overview of the disk shelves connected to the system including configuration and monitoring status information.

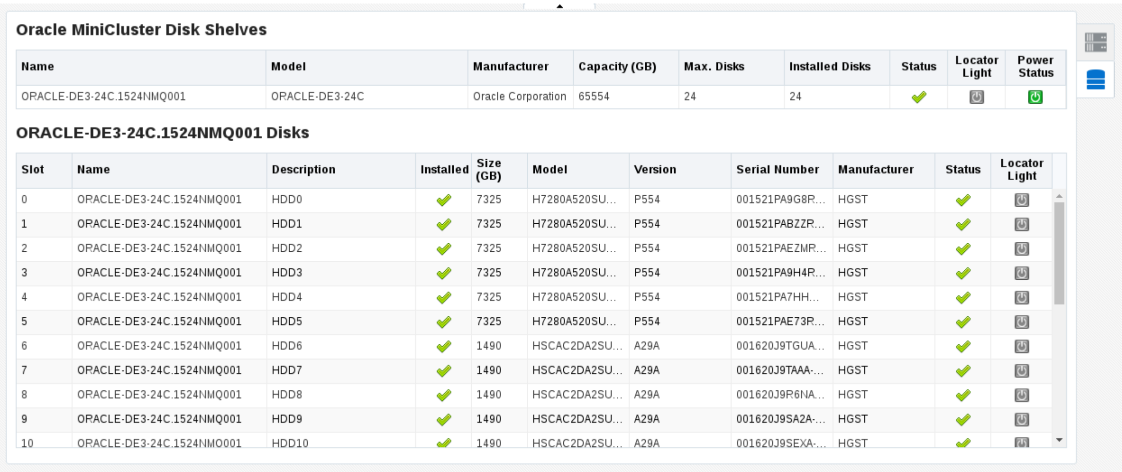

Oracle MiniCluster Disk Shelves

Name	Model	Manufacturer	Capacity (GB)	Max. Disks	Installed Disks	Status	Locator Light	Power Status
ORACLE-DE3-24C.1524NMQ001	ORACLE-DE3-24C	Oracle Corporation	65554	24	24	✓	⏻	⏻

ORACLE-DE3-24C.1524NMQ001 Disks

Slot	Name	Description	Installed	Size (GB)	Model	Version	Serial Number	Manufacturer	Status	Locator Light
0	ORACLE-DE3-24C.1524NMQ001	HDD0	✓	7325	H7280A520SU...	P554	001521PA9G8R...	HGST	✓	⏻
1	ORACLE-DE3-24C.1524NMQ001	HDD1	✓	7325	H7280A520SU...	P554	001521PABZZR...	HGST	✓	⏻
2	ORACLE-DE3-24C.1524NMQ001	HDD2	✓	7325	H7280A520SU...	P554	001521PAEZMR...	HGST	✓	⏻
3	ORACLE-DE3-24C.1524NMQ001	HDD3	✓	7325	H7280A520SU...	P554	001521PA9H4R...	HGST	✓	⏻
4	ORACLE-DE3-24C.1524NMQ001	HDD4	✓	7325	H7280A520SU...	P554	001521PA7HH...	HGST	✓	⏻
5	ORACLE-DE3-24C.1524NMQ001	HDD5	✓	7325	H7280A520SU...	P554	001521PAE73R...	HGST	✓	⏻
6	ORACLE-DE3-24C.1524NMQ001	HDD6	✓	1490	HSCAC2DA2SU...	A29A	001620J9TGUA...	HGST	✓	⏻
7	ORACLE-DE3-24C.1524NMQ001	HDD7	✓	1490	HSCAC2DA2SU...	A29A	001620J9TAAA-...	HGST	✓	⏻
8	ORACLE-DE3-24C.1524NMQ001	HDD8	✓	1490	HSCAC2DA2SU...	A29A	001620J9R6NA...	HGST	✓	⏻
9	ORACLE-DE3-24C.1524NMQ001	HDD9	✓	1490	HSCAC2DA2SU...	A29A	001620J9SA2A-...	HGST	✓	⏻
10	ORACLE-DE3-24C.1524NMQ001	HDD10	✓	1490	HSCAC2DA2SU...	A29A	001620J9SEXA-...	HGST	✓	⏻

Virtualization Management on the Oracle MiniCluster System

For Oracle MiniCluster virtualization management, the Enterprise Manager Agent has to be deployed to each of the virtualization platforms that is planned to be managed. You can install EM Agents into Oracle Solaris Global and Non-Global Zones using Add Host Targets wizard.

Note:

To enable monitoring of Oracle VM for SPARC the non-privileged user used to install and run the Enterprise Manager agent must be granted the `solaris.ldoms.read` and `solaris.ldoms.ldmpower` authorizations and be assigned the LDoms Power Mgmt Observability rights profile.

For example:

```
/usr/sbin/usermod -A solaris.ldoms.read,solaris.ldoms.ldmpower
oracle
```

```
/usr/sbin/usermod -P 'LDoms Power Mgmt Observability' oracle
```

Related Resources for Oracle MiniCluster

See the following for more information:

- Discovering and Adding Host and Non-Host Targets
- Discovering, Promoting, and Adding System Infrastructure Targets
- Using Incident Management
- Enterprise Monitoring

26
Managing Oracle SuperCluster

The following information is included in this chapter:

Getting Started with Oracle SuperCluster

Oracle SuperCluster is an Oracle Engineered System that integrates SPARC compute nodes, an Oracle ZFS Storage Appliance, InfiniBand and Cisco switches, PDUs, and Exadata Storage Servers into a single or multi-rack system.

Oracle SuperCluster is supported in the following configurations:

- SPARC SuperCluster T4-4
- Oracle SuperCluster T5-8
- Oracle SuperCluster M6-32
- Oracle SuperCluster M7-8
- Oracle SuperCluster M8-8

To be able to start target monitoring, it is necessary to discover it. The discovery is made quick and easy with the Guided Discovery Wizard that guides you through the whole process, requests only for necessary information, and helps you solve possible issues in order to successfully complete the discovery.

Actions for Oracle SuperCluster

You can perform the following actions, depending on the requirements.

- Discover Oracle SuperCluster
- View the Oracle SuperCluster system
- Monitor the Oracle SuperCluster system

Target Navigation for Oracle SuperCluster

The target navigation tree helps you to navigate between targets in the Oracle SuperCluster system. The first level under the Oracle SuperCluster consists of Racks and M-series servers. Subsequent levels consist of other hardware targets in the system which are logically grouped by their type. Figure 26-1 displays the hardware structure of the Oracle SuperCluster system. All targets are logically grouped by their type, except for the ZFS Appliance. Disk shelves are not listed in the navigation tree, see Managing Storage for more information.

Figure 26-1 Target Navigation of Oracle SuperCluster

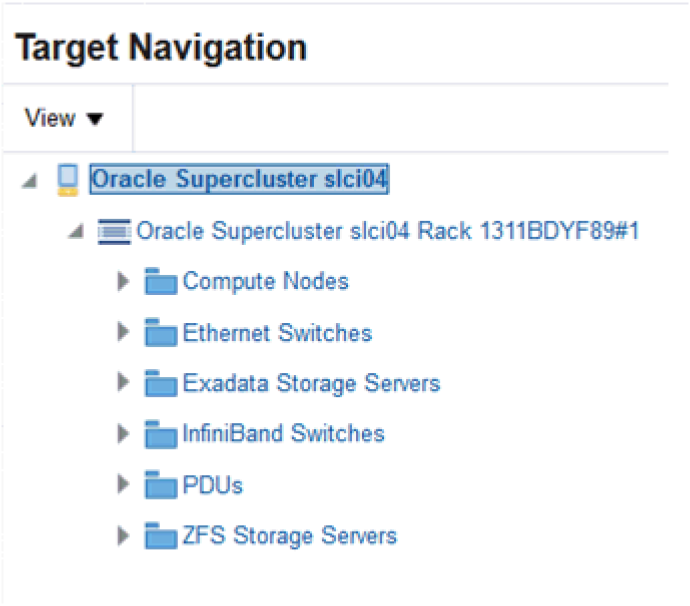

Target Navigation

View ▾

- Oracle Supercluster slci04
  - Oracle Supercluster slci04 Rack 1311BDYF89#1
    - Compute Nodes
    - Ethernet Switches
    - Exadata Storage Servers
    - InfiniBand Switches
    - PDUs
    - ZFS Storage Servers

After the Enterprise Manager Agents are deployed to the Oracle SuperCluster VM Servers for SPARC and Solaris zones, you can view the Virtualization stack for each compute node as seen in Figure 26-2. See Monitoring Oracle Solaris Zones and Monitoring Oracle VM Server for SPARC for details about the Virtualization stack.

Figure 26-2 Oracle SuperCluster Target Navigation with Virtualization Stack

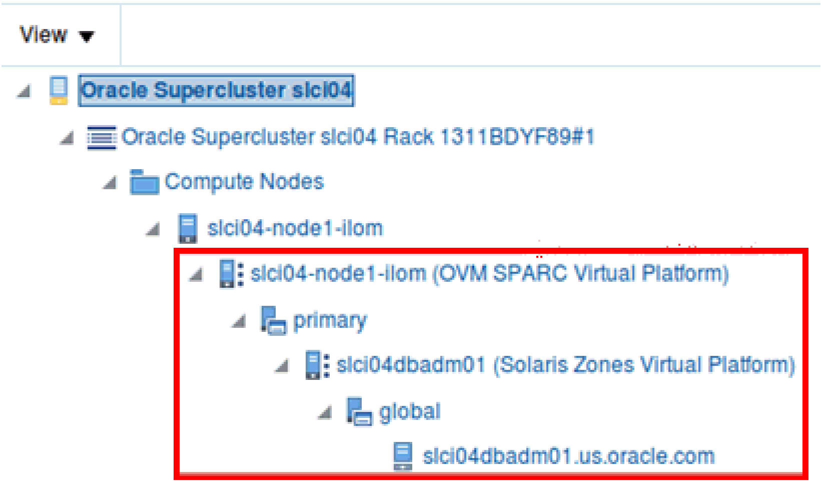

View ▾

- Oracle Supercluster slci04
  - Oracle Supercluster slci04 Rack 1311BDYF89#1
    - Compute Nodes
      - slci04-node1-ilom
        - slci04-node1-ilom (OVM SPARC Virtual Platform)
          - primary
            - slci04dbadm01 (Solaris Zones Virtual Platform)
              - global
                - slci04dbadm01.us.oracle.com

Viewing the Oracle SuperCluster System

The Oracle SuperCluster landing page provides information about the system and aggregated information about the targets it contains. The page is divided into two parts. The top region consists of dashlets which provide you the general overview of the system.

The first dashlet is the Summary dashlet. It displays the summary of the system, that is, total number of most important target types (racks, switches, servers, Exadata cells), amount of available memory, and number CPU cores.

The second dashlet displays the number of open incidents grouped by severity. Oracle SuperCluster collects incidents from all associated targets. You can get more details about incidents by clicking on the numbers in the dashlet. The details are displayed in a table with information such as Target, Summary, Last Updated, Acknowledged, and Status. You can click on a target to view more information of the particular target or click on the summary to view the details of the incident.

The third dashlet displays the time of the last configuration change and last reported incident.

The main page consists of one tab that displays the physical view of Oracle SuperCluster. Click on any target to view important information and to open the target's landing page. You can also switch to a different rack or M-series server by selecting it in the left menu (available only if there is more than one rack or M-series server in the system.)

Physical View of Oracle SuperCluster

Physical view of the Oracle SuperCluster system provides an overview of how the Oracle SuperCluster system is physically structured in the rack.If there is more than one rack or M-series server in the system, you can select a target from the selector available on the left side to view its detailed physical layout. (The selector is hidden if only one rack or server is present). You can click any target or component to get more information about it.If there are any incidents on a target, that particular target is highlighted by a red border to indicate it needs attention. Click on the target to view incidents grouped by severity. You can then click the severity to open the incident manager where you can further interact with the incident.

You can switch between several representations of the physical view.

- Photorealistic view provides a realistic picture of the system, providing detailed graphics of all the components.
- Schematic view is data oriented and displays the most important information such as locator light, status, temperature, host name. Each component in the view has its color based on the type for easy identification of the component.
- Table view is a tabular representation of the physical view, providing a list of displayed components with the most important information.

Figure 26-3 displays the front and rear views of the Oracle SuperCluster system.

Figure 26-3   Oracle SuperCluster Rack View

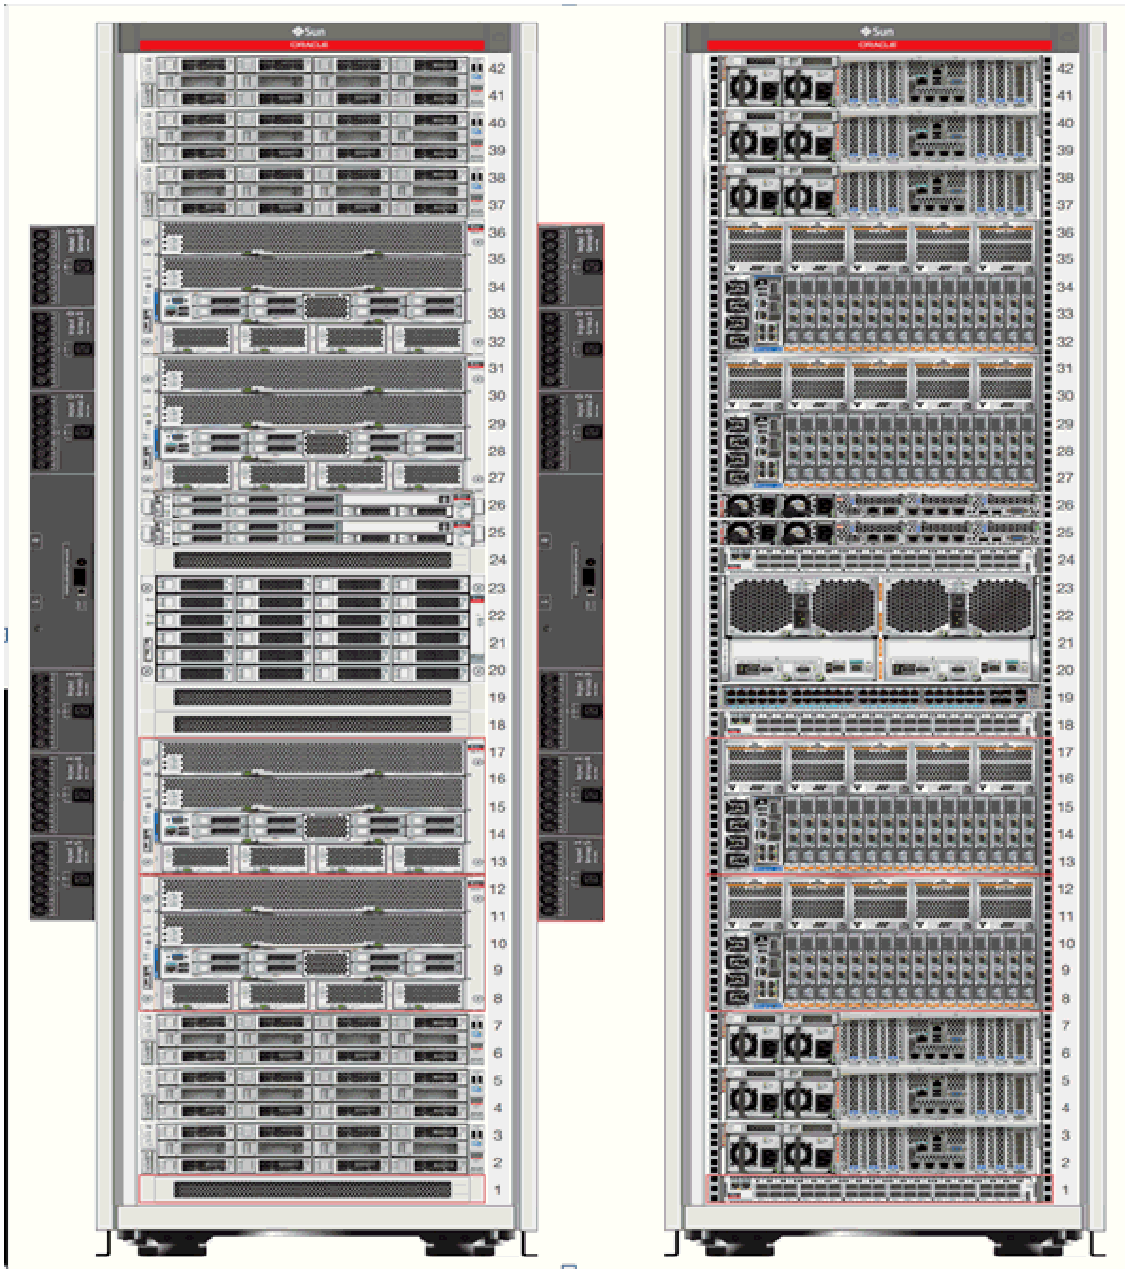

Virtualization Management on the Oracle SuperCluster System

For Oracle Supercluster virtualization management, the Enterprise Manager Agent has to be deployed to each of the virtualization platforms that is planned to be managed.

To monitor Oracle VM Server for SPARC, deploy the EM Agent for Host targets to the Control Domain Operating System. This agent monitors all the Oracle VM Server for SPARC resources as well as the Solaris zones configured on the Control Domain. After the Enterprise manager Agent is deployed on the guest domain operating system, Solaris zones configured in guest domains are monitored. You can install EM Agents into Control Domain and Oracle Solaris Zones using Add Host Targets wizard.

Deleting Oracle SuperCluster System

The Oracle SuperCluster can be removed in Enterprise Manager without SuperCluster members. So If the rack Oracle SuperCluster is deleted, the targets which are part of the Oracle SuperCluster System are not deleted from the Enterprise Manager and can be still accessed in the All Targets view. Or Oracle SuperCluster can be removed including targets which are part of the Oracle SuperCluster System. In this case targets placed in racks which are part of the Oracle SuperCluster System will be removed too.

Perform the following steps to delete Oracle Supercluster:

1. In the rack landing page, select the **Target** menu drop-down list.
2. Select Target Setup, then click **Remove Target**.
3. If you want remove targets which are part of the Oracle SuperCluster System, select “Remove all targets associated to target Oracle Supercluster”. A list of targets which are part of the Oracle SuperCluster system will be shown.
4. Some targets cannot be removed. Such targets have a red cross in the column Removal Details including a reason of why target cannot be removed. Unremovable targets will not be removed when Oracle SuperCluster target is removed.
5. Some targets provide a custom removal flow (like SNMP monitoring unsubscription). Such targets have a yellow exclamation mark in the column Removal Details. If you want to remove target with a custom removal flow, go to the target’s landing page and remove the target from there. If you don't remove targets with custom removal flow, they will be removed along with Oracle SuperCluster without custom removal options available in their removal flows.
6. Click **Yes** to confirm. Oracle SuperCluster is deleted.

Related Resources for Oracle SuperCluster

See the following for more information:

- Discovering and Adding Host and Non-Host Targets
- Discovering, Promoting, and Adding System Infrastructure Targets
- Using Incident Management
- Enterprise Monitoring

27 Monitoring Oracle Operating Systems

This chapter contains information about monitoring Oracle Solaris and Linux operating systems, or hosts.

The following topics are covered:

Get Started with Monitoring Oracle Operating Systems

The operating system is part of the core platform and its metrics are leveraged across other targets in Enterprise Manager. This chapter only covers Oracle Solaris and Linux monitoring.

To view all the hosts monitored by Oracle Enterprise Manager, select **Hosts** in the Targets menu of the Enterprise Manager Cloud Control. The Hosts section of the user interface displays Oracle Solaris and Linux operating system configuration, resource and process metrics.

Performance and configuration metrics provide you with a unified view of an operating system's CPU, memory, and process resource usage, enabling you to manage and optimize resources.

Monitoring is activated when you discover and manage the operating system. A series of escalating status levels notifies you when something is not operating as expected. The lowest level incident status is warning, then critical, and the highest level is a fatal incident.

The information on an operating system is designed to help you maximize performance and utilization. The Host Summary provides a high-level overview of the host. You can drill down to view CPU and resource usage and deeper to view the processes that are running on the host, and real time display of current top processes on a host.

The following features are available for operating systems:

- OS details: Displays operating system resource and processes information, Oracle Solaris Zones and boot environments.
- Incidents: Notification of incidents with links to view details. A series of monitoring rules and parameters monitor your managed assets. Events and incidents are raised for resources that are not performing as expected.
- Performance: Analytics: Provides a detailed view into operating system performance and resource usage.
- Configuration: Indicates configuration changes to the operating system.

Location of Oracle Operating System Information in the UI

Oracle Solaris and Linux operating system information is located in the following locations in the user interface.

Table 27-1 Location of Operating System Information in the UI

Object	Location
All Oracle Solaris and Linux operating systems	Select **Hosts** in the Targets selector. If the operating system parameter does not appear in the table, click **View**, click **Columns**, then select the Operating System parameter. To view the version, select **Target Version**.
A specific Oracle Solaris or Linux operating system	Select **Hosts** in the Targets selector. Click the host to display the Home page.
System resources for a specific Oracle Solaris or Linux operating system	Select **Hosts** in the Targets selector. Click the host to display the Summary page. To view more details, click one of the following tabs on the right side of the Host page: **CPU**, **Host Memory**, **Storage**, **Network Connectivity**, **Host Processes**, or **Host Services**.
Monitoring metrics and details, Program Resource Utilization, Metric and Collection Settings, Metric Collection Errors, Status History, Incident Manager, Alert History, and Blackouts and Brownouts.	Select **Hosts** in the Targets selector. Click the host to display the Home page. Click the **Host** drop down menu, then click **Monitoring**.

Features of Operating Systems

You can perform the following actions:

- View the configuration and status of hosts
- View host details
- View the platform, physical or virtual, on which the operating system is deployed
- View CPU and memory resource utilization and the top process utilization
- Diagnose problems using incidents and performance metrics

About the Dashboard for all Hosts

The Hosts dashboard is a sortable table that contains details about the managed hosts, including incidents, the type of operating system, the operating system version, CPU and memory utilization. More than 40 host parameters are available. A few parameters are selected by default, others are hidden. You can change the parameters that appear in the dashboard and the order in which they appear.

Viewing the Dashboard of all Hosts

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

   The Hosts page displays all managed hosts, including operating system details. You can sort most of the columns in either ascending or descending order.

2. Click **View**, then **Columns** to view or edit the parameters to display in the dashboard.
3. To add or remove parameters, select or deselect the parameter in the View menu. For example, you might want to display the Incidents, Operating System, and Target Version.

Figure 27-1 Hosts Dashboard

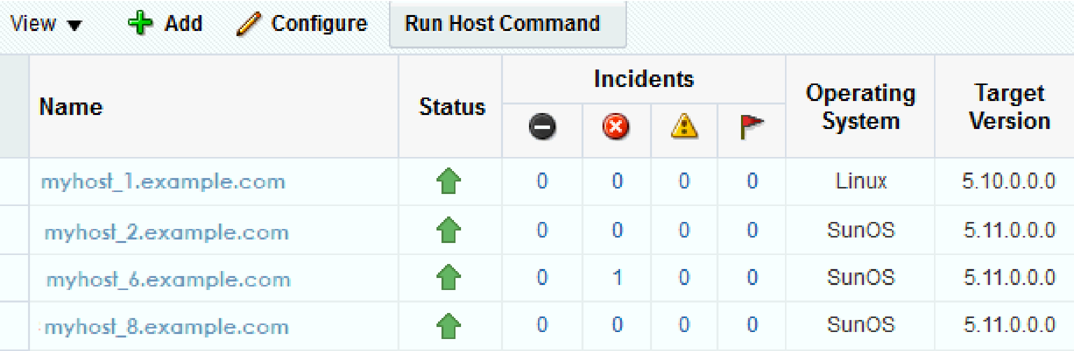

Name	Status	Incidents				Operating System	Target Version
		⊖	⊗	⚠	🚩		
myhost_1.example.com	↑	0	0	0	0	Linux	5.10.0.0.0
myhost_2.example.com	↑	0	0	0	0	SunOS	5.11.0.0.0
myhost_6.example.com	↑	0	1	0	0	SunOS	5.11.0.0.0
myhost_8.example.com	↑	0	0	0	0	SunOS	5.11.0.0.0

> **Note:**
>
> The Oracle Solaris operating system might appear as SunOS in the user interface and the 11.0.0.0 release might appear as Target Version 5.11.0.0.0.

4. (Optional) You can reorder the columns that appear on the dashboard. Click **View**, then **Reorder Columns**. Select a column and use the arrow button to move the selected item to a different position. Click **OK** to save the change.

How to Get Information About a Specific Host

The Home page displays details and metrics for the selected operating system:

- Dashlets: A series of dashlets at the top of the page contains summary information and might be associated with a more detailed information that is in a tab.
- Tabs: A series of tabs on the right side of the page link to more detailed information.

- Host menu: Links to detailed information about the host, including performance and configuration metrics, metric collection settings, status history, incidents, alert history, and blackout and brownouts.

Viewing the Host Target Home Page

You can view a summary of managed operating system targets.

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

   The Hosts page appears with a list of all managed hosts. You can sort the list.

2. Click the host name from the list of managed hosts to display the Home page for that host.

   The Home page contains a series of dashlets that provide useful information about the host and a summary of the host details and utilization.

About Dashlets for Hosts

The top of the host page contains a series of dashlets that provide a quick view of top statistics. Click the small button below the row of dashlets to toggle to the next series of dashlets.

Figure 27-2 Dashlets

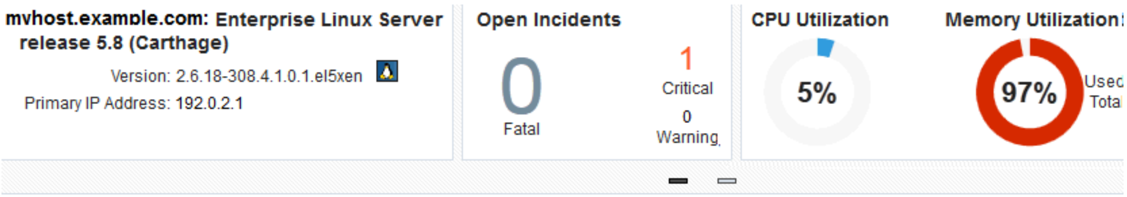

mvhost.example.com: Enterprise Linux Server release 5.8 (Carthage)

Version: 2.6.18-308.4.1.0.1.el5xen

Primary IP Address: 192.0.2.1

The following dashlets are available:

- Host details: Provides a short summary, including the host name, type of operating system, version and release, primary IP address, and the length of up time for the operating system.
- Open Incidents: Shows the number of Fatal, Critical, and Warning incidents. Mouse over the number to see a small snapshot of the incident. Click a number or a summary link to navigate to the Incident Manager console for incident details.
- CPU Utilization and Memory Utilization: Shows a graphical and number percentage for CPU and Memory utilization.
- OS Services State: A pie chart shows the percentage of services that are running, stopped, no state, and other states.
- Configuration changes: Shows the data and time of the last configuration change and last reported incident.

About Tabs for Hosts

The following tabs appear on the host home page and are represented with icons on its right side, click a tab to display more information:

- **Summary**: Displays host details, swap, CPU, CPU threads, memory utilization, filesystem distributions, network usage.

- **CPU**: Displays performance and resource metrics. The following performance and process utilization graphs are available: CPU Utilization, System Loads, CPU Threads Utilization (including Processor Group Threads Utilization for Oracle Solaris, and CPU Frequency State for Oracle Solaris on SPARC.)
- **Host Memory**: View of an operating system's memory utilization, IPCS and swap details.
- **Storage**: Links to detailed storage information, including the disks, filesystem, volume group and SAN configuration.
- **Network Connectivity**: Shows the network interface and subnet.
- **Boot Environments**: Displays the available alternate boot environments and boot environment snapshots for Oracle Solaris operating systems.
- **Host Processes**: View of an operating system's top CPU, memory utilization, and process resource usage.
- **Host Services**: View the services that are managed by the operating system and the state of the services.

About the Host Menu

The Host menu contains links to detailed information about the host, including monitoring and configuration information.

The following information is available in the Host Monitoring menu:

- CPU Details
- Memory Details
- Disk Details
- Program Resource Utilization
- All Metrics
- Metric and Collection Settings
- Metric Collection Errors
- Status History
- Incident Manager
- Alert History
- Blackouts

Viewing the Host Monitoring Menu

1. Click **Hosts** from the Targets page.
2. Click the target name to open the home page.
3. Click **Host** in the upper left corner of the page, then click **Monitoring**.
4. Click an option to view greater detail. For example, click **All Metrics** to view all metrics collected.

About Open Incidents

You can view all open incidents on the Hosts dashboard, or you can view open incidents for a specific host on the open incidents dashlet of the Host home page.

The Hosts Dashboard lists all managed hosts and displays open incidents. You can sort the columns, or click a number in a column to navigate directly to the Incident Manager page for details. Alternatively, you can click the host name to view more about the host and the incident from the host's home page. The Open Incident dashlet on the Host home page displays the number of Fatal, Critical, and Warning incidents.

Viewing Open Incidents

To view an open incident from the host's Home page:

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

   The Hosts page appears with a list of all managed hosts. You can sort the list.
2. Click the host name from the list of managed hosts to display the Home page for that host.
3. Click the number to view a summary of the open incidents.
4. Click the summary text to navigate to the Incident Manager.

Figure 27-3 Open Incidents Dashlet for a Host

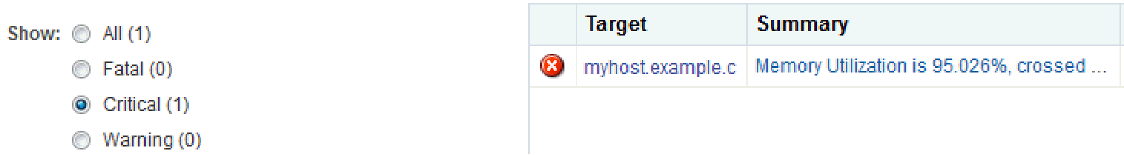

Show: ○ All (1)
○ Fatal (0)
● Critical (1)
○ Warning (0)

	Target	Summary
✖	myhost.example.c	Memory Utilization is 95.026%, crossed ...

The Incident Manager provides incident details and the events that led to the incident. Events, Notifications, My Oracle Support Knowledge tabs are located on the individual bookmarks of an Incident Manager page (horizontally from left to right). If you are online, a link will take you to My Oracle Support.

You can acknowledge the incident, add comments, or manage the incident from the Incident Manager page.

Identifying Changes in an OS Configuration

When an administrator changes a host configuration, it can be helpful to know when the configuration was last changed. This information appears in the configuration dashlet on the Home page. Detailed configuration information is available, including the ability to compare.

1. Click **Hosts** from the Targets page.
2. Click the target name to open the home page.
3. Click **Host** in the upper left corner of the page. Click **Configuration**.
4. Click the option to view **Last Collected**, **Comparison & Drift Management**, **Compare**, **Search**, **History**, **Save**, **Saved**, or **Topology**.

Overview of Performance and Resource Metrics

Performance and resource metrics provide details on the kernel configuration and performance, helping you to identify issues. The CPU Load chart and Free Memory chart enable you to easily view the status. High level CPU and Memory usage are available in the Summary tab with more details in the CPU and Host Memory tabs.

CPU data for the following metrics is collected every 15 minutes and appears in the CPU tab:

- About CPU Utilization
- About CPU Threads Utilization
- About Processor Group Utilization for Oracle Solaris 11

The following options are available for you to view kernel information:

- Shared memory
- CPU I/O wait and buffer cache read/write details
- Physical I/O read/write, disk and disk block read/writes
- Run queue length and paging activity
- Tunable kernel parameters

Resource metrics provides details on the operating system, the available resources, and the load on the operating system or zone.

The following details are available:

- Memory, total and available
- Swap, total configured and available
- CPU details, including the vendor name, number, frequency, revision and mask
- Number of cores and threads per CPU
- Bar chart showing the utilization percentage per CPU thread
- The amount of time spent by all CPUs in different frequencies for Oracle Solaris
- CPU and memory usage over time

About CPU Utilization

The CPU Utilization metric displays the percentage utilization of a CPU over time for Oracle Solaris and Linux targets. An abnormally high value indicates that the system is under heavy load. If the value is consistently high, consider reducing the load on the system.

CPU data is collected every 15 minutes. The default display is a graphical representation of the Run Queue Length with a red line for the 1 minute average, a green line for the 5 minute average, and a yellow line for the 15 minute average. To display the information in table format, click **Table View**.

Viewing CPU Metrics

To view the kernel and performance metrics for an operating system:

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

The Hosts page appears with a list of all managed hosts. You can sort the list.

2. Click the host name from the list of managed hosts.
3. Click the **CPU** tab to view the metrics and charts.

   The information appears in a graphical format. Click **Table View** to change the format. You can adjust the time frame to display historical data for the last two (2) hours, four (4) hours, 10 hours, one day, or one week. By default, CPU and System Load appear. Deselect to remove the information from the graphs.

About CPU Threads Utilization

The CPU Threads Utilization metric collects CPU thread diagnostics for Oracle Solaris and Linux targets, useful for analysis of multi-threaded CPUs. You can view the efficiency and the metrics for each CPU thread.

To help you to gauge the efficiency, the following charts are available:

- Bar chart showing the number of CPU threads at each frequency
- Historical charts showing the percentage of time spent at different frequencies

About Processor Group Utilization for Oracle Solaris 11

In addition to the CPU Utilization, you can view the following processor group utilization details for Oracle Solaris 11 operating systems:

- List of processor groups and the number of threads per group
- Bar chart of CPU utilization per processor group
- Type of group, such as integer pipeline

About Host Memory

The Host Memory page provides you with a unified view of an operating system's memory utilization.

The page displays the following information:

- Memory Utilization: Displays overall memory utilization. The default view is a graphical representation of the percentage of memory used over time. You can change the view to represent the MB of memory used over time. If you prefer, you can view the overall memory utilization in a table instead of a chart.
- Virtual Memory: Displays overall virtual memory utilization. The default view is a graphical representation of the percentage of swap space used over time. The default view is a graphical representation of the MB of memory used over time. You can change the view to represent the percentage of swap space used over time. If you prefer, you can view the overall virtual memory utilization in a table instead of a chart.
- Page Activity: Displays a paging statistics activity in a color-coded graph format. The chart shows the following page activity: Address Translation Page Faults appear as a blue line, Pages Paged-in appear as a green line, Pages Paged-out appear as an orange line, Active Pages appear as a blue line, and the Pages Scanned by Page Stealing Daemons appear as a violet line. All activity is on a per second basis for a single day. If you prefer, you can view the data in a table instead of a chart.
- Memory Details: A pie chart shows the memory details. The chart displays the Free Memory, Used Memory, and Other Shared Memory as a percentage of the entire memory.

- Swap File: Displays the swap file and amount of space used. A graphical representation shows at a glance the amount of used and free swap space.
- ZFS ARC cache usage: Displays ZFS ARC (Adaptive Replacement Cache) usage for Oracle Solaris operating systems.

Viewing Host Memory Utilization

To view host memory charts:

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

   The Hosts page appears with a list of all managed hosts. You can sort the list.
2. Click the host name from the list of managed hosts.
3. Click the **Host Memory** tab to view the metrics and charts.

   In some cases, you can click **Table View** next to the chart to view the information in a table format.
4. To change the y-axis of the Memory Utilization and Virtual Memory Utilization charts to display as a percentage instead of in MB, select **By Percentage**. Select the time frame from the Time Range menu to change the default from two hours.

Viewing Memory and Swap File Details

To view details about memory and swap file:

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

   The Hosts page appears with a list of all managed hosts. You can sort the list.
2. Click the host name from the list of managed hosts.
3. Click the **Host Memory** tab, then click **IPCS & Swap Details** in the center pane.

Viewing Memory Details for a Host

1. Click **Hosts** from the Targets menu.
2. Click the target name to the home page.
3. Click **Host** in the upper left corner of the page. Click **Monitoring**, then click **Memory Details**.

Viewing Host Storage

Host Storage contains links to detailed storage information, including the disks, filesystems, volume group and SAN configuration. You can view all storage, or filter by volume name, or select to view only local or remote storage.

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

   The Hosts page appears with a list of all managed hosts. You can sort the list.
2. Click the host name from the list of managed hosts.
3. Click the **Host Storage** tab.
4. Select **Local** or **Remote** to view a subset of storage. Enter a name in the Volume Name field to filter your results.

5. Click the icon for the storage details you want to view. For example, for Linux the options are Disks, Filesystems, Linux LVM Volume Group(s), and SAN Configuration.

Viewing Network Connectivity

Network Connectivity shows the network interface and subnet to associate with the host.

You can view different layers of the network:

- Network interface: View the network state, subnet and flag details for each network interface
- Data link: View the data link state, and physical address, and the type of media, such as ethernet, for each data link.

1. From the **All Targets** or **Hosts**menu, select **Hosts**.

   The Hosts page appears with a list of all managed hosts. You can sort the list.
2. Click the host name from the list of managed hosts.
3. Click the **Network Connectivity** tab.
4. Click the icon to display the **Interfaces** or **Data links** layer.

About Boot Environments

Oracle Solaris 11 Boot Environments use the `beadm` utility and ZFS file systems to create and manage boot environments. The Oracle Solaris 11 software automatically creates boot environments.

A boot environment is an instance of a bootable Oracle Solaris image plus additional software packages that are installed onto the image, and the set of all file systems and devices (disk slices and mount points) that are required to operate an Oracle Solaris OS instance. A system can have only one active boot environment, which is the booted environment. An alternate boot environment is an inactive environment that is not currently booted. A system can have many inactive boot environments.

A dual boot environment is often used to manage updates because it can significantly reduce the service outage time that is usually associated with patching. Maintaining multiple boot environments also enables quick and easy rollback to a version before the patches were applied, if needed.

The Boot Environment tab of the Oracle Solaris operating system page displays Oracle Solaris boot environment and file system details, including all available boot environments, the size, and the synchronization date. For a selected boot environment, you can view snapshot details, file system details, and any associated zone boot environments. This tab is only available for Oracle Solaris operating systems

The Boot Environment tab of the Oracle Solaris operating system page displays Oracle Solaris boot environment and file system details, including all available boot environments, the size, and the date the environment was created or synchronized. The Boot Environment tab is only available for Oracle Solaris operating systems

Viewing Oracle Solaris Boot Environments

1. From the **Targets** menu, select **Hosts**.
2. Select an Oracle Solaris operating system from the list of managed hosts.

3. Click the **Boot Environments** tab to view the boot environment snapshot and file system details. The file system details are at the bottom of the page, after the boot environments.
4. Expand the operating system to display snapshots of the boot environments.

Viewing Running Host Processes

The Host Processes page provides you with a unified view of an operating system's top processes, including the CPU and memory utilization of each process.

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

   The Hosts page appears with a list of all managed hosts. You can sort the list.
2. Click the host name from the list of managed hosts.
3. Click the **Host Processes** tab to view the processes.

Viewing Managed Host Services

With Host Services, you can see which services are managed by the operating system and the state of the services. This is useful when you want to quickly identify which services are in need of attention and the state of services that are important to you.

Host Services monitors and displays the services running on a host. You can view the current state of a service. However, you cannot create, delete, or modify the properties of a service. Fault Management Resource Identifier (FMRI) identifies each service on the system.

The following are the service states:

- Running: The service is running.
- Stopped. The service is either disabled or offline and the service is not running.

The Host Services page displays the service state, the number of spawned process identifiers, and the spawned process identifiers (PIDs).

1. Click **Hosts** from the **Targets** menu.
2. Click the host name from the list of managed hosts to display the Summary page for that host.
3. Click the **Second dashlet series** button below the dashlets to view a chart summarizing the current status of services.
4. Click the **Host Services** tab on the right side of the user interface to view details of the services.
5. The default view is to display stopped services. Click the radio button to change the view. For Oracle Solaris, the options are: **Offline**, **Online**, or **All**. For Linux, the options are: **Stopped**, **Running**, or **All**.
6. Click a number to view the spawned process identifiers (PIDs.)

Working with Host Metrics

More detailed host metrics are available from the Hosts menu, including the following:

- Viewing CPU, Memory, and Disk Details for a Host
- Viewing a Host's Program Resource Utilization

Viewing CPU, Memory, and Disk Details for a Host

1. Click **Hosts** from the Targets menu.
2. Click the target name to the home page.
3. Click **Host** in the upper left corner of the page. Click **Monitoring**, then click **CPU Details**, **Memory Details**, or **Disk Details**.

Viewing a Host's Program Resource Utilization

1. Click **Hosts** from the Targets menu.
2. Click the target name to open the home page.
3. Click **Host** in the upper left corner of the page. Click **Monitoring**, then click **Program Resource Utilization**.

Viewing All Metrics

1. Click **Hosts** from the Targets menu.
2. Click the target name to open the home page.
3. Click **Host** in the upper left corner of the page. Click **Monitoring**, then click **All Metrics**.
4. (Optional) To view by category instead of by metric, click **View**, then click **By Metric Category**.
5. Click a metric to view details, collection schedule, upload interval and other details.

Managing Metrics and Incident Notifications for Hosts

You can perform the following tasks to manage monitoring and incident notification:

- Viewing Host Metric Collection Error
- Editing Metric and Collection Settings for Hosts

Viewing Host Metric Collection Error

Metric collection errors are usually caused by installation or configuration issues.

1. Click **Hosts** from the Targets menu.
2. Click the target name to open the home page.
3. Click **Host** in the upper left corner of the page. Click **Monitoring**, then click **Metric Collection Errors**.

Editing Metric and Collection Settings for Hosts

The Metrics tab contains displays all of the monitored attributes. The default view is metrics with thresholds. For these types of monitored attributes, you can modify the comparison operator, the threshold limits, the corrective action, and the collection schedule.

1. Click **Hosts** from the Targets menu.
2. Click the target name to open the home page.

3. Click **Host** in the upper left corner of the page. Click **Monitoring**, then click **Metric and Collection Settings**.
4. Modify threshold limits or collection schedule. When a threshold field is empty, the alert is disabled for that metric.
5. Click the **Edit** icon for advanced settings.

   Click the **Other Collected Items** tab to view non-threshold monitored attributes. You can modify the collection period for these attributes, or disable monitoring.

About Host Compliance

Host compliance provides you information on the compliance frameworks, standards, and the targets that are associated with the compliance standard selected in the Compliance Standard Library.

Cloud Control displays the evaluation results and level of compliance of a target against a compliance framework. The Compliance Frameworks evaluation results provide an overview of the state of the framework, the level of compliance (Critical, Warning, or Compliant.) and the criticality of any violations (Critical, Warning, or Minor Warning.) You can view the average score, as a percentage, and the Author.

When Compliance Framework errors are detected, the following information is available:

- Root Compliance Standard
- Root Compliance Standard State
- Parent Compliance Standard
- Rule
- Root Target Information
- Target Information
- Error Date
- Error Message

Viewing Compliance Frameworks

The Compliance Framework tab displays the evaluation results and errors, if any.

1. Click **Hosts** from the Targets menu.
2. Click the target name to open the home page.
3. Click **Host** in the upper left corner of the page. Click **Compliance**, then click **Results**.
4. Click **Compliance Frameworks** to display the Evaluation Results tab.
5. Click the **Errors** tab to see if there are any Compliance Framework errors.

Viewing Compliance Standards

You can search for a specific compliance standard for a host and view the evaluation results and errors.

1. Click **Hosts** from the Targets menu.
2. Click the target name to open the home page.

3. Click **Host** in the upper left corner of the page. Click **Compliance**, then click **Results**.
4. Click **Compliance Standards** to display the Evaluation Results tab.
5. Click the **Errors** tab to see if there are any Compliance Framework errors.

Viewing Target Compliance

The Target Compliance table lists the targets that are associated with the compliance standard selected in the Compliance Standard Library.

1. Click **Hosts** from the Targets menu.
2. Click the target name to open the home page.
3. Click **Host** in the upper left corner of the page. Click **Compliance**, then click **Results**.
4. Click **Target Compliance** to display the targets that are in compliance with the standards.

Related Resources for Operating Systems

See the following for more information:

- Discovering and Adding Host and Non-Host Targets for the concepts around discovering and adding targets.
- Discovering, Promoting, and Adding System Infrastructure Targets for how to discover operating systems and other system infrastructure targets.
- Managing Storage for more details about viewing a host's storage filesystems and viewing storage information.
- View Network Details of a Host Target for more information about the types of information that are available.
- Using Incident Management for details on managing incidents.

Go to `http://docs.oracle.com/en/operating-systems/` for the following:

- Oracle Solaris documentation
- Oracle Linux documentation

28
Monitoring Oracle Solaris Zones

This chapter contains information about monitoring Oracle Solaris Zones.

The following topics are covered:

Get Started with Monitoring Oracle Solaris Zones

Oracle Solaris Zones, also known as Oracle Solaris Containers, are used to virtualize operating systems and provide an isolated and secure environment for running software applications. A zone is a virtualized operating system environment created within a single instance of the Oracle Solaris operating system.

Think of a zone as a box with flexible, software-defined walls. One or more applications can run in this box without interacting with the rest of the system. Because zones isolate software applications or services, applications that are running in the same instance of the Oracle Solaris OS are managed independently of each other. For example, you can run different versions of the same application in separate zones. Zones require a machine that is running at least an Oracle Solaris 10 operating system. Solaris 11 global zone and Solaris 10 update 11 global zones are supported.

The **global zone** is the default operating system and has control over all of the processes and has system-wide administrative control. The global zone oversees the CPU, memory, and network resource allocation of all of the non-global zones. A global zone always exists, even when no other zones are configured.

Non-global zones, or simply zones, are configured inside the global zone. Zones are isolated from the physical hardware by the virtual platform layer. A zone cannot detect the existence of other zones.

Kernel zones are zones that implement virtualization from within the global zone's operating system kernel. Each kernel zone has a separate kernel from the global zone, its own file systems and user space. Configuration of each zone (including the global zone) puts limits on the CPU, memory and I/O resources available to the zone. Kernel zones are supported beginning with Solaris 11.2. A kernel zone enables you to deploy a non-global zone with its own operating system kernel instance. The non-global zone has a different kernel version to the global zone. You can create one level of non-kernel zones inside kernel zones.

The following types of non-global zones are available with Oracle Solaris:

- Native zone: A separate Solaris 10 or Solaris 11 instance with the same version of Solaris as the global zone. You cannot create nested zones.
- Solaris 10 branded zone: An independent Solaris 10 instance running inside a Solaris 11 global zone, providing a migration path for existing Solaris 10 deployments. Nested non-global zones are not supported.
- Kernel Zone: Runs a separate kernel version inside the non-global zone. A kernel zone is fully independent operating system instance, which enables you to create nested (non-kernel) zones within the kernel zone. Kernel zones are available beginning with the Oracle Solaris 11.2 release.

Zones are represented by an icon in the user interface. Different types of zones, such as global zone, kernel zone, and non-global zones, have different icons.

You can monitor the following Solaris 10 and Solaris 11 global zones and their non-global zones through the Enterprise Manager user interface.

- Solaris 10 global zones running native Solaris 10 non-global zones
- Solaris 11 global zones running branded Solaris 10 non-global zones
- Solaris 11 global zones running native Solaris 11 non-global zones
- Solaris 11 global zones running Solaris 11 kernel zones

Enterprise Manager supports the following types of virtualization:

- Oracle Solaris Zones: operating system virtualization
- Oracle VM Server for SPARC: hardware virtualization on a SPARC platform

You can view zones within any type of logical domain on a SPARC platform.

The hypervisor is responsible for managing one or more non-global zones. A non-global zone is represented as its operating system instance deployed on a virtual server which is given a subset of the CPU, memory and I/O resources which are available from the physical server, and/or some virtual resources (such as virtual disks or networks) which are backed by configured resources from the global zone. The global zone always exists and is the controlling zone for the non-global zones.

Location of Oracle Solaris Zone Information in the UI

You can select any target that is a child or a parent of the zone (virtual platform, host, or server) and the zone will appear in the Navigation pane of the target.

The Virtualization Platform is the container on which zones are running. A global zone and all associated zones appear in a Zone Virtualization Platform page. A global zone is represented by a virtual server as well. Each non-global zone (virtual server) and kernel zone has its own Virtual Server page that displays details specific to that zone.

The following options are available on the All Targets page:

- Click **Virtualization Platform** to see the list of virtual platforms. You can click any virtual platform in this list to open the virtual platform's home page.
- Click **Virtual Server** to see the list of virtual servers.You can click any virtual server in this list to open the virtual server's home page.

Table 28-1 Location of Zone Information in the UI

Object	Location
Virtualization Platform	A Virtual Platform target home page is accessible from the All Targets page, or from any of its parent or child targets through the Target Navigation pane. Click the **All Targets** selector at the top of the page. In the Refine Search section, select **Virtualization Platform**. The target can be a host, a server, or a zone target running on this Virtual Platform. A Target Navigation pane is located in the top left corner of a parent or child home page. Click the Target Navigation pane to expand it and see the Virtual Platform in the Navigation pane.
Zone-specific (virtual server) page	A zone virtual server target home page is accessible from the Solaris Zone Virtualization Platform page or the All Targets page. Click the **All Targets** selector at the top of the page. In the Refine Search section, select **Virtual Server**. If you are on the Solaris Zone Virtualization Platform page, click a zone to navigate to the virtual server page for that zone.

Actions for Zones

The virtualization platform is automatically promoted when you discover Solaris 11 or Solaris 10 update 11 host operating systems.

You can perform the following actions:

- View the configuration and status of zones
- View CPU and memory resource utilization and the distribution of the CPU and memory consumers
- Diagnose problems using incidents and performance metrics
- Disable and enable monitoring notifications and monitoring

Target Navigation for Zones

The target navigation tree helps you to navigate between targets that are in a Solaris Zones Virtual Platform. When all resources are discovered and monitored through agent deployment on the global zone, you can navigate from the server down to the zones.

The top level, or node, of the target navigation tree is the physical server and the second node is the Oracle Solaris Zones Virtual Platform that is hosting the global zone. The global zone and all other types of zones appear under the Oracle Solaris Zones Virtual Platform. Expand the nodes to drill down the target hierarchy or click a specific target to open that target's landing page.

Figure 28-1 is an example of a target navigation tree that shows a global zone and two kernel zones. The first node is the physical server (smx42-1). The second node is the Solaris Zones Virtual Platform that is hosting the global zone and 2 kernel zones (kernelzone1 and kernelzone2.) The global zone is running the operating system (smx42-1-n17 host.) When the agent was deployed on kernelzone1, it triggered the discovery and promotion of the Solaris Zones Virtual Platform (smvt-175-10.) The Solaris Zones Virtual Platform in kernelzone1 is running a global zone and operating system (smvt-175-10 host) and 2 local zones (zone 1 and zone 2).

Figure 28-1 Target Navigation for Zones

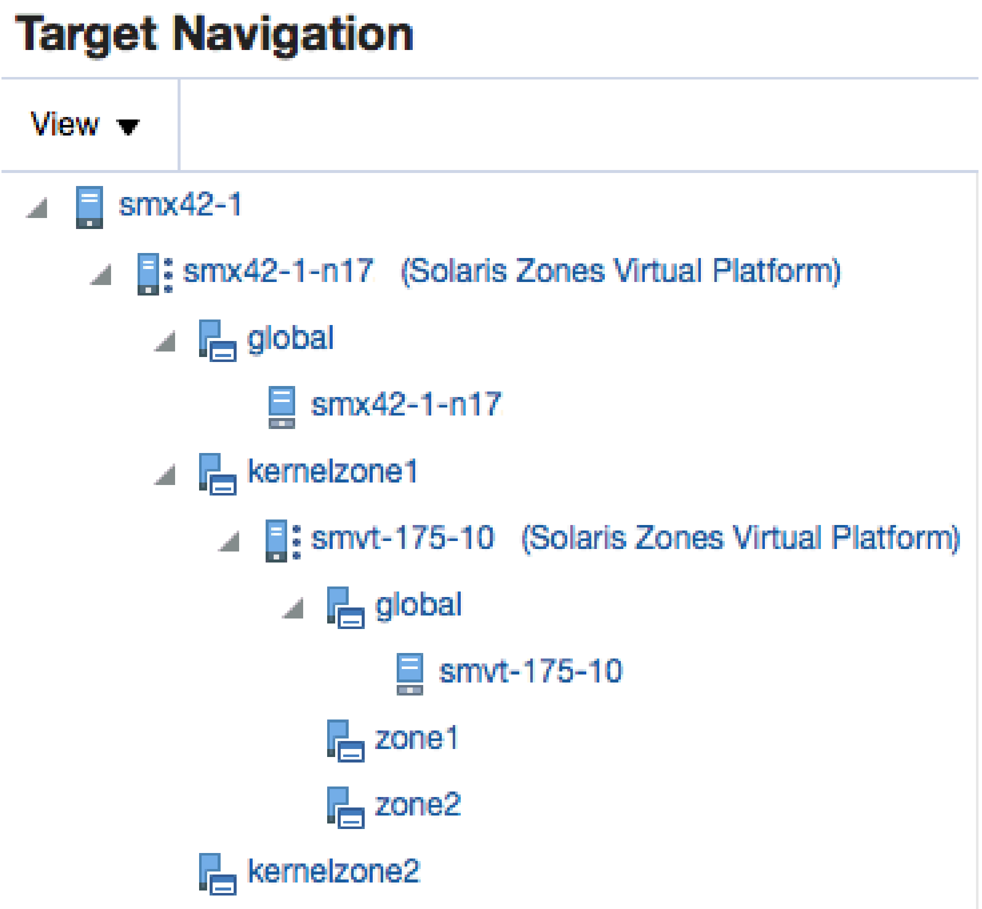

How to Get Information About a Zone

You can view information about the zone from a platform perspective, or for a specific zone. Zones appear in a Solaris Zone Virtualization Platform page and each global zone, kernel zone, and non-global zone has its own virtual server page.

The Solaris Zone Virtualization Platform page provides general information about the zone, such as open incidents, CPU usage, memory resources, number of running and configured zones, and a list of all zones for the global zone and their states.

The Virtualization Platform is the container on which zones run. Go to the Virtualization Platform page to see the list of zones (virtual servers) running on a virtual platform. Click a zone to see the zone's home page.

The Summary displays details and metrics for the selected zone:

- Dashlets: A series of dashlets at the top of the page contains summary information and might be associated with a more detailed information that is in a tab.

- Tabs: A series of tabs on the right side of the page link to more detailed information, including CPU and memory performance.

The top of the page contains a series of dashlets that provide a quick view of top statistics. Click the small button below the row of dashlets to toggle to the next series of dashlets. Click the icon below the toggle icons to minimize the dashlets.

The Virtualization Platform dashlets provide the following types of information:

- Platform status
- Time that the zone platform has been up and running
- Open incidents
- Guest configurations
- Distribution of the CPU and memory consumers per zone
- Date and time of last configuration changes and incidents

See Working with Zone Platform Metrics for more details.

The Virtual Server dashlets provide the following types of information about a specific zone:

- Zone status
- CPU and memory utilization
- Guest configurations
- Date and time of last configuration changes and incidents

See Working with Zone-Specific Metrics for more details.

Working with Zone Platform Metrics

Virtualization Server platform metrics enable you to monitor the performance and resource usage of the virtualization server in your data center. You can use this information to balance resources or plan ahead to add resources to improve future performance.

The distribution of CPU and memory consumers is useful in managing the most heavily loaded zones, enabling you to be proactive in identifying potential issues. The last dashlet on the Summary page shows the date and time stamp for the last configuration change and the last incident.

The following zone metrics are available in the dashlets across the top of the page:

- Platform Status
  - Oracle Solaris version.
  - Current state: View the current health status and the time that the zone has been up and running.
  - Guest count. The guest count shows the total number of zones. If zones are still in the process of being discovered and promoted, the guest count will indicate the total number of guests and the number of guests that are still in the queue for auto promotion.
  - Open Incidents: View the number of Fatal, Critical, and Warning incidents. Click a number to view a synopsis. Click the synopsis to navigate to the Incident Manager console for incident details.
  - CPU Usage and Memory Usage: This section displays circular gauges showing the current platform CPU and memory as a percentage of the maximum values.

  - Up and Running Guests: This section displays the number of running zones out of the number of configured zones.
  - Virtual CPUs and Memory: This section displays pie charts showing the current number of allocated and available virtual CPUs and amount of allocated and available memory
  - Distribution of the CPU consumers per zone: View the total CPU consumption and the number of zones, or guests, that are consuming CPUs in the following ranges: 0-20, 20-40, 40-60, 60-80, and 80-100 percent.
  - Distribution of the Memory consumers per zone: View the total memory consumption and the number of zones, or guests, that are consuming memory in the following ranges: 0-20, 20-40, 40-60, 60-80, and 80-100 percent.
- Last changes
  - Configuration change: View the date and time stamp for the last change in configuration.
  - Incident: View the date and time stamp for the last incident.

The main body is made of 2 sections. The top one provides graphs showing the vCPU distribution per resource pool and zones, as well as the memory distribution per zones. The second section is the list of the guests. The Virtual Platform's Guests contains a list of the zones and zone details including the resources (vCPU, core, socket, share, memory) configured for the zone and those actually allocated to running zones. You can display the page as a list or a table. The default view is the List view, sorted by Incident count. The List view has sorting and filtering options that enable you to sort by incident and by allocated resources. Alternatively, you can view the Virtual Platform's Guests page as a sortable table. You can sort by the zone type, zone state, and by zone resource pool. You can also sort by the Fatal, Critical, or Warning incident columns.

Viewing Zone Platform Metrics

Some metrics appear in the dashlets, for details on specific zone metrics, view the Summary page.

1. From the Targets list, select **All Targets**.
2. From Servers, Storage, and Network, select **Virtualization Platform**.
3. Click the target name to open the Summary page for the Solaris Virtualization platform.
4. The dashlets display the metric information. The Summary page shows details for each zone. Details include the zone name, type of zone (global, non-global, kernel), the status, the number of vCPUs, memory, and the incidents.
5. The main body of the page contains details about the virtualization platform's guests, or zones.

Working with Zone-Specific Metrics

Virtual Server metrics enable you to monitor the performance and resource usage of a specific zone.

The Virtualization Platform shows the distribution of the CPU and memory consumers per zone. The Virtual Server shows the load imposed by a single zone (global, non-global, or kernel.) The last dashlet on the Virtual Server Summary page shows the date and time stamp for the last configuration change and the last incident.

The following metrics are available in the dashlets across the top of the page:

- Virtual Server Status
  - Current state: View the current health status and the type of zone (global, non-global, or kernel.)
  - Open Incidents: View the number of Fatal, Critical, and Warning incidents. Click a number to view a synopsis. Click the synopsis to navigate to the Incident Manager console for incident details.
- CPU and Memory Utilization
  - CPU Usage: View the percentage CPU used by the virtual server.
  - Memory Usage: View the percentage of memory used by the virtual server.
- Virtual Server Configuration
  - Guest OS Information: View the host name, IP address, and time that the guest (zone) has been up and running. To receive guest OS information, you must have an agent deployed on the zone's operating system.
  - Guest Configuration: View the guest UUID, host ID, and whether automatic boot is enabled.
- Last Changes
  - Configuration change: View the date and time stamp for the last change in configuration.
  - Incident: View the date and time stamp for the last incident.

The main body provides details about the CPU Resources (Shares, vCPUs, cores, sockets) configured and also being currently allocated to the zone and memory configured for the zone and currently allocated to the zone. The page also contains two graphs that show the CPU usage profile and the Memory usage profile. Each graph shows the CPU or memory usage during the past hour, day or week. More detailed CPU and Memory graphs appear in the Virtual Server Usages tab.

Viewing a Summary of Zone Metrics

Zone-specific metrics and resource usage charts appear on the zone's virtual platform Summary page.

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, click **Systems Infrastructure Virtualization Platform**. The Virtualization Platform provides the ability to drill down to see the virtual server that has the open incident.
3. Click the target name to open the Summary page for the Solaris Virtualization Platform.

Viewing Zone CPU and Memory Metrics

Zone-specific CPU and memory usage metrics and resource usage charts appear on the zone's virtual server page. More detailed metrics are available in the All Metrics page.

1. Click **Systems Infrastructure Virtual Server** from the All Targets page.
2. Click the target name to open the Summary page for the virtual server. The CPU and Memory metrics appear on the zone's home page and in the home page dashboard. The usage profiles appear for the last hour, day, or week.

3. Click the **Virtual server usages** icon on the right side of the page to view CPU and Memory Resource Usages Charts. Move the slider icon on the chart page to change the number of days to display in the chart. You can view up to seven days.

Viewing All Metrics

Platform and virtual server metrics are available in the All Metrics page.

1. Click **Systems Infrastructure Virtual Server** from the All Targets page.
2. Click the target name to open the **Solaris Virtualization Platform** or **Virtual Server** page, depending on your target.
3. Click **Solaris Virtualization Platform** or **Virtual Server** in the upper left corner of the page. Click **Monitoring**, then click **All Metrics**.Click a metric to view details, collection schedule, upload interval and other details.

> **Note:**
>
> Although it appears in the list of configuration metrics, the Component Faults metric is not supported for a Virtual Server.
>
> The `CoreUsage` metric is only reported for guest domains that have full cores allocated.

Working with Incidents for Zones

The following information will help when working with incident information for zones:

- About Incidents for Zones
- Viewing Open Incidents for Zones

About Incidents for Zones

The Virtualization Platform page shows all incidents that are open on the virtualization platform, enabling you to quickly see if there are any issues on the platform.

For zones, the number of incidents that appear on the Virtualization Platform page includes incidents for the global zone and any zones associated with the global zone. Click the number in the Open Incidents dashlet to display the list of incidents and the target on which the incident occurred.

In the list of incidents, you can click the target (global zone or zone) link to go to the corresponding target home page or you can click the incident to be redirected to the corresponding target Incident Manager page.

> **Note:**
>
> The number of open incidents on the Virtualization Platform page indicates all open incidents for the platform - for the global zone and all zones. By default, the Incident Manager page for the global zone shows the open incidents for the global zone, not the associated zones. You can change the display in the Incident Manager to show incidents for the global zone and zones. In Incident Manager, select Search, then select **Target and all members** for the Include Members search criteria.

Viewing Open Incidents for Zones

The Virtualization Platform and Virtual Server pages have an Open Incident dashlet, which displays the number of Fatal, Critical, and Warning incidents. The Virtualization Platform page shows open incidents for the global zone and all associated zones (non-global and kernel.) The Virtual Server page only shows incidents for the selected zone.

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform**. The Virtualization Platform does provide the ability to drill down to see the virtual server that has the open incident.
3. Click the target name to open the Summary page for the Solaris Virtualization Platform.
4. The Open Incidents dashlet displays a number to indicate how many open incidents are associated with the global zone and non-global zones.

   When a zone has an open incident, a number appears by the zone name in the Virtual Platform's Guests section of the Summary pane.
5. Click the target name in the Open Incident dashlet to go to the Virtual Server page for the zone.
6. Click the summary text to navigate to the Incident Manager page.

   The Incident Manager provides incident details and the events that led to the incident. You can drill down to get details on the events and notifications.You can acknowledge the incident, add comments, or manage the incident from the Incident Manager page.

   > **Note:**
   >
   > The Virtualization Platform page shows all incidents for the global zone and its zones. By default, the Incident Manager page only displays incidents for the global zone. You can change the setting in Incident Manager to display the target and all members.

Managing Metrics and Incident Notifications for Zones

You can perform the following tasks to manage monitoring and incident notification:

- Viewing Zone Metric Collection Errors
- Editing a Zone's Monitoring Configuration
- Suspending Monitoring Notifications for Zones

- Suspending Zone Monitoring for Maintenance
- Ending a Monitoring Brownout or Blackout for Zones

Viewing Zone Metric Collection Errors

Metric collection errors are usually caused by installation or configuration issues. You can view errors for a virtual server or for the virtualization platform.

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Monitoring**, then click **Metric Collection Errors**.

Editing Metric and Collection Settings for Zones

The Metrics tab contains displays all of the monitored attributes. The default view is metrics with thresholds. For these types of monitored attributes, you can modify the comparison operator, the threshold limits, the corrective action, and the collection schedule.

To edit the metrics and settings for a zone, navigate to the Virtual Server. To edit the settings for a virtualization platform, navigate to the Virtualization Platform. The parameters are different for each.

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** in the upper left corner of the page. Click **Monitoring**, then click **Metric and Collection Settings**.
4. Modify threshold limits or collection schedule. When a threshold field is empty, the alert is disabled for that metric.
5. Click the **Edit** icon for advanced settings.

   Click the **Other Collected Items** tab to view non-threshold monitored attributes. You can modify the collection period for these attributes, or disable monitoring.

Editing a Zone's Monitoring Configuration

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Target Setup**.
4. Click **Monitoring Configuration**.

Suspending Monitoring Notifications for Zones

Brownouts enable you to temporarily suppress notifications on a target. The Agent continues to monitor the target under brownout. You can view the actual target status along with an indication that the target is currently under brownout.

You can create a brownout for a virtual server or for the virtualization platform.

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Control**.
4. Click **Create Brownout**.
5. Enter a name for the brownout event.
6. Select a reason from the menu and add comments, as needed.
7. Click the options to define how jobs will run and the maintenance window.
8. Click **Submit**.

Suspending Zone Monitoring for Maintenance

Blackouts enable you to suspend monitoring on one or more targets in order to perform maintenance operations. To place a target under blackout, you must have at least the Blackout Target privilege on the target. If you select a host, then by default all the targets on that host are included in the blackout. Similarly, if you select a target that has members, then by default all the members are included in the blackout.

You can create a blackout for a virtual server or for the virtualization platform.

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Control**.
4. Click **Create Blackout**.
5. Select a reason from the menu.
6. Add comments, as needed.
7. Click **Submit**.

Ending a Monitoring Brownout or Blackout for Zones

You can end a blackout or brownout for a virtual server or for the virtualization platform.

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Control**.
4. Click **End Blackout** or **End Brownout**.

Administering Zones

You can perform the following tasks to manage and administer zones:

- Viewing Zone Compliance
- Identifying Changes in a Zone Configuration
- Editing Zone Administrator Access
- Adding a Zone to a Group
- Editing Zone Properties

Viewing Zone Compliance

The Compliance pages enable you to view the compliance framework, standards, and the zone's compliance.

You can view compliance for a virtual server or for the virtualization platform.

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Compliance**.
4. Click the option to view **Results**, **Standard Associations**, or **Real-time Observations**.

Identifying Changes in a Zone Configuration

When an administrator changes a zone configuration, it can be helpful to know the when the configuration was last changed. This information appears in the configuration dashlet on the Summary page. The zone configuration information does not require an agent to be installed on the zone.

To view more detailed information for a virtual server or virtualization platform:

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Configuration**.
4. Click the option to view **Last Collected**, **Comparison and Drift Management**, **Compare**, **Search**, **History**, **Save**, **Saved**, or **Topology**.

Editing Zone Administrator Access

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Target Setup**.
4. Click **Administrator Access**.

Adding a Zone to a Group

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Target Setup**.
4. Click **Add to Group**.

Editing Zone Properties

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Target Setup**.
4. Click **Properties**.

Additional Resources for Oracle Solaris Zones

See the following for more information:

- Discovering and Adding Host and Non-Host Targets
- Discovering, Promoting, and Adding System Infrastructure Targets
- Using Incident Management
- Using Blackouts

For in-depth information about Oracle Solaris and Oracle Solaris Zones, go to `http://docs.oracle.com/en/operating-systems/`.

29 Monitoring Oracle VM Server for SPARC

The following features and topics are covered in this chapter:

- Getting Started With Oracle VM Server for SPARC Virtualization
- Location of Oracle VM Server for SPARC Information in the UI
- Actions for Oracle VM Server for SPARC
- Target Navigation for Oracle VM Server for SPARC
- Supported Versions
- Viewing all Oracle VM Server for SPARC Virtualization Platforms
- About Virtualization Platform Information
- Zones within a Logical Domain
- About Logical Domain Information
- Related Resources for Oracle VM Server for SPARC

Getting Started With Oracle VM Server for SPARC Virtualization

Oracle VM Server for SPARC technology enables server virtualization on SPARC platforms. You can create and manage multiple virtual machine instances simultaneously on a single SPARC machine. Each virtual machine, or guest, can run a separate Oracle Solaris 10 or Oracle Solaris 11 operating system.

Oracle VM Server for SPARC technology is virtualization of SPARC servers. This technology is part of a suite of methodologies for consolidation and resource management for SPARC Chip Multi Threading (CMT) systems. Using this technology, you can allocate the various resources of the system such as memory, CPU threads, and devices, into logical groupings and create multiple discrete systems. These discrete systems have their own operating system, resources, and identity within a single system. By careful architecture, an Oracle VM Server for SPARC environment can help you achieve greater resource usage, better scaling, and increased security and isolation.

Terminology

In Oracle Enterprise Manager, a virtualization platform is a virtualization technology, such as an Oracle VM Server for SPARC control domain, that can host guests. A virtual server is a guest, such as a logical domain.

Logical Domains

When Oracle VM Server for SPARC software is installed, a domain called the control domain is created. From this control domain, you create virtual machines called logical domains that each run an independent OS. A logical domain is a virtual machine with resources, such as CPU threads, memory, I/O devices, and its own operating system. The control domain manages the logical domains. Each logical domain can be created, destroyed, reconfigured, and rebooted independently of other logical domains.

Location of Oracle VM Server for SPARC Information in the UI

You can select any target that is a child or a parent of the logical domain (virtual platform, host, or server) and the logical domain will appear in the Navigation pane of the target.

The Virtualization Platform is the container on which logical domain virtual servers or zones are running.

The following options are available on the All Targets page:

- Click **Systems Infrastructure Virtualization Platform** to see the list of virtual platforms. You can click any virtual platform in this list to open the virtual platform's home page.
- Click **Systems Infrastructure Virtual Server** to see the list of virtual servers.You can click any virtual server in this list to open the virtual server's home page.

Table 29-1 Location of Oracle VM Server for SPARC Information in the UI

Object	Location
Virtualization Platform	A Virtual Platform target home page is accessible from the All Targets page, or from any of its parent or child targets through the Target Navigation pane. Click the **All Targets** selector at the top of the page. In the Refine Search section, select **Systems Infrastructure Virtualization Platform**. Alternatively, select any target that is a parent or a child of the Virtual Platform in the All Targets selector. The target can be a host, a server, or a logical domain target running on this Virtual Platform. A Target Navigation pane is located in the top left corner of a parent or child home page. Click the Target Navigation pane to expand it and see the Virtual Platform in the Navigation pane.
Virtual Server	A logical domain virtual server target home page is accessible from the zone Virtualization Platform page or the Virtual Server page. Click the All Targets selector at the top of the page. In the Refine Search section, select **Systems Infrastructure Virtual Server**. If you are on the Virtualization Platform page, click a zone to navigate to the virtual server page for that zone.

Actions for Oracle VM Server for SPARC

You can discover and promote Oracle VM Servers for SPARC by deploying an Em Agent on the Control Domain operating system. You can then monitor the Oracle VM Server for SPARC and its guests. If the Control Domain or its guests have Oracle Solaris Zones installed, they are displayed in the target navigation tree, and you can monitor them if an EM Agent is installed in the guest domain.

You can perform the following actions:

- View the configuration and status of domains.
- View resource utilization and the top consumers of resources.
- Diagnose problems using incidents and performance metrics.

Target Navigation for Oracle VM Server for SPARC

The target navigation tree helps you to navigate between targets that are in an Oracle VM Server for SPARC Virtual Platform (OVM SPARC Virtual Platform.) When all resources are discovered and monitored through agent deployment on the primary domain, you can navigate from the server down.

The top level, or node, is the physical server and the second node is the OVM SPARC Virtual Platform that is hosting the primary domain and all other logical domains. Expand the nodes to drill down the target hierarchy or click a specific target to open that target's landing page.

Figure 29-1 is an example of a target navigation tree for an OVM SPARC Virtual Platform. The first node is the physical server (slci04-node4-ilom), which is running the Oracle VM Server for SPARC hypervisor (slci04dbadm07 (OVM SPARC Virtual Platform)). The hypervisor appears on the second node. The third node displays four logical domain guests: primary, ssccn4-app1, ssccn4-app2, and ssccn4-app3. The primary domain is also a Solaris Zones Virtual Platform (slci04dbadm07 (Solaris Zones Virtual Platform)) and that has one global zone that is running the operating system (slci04dbadm07 host.) The fourth logical domain (ssccn4-app3) is expanded. When an agent was deployed on the logical domain, the domain was promoted as a zone virtual platform that is running the global zone and host slci04dbadm08.

Figure 29-1 Target Navigation for Oracle VM Server for SPARC

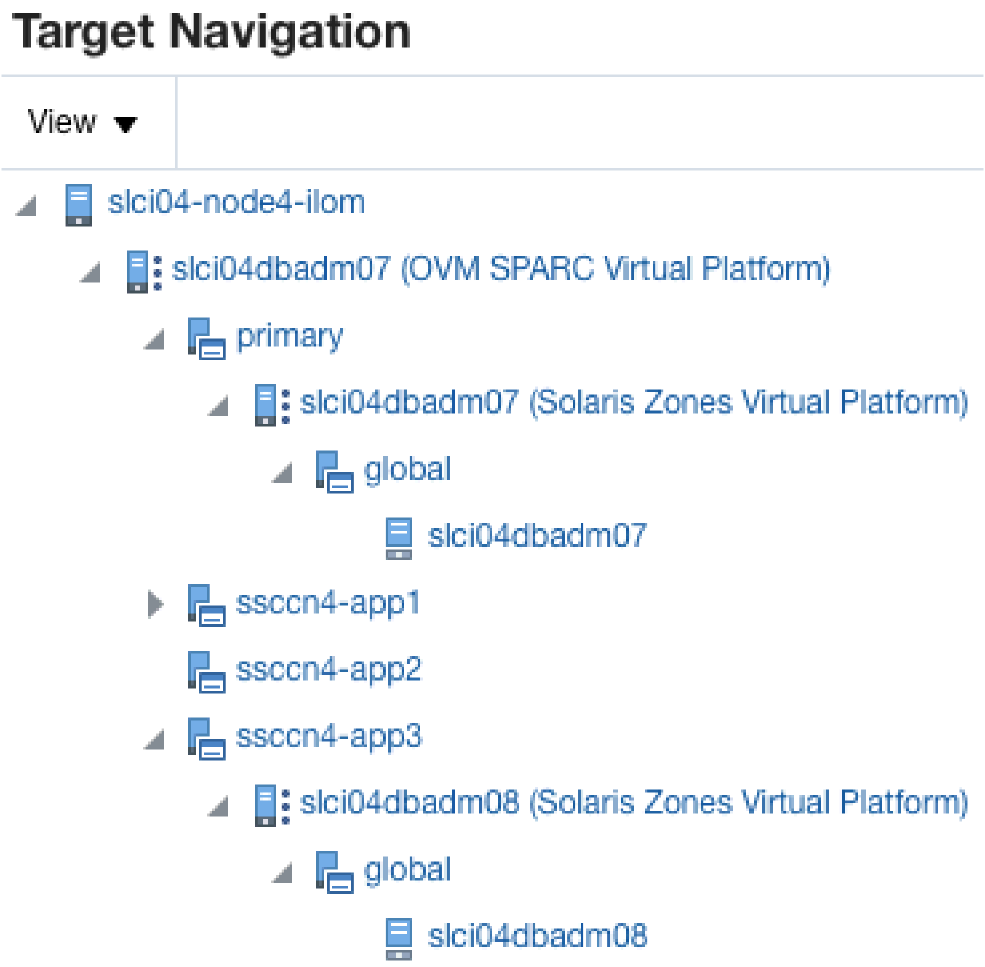

Target Navigation

- slci04-node4-ilom
  - slci04dbadm07 (OVM SPARC Virtual Platform)
    - primary
      - slci04dbadm07 (Solaris Zones Virtual Platform)
        - global
          - slci04dbadm07
    - ssccn4-app1
    - ssccn4-app2
    - ssccn4-app3
      - slci04dbadm08 (Solaris Zones Virtual Platform)
        - global
          - slci04dbadm08

Supported Versions

Oracle VM Server for SPARC Control Domain:

- Oracle Solaris 11.1 and later
- Oracle VM Server for SPARC 3.1 and later

Oracle VM Server for SPARC Guest Domain Operating System:

- Oracle Solaris 10 1/13
- Oracle Solaris 11.1 and later

Viewing all Oracle VM Server for SPARC Virtualization Platforms

You can view all virtualization platforms, including all Oracle VM Server for SPARC virtualization platforms, from the targets list.

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Virtualization Platform**.

   A list of the target virtualization platforms is displayed.

About Virtualization Platform Information

You can select a virtualization platform to view information about the virtualization platform and its guests.

The top section of the UI displays basic information about the selected virtualization platform. You can click the scroll icons to move left or right.

- Platform: This section displays the platform type, its version, and its uptime.
- Open Incidents: This section displays the number of fatal, critical, and warning incidents on the system. You can click on a category to bring up a detailed list of incidents within that category.

  > **Note:**
  >
  > The number of open incidents on the Virtualization Platform page indicates all open incidents for the platform - for the control domain and all logical domains. By default, the Incident Manager page for the control domain shows the open incidents for the control domain, not the associated logical domains. You can change the display in the Incident Manager to show incidents for the control domain and logical domains. In Incident Manager, select Search, then select Target and all members for the Include Members search criteria.

- CPU and Memory: This section displays circular gauges showing the current CPU and memory as a percentage of the maximum values. For more details, click the third tab. Beginning with Oracle VM Server for SPARC version 3.2, the Resource Group feature is available for eligible platforms with the proper service processor firmware. The Resource Group feature provides physical CPU information for logical domain guests without requiring a discovered ILOM.

> **Note:**
>
> On older hardware and Oracle VM Server for SPARC versions earlier than 3.2, this information is displayed only if the ILOM of the hardware has been discovered.

- Up and Running Guests: This section displays the number of running guests out of the number of configured guests.
- Virtual CPUs, Memory, and Cores: This section displays pie charts showing the current number of allocated and available virtual CPUs, amount of allocated and available memory, and number of allocated and available cores.
- Total CPU Consumption and its Distribution per Guest: This section shows the total CPU consumption and a graph of the guest count for each cpu range.
- Total Memory Consumption and its Distribution per Guest: This section shows the total memory consumption for the virtualization platform and a graph of the guest count for each memory range.
- Total Power Consumption and its Distribution per Guest: This section shows the total power consumption and a graph of the guest count for each power range.
- Last Configuration Change and Incident: This section shows the date of the last configuration change on the system, and the date of the last reported incident.

Viewing the Virtualization Platform Basic Information

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Virtualization Platform**.

   A list of the target virtualization platforms is displayed.
3. Click the target name to open the Summary page for the virtualization platform. The dashlets appear on the top of the page and provide the summary information.

About the Virtualization Platform's Guest Summary

The Guest Summary section displays details about the guests that are managed by the virtualization platform. Select the Summary tab to view this information.

The top section displays bar graphs showing the virtual CPUs and memory configured for each guest, as a portion of the total available.

The bottom section displays a table, or list, of the guests. You can select list or table for the guest display. In list mode, use the sort options in the upper left to sort the guests by type, incidents, memory, or vCPUs.

By default, the domains are displayed in descending order based on the number of incidents.

For each guest, the following information is displayed:

- Target Status Icon: This icon indicates whether the guest is monitored.
- Type Icons: These icons identify the type of guest. Separate icons identify control domains, guest domains, root domains, IO domains, and service domains.
- Name: The guest's name.

- CPU Information: Displays the number of CPU cores and vCPUs allocated to the guest
- Memory: Displays the memory available to the guest.
- Operational Status Icon: This icon indicates whether the guest is unbound, started, or stopped.
- Incidents: The numbers of open fatal, critical, and warning incidents for the guest.
- Cores: Displays the number of cores allocated to the guest.
- CPU Usage Graph: Displays the current CPU usage and the CPU usage over the past five hours. This information is displayed only if the guest is started, and is only displayed in list mode.

You can use the search field at the top of the guest table to search for specific guests.

Viewing the Virtualization Platform Guest Summary

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Virtualization Platform**.

   A list of the target virtualization platforms is displayed.
3. Click the target name to open the Summary page for the virtualization platform.
4. Click the **Summary** tab.

About the Virtualization Platform's Services

You can view details about the I/O and network services available on the virtualization platform, and view the topology of a virtualization platform, showing what services are provided and consumed. Select the Services tab to view this information.

The top section displays a table of I/O and network services. For each resource, the following information is displayed:

- Name
- Type
- Operational Status icon
- Number of guests using the resource

You can use the search field at the top of the table to search for specific services.

For each resource, you can click More to display additional information.

The bottom section displays a topology diagram, which shows each guest and the network resources provided or consumed by each guest. You can select the only guests option to display only the guests.

In the guests and services display, you can hover over a guest to highlight the network services it is using, or hover over a network resource to highlight the guests using it.

In the only guests display, you can hover over a guest to show the network resources it shares with other guests.

Click the Control Panel button to access zoom controls. You can use these controls to zoom in, zoom out, or zoom to fit the current window.

Viewing the Virtualization Platform Services

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Virtualization Platform**.

   A list of the target virtualization platforms is displayed.
3. Click the target name to open the Summary page for the virtualization platform.
4. Click the **Services** tab.

About the Virtualization Platform's vCPU and Core Allocation

You can view details about vCPU and core allocation. Select the Core Distribution tab to view this information.

This section displays a pie chart, showing which CPUs and cores are allocated to which guests, and which are unallocated. You can click on a guest in the outer layer to view detailed information about that guest's resource consumption. The following fields are displayed:

- Name: The guest's user-friendly name.
- Type Icons: These icons identify the type of guest. Separate icons identify control domains, guest domains, root domains, IO domains, and service domains.
- Operational Status Icon
- Number of vCPUs: The number of virtual CPUs assigned to the guest.
- CPU Usage: Displays the guest's CPU usage over the last hour, day, or week.

Viewing the Virtualization Platform vCPU and Core Allocation

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Virtualization Platform**.

   A list of the target virtualization platforms is displayed.
3. Click the target name to open the Summary page for the virtualization platform.
4. Click the **Core Distribution** tab.

About Virtualization Platform Metrics

You can view a complete list of the metrics for a selected virtualization platform.

Viewing Platform Metrics

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Virtualization Platform**.

   A list of the target virtualization platforms is displayed.
3. Click the target name to open the Summary page for the virtualization platform.

4. Click **Virtualization Platform** in the upper left corner of the page. Click **Monitoring**, then click **All Metrics**.
5. Click a metric to view details, collection schedule, upload interval and other details.

Zones within a Logical Domain

You can discover and monitor zones installed on the operating system of a logical domain, including the control domain, if you have installed an EM Agent on the logical domain operating system.

Viewing Zones in a Logical Domain

The zones view for zones installed on a logical domain is shown beneath the logical domain in the target navigation pane. You can select the zones view or select an individual zone to view more information.

About Logical Domain Information

Oracle Enterprise Manager collects detailed information about monitored logical domains, including CPU and memory usage, logical domain status, and incidents. You can view these metrics by selecting a logical domain.

The top section of the UI displays basic information about the selected logical domain. You can click the scroll icons to move left or right.

The following sections are displayed:

- Logical Domain: This section displays the logical domain name, its types, and its uptime.
- Open Incidents: This section displays the number of fatal, critical, and warning incidents on the system. You can click on a category to bring up a detailed list of incidents within that category.
- CPU and Power: This section displays a circular gauge showing the current CPU usage as a percentage of the maximum value and the current power usage in Watts.
- Guest OS Information: This section displays information about the guest's operating system, if it has been discovered. This information includes the hostname, IP address, and uptime.
- Guest Configuration: This section shows guest configuration information, including the guest UUID, an icon indicating whether the guest is set to boot automatically, and the console port.
- Last Configuration Change and Incident: This section shows the date of the last configuration change on the system, and the date of the last reported incident.

Viewing the Logical Domain's Basic Information

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Virtual Server**.

   A list of the target virtual servers is displayed.
3. Click the target name to open the Summary page for the virtual server.

About the Virtual Server Summary Information

The virtual server summary tab displays basic information about the virtual server, including status, virtual CPU, and memory usage.

The following fields are displayed:

- Number of vCPUS or Cores
- Amount of Memory
- Automatic Boot Configuration
- UUID
- Graphs of the historical CPU usage percentage over the past hour, day, and week
- Graphs of the historical Power usage in Watts over the past hour, day, and week

Viewing the Virtual Server Summary Information

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Virtual Server**.

   A list of the target virtual servers is displayed.
3. Click the target name to open the Summary page for the virtual server.
4. Click the **Summary** tab.

About the Virtual Server Power and CPU Usage Charts

The Virtual Server Usages tab displays charts of the guest's CPU and power usage. Select one of the following icons to display the relevant chart:

- CPU (%): Displays a chart of the virtual server's CPU usage, as a percentage of the total.
- Power (Watts): Displays a chart of the power in Watts consumed by the CPU and by the memory.

For each chart, you can select a number of days to display using the slider in the upper right.

Viewing the Virtual Server Power and CPU Usage Charts

1. Select **All Targets** from the Targets list.
2. Under the Servers, Storage, and Network heading, select **Systems Infrastructure Virtual Server**.

   A list of the target virtual servers is displayed.
3. Click the target name to open the Summary page for the virtual server.
4. Click the **Virtual Server Usages** tab.

Managing Metrics and Incident Notifications

You can perform the following tasks to manage monitoring and incident notification:

Viewing Metric Collection Errors

Metric collection errors are usually caused by installation or configuration issues. You can view errors for a virtual server or for the virtualization platform.

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Monitoring**, then click **Metric Collection Errors**.

Editing Metric and Collection Settings

The Metrics tab contains displays all of the monitored attributes. The default view is metrics with thresholds. For these types of monitored attributes, you can modify the comparison operator, the threshold limits, the corrective action, and the collection schedule.

To edit the metrics and settings for a virtual server, navigate to the Virtual Server. To edit the settings for a virtualization platform, navigate to the Virtualization Platform. The parameters are different for each.

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Monitoring**, then click **Metric and Collection Settings**.
4. Modify threshold limits or collection schedule. When a threshold field is empty, the alert is disabled for that metric.
5. Click the **Edit** icon for advanced settings.

   Click the **Other Collected Items** tab to view non-threshold monitored attributes. You can modify the collection period for these attributes, or disable monitoring.
6. Click **OK** to save your changes.

Editing a Monitoring Configuration

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Target Setup**.
4. Click **Monitoring Configuration**.

Suspending Monitoring Notifications

Brownouts enable you to temporarily suppress notifications on a target. The Agent continues to monitor the target under brownout. You can view the actual target status along with an indication that the target is currently under brownout.

You can create a brownout for a virtual server or for the virtualization platform.

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Control**.
4. Click **Create Brownout**.
5. Enter a name for the brownout event.
6. Select a reason from the menu and add comments, as needed.
7. Click the options to define how jobs will run and the maintenance window.
8. Click **Submit**.

Suspending Monitoring for Maintenance

Blackouts enable you to suspend monitoring on one or more targets in order to perform maintenance operations. To place a target under blackout, you must have at least the Blackout Target privilege on the target. If you select a host, then by default all the targets on that host are included in the blackout. Similarly, if you select a target that has members, then by default all the members are included in the blackout.

You can create a blackout for a virtual server or for the virtualization platform.

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Control**.
4. Click **Create Blackout**.
5. Select a reason from the menu.
6. Add comments, as needed.
7. Click **Submit**.

Ending a Monitoring Brownout or Blackout

You can end a blackout or brownout for a virtual server or for the virtualization platform.

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Control**.

4. Click **End Blackout** or **End Brownout**.

Administering Oracle VM Server for SPARC

You can perform the following tasks to manage and administer Oracle VM Server for SPARC:

- Viewing Compliance
- Identifying Changes in a Virtual Server Configuration
- Editing Virtual Server Administrator Access
- Adding a Virtual Server to a Group
- Editing Virtual Server Properties

Viewing Compliance

The Compliance pages enable you to view the compliance framework, standards, and the virtual server or virtual platform's compliance.

You can view compliance for a virtual server or for the virtualization platform.

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Compliance**.
4. Click the option to view **Results**, **Standard Associations**, or **Real-time Observations**.

Identifying Changes in a Virtual Server Configuration

When an administrator changes a virtual server or virtualization platform configuration, it can be helpful to know the when the configuration was last changed. This information appears in the configuration dashlet on the Summary page.

To view more detailed information for a virtual server or virtualization platform:

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Configuration**.
4. Click the option to view **Last Collected**, **Comparison and Drift Management**, **Compare**, **Search**, **History**, **Save**, **Saved**, or **Topology**.

Editing Virtual Server Administrator Access

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Target Setup**.

4. Click **Administrator Access**.

Adding a Virtual Server to a Group

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Target Setup**.
4. Click **Add to Group**.

Editing Virtual Server Properties

1. Click **Systems Infrastructure Virtual Server** or **Systems Infrastructure Virtualization Platform** from the All Targets page.
2. Click the target name to open the home page.
3. Click **Virtual Server** or **Virtualization Platform** in the upper left corner of the page. Click **Target Setup**.
4. Click **Properties**.

Related Resources for Oracle VM Server for SPARC

See the *Oracle Enterprise Manager Cloud Control Administrator's Guide* for information about the following:

- Discovering and Adding Host and Non-Host Targets
- Discovering, Promoting, and Adding System Infrastructure Targets
- Using Incident Management

30 Provisioning Zones with Oracle Database on Database Domains

This section describes how to create Solaris Zones in Database Domains and deploy Oracle RAC (Real Application Clusters) databases in these Solaris zones.

Topics:

Prerequisites

Deploy an EM agent on SuperCluster Control Domains and Database Domains

An EM agent must be deployed on the SuperCluster Control Domains, as a user who is granted Role-Based Access Control privileges to monitor Oracle VM Server for SPARC. See Discovering and Promoting Oracle VM Server for SPARC .

An EM agent must also be deployed on Database Domains on which Solaris Zones will be created.

Discover the SuperCluster

See Discovering and Promoting Oracle SuperCluster.

Discover the Exadata Database Machine

See Exadata Database Machine Discovery.

Import Grid Infrastructure and Database Software gold images in Software Library

There are two options to import Grid Infrastructure and Database gold images in the Enterprise Manager Software Library. Either you can clone an existing deployment, or you can import Provisioning Archive (PAR) files stored on a host.

- **Option 1:**

  Once you have discovered a Host on which Oracle Clusterware and a Database is deployed, you can create an Oracle Clusterware Clone and Oracle Database Software Clone component in the Software Library.

  Once done for both Component sub-type "Oracle Clusterware Clone" or "Oracle Database Software Clone", you can create a Database Zones Cluster.

  You can then also export these clones as Provisioning Archives (PAR files) to any host, so that you can later import these PAR files on any other Enteprise Manager Software Library to have "Oracle Clusterware Clone" and "Oracle Database Software Clone" imported in this library, this is the option 2 described below.

To create such PAR file, from the Software Library Home Page, select **Actions**, then click **Export**.

- **Option 2:**

  If you have already Provisioning Archive (PAR) files stored on a host (files created as described above in option 1), then from any Enterprise Manager you can deploy an agent on this host and import the par files on this host to the Software Library.

  It can be done from the Software Library Home Page through selecting **Actions**, then clicking **Import**.

  Once you have imported a Provisioning Archive for Oracle Clusterware and imported another PAR file for the Database, you are ready to create a Database Zones Cluster.

Create a DB Zones Cluster

1. Create a Cluster Definition.
   a. Select a Database Machine target and once on the Database Machine Home Page, select **Database Machine**, then **Provisioning**, and finally click **Create cluster**.
   b. Enter a cluster name, the list of available Domains provided are Domains on which an Enterprise Manager agent was deployed.
   c. Select one or several Database Domains on which a Solaris zone will be created.
   d. Select at least 3 Exadata Storage Servers:

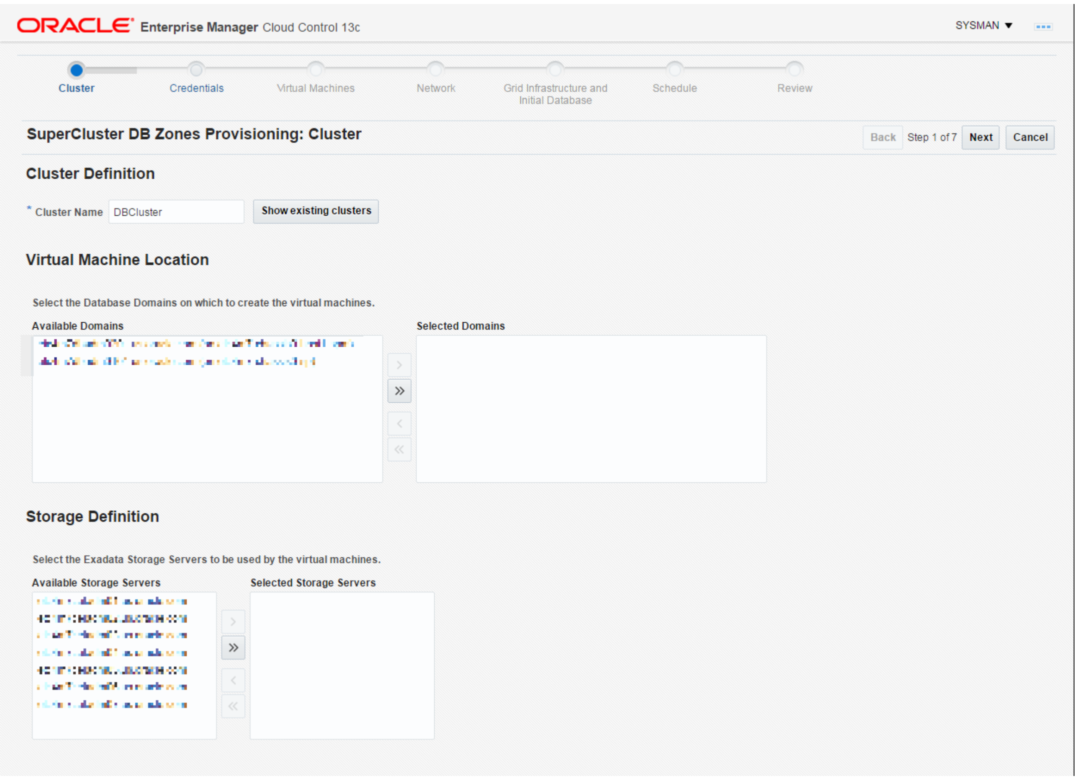

SuperCluster DB Zones Provisioning: Cluster

Cluster Definition

* Cluster Name: DBCluster

Virtual Machine Location

Select the Database Domains on which to create the virtual machines.

Available Domains: [redacted]

Selected Domains

Storage Definition

Select the Exadata Storage Servers to be used by the virtual machines.

Available Storage Servers: [redacted]

Selected Storage Servers

2. Create Credentials by providing root credentials of Database Domains and administrator credentials of Exadata Storage Servers that were selected:

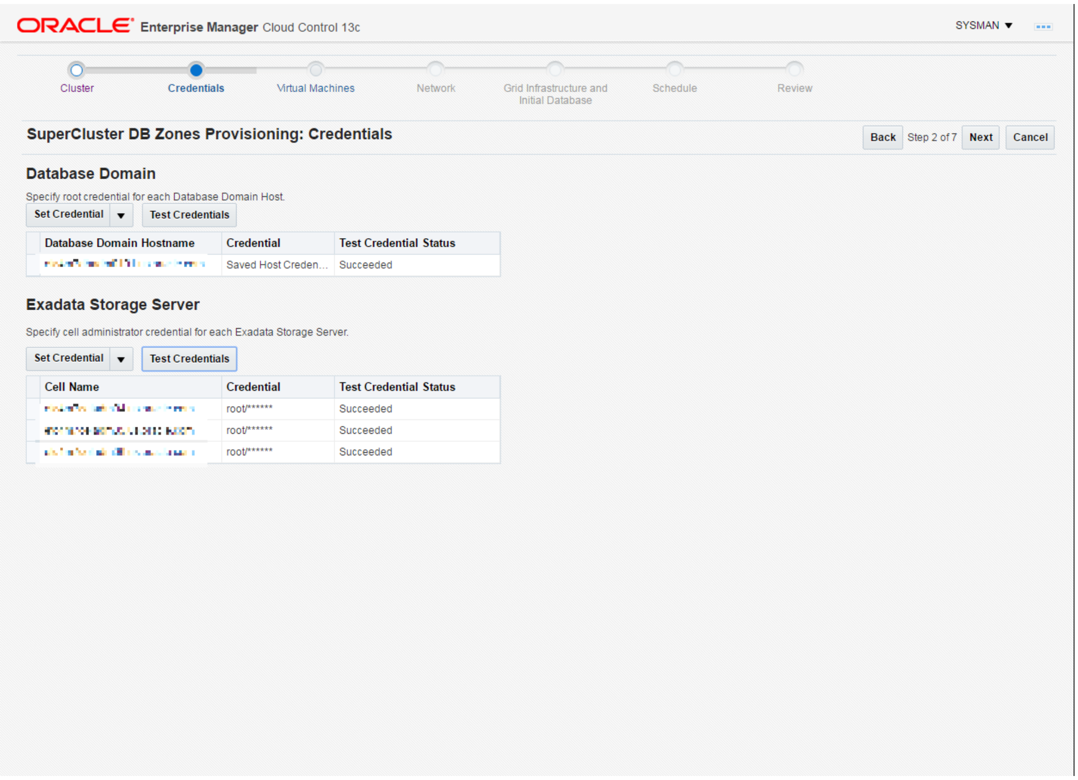

SuperCluster DB Zones Provisioning: Credentials

Database Domain

Specify root credential for each Database Domain Host.

Set Credential ▼ | Test Credentials

Database Domain Hostname	Credential	Test Credential Status
[redacted]	Saved Host Creden...	Succeeded

Exadata Storage Server

Specify cell administrator credential for each Exadata Storage Server.

Set Credential ▼ | Test Credentials

Cell Name	Credential	Test Credential Status
[redacted]	root/******	Succeeded
[redacted]	root/******	Succeeded
[redacted]	root/******	Succeeded

3. Create a Virtual Machine Definition.

   One Solaris Zone will be created on each selected Database Domain.

   By default the **Small** Virtual Machine size is selected. This will allocate 1 CPU core to each Solaris Zone to create, and a quota of 50GB will be set on the ZFS filesystem mounted on /u01 on the Solaris Zone and used to store the Grid Infrastructure and Database. You can select another Virtual Machine size (**Medium** or **Large**) or click on **Customize...** to adjust the number of CPU cores and storage size. You can click on **Show Details...** to see the number of available CPU cores on each of the Database Domains selected. A certain number of cores should be set aside for the global zone (the Operating System instance on the Database Domain) depending on the size of the Domain: 2 cores could be set aside for the global zone on a small Domain of less than 32 cores, and 4 cores could be set aside for the global zone on Domains with more cores. The remaining cores can be allocated to new Solaris zones that will be created on this Domain.

   a. Provide the root password for the root account that will be created on each Solaris Zone.

   b. DNS IP addresses are pre-populated. Enter the IP address of a NTP (Network Time Protocol) Server.

   c. Grid Infrastructure and Database Software versions are populated from the Gold Images in Software Library (see the Prerequisite section above).

   d. Provide the password for new accounts that will be created on each Solaris zone.

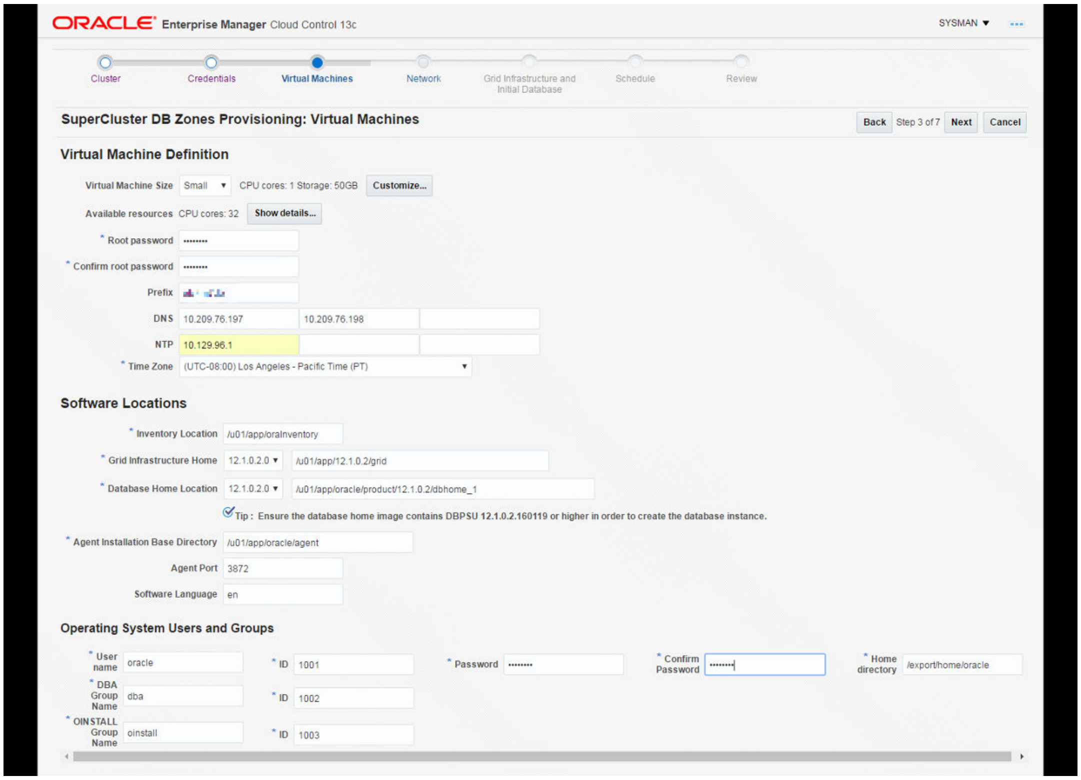

SuperCluster DB Zones Provisioning: Virtual Machines

Virtual Machine Definition

Virtual Machine Size	Small ▼	CPU cores: 1 Storage: 50GB	Customize...
Available resources	CPU cores: 32	Show details...	
* Root password	••••••••		
* Confirm root password	••••••••		
Prefix	[redacted]		
DNS	10.209.76.197	10.209.76.198	
NTP	10.129.96.1		
* Time Zone	(UTC-08:00) Los Angeles - Pacific Time (PT) ▼		

Software Locations

* Inventory Location	/u01/app/oraInventory	
* Grid Infrastructure Home	12.1.0.2.0 ▼	/u01/app/12.1.0.2/grid
* Database Home Location	12.1.0.2.0 ▼	/u01/app/oracle/product/12.1.0.2/dbhome_1

✓Tip : Ensure the database home image contains DBPSU 12.1.0.2.160119 or higher in order to create the database instance.

* Agent Installation Base Directory	/u01/app/oracle/agent
Agent Port	3872
Software Language	en

Operating System Users and Groups

* User name	oracle	* ID	1001	* Password	••••••••	* Confirm Password	••••••••	* Home directory	/export/home/oracle
* DBA Group Name	dba	* ID	1002						
* OINSTALL Group Name	oinstall	* ID	1003						

4. Configure the Network.

   Network Domains and Subnet masks are pre-populated.

   Provide IP addresses for Admin Network Gateway IP and and Client Network Gateway IP.

   Provide IP addresses and hostnames for each network (admin network, client network, Virtual IP on client network, private network) for the Solaris Zone on each Database Domain selected:

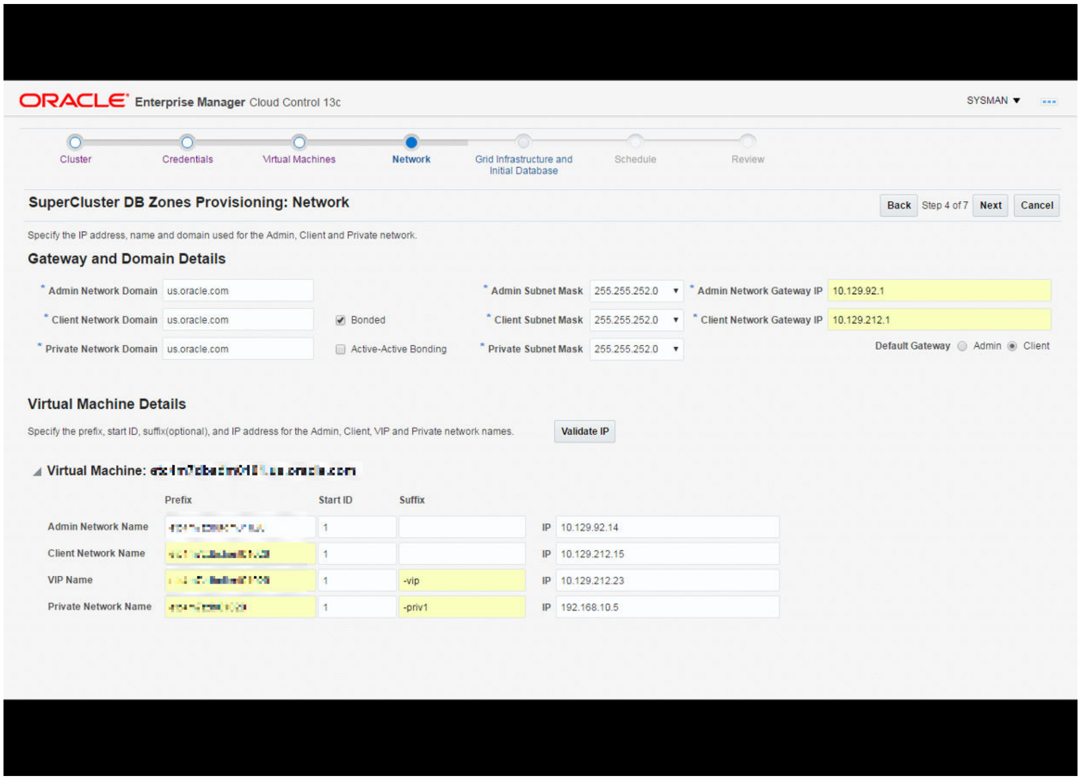

SuperCluster DB Zones Provisioning: Network

Specify the IP address, name and domain used for the Admin, Client and Private network.

Gateway and Domain Details

* Admin Network Domain	us.oracle.com		* Admin Subnet Mask	255.255.252.0 ▾	* Admin Network Gateway IP	10.129.92.1
* Client Network Domain	us.oracle.com	☑ Bonded	* Client Subnet Mask	255.255.252.0 ▾	* Client Network Gateway IP	10.129.212.1
* Private Network Domain	us.oracle.com	□ Active-Active Bonding	* Private Subnet Mask	255.255.252.0 ▾	Default Gateway	○ Admin ● Client

Virtual Machine Details

Specify the prefix, start ID, suffix(optional), and IP address for the Admin, Client, VIP and Private network names.

Virtual Machine: [redacted]

	Prefix	Start ID	Suffix		
Admin Network Name	[redacted]	1		IP	10.129.92.14
Client Network Name	[redacted]	1		IP	10.129.212.15
VIP Name	[redacted]	1	-vip	IP	10.129.212.23
Private Network Name	[redacted]	1	-priv1	IP	192.168.10.5

5. Create a Grid Infrastructure and Initial Database.

   Provide the SCAN hostname (Single Client Access Name that provides a single hostname for clients to access Oracle Databases running in the cluster), ASM password, Disk group names (must be unique, a check will be performed to verify no such disk group already exists), global database name and SID, and Administrative users passwords.

Click the check box to create the initial database and provide additional information for Database Identification and Administrator Credentials:

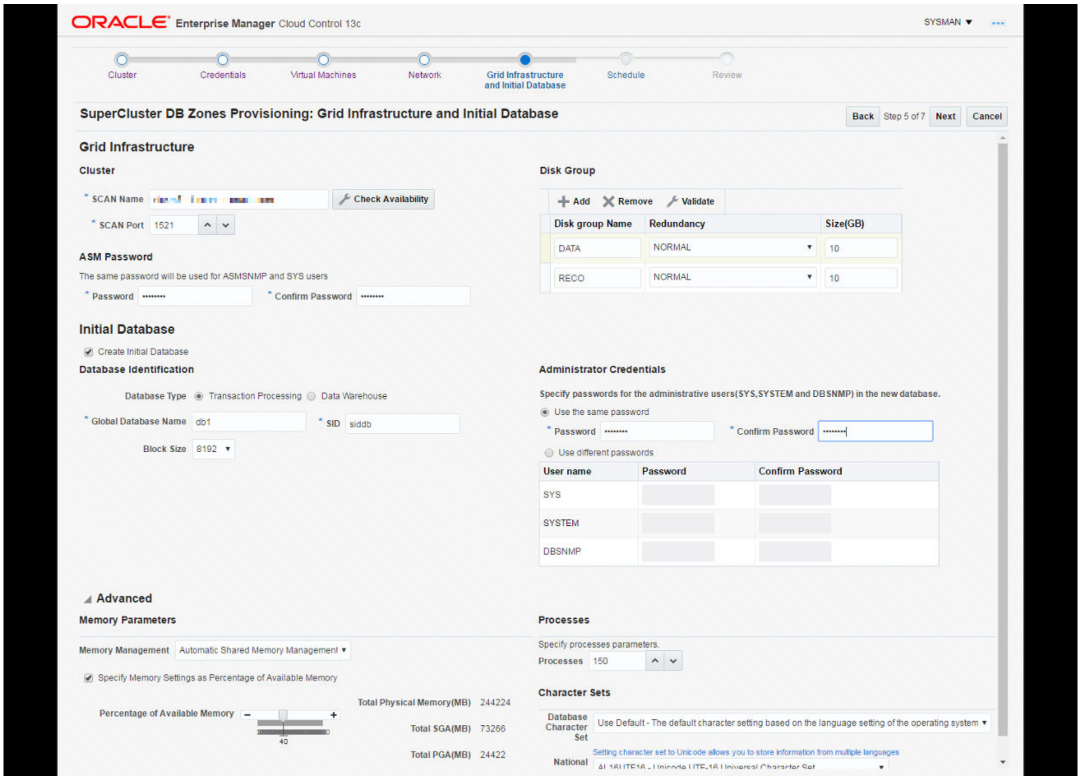

ORACLE Enterprise Manager Cloud Control 13c — SYSMAN

SuperCluster DB Zones Provisioning: Grid Infrastructure and Initial Database

Grid Infrastructure

Cluster

* SCAN Name [redacted] Check Availability

* SCAN Port 1521

ASM Password

The same password will be used for ASMSNMP and SYS users

* Password •••••••• * Confirm Password ••••••••

Disk Group

Disk group Name	Redundancy	Size(GB)
DATA	NORMAL	10
RECO	NORMAL	10

Initial Database

☑ Create Initial Database

Database Identification

Database Type ● Transaction Processing ○ Data Warehouse

* Global Database Name db1 * SID siddb

Block Size 8192

Administrator Credentials

Specify passwords for the administrative users(SYS,SYSTEM and DBSNMP) in the new database.

● Use the same password

* Password •••••••• * Confirm Password ••••••••

○ Use different passwords

User name	Password	Confirm Password
SYS		
SYSTEM		
DBSNMP		

Advanced

Memory Parameters

Memory Management Automatic Shared Memory Management

☑ Specify Memory Settings as Percentage of Available Memory

Percentage of Available Memory 40

Total Physical Memory(MB) 244224

Total SGA(MB) 73266

Total PGA(MB) 24422

Processes

Specify processes parameters.

Processes 150

Character Sets

Database Character Set Use Default - The default character setting based on the language setting of the operating system

Setting character set to Unicode allows you to store information from multiple languages

National AL16UTF16 - Unicode UTF-16 Universal Character Set

6. Schedule a Deployment

   Schedule the deployment, by default this is started immediately.

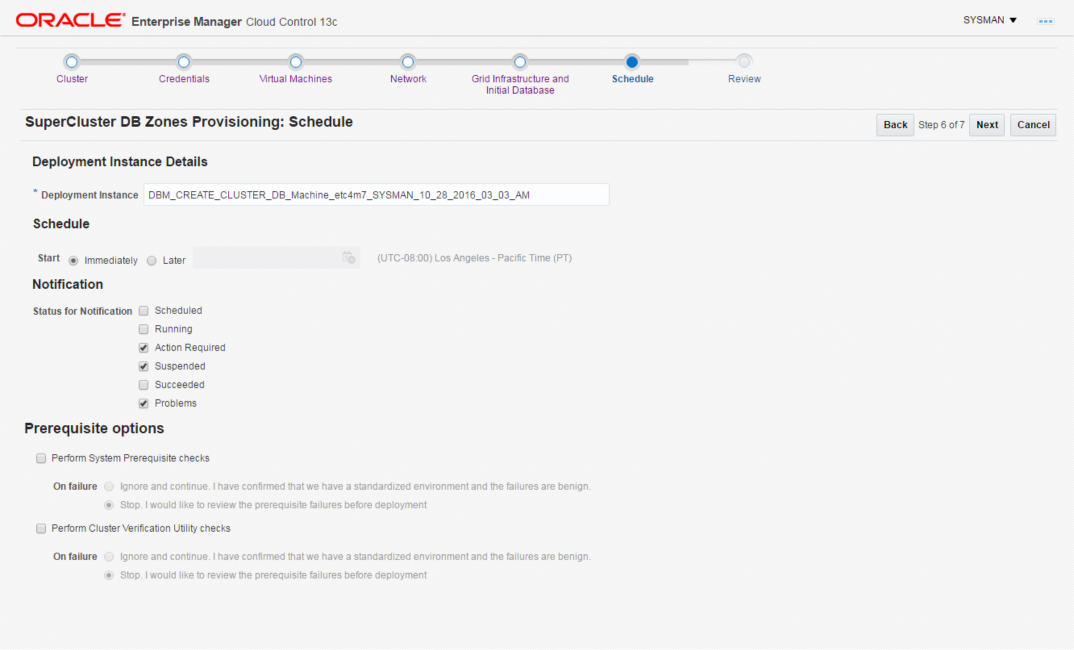

ORACLE Enterprise Manager Cloud Control 13c — SYSMAN

SuperCluster DB Zones Provisioning: Schedule

Deployment Instance Details

* Deployment Instance DBM_CREATE_CLUSTER_DB_Machine_etc4m7_SYSMAN_10_28_2016_03_03_AM

Schedule

Start ● Immediately ○ Later (UTC-08:00) Los Angeles - Pacific Time (PT)

Notification

Status for Notification
☐ Scheduled
☐ Running
☑ Action Required
☑ Suspended
☐ Succeeded
☑ Problems

Prerequisite options

☐ Perform System Prerequisite checks

On failure ○ Ignore and continue. I have confirmed that we have a standardized environment and the failures are benign.
● Stop. I would like to review the prerequisite failures before deployment

☐ Perform Cluster Verification Utility checks

On failure ○ Ignore and continue. I have confirmed that we have a standardized environment and the failures are benign.
● Stop. I would like to review the prerequisite failures before deployment

7. Review and submit. Once submitted, the deployment procedure activity can be followed from the menu selecting **Enterprise**, then selecting **Provisioning and Patching**, finally clicking **Procedure Activity**.

Scale Up Cluster

1. Select a cluster and Database Domains

   Select a Database Machine target and once on the Database Machine Home Page, select **Database Machine**, then select **Provisioning** , then click **Scale Up Cluster**.

   Select the **Cluster** to extend, then select **Database Domains** on which to create a new DB zone.

   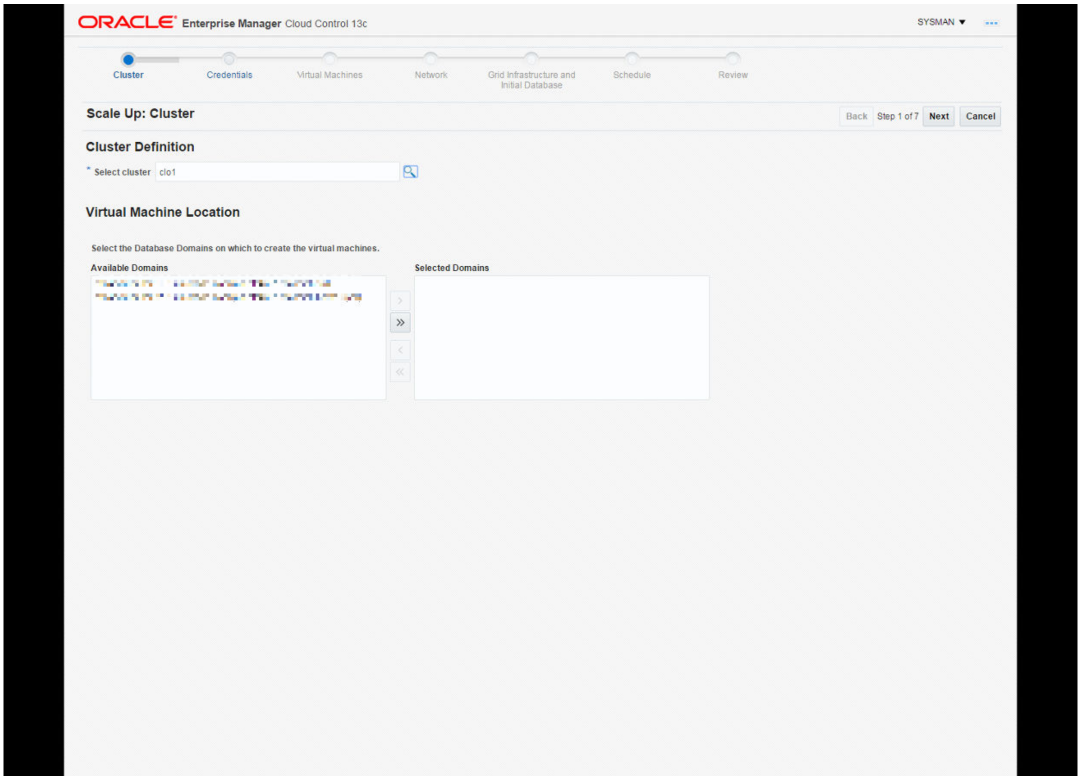

   Scale Up: Cluster

   Cluster Definition

   * Select cluster: clo1

   Virtual Machine Location

   Select the Database Domains on which to create the virtual machines.

   Available Domains: [redacted]

   Selected Domains

2. Provide the required credentials.

   Provide root credentials for Database Domains.

   Provide Solaris zone Host credentials for the Oracle Home (oracle user).

   Provide Solaris zone host root credentials.

   Provide Cluster ASM credentials (user sys, role sysdba).

   Provide Cluster Database Credential (user sys, role sysdba).

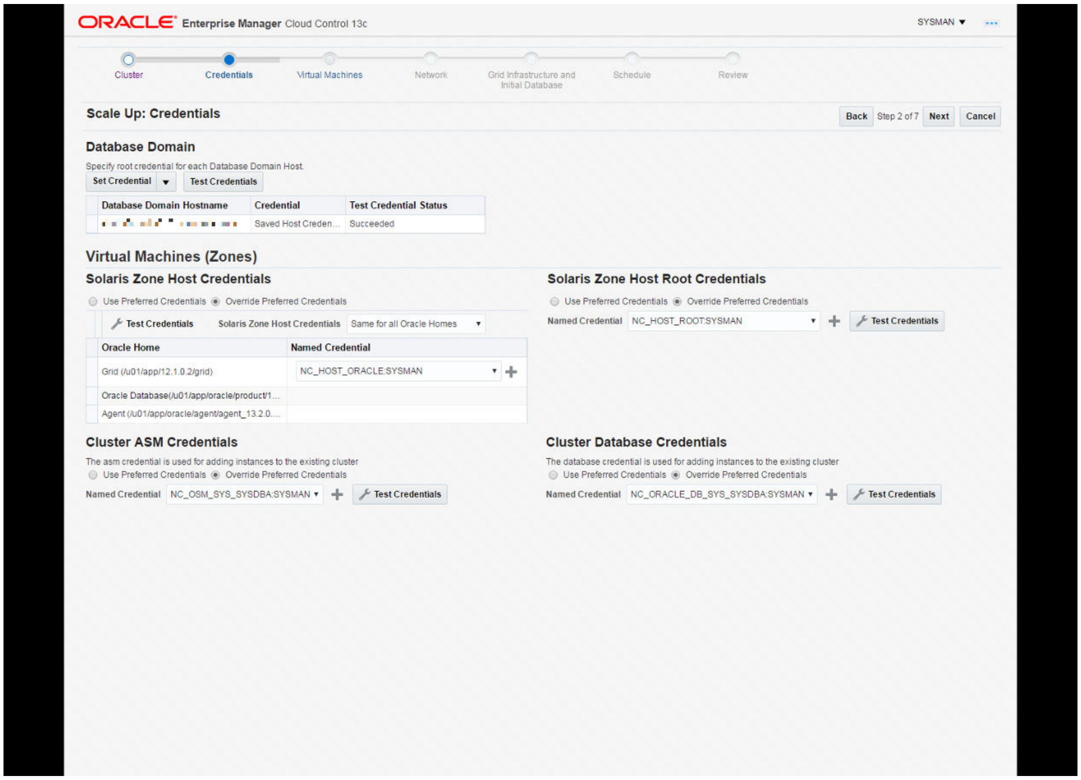

Scale Up: Credentials

Database Domain

Specify root credential for each Database Domain Host.

Database Domain Hostname	Credential	Test Credential Status
[redacted]	Saved Host Creden...	Succeeded

Virtual Machines (Zones)

Solaris Zone Host Credentials

○ Use Preferred Credentials ● Override Preferred Credentials

Solaris Zone Host Credentials: Same for all Oracle Homes

Oracle Home	Named Credential
Grid (/u01/app/12.1.0.2/grid)	NC_HOST_ORACLE:SYSMAN
Oracle Database(/u01/app/oracle/product/1...	
Agent (/u01/app/oracle/agent/agent_13.2.0....	

Solaris Zone Host Root Credentials

○ Use Preferred Credentials ● Override Preferred Credentials

Named Credential: NC_HOST_ROOT:SYSMAN

Cluster ASM Credentials

The asm credential is used for adding instances to the existing cluster

○ Use Preferred Credentials ● Override Preferred Credentials

Named Credential: NC_OSM_SYS_SYSDBA:SYSMAN

Cluster Database Credentials

The database credential is used for adding instances to the existing cluster

○ Use Preferred Credentials ● Override Preferred Credentials

Named Credential: NC_ORACLE_DB_SYS_SYSDBA:SYSMAN

3. Verify the info of the Virtual Machines Definition

   No input to provide here, values are taken from DB zones already in the cluster.

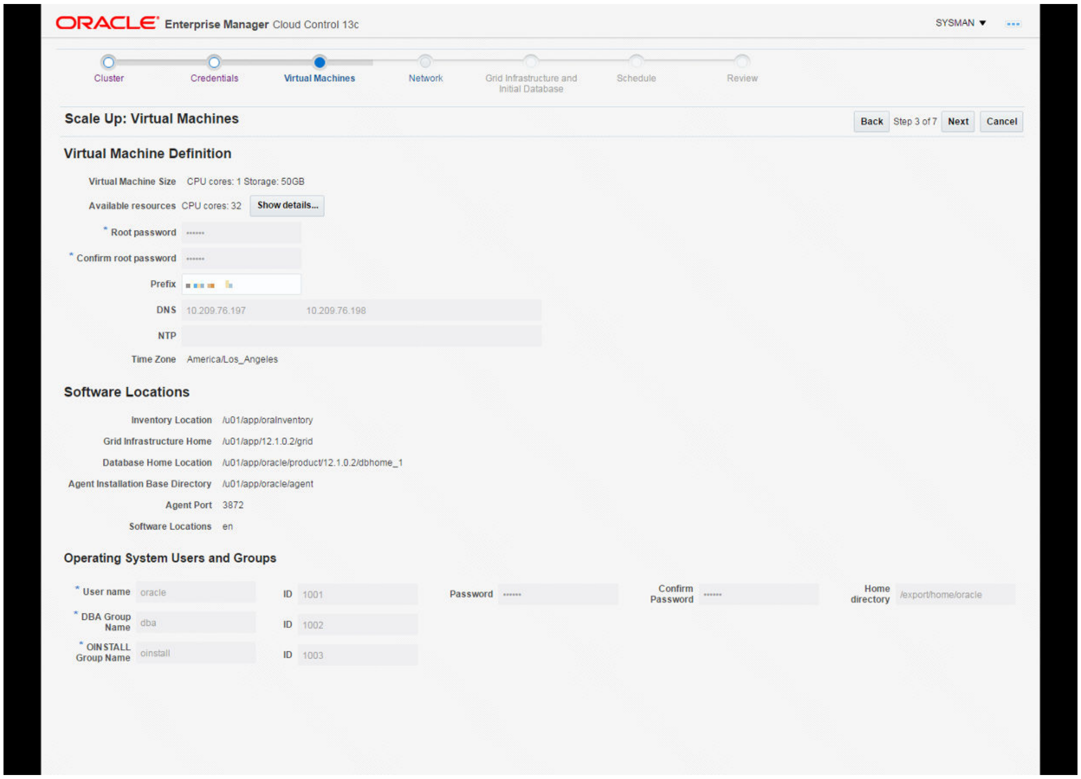

Scale Up: Virtual Machines

Virtual Machine Definition

Virtual Machine Size: CPU cores: 1 Storage: 50GB
Available resources: CPU cores: 32
* Root password: ••••••
* Confirm root password: ••••••
Prefix: [redacted]
DNS: 10.209.76.197 10.209.76.198
NTP:
Time Zone: America/Los_Angeles

Software Locations

Inventory Location: /u01/app/oraInventory
Grid Infrastructure Home: /u01/app/12.1.0.2/grid
Database Home Location: /u01/app/oracle/product/12.1.0.2/dbhome_1
Agent Installation Base Directory: /u01/app/oracle/agent
Agent Port: 3872
Software Locations: en

Operating System Users and Groups

* User name	oracle	ID	1001	Password	••••••	Confirm Password	••••••	Home directory	/export/home/oracle
* DBA Group Name	dba	ID	1002						
* OINSTALL Group Name	oinstall	ID	1003						

4. Define the Network

   Network Domains and Subnet masks are pre-populated.

   Provide Gateway IP addresses for Admin and Client Networks.

Provide IP addresses and hostnames for each network (admin network, client network, Virtual IP on client network, private network).

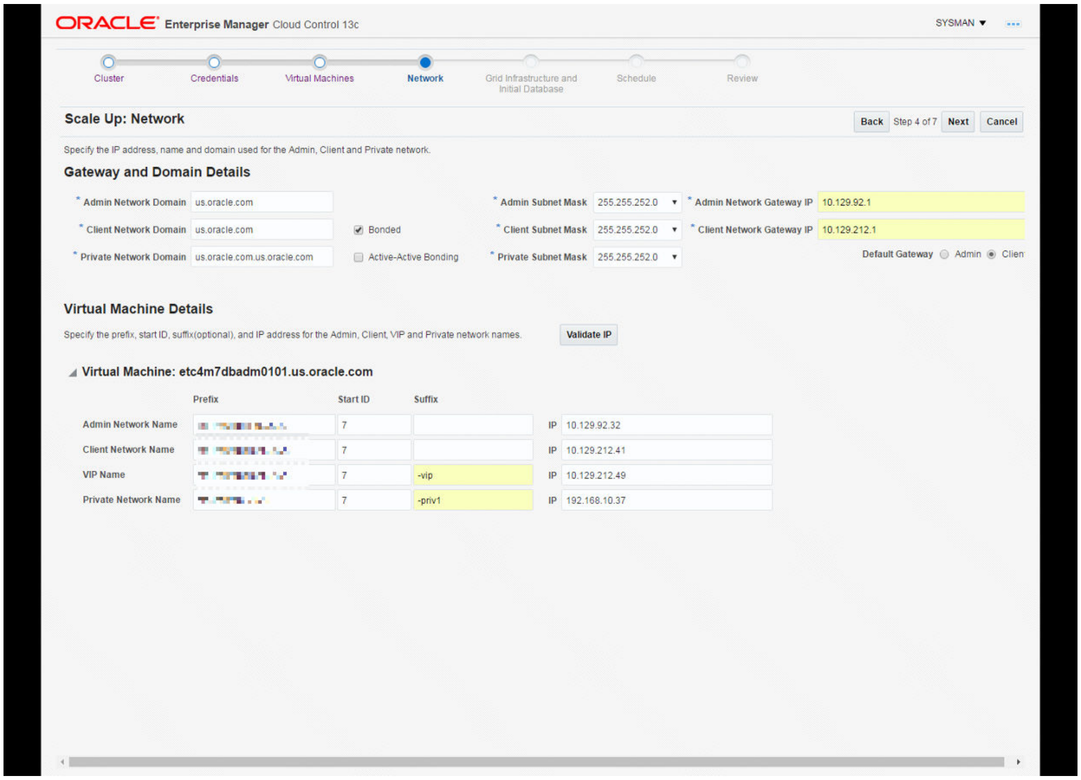

Scale Up: Network

Specify the IP address, name and domain used for the Admin, Client and Private network.

Gateway and Domain Details

* Admin Network Domain	us.oracle.com		* Admin Subnet Mask	255.255.252.0 ▾	* Admin Network Gateway IP	10.129.92.1
* Client Network Domain	us.oracle.com	☑ Bonded	* Client Subnet Mask	255.255.252.0 ▾	* Client Network Gateway IP	10.129.212.1
* Private Network Domain	us.oracle.com.us.oracle.com	☐ Active-Active Bonding	* Private Subnet Mask	255.255.252.0 ▾	Default Gateway	○ Admin ● Clien

Virtual Machine Details

Specify the prefix, start ID, suffix(optional), and IP address for the Admin, Client, VIP and Private network names.

◢ **Virtual Machine: etc4m7dbadm0101.us.oracle.com**

	Prefix	Start ID	Suffix		
Admin Network Name	[redacted]	7		IP	10.129.92.32
Client Network Name	[redacted]	7		IP	10.129.212.41
VIP Name	[redacted]	7	-vip	IP	10.129.212.49
Private Network Name	[redacted]	7	-priv1	IP	192.168.10.37

5. Deployment schedule

Schedule the Deployment, by default started immediately:

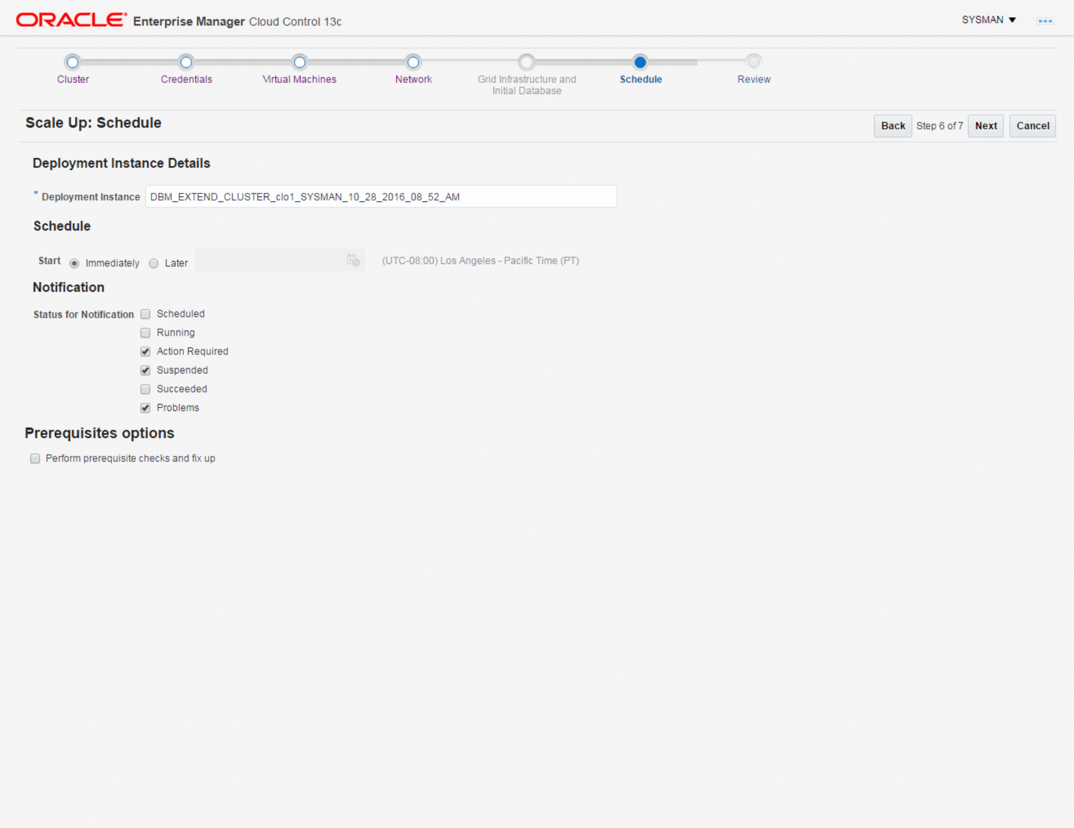

Scale Up: Schedule

Deployment Instance Details

* Deployment Instance: DBM_EXTEND_CLUSTER_clo1_SYSMAN_10_28_2016_08_52_AM

Schedule

Start ● Immediately ○ Later (UTC-08:00) Los Angeles - Pacific Time (PT)

Notification

Status for Notification
☐ Scheduled
☐ Running
☑ Action Required
☑ Suspended
☐ Succeeded
☑ Problems

Prerequisites options

☐ Perform prerequisite checks and fix up

Scale Down Cluster

1. Select the cluster and Solaris zone to delete

   From DB machine Home Page, select **Database Machine**, then selecting **Provisioning**, then click **Scale down cluster**.

   Select the Cluster to scale down. Once selected the list of DB zones in the Cluster appear.

   Select the Solaris Zones to delete

   Provide Solaris zone Host credentials for the Oracle Home (oracle user)

   Provide Solaris zone Host root credentials

   Provide Exadata Storage Server Credentials

   Provide Database Domains host root credentials

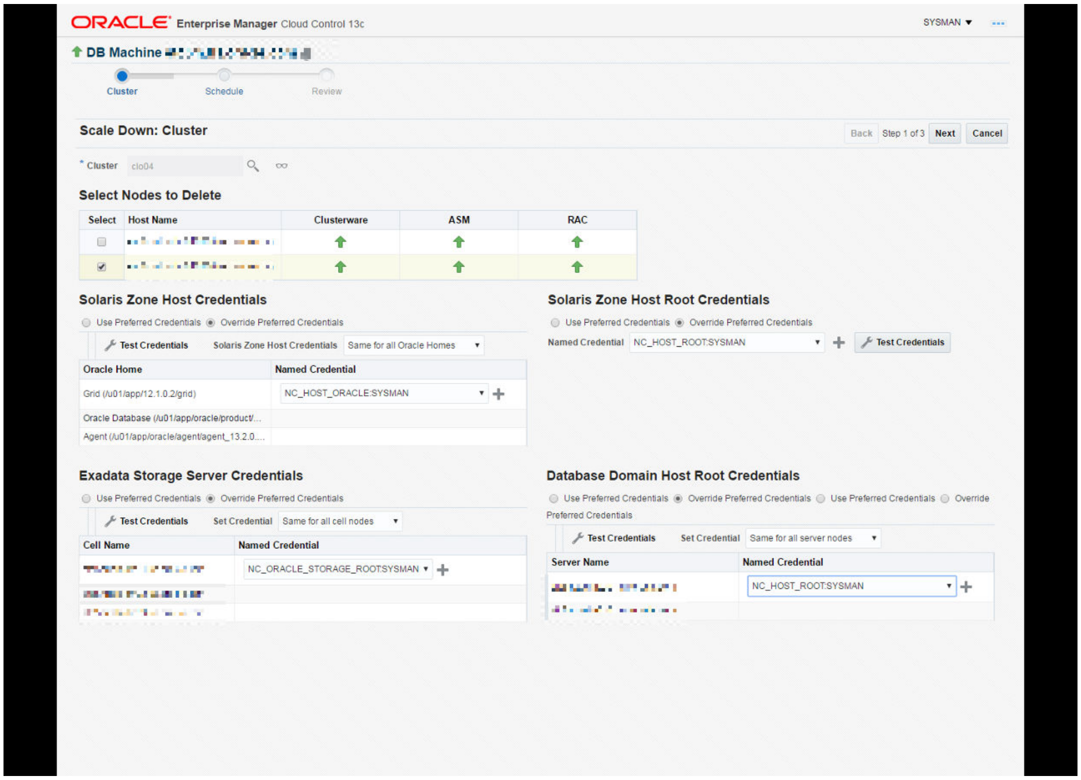

Scale Down: Cluster

* Cluster: clo04

Select Nodes to Delete

Select	Host Name	Clusterware	ASM	RAC
□	[redacted]	↑	↑	↑
☑	[redacted]	↑	↑	↑

Solaris Zone Host Credentials

○ Use Preferred Credentials ● Override Preferred Credentials

Solaris Zone Host Credentials: Same for all Oracle Homes ▾

Oracle Home	Named Credential
Grid (/u01/app/12.1.0.2/grid)	NC_HOST_ORACLE:SYSMAN ▾
Oracle Database (/u01/app/oracle/product/...	
Agent (/u01/app/oracle/agent/agent_13.2.0....	

Solaris Zone Host Root Credentials

○ Use Preferred Credentials ● Override Preferred Credentials

Named Credential: NC_HOST_ROOT:SYSMAN ▾

Exadata Storage Server Credentials

○ Use Preferred Credentials ● Override Preferred Credentials

Set Credential: Same for all cell nodes ▾

Cell Name	Named Credential
[redacted]	NC_ORACLE_STORAGE_ROOT:SYSMAN ▾
[redacted]	
[redacted]	

Database Domain Host Root Credentials

○ Use Preferred Credentials ● Override Preferred Credentials ○ Use Preferred Credentials ○ Override Preferred Credentials

Set Credential: Same for all server nodes ▾

Server Name	Named Credential
[redacted]	NC_HOST_ROOT:SYSMAN ▾
[redacted]	

2. Define a Deployment schedule

Schedule the Deployment, by default started immediately:

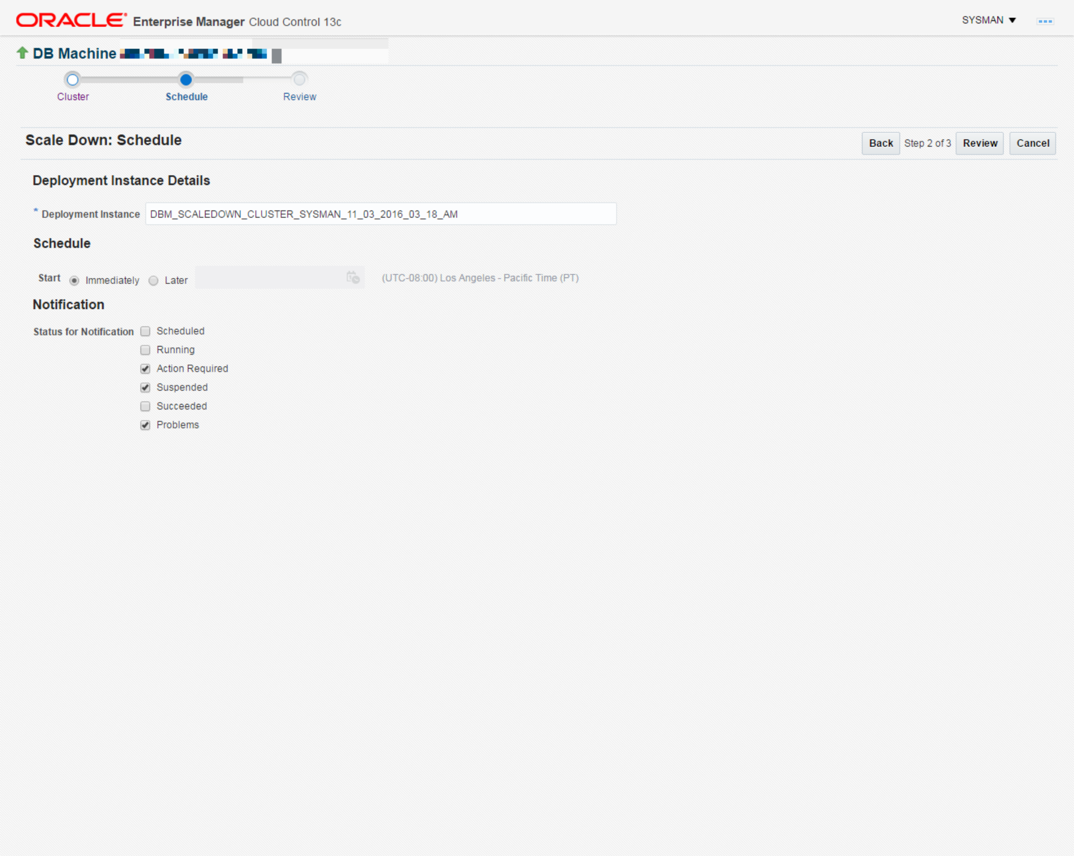

Scale Down: Schedule

Deployment Instance Details

* Deployment Instance: DBM_SCALEDOWN_CLUSTER_SYSMAN_11_03_2016_03_18_AM

Schedule

Start ● Immediately ○ Later (UTC-08:00) Los Angeles - Pacific Time (PT)

Notification

Status for Notification
- □ Scheduled
- □ Running
- ☑ Action Required
- ☑ Suspended
- □ Succeeded
- ☑ Problems

3. Review and submit

Delete Cluster

1. Select the cluster to delete

   From DB machine Home Page, select **Database Machine**, then select **Provisioning**, then click **Delete cluster**.

   Select the Cluster to delete. Once selected the list of DB zones in the Cluster appear.

   Provide Solaris zones Host credentials for the Oracle Home (oracle user).

   Provide Solaris zones Host root credentials.

   Provide Exadata Storage Server Credentials.

   Provide Database Domains host root credentials.

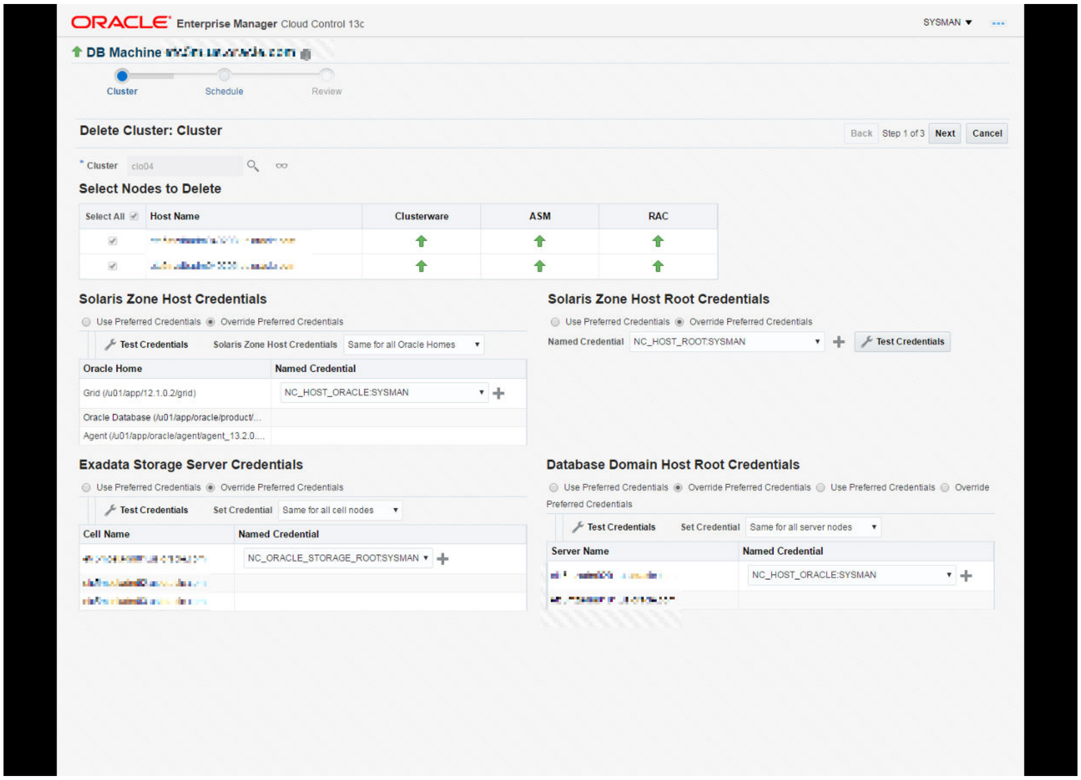

ORACLE Enterprise Manager Cloud Control 13c   SYSMAN ▾

DB Machine [redacted]

Cluster — Schedule — Review

Delete Cluster: Cluster   Back Step 1 of 3 Next Cancel

* Cluster: clo04

Select Nodes to Delete

Select All ☑	Host Name	Clusterware	ASM	RAC
☑	[redacted]	↑	↑	↑
☑	[redacted]	↑	↑	↑

Solaris Zone Host Credentials

○ Use Preferred Credentials ● Override Preferred Credentials

Test Credentials   Solaris Zone Host Credentials: Same for all Oracle Homes ▾

Oracle Home	Named Credential
Grid (/u01/app/12.1.0.2/grid)	NC_HOST_ORACLE:SYSMAN ▾ +
Oracle Database (/u01/app/oracle/product/...	
Agent (/u01/app/oracle/agent/agent_13.2.0....	

Solaris Zone Host Root Credentials

○ Use Preferred Credentials ● Override Preferred Credentials

Named Credential: NC_HOST_ROOT:SYSMAN ▾ +   Test Credentials

Exadata Storage Server Credentials

○ Use Preferred Credentials ● Override Preferred Credentials

Test Credentials   Set Credential: Same for all cell nodes ▾

Cell Name	Named Credential
[redacted]	NC_ORACLE_STORAGE_ROOT:SYSMAN ▾ +
[redacted]	
[redacted]	

Database Domain Host Root Credentials

○ Use Preferred Credentials ● Override Preferred Credentials ○ Use Preferred Credentials ○ Override Preferred Credentials

Test Credentials   Set Credential: Same for all server nodes ▾

Server Name	Named Credential
[redacted]	NC_HOST_ORACLE:SYSMAN ▾ +
[redacted]	

2. Define a Deployment schedule

Schedule the Deployment, by default it starts immediately:

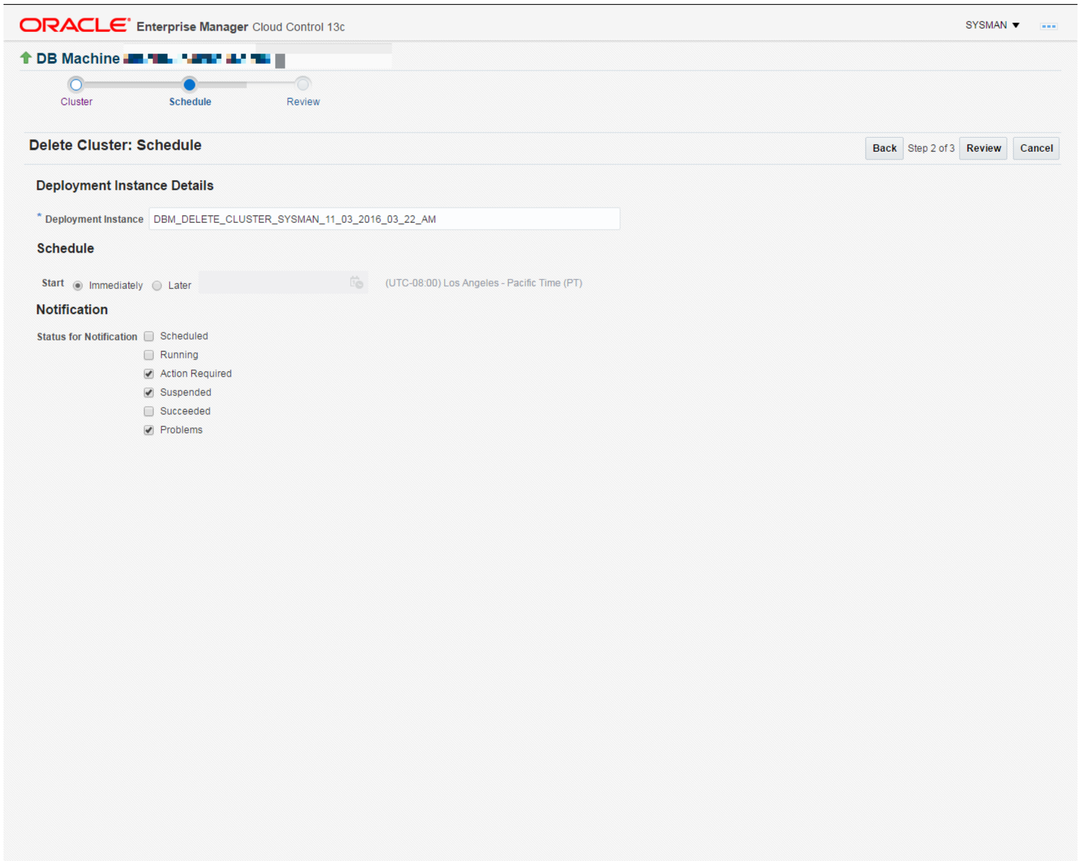

ORACLE Enterprise Manager Cloud Control 13c   SYSMAN ▾

DB Machine [redacted]

Cluster — Schedule — Review

Delete Cluster: Schedule   Back Step 2 of 3 Review Cancel

Deployment Instance Details

* Deployment Instance: DBM_DELETE_CLUSTER_SYSMAN_11_03_2016_03_22_AM

Schedule

Start: ● Immediately ○ Later   (UTC-08:00) Los Angeles - Pacific Time (PT)

Notification

Status for Notification:
- ☐ Scheduled
- ☐ Running
- ☑ Action Required
- ☑ Suspended
- ☐ Succeeded
- ☑ Problems

3. Review and submit

Part VI
Appendix

Overview of Target Availability States

The following sections summarize available states and how to set real-time target status updates.

Target Availability State Changes

Enterprise Manager displays a comprehensive array of target availability statuses in the form of informational icons. Various Cloud Control console pages display these icons to indicate the current status of targets in the repository.

The following table contains all available target availability status icons and their meaning.

Icon	**Availability State**	**Description**
N/A	N/A	Target availability state does not apply.
⬇	Down	Target is down. The target may be unreachable due to the fact that the Agent is down. If the Agent was brought down as part of planned maintenance, consider creating a blackout on the Agent.
⬆	Up	Target is up.
⊗	Availability Evaluation Error	An error occurred while attempting to determine target availability status. A target availability evaluation error can be caused by metric collection errors, the Agent being unreachable, or network problems.
⚙	Agent Down	The Agent monitoring the target is down. If an Agent was brought down in error it should be restarted. If Agent was brought down as part of planned maintenance, consider creating a blackout on the Agent.

Icon	Availability State	Description
	Agent Down, Target Up	The Agent monitoring the target is down, however, the target is currently up but not monitored. To troubleshoot, go to the Agent homepage and run the *Symptom Analysis* tool located next to the Status field.
	Agent Unreachable	The Agent is not reachable. Specifically, the Oracle Management Service (OMS) cannot communicate with the Agent. An Agent is generally unreachable when it is down, when it is blocked by the OMS, or when the Management Agent host is down. A Management Agent may also be unreachable due to network problems or certain other issues.
	Agent Unreachable (Under Migration)	The Agent is unreachable because it is in the process of being migrated.
	Agent Unreachable (Cannot Write to File System)	The Agent cannot write to the file system. Check the Agent file system for accessibility. To troubleshoot problems, navigate to the Agent home page from the Enterprise Manager console and run the *Symptom Analysis* tool (located next to the *Status* field).
	Agent Unreachable (Collections Disabled)	Agent metric collection has been disabled. Check that the Agent can upload to the OMS. To troubleshoot problems, navigate to the Agent home page from the Enterprise Manager console and run the *Symptom Analysis* tool (located next to the Status field).
	Agent Unreachable (Disk Full)	The Agent file system is full. Check the Agent file system for available space. To troubleshoot problems, navigate to the Agent home page from the Enterprise Manager console and run the *Symptom Analysis* tool (located next to the Status field).
	Agent Unreachable (Post Blackout)	The Agent is unreachable because the first alert condition has not yet occurred since the blackout period ended.
	Agent Blocked (Blocked Manually)	The Agent has been blocked manually. Unblock the Agent.
	Agent Blocked (Plug-in Mismatch)	The Agent has been blocked due to a plug-in mismatch. If the Agent has been restored from a backup, perform an Agent Resync.

Icon	Availability State	Description
N	Agent Blocked (Bounce Counter Mismatch)	The Agent has been blocked due to Bounce Counter mismatch. If the Agent has been restored from a backup, perform an Agent Resync.
⚙✖	Agent Unreachable (Agent Misconfigured)	The Agent is configured for communication with a different OMS. Check the Agent configuration to ensure the Agent is communicating with the correct OMS.
⚙	Agent Unreachable (Communication Broken)	The Agent is unreachable due to a communication break between the Agent and the OMS.
🕒↓	Blackout	The target is currently blacked out.
🕒	Status Pending	The target status is currently unknown.
🕒+	Status Pending (Target Addition in Progress)	The target status is currently unknown. Target addition is in progress.
🕒↓	Status Pending (Post Blackout)	The target status is currently unknown. Blackout has recently ended on this target and *Availability Status* is pending.
🕒✖	Status Pending (Post Metric Error)	A metric error has recently ended on the target and Availability Status is pending. To troubleshoot, refer to My Oracle Support article Enterprise Manager 12c: How to run the "Targets Status Diagnostics Report" to Troubleshoot Target Status Availability Issues (up, down, metric collection error, pending, unreachable) for all Targets (Doc ID 1546575.1).

Target Status Change Updates

Enterprise Manager can automatically update target information for specific target context UI pages without having to refresh the browser page or wait for the status change to be detected by the Response metric, where the collection interval delay may take anywhere from a few tenths of a second to a few minutes.

A *target context* page displays information about a particular target. It has a context header at the top showing information such as target name, target type, target status, or target menu.

Status change updates are available for the following target types:

- Agent
- Host
- Database Instance (Single Instance Database Only)
- Application Deployment
- WebLogic Server

As mentioned earlier, this feature allows target context pages to be updated automatically when that target's status changes (from *up* to *down*, for example). By default, automatic status change update is off. You can toggle this feature on and off using the *oracle.sysman.core.uifwk.realTimeUIEnabled* OMS property.

To enable status change updates, run the following emctl command:

```
emctl set property -name oracle.sysman.core.uifwk.realTimeUIEnabled -value true
```

Timeout Values for Enterprise Manager Components

Table 1 describes the timeout values for Enterprise Manager components.

Table 1 Time Out Values for Enterprise Manager Components

Component	Description	Timeout Value (in minutes)	Command

Apache timeout	Number of seconds that an Apache session is kept active.
	If Apache timeout is set beyond the operating system TCP timeout, it will cause unpredictable results. The operating system timeout is set to 2 hours by default.
		5 mins by default	Run the following command:
			`$ omsvfy show tcp`
			`Parameters`
			`Incoming`
			`Value`
			`--------------------`
			`tcp_keepalive_time`
			`7200`
			`tcp_keepalive_intvl`
			`75`
			`tcp_fin_timeout`
			`60`
			`--------------------`

Table 1 (Cont.) Time Out Values for Enterprise Manager Components

Component	Description	Timeout Value (in minutes)	Command
OMS timeout or Login timeout	This is the `oracle.sysman.eml.maxInactiveTime` parameter that can be set per OMS. To prevent unauthorized access to the Cloud Control, Enterprise Manager will automatically log you out of Cloud Control when there is no activity for a predefined period of time. For example, if you leave your browser open and leave your office. This default behavior prevents unauthorized users from using your Enterprise Manager administrator account. If you make changes to the login timeout value, be sure to consider the security implications of leaving your session open for other than the default timeout period. **Note:** The default timeout value does not apply when you restart the Web server or the OMS. In both of those cases, you will be asked to log in to the Cloud Control Console, regardless of the default timeout value.	45 min by default	Run the following command: `emctl set property -name oracle.sysman.eml.maxInactiveTime -value time_in_minutes -module emoms` Then, restart OMS for the value to take effect.
ADF timeout	This is controlled by the variable `oracle.adf.view.rich.poll.timeout`. The variable applies to pages that have auto poll. ADF pages may be enabled with automatic poll. After a page does not receive any keyboard or mouse event for duration of `oracle.adf.view.rich.poll.timeout` variable, then the poll stops. From that point on, the page participates in the standard server-side session timeout.	10 min	None

Executing SQL via REST API

Enterprise Manager has a rich set of monitoring data collected in its repository that can be extracted via REST API. You can use your own SQL scripts and Enterprise Manager's REST endpoints to extract repository data.

Enterprise Manager repository data can be extracted and used for a variety of purposes such as building custom dashboards, or Key Performance Indicator (KPI) reports. You can easily extract repository information via SQL by using Enterprise Manager's HTTP-based REST endpoints.

In addition to extracting data from the Enterprise Manager repository, the REST API also allows you to use some of the REST endpoints to extract data from any database target that is monitored/managed by Enterprise Manager.

With Enterprise Manager you can run a *SQLScript* job against a database target to automate data extraction. This job type requires both host and database credentials for job execution. For situations where the use of host credentials are not permitted, you can run an *Execute SQL* job, which requires only database credentials.

Refer to the following tables for REST endpoints:

- Table 2: REST Endpoints for Repository Operations
- Table 3: REST Endpoints for Target Database Operations

REST Endpoints Accessibility

Because the repository is a critical component of the Enterprise Manager framework, specific protections must be implemented to ensure that the repository database is secure.

Repository-related REST endpoints are protected by the following:

- The OMS property **oracle.sysman.db.restfulapi.executesql.repository.query.enable** must be set to *true* using `emctl`.
  Example:

  ```
  emctl set property -name
  oracle.sysman.db.restfulapi.executesql.repository.query.enable -value true
  -sysman_pwd “sysman”
  ```

- **Authorization Header:** Enterprise Manager User Credentials need to be passed as part of header.

For specific repository-related REST operations:

In addition to the above endpoint protection settings, you also need to set the following:

- Set the **oracle.sysman.db.restfulapi.executesql.repository.update.enable** OMS property to *true* using `emctl` for `/repository/update` REST method invocation on the repository database.

  Example:

  ```
  emctl set property -name
  oracle.sysman.db.restfulapi.executesql.repository.update.enable -value
  true -sysman_pwd “sysman”
  ```

- Set the **oracle.sysman.db.restfulapi.executesql.repository.plsql.enable** OMS property to *true* using `emctl` for `/repository/plsql` method invocation on the repository database.

  ```
  emctl set property -name
  oracle.sysman.db.restfulapi.executesql.repository.plsql.enable -value true
  -sysman_pwd “sysman”
  ```

- The Enterprise Manager user running a SELECT query REST operation (i.e., /repository/query REST) must be granted any out-of-box Enterprise Manager roles (in addition to the

EM_USER role), otherwise Fine Grained Auditing (FGA) will restrict the result of the SQL query to *no rows*.

Database target-related REST endpoints are protected by the following:

In addition to the aforementioned roles, the Enterprise Manager user should also have the following Target Privileges:

- **Connect target**
- **Run any sql on Database**

The **oracle.sysman.db.restfulapi.executesql.target.query.enable** OMS property must be set to *true* using `emctl` to enable REST operations on the database target.

```
emctl set property -name
oracle.sysman.db.restfulapi.executesql.target.query.enable -value true -
sysman_pwd "sysman"
```

For specific database target-related REST operations:

In addition to the above settings, the following OMS properties must be set to *true*.

- The **oracle.sysman.db.restfulapi.executesql.target.update.enable** OMS property must be set to *true* using `emctl` for `/target/update` REST method invocation on the target database:
  Example:

  ```
  emctl set property -name
  oracle.sysman.db.restfulapi.executesql.target.update.enable -value true -
  sysman_pwd "sysman"
  ```

- The **oracle.sysman.db.restfulapi.executesql.target.plsql.enable** OMS property must be set to *true* using `emctl` for `/target/plsql` REST method invocation on the target database:
  Example:

  ```
  emctl set property -name
  oracle.sysman.db.restfulapi.executesql.target.plsql.enable -value true -
  sysman_pwd "sysman"
  ```

Query Result Limitation

To prevent the repository/target database from being overloaded, flood control mechanisms that limit the number of returned rows and columns by REST SQL queries have been implemented via the following OMS properties:

- **oracle.sysman.db.httpsql.numrows** : If no value has been specified, then by default 1000 rows will be returned in the resultset.
- **oracle.sysman.db.httpsql.numcols** : If no value has been specified, then by default only 20 columns will be returned in the resultset.

REST Endpoints

The following tables list available Enterprise Manager repository and target database REST endpoints.

Table 2 REST Endpoints for Repository Operations

REST Endpoint	Sample Payload	HTTP Method	Comments	Sample Output
https://<EM_HOST>:<PORT>/em/websvcs/restful/emws/oracle.sysman.db/**executesql/repository/query/**v1	{ "sqlStatement": "**SELECT * FROM sysman.MGMT$TARGET_METRIC_SETTINGS**", "maxRowLimit": **2**, "maxColumnLimit": **4** }	POST	Executes the given SELECT query on the repository DB. **maxRowLimit and maxColumnLimit are optional.**	{"Result" : [{"target_name":"Management_Servers","target_type":"oracle_emsvrs_sys","target_guid":"[B@2f99ae59","metric_name":"Response"},{"target_name":"\/EMGC_EMGC_DOMAIN\/EMGC_DOMAIN\/EMGC_ADMINSERVER\/mds-owsm","target_type":"metadata_repository","target_guid":"[B@12856d79","metric_name":"MDS_REPOSITORY_ROLLUP"}]}
https://<EM_HOST>:<PORT>/em/websvcs/restful/emws/oracle.sysman.db/**executesql/repository/plsql/**v1	**Example #1: executing a PL/SQL block which does not have any data to return.** { "sqlStatement": " begin execute immediate 'create table test_sql_t1(col1 number(10))'; end; " }	POST	Executes the given PL/SQL block on the repository DB.	{ "Result":"PLSQL block executed successfully" }

Table 2 (Cont.) REST Endpoints for Repository Operations

REST Endpoint	Sample Payload	HTTP Method	Comments	Sample Output
.	**Example #2: executing a PL/SQL block which have multiple outputs.** `{ "sqlStatement": " DECLARE v_target_guid RAW(16); v_status NUMBER; v_sub_status NUMBER; BEGIN SELECT target_guid INTO v_target_guid FROM sysman.EM_TARGETS WHERE target_type = 'oracle_database' AND target_name = 'Oemrep_Database'; SYSMAN.MGMT_AVAIL.get_avail_state(v_target_guid,?,?); SYSMAN.EMD_MAINT_UTIL.get_em_db_sessions_cursor(?); END; ", "sqlParameters": [` **{"type":"INTEGER","isOutparameter":true}, {"type":"INTEGER","isOutparameter":true}, {"type":"CURSOR","isOutparameter":true}** `], "maxRowLimit": 2, "maxColumnLimit": 2 }`	POST	Executes the given PL/SQL block on the target DB using the DB user ID and password.	`{"Result":[{"OutParameter3":[{"INST_ID":1,"SID":10},{"INST_ID":1,"SID":11}]},{"OutParameter1":1},{"OutParameter2":99}]}`

Table 2   (Cont.) REST Endpoints for Repository Operations

REST Endpoint	Sample Payload	HTTP Method	Comments	Sample Output
.	**Example #3: excuting PL/SQL blocks having both input and output parameters.** { "sqlStatement": " DECLARE BEGIN SYSMAN.EM_TARGET.GET_AGENT_VERSION_FOR_TARGET(?,?,?); END; ", "sqlParameters": [**{"type":"STRING","value":"Oemrep_Database"}** , **{"type":"STRING","value":"oracle_database"}, {"type":"STRING","isOutparameter":true}**] }	POST	Executes the given PLSQL block on the target and returns results, if any.	{"Result":[{"OutParameter3":"13.4.0.0.0"}]}
https://<EM_HOST>:<PORT>/em/websvcs/restful/emws/oracle.sysman.db/**executesql/repository/update/**v1	{ "sqlStatement": "CREATE TABLE test_SQL_T3 AS (SELECT * FROM sysman.ADP_EVENT_J2EE where 1<>1)" }	POST	Executes the given DML query on the repository DB.	{ "Result":0 }

Table 3   REST Endpoints for Target Database Operations

REST Endpoint	Sample Payload	HTTP Method	Comments	Sample Output
`https://<EM_HOST>:<PORT>/em/websvcs/restful/emws/oracle.sysman.db/`**`executesql/target/query/`**`v1`	`{"targetName":"Oemrep_Database", "targetType":"oracle_database", "sqlStatement":"SELECT * FROM sysman.MGMT$TARGET_METRIC_SETTINGS", "credential":{"`**`DBCredsMonitoring`**`":"`**`testcred`**`"}, "maxRowLimit":3, "maxColumnLimit": 2}`	POST	Executes the given SELECT query on a target DB using named DB credentials referred by property *DBCredsMonitoring* in the payload. *maxRowLimit* and *maxColumnLimit* are optional. If *maxRowLimit*/ *maxColumnLimit* value is -1, then all rows/columns will be returned in the result set.	`{"Result" : [{"target_name":"Management_Servers", "target_type":"oracle_emsvrs_sys"}, {"target_name":"\/EMGC_EMGC_DOMAIN\/EMGC_DOMAIN\/EMGC_ADMINSERVER\/mds-owsm", "target_type":"metadata_repository"}, {"target_name":"\/EMGC_EMGC_DOMAIN\/EMGC_DOMAIN\/EMGC_ADMINSERVER\/mds-owsm", "target_type":"metadata_repository"}]}`

Table 3 (Cont.) REST Endpoints for Target Database Operations

REST Endpoint	Sample Payload	HTTP Method	Comments	Sample Output
`https://<EM_HOST>:<PORT>/em/websvcs/restful/emws/oracle.sysman.db/`**`executesql/target/query/`**`v1`	`{"targetName":"Oemrep_Database", "targetType":"oracle_database", "sqlStatement":"SELECT * FROM sysman.MGMT$TARGET_METRIC_SETTINGS where target_name=? and warning_operator=?", "sqlParameters": [ {"type":"STRING","value":"Management_Servers"}, {"type":"INTEGER","value":"1"} ], "credential": {` **`"DBCredsMonitoring":"testcred"`** `}, "maxRowLimit": 2, "maxColumnLimit": 2 }` `sqlParameters` is required only when the `sqlStatement` contains a question mark "?". Type can be one of the following • STRING • INTEGER • DATE • BYTE • BOOLEAN • CURSOR (only in case of PL/SQL kind of statements)	POST	Executes the given SELECT query on a target DB using specified named DB credentials (refer to property *DBCredsMonitoring* in the payload). **maxRowLimit and maxColumnLimit are optional.**	`{"Result" : [{"target_name":"Management_Servers","target_type":"oracle_emsvrs_sys"}]}`

Table 3 (Cont.) REST Endpoints for Target Database Operations

REST Endpoint	Sample Payload	HTTP Method	Comments	Sample Output
https://<EM_HOST>:<PORT>/em/websvcs/restful/emws/oracle.sysman.db/**executesql/target/plsql/**v1	**Example #1: executing a PL/SQL block which does not have any data to return.** {"targetName":"Oemrep_Database", "targetType":"oracle_database", "sqlStatement":" begin execute immediate 'create table test_sql_t1(col1 number(10))'; end; ", "credential": { "**DBCredsMonitoring**":"**testcred**" } }	POST	Executes the given PL/SQL block on a target DB using specifed named DB credentials (refer to property *DBCredsMonitoring* in the payload).	{""Result"": "PLSQL Block executed successfully"}

Table 3 (Cont.) REST Endpoints for Target Database Operations

REST Endpoint	Sample Payload	HTTP Method	Comments	Sample Output
.	**Example #2: executing a PL/SQL block which have multiple outputs.** `{"targetName":"Oemrep_Database", "targetType": "oracle_database", "sqlStatement": " DECLARE v_target_guid RAW(16); v_status NUMBER; v_sub_status NUMBER; BEGIN SELECT target_guid INTO v_target_guid FROM sysman.EM_TARGETS WHERE target_type = 'oracle_database' AND target_name = 'Oemrep_Database'; SYSMAN.MGMT_AVAIL.get_avail_state(v_target_guid,?,?); SYSMAN.EMD_MAINT_UTIL.get_em_db_sessions_cursor(?); END; ", "sqlParameters":` **`[{"type":"INTEGER","isOutparameter":true}, {"type":"INTEGER","isOutparameter":true}, {"type":"CURSOR","isOutparameter":true}],`** `"credential": {"DBCredsMonito`	POST	Executes the given PL/SQL block on a target DB using specified named DB credentials (refer property *DBCredsMonitoring* in the payload). Executes the given SELECT query on a target DB using named DB credentials referred by property *DBCredsMonitoring* in the payload.	`{"Result": [{"OutParameter3": [{"INST_ID":1,"SID":10}, {"INST_ID":1,"SID":11}]}, {"OutParameter1":1}, {"OutParameter2":99}]}`

Table 3 (Cont.) REST Endpoints for Target Database Operations

REST Endpoint	Sample Payload	HTTP Method	Comments	Sample Output
	`ring":"testcred"}, "maxRowLimit": 2, "maxColumnLimit": 2 }`			
.	**Example #3: excuting PL/SQL blocks having both input and output parameters.** `{"targetName":"Oemrep_Database", "targetType": "oracle_database", "sqlStatement": " DECLARE BEGIN SYSMAN.EM_TARGET.GET_AGENT_VERSION_FOR_TARGET(?,?,?); END; ", "credential": {"DBCredsMonitoring":"testcred"}, "sqlParameters": [` **`{"type":"STRING","value":"Oemrep_Database"}, {"type":"STRING","value":"oracle_database"}, {"type":"STRING","isOutparameter":true}`** `] }`	POST	Executes the given PL/SQL block on a target DB using specified named DB credentials (refer to property *DBCredsMonitoring* in the payload) and returns results, if any.	`{"Result": [{"OutParameter3":"13.4.0.0.0"}]}`

Table 3 (Cont.) REST Endpoints for Target Database Operations

REST Endpoint	Sample Payload	HTTP Method	Comments	Sample Output
`https://<EM_HOST>:<PORT>/em/websvcs/restful/emws/oracle.sysman.db/`**`executesql/target/update/`**`v1`	`{"targetName":"Oemrep_Database", "targetType":"oracle_database", "sqlStatement":"DELETE FROM mytable where empId<2000 and empId > 1998", "credential":{ `**`"DBCredsMonitoring":"testcred"`**` } }`	POST	Executes the given DML query on a target DB using specified named DB credentials (refer to the property *DBCredsMonitoring* in the payload). For most of the DML statement, the result will have a number of rows affected by that DML.	`{ "Result": 2 }`

Note:

Only Named Credentials are accepted in the payload for *target database* endpoints. Ensure that Enterprise Manager administrators executing the REST endpoint operations have already created a valid named credential for the database target specified in the payload.

Managing Database Target Credentials

Password lifecycle management plays a crucial role in maintaining a secure database environment. This typically involves changing the password for a user and then updating all Enterprise Manager configurations that use this password for monitoring or managing a database.

Enterprise Manager lets you use the job system to automate database password change for both monitoring and non-monitoring database users.

For more information about automated password management, see the following:

- Automate Monitoring User Password Management
- Automate Non-monitoring User Password Management

Index