ORACLE

Strategies for Scalable, Smarter Monitoring using Oracle Enterprise Manager Cloud Control 13c

PURPOSE STATEMENT

This document provides an overview of features and enhancements included in Oracle Enterprise Manager 13c. It is intended solely to help you assess the business benefits of upgrading to Oracle Enterprise Manager 13c and to plan your I.T. projects.

DISCLAIMER

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle.

Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

TABLE OF CONTENTS

PRODUCT OVERVIEW

Oracle Enterprise Manager is Oracle's integrated enterprise IT management product line and provides the industry's first complete cloud lifecycle management solution. Oracle Enterprise Manager's Business-Driven IT Management capabilities allow you to quickly set up, manage and support enterprise clouds and traditional Oracle IT environments from applications to disk. Enterprise Manager allows customers to achieve:

- *Best service levels for traditional and cloud applications* through management from a business perspective
- *Maximum return on IT management* investment through the best solutions for intelligent management of the Oracle stack and engineered systems
- *Unmatched customer support experience* through real-time integration of Oracle's knowledge base with each customer environment

INTRODUCTION

Whether you're supporting enterprise clouds or traditional IT applications, the need for the proactive and complete monitoring of your business applications and their underlying IT infrastructure, on-premises or in the cloud, continues to be a critical requirement for any datacenter. Effective monitoring of today's rapidly changing datacenter requires a management tool that can scale dynamically as the enterprise grows and IT staff that can use the tool in conjunction with best practice standards and processes. Enterprise monitoring has always been built into Enterprise Manager's DNA since its initial release. Over time it has evolved to work seamlessly with new Oracle technologies and meet the ever-changing requirements of IT staff who use it to manage the Oracle footprint in their datacenters. In the years that we've worked with customers in their Enterprise Manager deployments, we've sought to understand their product requirements, processes and strategies used to monitor their own datacenters. This white paper is a consolidation of these monitoring best practice strategies used in conjunction with product capabilities in Enterprise Manager 13c.   These strategies are meant to provide high level guidance in using Enterprise Manager's monitoring features to:

- Meet your datacenter's monitoring requirements
- Comply with security best practices
- Provide a solution that is easy to set up and manage
- Provide a solution that scales as your enterprise grows

These monitoring strategies are divided into three sections:

- Laying the Groundwork for Monitoring
- Managing Incidents
- Additional Monitoring Requirements and Recommendations

LAYING THE GROUNDWORK FOR MONITORING

Setting up your enterprise for monitoring requires a combination of planning and implementation. In order to manage and monitor any large datacenter to scale while meeting SLAs and other business-related requirements, it is important to have some uniformity in the way managed entities (called targets) are monitored and standard IT management procedures to handle events and incidents raised on those targets.   At a minimum, there are three major areas to address when it comes to setting up target monitoring:

- Defining and setting up a standard set of monitoring settings for your targets

- Defining and granting the appropriate level of privileges in Enterprise Manager to administrators who are responsible for managing these targets
- Setting up rules to automate IT operational processes such as sending email notifications for events, opening helpdesk tickets, escalating long running events, etc.

For each area, the IT staff has to determine the specific requirements for their datacenter based on business needs. For example, they will need to determine the appropriate set of metrics and thresholds to monitor their production servers in order to meet SLAs.   They will need to determine which administrators are responsible for operations such as defining monitoring settings, responding to events, etc. and then grant them the appropriate privileges in Enterprise Manager in support of those operations.   It is important to take time to plan and define these in advance before anything is implemented in Enterprise Manager. Assuming such planning has already taken place, this paper will recommend strategies for implementing these plans in Enterprise Manager in a scalable way.   This means you will only need to set it up once, and as your enterprise grows and more targets and/or administrators are added, the additional setup needed to accommodate the growth will be kept to a minimum because your monitoring setup will be automatically leveraged. Each of these areas will now be discussed in more detail.

Define and Set Up Standards for Monitoring

This first phase involves defining how your targets should be monitored in Enterprise Manager. Several steps are involved and each of these steps is discussed in detail below.

Organize Targets Based on How They are Monitored

Not all targets are alike. Some support mission critical applications, some support test environments and others support development environments. Based on these different usage profiles, most datacenters are likely to have different sets of monitoring settings (i.e. metrics, thresholds, collection schedules, corrective actions) for each of these different usage profiles. (The details of specific metrics to monitor, the threshold values to use, etc. are outside the scope of this paper but should be determined by the appropriate administrators in your IT organization).   For example, mission critical targets might be monitored more comprehensively for availability, performance and space usage while test targets might only be monitored for basic availability. It is first important to identify these different monitoring profiles in your datacenter and then plan on how you can group your targets based on these monitoring profiles.

As a simple example, you might have this type of grouping:

Figure 1. *Monitored targets are divided into Production targets and Non-Production targets*

In Figure 1 above, the set of targets monitored in Enterprise Manager are divided into 2 groups: all production targets in one group because they are all monitored using production monitoring settings

and all non-production targets such as test and development targets are put in another group because they are all monitored using non-production settings.

Suppose there are additional monitoring requirements within the group of production targets. For example, the targets supporting the production applications in the Sales Line of Business might have, in addition to the production monitoring settings, additional metrics that need to be monitored. The targets supporting the applications in the Finance Line of Business might also have additional metrics and/or metric thresholds that are slightly stricter than the thresholds used for the other production targets. This scenario leads to this type of grouping:
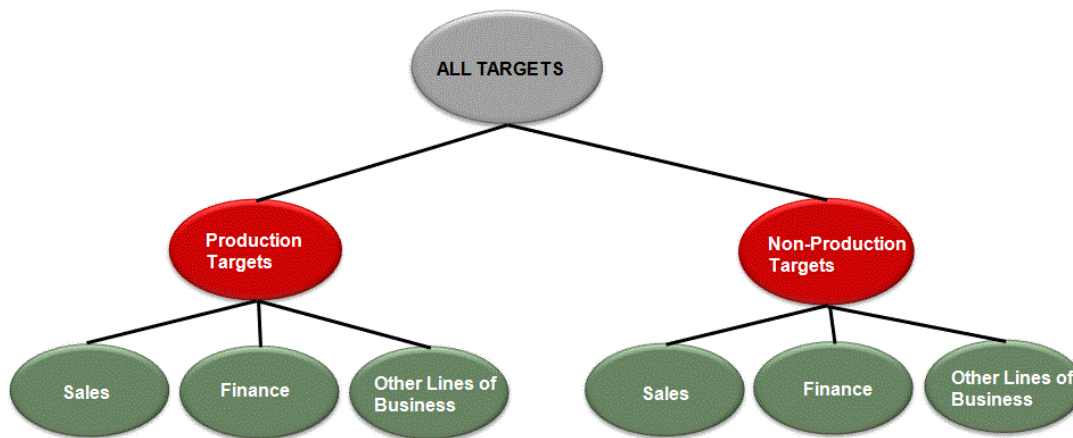

Figure 2. *Monitored targets further broken down by Line of Business*

In Figure 2 above, you will notice there is a further breakdown of the Production Targets group into different subgroups based on Lines-of-Business. The grouping at this level is again based on how targets are monitored, where targets monitored in the same way are put together in the same group. So, under the Production targets group, the Sales targets have additional monitoring settings, the Finance group has additional monitoring settings and all other Lines of Business (HCM, Manufacturing, and Marketing) that don't have these additional monitoring settings are put together in another group. A similar scenario can happen for the targets in the Non-Production group.

When designing your Administration Group hierarchy, it is important to remember that the primary goal here is to define the group hierarchy based on how targets are monitored. While you can re-use the same hierarchy for other purposes (e.g. reporting, job submission), designing the hierarchy with these other applications in mind might result in a hierarchy that is bigger and more complicated. You will need to weigh the benefits that such a hierarchy may provide versus the cost of managing any added complexity.

Once this grouping has been defined, it is important to understand how this group hierarchy is specified and implemented in Enterprise Manager. In Enterprise Manager, each node in the group hierarchy is a (target) group and each level in the group hierarchy is identified by a *target property*. (A *target property* is an attribute associated with all targets in Enterprise Manager and is used to annotate operational characteristics of the target, e.g. Line of Business, Owner, etc.). Within a level in the group hierarchy, specific values of the target property determine the membership criteria of groups at that level. In Figure 3 below, the target property *Lifecycle Status* is used to identify the two groups' membership criteria. The membership criterion for the Production Targets group is that the target's Lifecycle Status property should be 'Production.' The membership criteria of the Non-Production

Targets group (consisting of test targets, development targets or staging targets) is that the target's Lifecycle Status property should be Development or Test or Staging. A target cannot be added to these groups directly, rather, its target properties need to be set such that it matches the group's membership criteria. Once that happens, Enterprise Manager will automatically add the target to the appropriate group. For example, if you set a target's Lifecycle Status property to 'Production', it will automatically be added by Enterprise Manager to the Production Targets group. These types of groups are called Administration Groups and together they form an Administration Group hierarchy.
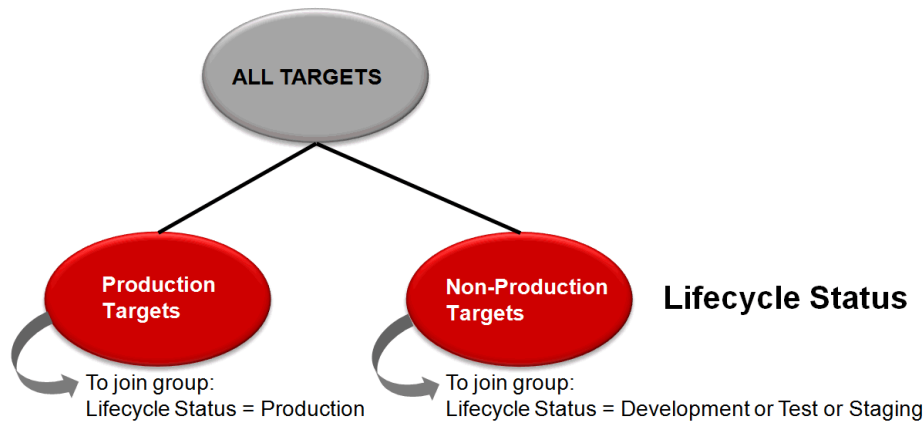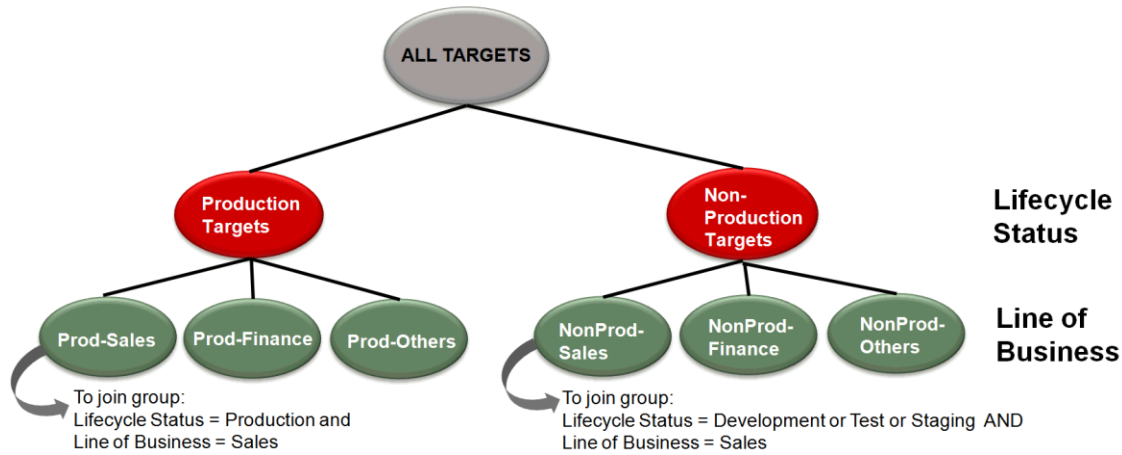

Figure 3. In this Administration Group hierarchy, the Lifecycle Status target property is used to define the membership criteria of the 2 groups.

Defining Your Administration Group Hierarchy

Administration Groups are a special type of group designed primarily to deploy monitoring settings to targets as they join the group[1]. (Details of how this occurs will be discussed later.) Defining your Administration Group hierarchy involves defining the target properties that make up each level of the hierarchy and the values of the target properties that determine the different groups at each level. In Figure 4 below, the first level of Administration Group is based on Lifecycle Status target property and the second level is based on Lines-of-Business. The membership criteria for the Prod-Sales group are *Lifecycle Status = Production* and *Line of Business = Sales*. The membership criteria for Prod-Finance are *Lifecycle Status = Production* and *Line of Business = Finance*. The membership criteria for Prod-Others is *Lifecycle Status = Production* and *Line of Business = HCM or Manufacturing or Marketing*. You will need to specify the exact list of values for each target property used as membership criteria. In Enterprise Manager, the creation of the Administration Group hierarchy involves choosing the target property that defines the level and for that level, specifying the values of the target properties that define each group for that level. For a complete list of target properties supported for Administration Groups and for more details on creating the Administration Group hierarchy, refer to the "Implementing Administration Groups and Template Collections" chapter of the *Oracle Enterprise*

---

[1] You can also use Administration Groups to automatically deploy Compliance Standards and Cloud Policies to targets as they join the group. However, that is outside the scope of this paper.

Manager Cloud Control Administrator's Guide 13c. Figure 4 below shows the creation of the Administration Group hierarchy in Enterprise Manager.
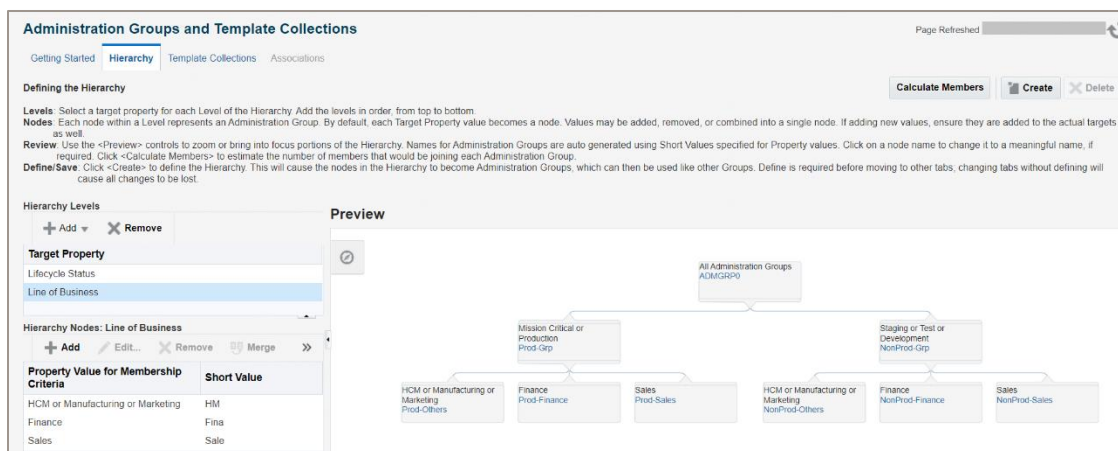

Figure 4. In this Administration Group hierarchy, the Lifecycle Status and Line of Business target properties define the membership criteria of the groups.

Regardless of the number of levels you define in your Administration Group hierarchy, it is important to note three points:

- Targets are always added to the leaf-level groups of the hierarchy.

- A target in the hierarchy can directly belong to at most one Administration Group. (It can directly belong to any number of regular, non-Administration Groups). As you'll see in the later sections of this paper, this is to prevent the potential ambiguity that can arise from different monitoring settings if a target is part of multiple Administration Groups associated with different monitoring templates.

- A target must match all membership criteria defined by the levels of the hierarchy in order to join an Administration Group. For example, in the Administration Group hierarchy in Figure 4, a target's properties must be set for both *Lifecycle Status* and *Line of Business*. In addition, the values set for these properties must match the membership criteria of one of the Administration Groups. If a target only has its *Lifecycle Status* property set, it will not join the Administration Group because its *Line of Business* has not been set. If a target has *Lifecycle Status=Production* and *Line of Business=IT,* then it will not join the Administration Group because there is no Administration Group with that membership criteria.

Figure 5. *When creating the Administration Group hierarchy, specify the target property for each level and the values of the target property that define each group in that level.*

Setting Target Property Values

As mentioned, for a target to join an Administration Group, its target properties need to be set such that it matches the membership criteria of the Administration Group.

You can set the value of a target's properties as part of its process of being added to Enterprise Manager. If the target is added via the console using either the manual target addition or target promotion workflow, there are steps in the workflow that enable the user to specify the target properties of the target that is being added.

Enterprise Manager administrator accounts also have attributes. These include Contact, Location, Department, Cost Center, and Line of Business. If you set values to any of these attributes, then any target added by the administrator will automatically have the same values set to its target properties. For example, if the administrator's Line of Business is set to Finance, any targets added by that administrator will have its Line of Business also set to Finance by default (this can always be overwritten). Hence if any of these properties are used as Administration Group criteria, to facilitate setting of target properties, consider setting the values at the administrator level and use that administrator account to add targets to Enterprise Manager.

If the target is added using Enterprise Manager Command Line Interface (EM CLI[2]) *add_target* verb and the attributes of the administrator who added the target was not set or did not match the Administration Group criteria, you can follow that with another verb *set_target_property_value* to specify the values of the target properties. In the following example, this EM CLI command sets the Lifecycle Status and Line of Business properties for a database target that has been added to Enterprise Manager:

```
$ emcli set_target_property_value
-property_records="MyDB:oracle_database:LifeCycle
Status:MissionCritical;MyDB:oracle_database:Line of Business:Finance"
```

---

[2] EM CLI is Enterprise Manager's command line utility that enables you access to Enterprise Manager functionality within scripts.

The EM CLI verb *set_target_property_value* is the recommended way to set target properties in bulk across many different targets at a time. For more details on the EM CLI verbs, refer to the *Oracle Enterprise Manager Command Line Interface* documentation.

To specify target properties of a target via the console, go to the target's Target Properties page which is accessible via Target menu → Target Setup → Properties. The Target Properties page that appears will allow you to set properties for the target. As a reference, it will also display the membership criteria of the Administration Group hierarchy.
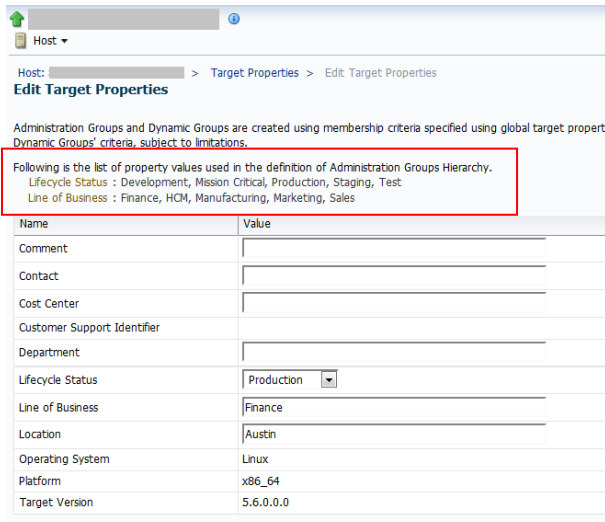
Figure 6. *The Edit Target Properties page shows you the criteria used for the Administration Group hierarchy.*

The All Targets page (accessible from Targets menu → All Targets) allows you to select additional columns for the All Targets table. The columns that you can add to the table include target properties.
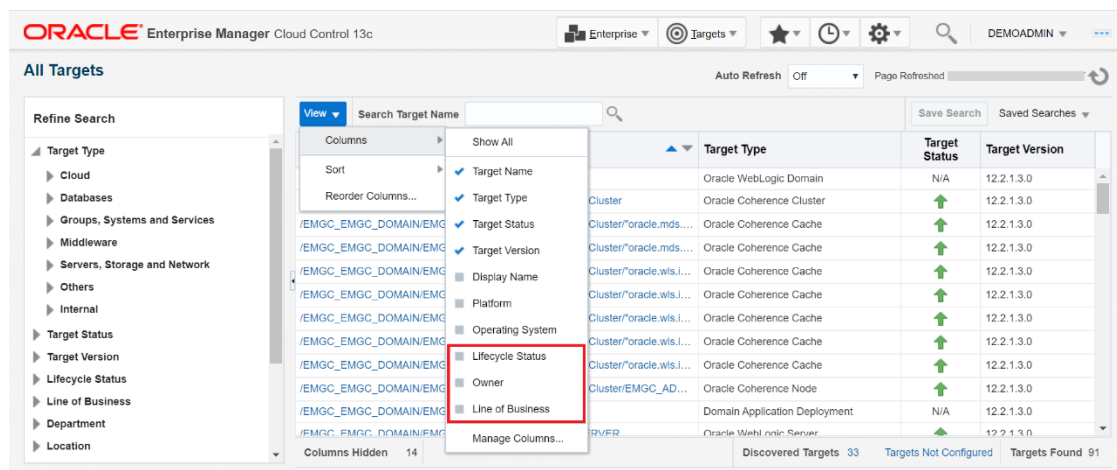

Figure 7. *The All Targets page allows you to select target properties as additional columns to be displayed in the targets table.*

If you customize the All Targets page by adding the target properties used as Administration Group criteria, you can easily see the target properties defined for the targets and also find out which targets do not have their target properties set. Note that a maximum of 2000 targets are shown in this page. If you're looking for a target and it's not in the list, you can use the Search Target Name field to filter the list to show targets matching the name specified. To set a target's properties, select the target from the table, right-mouse click to open up the target menu, and select Target Setup → Properties to access the Target Properties page.

One important note about setting target property values for aggregate targets, i.e. targets that contain other targets (e.g. Oracle WebLogic Domain target containing Oracle WebLogic Servers): if the intent is to set the target properties for the aggregate target itself as well as its member targets, then you will need to use the EM CLI verb *set_target_property_value* on the aggregate target and use the option *–propagate_to_members*. This will set the target property values on the aggregate target itself as well as all of its *current* member targets. Any new member targets that will be added in the future will NOT have its target properties automatically set; hence you will need to set the appropriate target property values for any members added in the future. The reason for not automatically setting the target properties for members of these aggregates is because these member targets could potentially be part of other aggregate targets with different values specified for their target properties. Hence, the administrator must decide and set the appropriate values of target properties for such targets.

The following EM CLI example sets the Location target property of a database system (aggregate target) and all its members using the *–propagate_to_members* option:

```
$ emcli set_target_property_value
-property_records="dbrac_sys:oracle_dbsys:Location:Bangalore"
-propagate_to_members
```

Cluster targets on the other hand are aggregate targets that have strong membership semantics, i.e. members of clusters can only belong to one cluster aggregate. Examples of these targets types are Redundancy System, Database Cluster, Host Cluster, etc. If you set a target property for a cluster target, then the same target property automatically applies to all members of the cluster target including any new members that will be added in the future.

Define Standards for Monitoring

The next step is to define the monitoring settings for the groups in the Administration Group hierarchy. For example, in our Administration Group hierarchy, you'll need to define the monitoring settings (metrics and associated thresholds) for the targets in the production group and the set of monitoring settings for the targets in the non-production group. While Enterprise Manager's out-of-box monitoring settings for targets might provide some initial guidance, IT organizations should give some thought in determining what's appropriate for their targets. Determining this appropriate set of metrics and thresholds for targets is outside the scope of this paper, but there some general guidelines provided in the "Additional Monitoring Requirements and Recommendations" section that is available in the latter part of this paper (refer to the subsection "Too Many Alerts"). Once you've defined these monitoring settings, create monitoring templates in Enterprise Manager containing these monitoring settings. (Monitoring templates are named collections of metric settings for a target type. These include metrics, their associated collection schedules, thresholds, and optionally, corrective actions).

Monitoring templates are defined per target type so create one monitoring template for each of the target types in your Administration Group. For example, if your Production Targets group contains hosts, databases and listeners, create three monitoring templates -- one for host targets, another for database targets and another for listener targets. See Figure 8 on the next page.

**Figure 8.** *Sample host monitoring template*

Once you've created the monitoring templates, combine these together into a container called a Template Collection.   A Template Collection is a collection of different monitoring templates designed to specifically contain the monitoring templates for an Administration Group. For our Administration Group in Figure 3, you will end up defining a Template Collection (containing monitoring templates) for the Production Targets group and another Template Collection for the Non-Production Targets group. Refer to the "Implementing Administration Groups and Template Collections" chapter of the *Oracle Enterprise Manager Cloud Control Administrator's Guide 13c* for specific steps in creating Template Collections.

Note: In all our diagrams, the top-level node of the Administration Group hierarchy is labeled "All Targets" because of the expectation that all targets in the Enterprise Manager site will be monitored using the Administration Group – Template Collection setup. However, it is possible that some datacenters may choose not to include all targets in the Administration Group hierarchy. If this is the case, then the top-level node really represents the group of all targets that are part of the Administration Group hierarchy and not all the targets in the Enterprise Manager site. It is a good practice to periodically check the Unassigned Targets Report (accessible from the Administration Group UI in the console) to ensure that no target that is meant to be part of the Administration Group hierarchy has been missed. Such targets could be missed because its target properties have not been set to match the membership criteria of any Administration Group.

Associate Template Collections with Administration Groups

After you've created the Template Collections, the next step is to associate these with the appropriate Administration Group. In Enterprise Manager, this is done by selecting the Administration Group, clicking the associate button and choosing the Template Collection (see Figure 9). Refer to the documentation ("Implementing Administration Groups and Template Collections" chapter of the Oracle Enterprise Manager Cloud Control Administrator's Guide 13c) for details on this step. Once this

association is done, then any target added to the Administration Group will automatically be applied with the monitoring settings from the associated Template Collection.
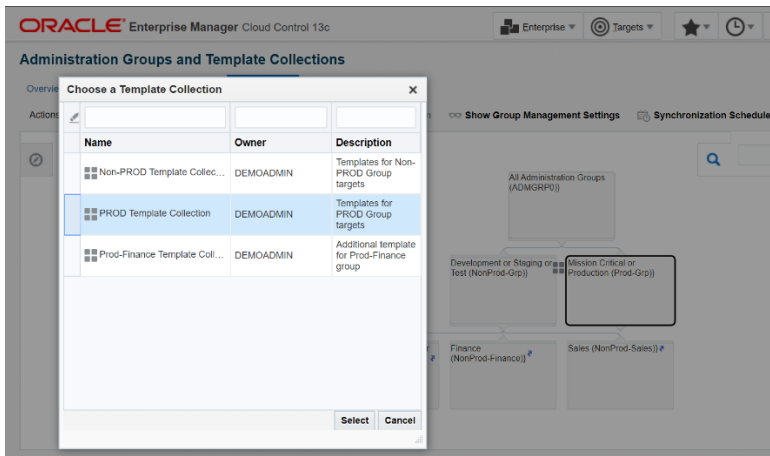
Figure 9. *Associate the PROD Template Collection with the PROD Administration Group.*

For the 3 level Administration Group hierarchy in Figure 4, all the groups under the Production Targets group start off with the production monitoring settings. Hence a Template Collection (e.g. Prod Template Collection) can be created containing monitoring templates with the production monitoring settings. This Template Collection should then be associated with the Production Targets group. All targets in subgroups Prod-Sales and Prod-Others will be applied with the monitoring settings. The Prod-Finance Group had monitoring settings required in addition to the production monitoring settings. Create monitoring template(s) containing these additional monitoring settings and put them in another Template Collection. Associate this new Template Collection with the Prod-Finance Group (see Figure 10 below). Targets in the Prod-Finance Group will be applied with a union of monitoring settings from the Prod Template Collection and Prod-Finance Template Collection. If there are any metrics in common between the two templates, the metric settings from Prod-Finance Template will take precedence.
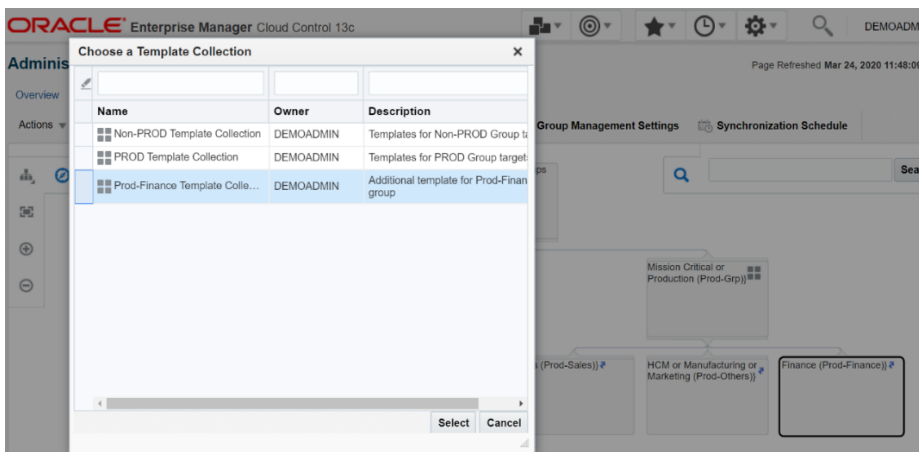
Figure 10. *Associate the PROD-Finance Template Collection with the Prod-Finance group. This template collection contains additional templates for the Prod-Finance group.*

To see the combined (final) set of monitoring settings that will be applied to targets in the Prod-Finance group (or any group), select the group and then click on the Show Group Management Settings option. This is especially useful in scenarios where monitoring settings to be applied to targets are based on multiple monitoring templates. In the subsequent page that appears, you can review the combined (aggregate) monitoring settings that will be applied to each target type in the group, check the synchronization status of members (if any) in the group and quickly go to any of its parent groups to perform similar actions. Refer to Figure 11 for more details.
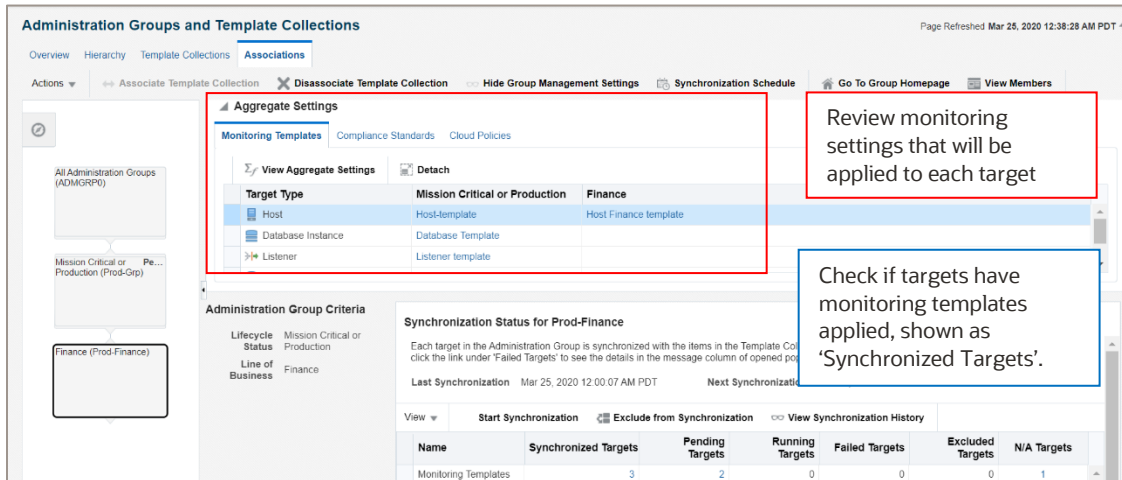
Figure 11. *Group Management Settings page enable you to review monitoring settings for the group and check the synchronization status of its members.*

Defining Monitoring Templates for Primary and Standby Databases

When monitoring primary and standby databases, there are a set of metrics and thresholds that apply to the database based on its current role – primary or standby. If there is a role change, i.e. the primary becomes the standby and the standby becomes the primary, you want to ensure that the applicable metrics and thresholds that pertain to the new standby and new primary are correspondingly changed. For example, the *Redo Generation Rate (KB/second)* metric applies to the primary database and the *Apply Lag (seconds)* metric applies to the standby database. Let's say you have DB1 (primary database) and DB2 (standby database) with the following metric settings:

- DB1 (primary database)
  - *Redo Generation Rate (KB/second)* metric: Warning = 25, Critical = 150 thresholds
- DB2 (standby database)
  - *Apply Lag (seconds)* metric: Warning = 7 and Critical = 17 thresholds

If there is a role change and DB1 becomes the standby and DB2 becomes the primary, you would like to have the following metric settings:

- DB1 (standby database)
  - *Apply Lag (seconds) metric*: Warning = 7 and Critical = 17 thresholds
- DB2 (primary database)
  - *Redo Generation Rate (KB/second)* metric: Warning = 25, Critical = 150 thresholds

To accomplish this, first make sure the database monitoring template contains the union of database metrics and thresholds that apply to both the primary and standby database. In our example above, make sure the monitoring template has:

- *Redo Generation Rate (KB/second)* metric: Warning = 25, Critical = 150 thresholds
- *Apply Lag (seconds)* metric: Warning = 7 and Critical = 17 thresholds

Set the target properties of both the primary and standby databases so that they join the same Administration Group. Add the monitoring template to the template collection that will be associated with the Administration Group (or parent Administration Group) to which the databases belong. When the template settings are initially applied, the appropriate metrics and thresholds will be applied to the database based on its role. In our example, DB1 (primary) will have *Redo Generation Rate (KB/second)* metric with the corresponding thresholds and DB2 (standby) will have *Apply Lag (seconds)* metrc with the corresponding thresholds. If there is a switchover later between the primary and standby, the corresponding database roles for the databases will be updated (i.e. DB1 becomes the standby and DB2 becomes the primary) and the associated monitoring template will be re-synced with the targets. This wil cause the *Redo Generation Rate (KB/second)* metric thresholds to be applied to the new primary (DB2) and the *Apply Lag (seconds)* metric thresholds to be applied to the new standby (DB1).

If primary and standby databases need to be in different Administration Groups, ensure that the monitoring templates associated with their respective Administration Groups contain a union of the database metrics that apply to primary and standby databases. The monitoring templates associated with these different Administration Groups can have different threshold values for the database metrics.[3] When the primary and standby databases join their respective Administration Groups, the appropriate metric threshold settings will be applied based on their current database roles. Later, when a switchover occurs, the database roles will be updated for the databases and the associated monitoring templates will be automatically re-synced (i.e. re-applied) with the appropriate metric thresholds based on their new database roles.

Set the Synchronization Schedule

The synchronization schedule determines when sync operations are executed by Enterprise Manager. Sync operations refer to the application of monitoring templates within a Template Collection to the relevant targets in the associated Administration Group. When a target is added to an Administration Group, then the associated monitoring template is automatically applied to the target. However, there are other conditions under which templates need to be applied to a target in an Administration Group:

- When a Template Collection containing monitoring templates is initially associated with an Administration Group that already has member targets or
- When there are changes made to a monitoring template that is part of a Template Collection associated with an Administration Group or
- When there are changes made to the monitoring settings of a specific target and the option to prevent 'Template Override' has not been specified

In all these cases, the targets in the Administration Group need to be made 'in sync' with the associated monitoring template by applying the monitoring template to the targets. These sync operations (i.e.

---

[3] In this scenario, if you are using different monitoring templates that contain different thresholds for the metrics, after a database role change and templates are re-synced, you will not see a 'swap' of thresholds between the primary and standby databases because they are separately associated with monitoring templates that have different threshold values for the metrics.

template apply operations) are not done right away but are scheduled based on the synchronization schedule.   At the day/time specified by the synchronization schedule, all pending sync operations are executed. Thus, you might want to think about specifying a synchronization schedule during off peak hours (e.g. Saturday at 11 pm), when there is little or no impact to production operations.   The synchronization schedule allows you to specify both a start date/time as well as a frequency in days.
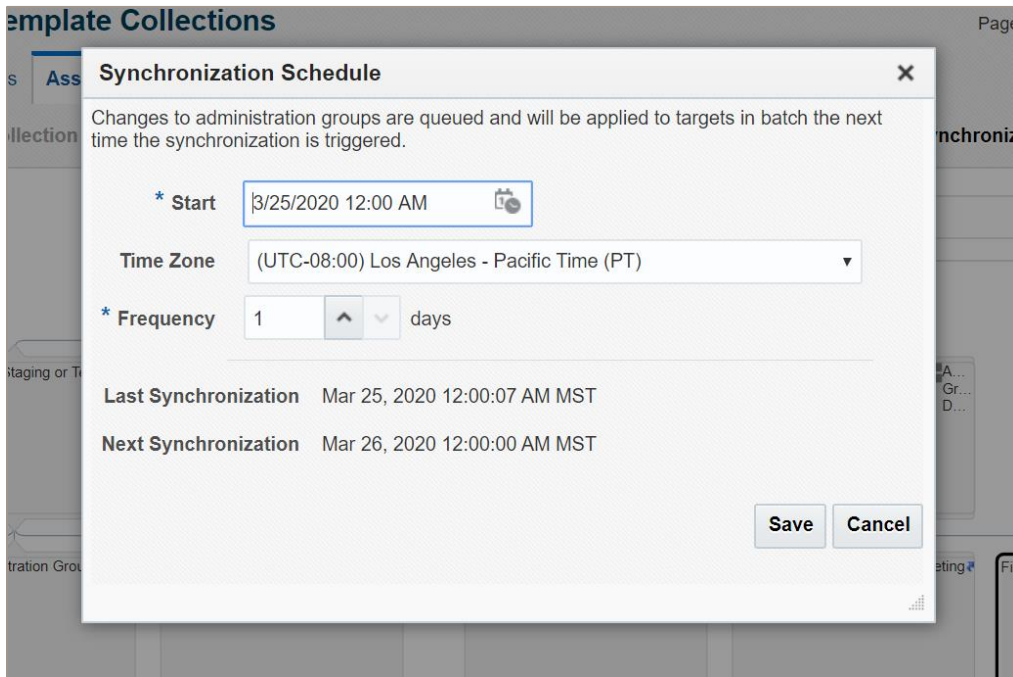
Figure 12. *The Synchronization Schedule determines when scheduled sync operations (i.e. template apply operations) will be performed by Enterprise Manager.*

For example, to specify a schedule where sync operations occur only on Saturdays, specify a start date that occurs on a Saturday and a frequency of 7 days.   Because the synchronization schedule impacts all sync operations, only an Enterprise Manager super administrator can specify or change the synchronization schedule.   Note that sync operations are scheduled only as needed, i.e. when the conditions previously described occur.   For more details on the synchronization schedule, refer to the Implementing Administration Groups and Template Collections chapter of the *Oracle Enterprise Manager Cloud Control Administrator's Guide 13c*.

Define Roles for Different Job Responsibilities

Plan Job Responsibilities (*Who* can do *what* operations?)

In order to manage the group, you will need to think about the different job responsibilities as it pertains to managing the group and its member targets. Here are some things to consider:

- Who can define group membership?
- Who can grant privileges on the group to other administrators?
- Who can do the following operations on the member targets:

- o Define and apply monitoring settings
- o Define and apply notification settings (e.g. who gets notified on events, etc.)
- o Determine which events should have incidents created for them
- o View and receive notifications for events/incidents
- o Acknowledge and work on incidents
- o Perform target blackouts for planned downtime activities

As mentioned earlier, there is a good amount of planning needed outside of Enterprise Manager to capture your datacenter requirements and determine how these can be implemented in Enterprise Manager. Planning discussions should also include defining job responsibilities.

For illustration purposes, here are some examples of some job responsibilities:

- Group Administrator
  - o Responsible for defining group membership and for granting privileges on groups to other administrators
- Senior Administrator
  - o Responsible for adding and removing targets in Enterprise Manager, and for planning and setting up monitoring settings for targets in Enterprise Manager. He is also responsible for setting up rules related to creating incidents for events and sending notifications based on the agreed upon plans.
- Target Owner
  - o For the targets he owns, he is responsible for setting monitoring settings, responding to events/incidents, and for performing maintenance operations
- First Level Support
  - o Responsible for responding to events/incidents on targets

After you've defined the different job responsibilities, you will next need to understand the privileges in Enterprise Manager required to support the various job responsibilities.

Understand Enterprise Manager Privileges to support job responsibilities

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.

Table 1. Enterprise Manager Privileges to Support Job Responsibilities

JOB RESPONSIBILITY	ENTERPRISE MANAGER PRIVILEGE
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)

JOB RESPONSIBILITY	ENTERPRISE MANAGER PRIVILEGE
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	• *Create Monitoring Template* resource privilege is required to create new monitoring templates. • If you are using monitoring templates created by another user, then you will need at least *View* privileges on those monitoring templates. • If you are using corrective actions, you will need the *Create Corrective Actons* resource privilege to create corrective actions or use any corrective action with your monitoring template.
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 user • *View Any Template Collection* to view/associate any Template Collection • *Full 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 (e.g. grant to other users) • Add a target to an Administration Group by setting its target properties • Perform a manual sync of the group with the associated Template Collection	• *Group Administration* on the group • *Configure Target* (on the target to be added to the Administration Group) • *Manage Template Collection Operations* on the group

JOB RESPONSIBILITY	ENTERPRISE MANAGER PRIVILEGE

Operations on the members of the Administration Group
 	
• Delete the target from Enterprise Manager
 	
• 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
	• *Full* on the target (Full also contains the privileges enumerated below)
	• *Operator* on the target also contains all the privileges enumerated below
	   o *Blackout Target* on the target
	   o *Manage Target Metrics* on the target
	   o *Configure Target* on the target
	   o *Manage Target Events* on the target
	   o *View* on the target

Create rules for managing events and incidents on targets	• *Create Enterprise Rule Set* resource privilege • *Manage Target Events* on the target on which the rule set operates

For convenience, there is an out-of-box role called EM_TC_DESIGNER that contains the necessary privileges required for creating Template Collections. Based on the job responsibilities you've defined, you will then need to map the corresponding Enterprise Manager privileges needed to support the various job responsibilities.

Use roles to manage user privileges

Privileges are ultimately granted to administrators to enable them to manage targets in Enterprise Manager. While you can grant specific privileges to individual administrators, tracking and granting privileges on many targets across many administrators easily becomes error-prone and an administrative burden. Our recommendation is to define and use roles to manage the granting of privileges to administrators.   A role is a user-defined set of privileges typically containing the set of privileges that you want to grant to a team of users. A role can contain other roles as well. For example, you can create a First Line Support role containing the privileges needed for the administrators to view and manage incidents on targets. Once this role is created, you can grant this role to the appropriate administrators who will manage these incidents as part of their job responsibility.   If you need to change the set of privileges for your administrators, e.g. add new privileges or remove privileges, then all you need to do is update the role.   The updated set of privileges in the role is automatically enabled for the administrators to whom the role has been granted. Likewise, if new administrators are added, all you need to do is grant them the appropriate role(s) instead of granting them individual privileges.

Using roles is one big step towards managing privileges. However, there is still the challenge of having to keep the role updated with privileges on new targets as they are added to Enterprise Manager. Privilege-propagating groups are meant to address this challenge and will be discussed next.

Leverage the privilege-propagating nature of Administration Groups

Administration Groups are privilege-propagating in nature. This means that a privilege on the Administration Group that is granted to a user or a role automatically applies (i.e. "propagates") to all members of the group including any subgroups.   If a new target is added to an Administration Group, then because the Administration Group is privilege-propagating, the user or role that has privileges on the Administration Group automatically gets privileges on the newly added target by virtue of it joining the group. No additional work is needed for granting privileges on the new target. Thus granting target privileges is much simpler because all you need to do is a one-time setup of granting privileges on the group to a role.

 In Figure 13 below, we have an Administration Group hierarchy. A role called "Senior Administrator Role" has been granted *Full* on the PROD Administration Group. This means any user that has been granted this role has *Full* privileges not only on the PROD group itself but all subgroups and all the members of the subgroups Sales, Finance, Others under the PROD group. Similarly, a (lesser) privilege *Manage Target Events* on the PROD group has been granted to the "First Line Support" role. This means any administrator that has been granted this role has *Manage Target Events* on PROD group and all its subgroups.   In the future, if a new target is added to the Sales, Finance, Others subgroups under PROD group, the user who has been granted "Senior Administrator Role" will automatically get *Full* privileges on the newly added target and the user who has been granted "First Line Support" role will automatically get *Manage Target Events* on the newly added target.   Thus, roles and privilege propagating groups provide benefits in economies of scale when it comes to managing a large number of target privileges for a large number of users.
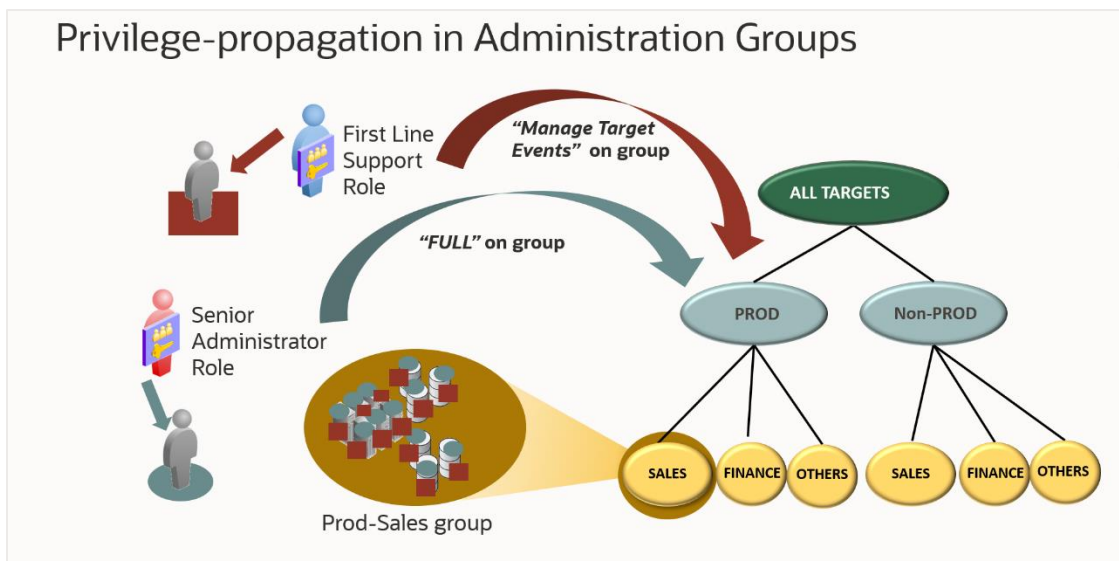
Figure 13. *This shows the privilege-propagating nature of Administration Groups.*

Create Roles for different job responsibilities

After you've planned the various job responsibilities and mapped these to the corresponding privileges in Enterprise Manager, the next step is to create roles in Enterprise Manager containing privileges required for each job responsibility. In our example below, here are the various roles that need to be

created for each job responsibility. Note that when it comes to privileges on targets in the Administration Group, the recommendation is to grant the privilege on the Administration Group and not on individual targets in order to leverage the privilege propagating nature of Administration Groups:

Table 2. EXAMPLES OF ROLES YOU CAN CREATE FOR DIFFERENT JOB RESPONSIBILITES*

JOB RESPONSIBILITY	ROLE IN ENTERPRISE MANAGER	*PRIVILEGES IN THE ROLE (MINIMUM SET)
Group Administrator Responsible for defining group membership and for granting privileges on the group to other administrators	GROUP_ADMIN_ROLE	*Group Administration* on the group
Senior Administrator Responsible for adding and removing targets in Enterprise Manager, and for planning and setting up monitoring settings for targets. He is also responsible for setting up rules related to creating incidents for events and sending notifications.	SENIOR_ADMIN_ROLE	• *Add Any Target* • *Create Enterprise Rule Set* • *Operator* on the group • *Create on Job System* • EM_TC_DESIGNER role
Target Owner For the targets he owns, he is responsible for setting monitoring settings, responding to events/incidents, and for performing maintenance operations	TARGET_OWNER_ROLE	• *Operator* on the Administration Group(s) that he is managing • *Create on Job System* • *View Any Monitoring Template* *View* on the Template Collection(s) associated with the group(s) he is managing
First Level Support Responsible for responding to events/incidents on targets. As part of operational procedures, he is allowed to blackout a target that is down.	FIRST_LEVEL_SUPPORT	• *Manage Target Events* on the appropriate Administration Group(s) • *Blackout Target* on the appropriate Administration Group(s)

* The privileges listed in Table 2 represent the minimum set of privileges in the role. Additional privileges can be added based on other responsibilities. Also note that you will need to have Super administrator privileges to create roles.

To create roles, log on as an Enterprise Manager super administrator. Go to Setup → Security → Roles. Select the Create button and follow the steps in the Create Role wizard to create the role with the necessary privileges.

There may be cases wherein the Administration Groups do not exactly fit the group requirements as far as managing privileges are concerned. For example, for an Administration Group, you might want to grant privileges for only a subset of the targets in the group. For recommendations on how to resolve this, refer to the section called *Using Administration Groups for Other Group Operations* in the latter part of this paper.

Assign Roles to your Enterprise Manager administrators

Once roles have been defined, you can now grant these roles to your Enterprise Manager administrators. This can be done in several ways:

- When creating/editing an Enterprise Manager administrator, you can assign role(s) as part of the "Create/Edit Administrator" wizard.

- As part of creating/editing a role, the "Create/Edit Role" wizard allows you to choose administrators to whom you would like to grant the role.

- When creating/editing administrators using the Enterprise Manager Command Line tool (EM CLI) create_user or modify_user, you can specify the roles granted to the user. Also EM CLI verb grant_roles also grants a role to a user (administrator).

# Set up Rules to Manage Events and Incidents

The last phase in setting up monitoring involves setting up rules in Enterprise Manager to automate the operational processes for managing events based on your requirements.   For example, one requirement might be to send page notifications for critical events and send email notifications for warning events.   Or another might be to open helpdesk tickets for target down events. Rule Sets enable the automation of these notification actions as well as other actions on events and incidents. This phase involves following:

- Understanding rule sets and rules

- Using groups in rule sets

- Auto-creating incidents using rules

- Planning your rules to leverage Enterprise Manager rule features

- Implementing the rules in Enterprise Manager

Each of these will be discussed in further detail. But before we get into these details, it is first important to get a basic understanding of events and incidents in Enterprise Manager.

Events and Incidents

An event is a significant occurrence, typically on a managed target, that has been detected and raised by Enterprise Manager. For example, a target that is down causes a 'target down' event to be raised. When disk usage is nearing its capacity, a 'disk full' event is raised. If a job fails, then a 'job failed' event is raised.   There are different types of events in Enterprise Manager[4] and sometimes you might see multiple events raised that pertain to the same underlying issue. For example, if a host target is under heavy load, then you could potentially see a CPU Utilization(%) event, Memory Utilization(%) event and Swap Utilization(%) event on the host all trigger within minutes of each other.   Administrators can get easily overwhelmed by the volume of events raised. Thus, it is desirable to be able to narrow down these set of events into the subset of actionable events that need to be addressed (because they impact your business applications) and also identify which of these pertain to the same issue and thus can be logically managed collectively as one unit. It is also desirable to be able to assign, prioritize and track the resolution of these important events. In Enterprise Manager 13c, we enable this through Incident Management features. The primarily goal of incident management is to be able to monitor and resolve disruptions to services as quickly and efficiently as possible. From the set of events raised, one should create incidents for significant, actionable events (e.g. target down event) or for a combination of significant related events that all pertain to the same issue (e.g. one incident for CPU Utilization(%), Memory Utilization(%), Swap Utilization(%) events on the same host). An incident is thus an object in Enterprise Manager containing either a significant actionable event or a combination of actionable related events that pertain to the same issue.
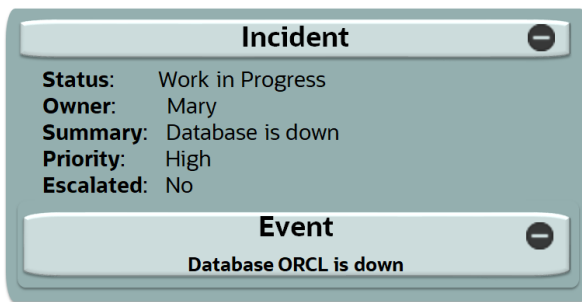
Figure 14. This is an example of an incident containing one event.

In Figure 14 above, the incident contains one target down event. In Figure 15 below, the incident contains multiple events. You can also combine multiple events into the same incident using the Incident Manager console or using event rules. The latter method will be discussed in the section *Minimizing Event Storms by Compressing Multiple Events into a Single Incident*. The severity of the incident is the worst-case severity of the events it contains.

---

[4] There are multiple event types supported. These include: Metric Alert, Target Availability, Job Status Change, Compliance Standard Score Violation, Service Level Agreement Alert, Metric Evaluation Error, JVM Diagnostics Threshold Violation, etc. For more details refer to the Incident Management chapter of the Oracle Enterprise Manager Cloud Control Administrator's Guide 13c.
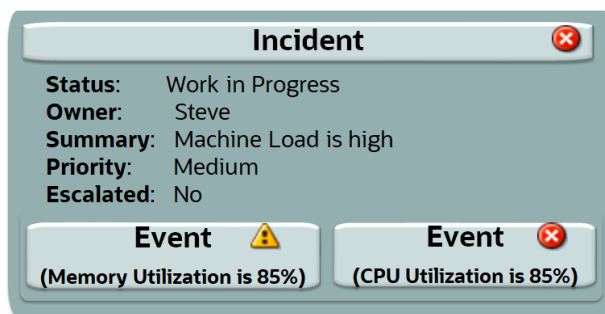
Figure 15. *This is an example of an incident containing two events.*

Once an incident is created, a rich set of incident lifecycle management features are available in Enterprise Manager to help manage the resolution of incidents. These include the sending of notifications, assigning ownership, acknowledging an incident, tracking its resolution status, prioritizing, escalating long running incidents, adding annotations, etc. Thus incidents (and not individual discrete events) are the primary way to monitor and manage events raised in Enterprise Manager. For more details on incidents refer to the "Using Incident Management" chapter of the *Oracle Enterprise Manager Cloud Control Administrator's Guide 13c*. Additional recommendations on managing incidents will be discussed in subsequent sections of this paper.

Understand Rule Sets and Rules

Automating actions on events and incidents in support of operational processes is an important part of any scalable monitoring solution. In Enterprise Manager, this is supported using rule sets and rules. It is important to get a good understanding of rule sets and rules before you can leverage them in your Enterprise Manager deployment. A rule is an instruction for Enterprise Manager to take action on an event or incident or problem. (A Problem is another type of object used in Enterprise Manager to manage critical errors in Oracle software. These will be discussed separately). Actions include sending of email notifications, opening of helpdesk tickets, creation of incidents for events, assignment of incidents, etc. A rule set is a set of rules that operate on a common object such as a group of targets. A rule consists of:

- Criteria: specified set of events or incidents on which the rule will operate. Examples are: all target down events on all targets in the group, specific metric alert (events) such as host Filesystem Space Available(%), database Tablespace Space Used (%), etc.

- Actions: one or more actions that Enterprise Manager should take on the specified set of events/incidents in the criteria. Actions in turn can be 'conditional', i.e. only execute the action if a specified condition is met. For example, your rule criteria could be "all incidents of fatal severity" and the action could be: if "the incident has been opened for more than 48 hours" (condition part), then set the escalation level to 2 (action part).

If you have several rules that apply to the same object (such as a group), then they should be combined in a rule set. Rules within rule sets are executed in a certain order. By default, they will be executed in the order in which they are defined, but they can be re-ordered as needed. Rule sets are also executed in a specified order. By default, they will be executed in the order in which they are defined, but they can also be re-ordered. Creation and use of rule sets are defined in detail in "Using Incident Management" chapter of the *Oracle Enterprise Manager Cloud Control Administrator's Guide 13c*. It is recommended that you refer to the documentation for details on rule set features. This paper will focus on how you can set up rule sets in support of your operational processes.

Use Groups in Rule Sets

For monitoring purposes, the best practice recommendation is to specify a group (or groups) of targets as the object of the rule set. In a hierarchy of groups, i.e. group containing other groups, specify the highest-level group for which the rules apply. Each rule in the rule set operates on the applicable member targets of the group including members of any subgroups. It is also recommended that any group you use for rule sets be included in only one rule set and not in multiple rule sets. Putting all the rules related to the same group into one rule set makes it much easier to track and manage all actions on the group because it's centrally defined in one place. The group you use in your rule set should contain targets that have common requirements for notification as well as common requirements for actions on events and incidents (such as assignment of incidents, etc.). As the group later expands and additional targets are added, the rule set will automatically apply to the newly added members without further modification of the rule set. If an Administration Group defined in the prior sections can serve the purpose of being used for rule sets (i.e. all members of the Administration Group have common requirements for notifications and event/incident management actions), then the Administration Group should be used as the target of the rule set.

Auto-create Incidents Using Rules

Since incidents are the recommended way of monitoring and managing events, it is recommended that rules be used to automatically create incidents for events that are important to be managed. In a rule set, the rules that create incidents should typically be the first set of rules in the rule set. There is an out-of-box rule set called *Incident management rule set for all targets* that automatically creates incidents for a subset of important events (such as target down availability events). However, it most likely needs to be adjusted based on your particular Enterprise Manager deployment needs. It is recommended that you review this out-of-box rule set to see if it meets your requirements. If it does not, you can do a 'create like' on the rule set and change the new rule set as needed. If you do so, remember to disable the original out-of-box rule set.

Minimizing Event Storms by Compressing Multiple Events into a Single Incident

When using rules to create incidents for events, you should consider using Event Compression Policies to compress or combine multiple related events into a single incident. Here are some examples of where you may use this feature:

- Compress or combine target down events for all RAC instances that belong to the same RAC into one incident

- Compress or combine configuration standard violation events into one incident

- Compress or combine all metric collection error events for a target into one incident

Event Compression Policies identify the conditions under which multiple correlated events are grouped together, or compressed, into one incident. You can use any Oracle-provided event compression policy or create your own. Once you've identified and enabled the policies to use, review your incident-creating event rules and make sure it has the option to "Use Event Compression Policies" when creating an incident. With this option in place, if an event occurs, Enterprise Manager will check to see if there is an applicable event compression policy that can be used to add the event to an incident instead of creating a new incident for each event.

For incidents with compressed events, the overall severity of the incident is the worst severity across all of the events it contains. For example, if the incident contains an event of critical severity and an event

of warning severity, the severity of the incident is 'critical'. If you need to send notifications for these events, send notifications for the incident, and not for the individual events. Using the compression feature in this way reduces possible 'event storms' when these events happen at about the same time, and enables you to work on these events in a more manageable way.

Plan Your Rules

Rules are meant to automate the operational processes for managing events and incidents. Before creating the rules in Enterprise Manager, it is important to take time to plan the structure of your rules.

For purposes of illustration, it is assumed that the set of rules that administrators would like to set up follow this type of pattern:

Rule requirements for the production group PROD:

1. For target down events, these should be treated as highest priority. Page the administrators. (Note: Target down events are of 'fatal severity'.)

2. For specific set of metric alert (events) on the databases in the group, send a notification based on the severity: if it is critical severity, page the administrators; if it is warning severity, email the administrators.

3. For specific set of metric alert (events) on hosts in the group, send a notification based on the severity: if it is critical severity, page the administrators; if it is warning severity, email the administrators.

4. For specific set of metric alert (events) on WebLogic Servers in the group, send a notification based on the severity: if it is critical severity, page the administrators; if it is warning severity, email the administrators.

5. If any job against the group fails, page the administrators. (Note: Job failure events are of 'critical' severity).

6. If the target down event isn't resolved within 24 hours, escalate to level 1 and email the manager.

7. If the target down event isn't resolved within 72 hours, escalate to level 2 and email the manager.

When transforming the above rule requirements into actual rules in Enterprise Manager, here are some guidelines you can use:

1. Put all the rules that pertain to the common target (e.g. group) into one rule set instead of across multiple rule sets. It will make it easier to manage the rules that pertain to the same target in one place and avoid potential duplication of actions across rule sets.

2. Incidents should typically be created for the events specified in the rule. Since there are operational requirements to send notifications for these events, then this means these events are actively managed and thus incidents should be created for these. Also think about which events can be logically compressed (or combined) into one incident. Rules that create incidents for events should typically be the first set of rules in the rule set.

3. Once incidents are created, subsequent actions in the rule set such as notifications, etc. should typically be performed on the incident instead of on the event. One exception would be in the case where interested parties are interested in receiving notifications for specific events. For example, business owners might want to be notified of 'target down events' for targets that impact their business applications. For these cases where the primary interest is to be notified if such events occur, you can create rules on the specific events of interest and specify notification actions for the interested parties.   Extending our example, if business owners wanted to be notified if any of their important targets went down, you could create a rule on the 'Target Availability' down event and the action of the rule would be to send email notification to the business owners.

4. If there are common actions across rules (e.g. for rules #1 through rule #5 in the rule requirements list above, actions are to send notifications), consider creating a separate rule for these actions and have this new rule be executed after the other rules. If it is set up this way, then should there be any future changes to the actions, you will only need to make the changes in one rule instead of multiple times across several rules. In our example, the actions that are common across rules #1 through rule #5 are to send page notifications if the severity is fatal or critical, and send email notifications if the severity is warning. Instead of specifying the notification actions for each of the five rules, you can create a separate rule for the notification actions. If there is any change needed in the notification action (such as a change in the recipient), then you will only need to make the change in this separate rule.

5. Consider combining rules that are similar in nature and only differ by different parameters (e.g. rules #6 and #7 above) by leveraging conditional action support in rule actions. In our example, rules #6 and #7 both apply to target down events. They differ in the length of time the event has been opened and subsequent actions based on this. This can be combined into one rule with two conditional actions: the rule will apply to target down events (or incidents containing target down events) and the conditions for the actions will be based on 'duration' i.e. how long the incident has been opened. The first action will have a condition based on the incident being open for at least 24 hours, and the second action will have a condition based on the incident being open for at least 72 hours.

Applying the above guidelines to the rule requirements, we now have these rules:

Rules for the production group PROD:

1. Create incidents for the following events: target down events, specific metric alert (events) for database, host, WebLogic targets, and job failure events. For target down events from RAC instances, compress these into a single incident, one incident per RAC target. Also set the priority of these incidents to Urgent.

2. Send notifications for incidents as follows: send page notifications to administrators for incidents of fatal severity or critical severity; send email notifications to administrators for incidents of warning severity. (Note: target down events are the (only) events with fatal severity; backup job failure events have critical severity). As a variation of this rule, you could have different recipients for the notifications based on target type, i.e. the DBA team gets

notified for incidents on database targets, the middleware team gets notified on incidents on WebLogic targets and the system administrators get notified for incidents on host targets.

3. For incidents of fatal severity (i.e. incidents for target down events):

   a. If the incident is still open after 24 hours, then escalate incident to level 1 and email the manager.

   b. If the incident is still open after 72 hours, then escalate the incident to level 2 and email the manager.

Which type of rule should I use?

Now that you have planned the rules that you would like to setup, when you define these rules in Enterprise Manager, there is a choice of the type of object on which the rule will operate: event or incident, or problem.   General recommendations on which object to use for rules are described in Table 3.

Table 3. RECOMMENDATIONS FOR USES OF RULES

TYPE OF RULE	PURPOSE
Rule on Events	- Create incidents based on one or more events - Create helpdesk tickets for incidents (create incident based on the event, then create the ticket based on the incident) - Send events to third party management systems - Send email for events you're interested in (e.g. sending email to users such as business owners for specific events of interest)
Rules on Incidents	- Automate incident workflow operations (e.g. assignment of incident, prioritization, etc.) - Send notifications on incidents - Create helpdesk tickets for incidents. You can optionally specify conditions for when the ticket should be created (e.g. create ticket if incident is escalated to level 2)
Rules on Problems	- Automate problem workflow operations (e.g. assignment, prioritization, etc.) - Send notifications on problems

The rule set for the production group PROD should thus be defined as follows in Enterprise Manager:

Rule Set name: Rule set for production group PROD
Target: PROD Group

Rules: (See Table 4 below)

Note the targets and events defined in the rule criteria in Table 4 below are scoped to the PROD group which is the target specified in this rule set.

Table 4. RULES IN THE RULE SET FOR PRODUCTION GROUP PROD

EXECUTION ORDER	RULE APPLIES TO...	RULE CRITERIA	RULE ACTION(S)
1	Events	*(All target down events from all database instances)* • Event type = Target Availability AND • Event = (Target Type = Database Instance, Availability State = 'Down'), (Target Type = Cluster Database, Availability State = 'Down')	• Create incident, compress events into an incident by 'Target' with same ancestor type 'Cluster Database' • Set incident priority = Urgent *Note: This compresses all target down events from RAC Instances into an incident, one incident per RAC. If the RAC target is also down, its event is also included in the same incident.* *Target down events from Single Instance databases (that have the same target type as RAC instance) will have individual incidents created for them.*
2	Events	*(All target down events across all WebLogic servers)* • Event type = Target Availability AND • Event = 'Target down' for all WebLogic servers	Create incident, Set incident priority = Urgent
3	Events	*(Specified set of metric alerts for databases, hosts, WebLogic servers)* • Event type = Metric Alert AND • Metrics = <specified set of metrics> with warning or critical severity	Create incident
4	Events	*(All job failure events)* • Event type = Job Status Change AND • Job Status = Problems	Create incident

EXECUTION ORDER	RULE APPLIES TO…	RULE CRITERIA	RULE ACTION(S)
5	Incidents	*(All incidents of fatal, critical or warning severity)* • Severity in (Fatal, Critical, Warning) AND • Target Type = Host	• If severity is critical, page <specified set of system administrators> • If severity is warning, email <specified set of system administrators>
6	Incidents	*(All incidents of fatal, critical or warning severity)* • Severity in (Fatal, Critical, Warning) AND • Target Type = Database	• If severity is fatal or critical, page <specified set of DBAs> • If severity is warning, email <specified set of DBAs>
7	Incidents	*(All incidents of fatal, critical or warning severity)* • Severity in (Fatal, Critical, Warning) AND • Target Type = WebLogic Server	• If severity is fatal or critical, page <specified set of middleware administrators> • If severity is warning, email <specified set of middleware administrators>
8	Incidents	*(All incidents of fatal severity)* Severity in (Fatal)	• If incident is open for at least 24 hours, set incident escalation level to 1 and email <specified manager>. • If incident is open for at least 72 hours, set incident escalation level to 2 and email <specified manager>.

Other Considerations for Planning Rules

Rules were designed to offer much flexibility in specifying the conditions and actions that the rule should take. As such, there could be possibly more than one way to define rules that meet your operational requirements. Here are additional options you might want to consider when defining rules:

1.   Which rule should create the incident?

An event can be part of at most one incident. Hence if there are multiple rules that create an incident for the same event, then the rule that is executed *first* will create the incident containing the event. Subsequent rules that contain actions to create an incident for the same event will not be executed because an incident has already been created for the event. However, other applicable actions (such as sending notifications) will still be executed. Hence to avoid surprises, it is important to plan which rule set will contain rules for creating incidents.

2.   Execution order matters if rules operate on the same set of events and incidents

Incident workflow attributes such as owner, priority, etc. can be changed via rules. Thus, it is technically possible for a set of rules in the same rule set to set the same incident attribute to different values.   If this happens, the last one wins, i.e. the final value of the incident attribute is based on the last rule that was executed. You might have requirements that require the setting of incident attributes to different values based on certain conditions. In this case, think about putting rules that apply to broader criteria first, and rules that apply to more specific criteria afterwards. For example, if your requirements are to assign incidents on database targets to the DBA-TEAM user and all other incidents to the SUPPORT-TEAM, then you can set the rules up this way:

Rule1: For all incidents, set owner to SUPPORT-TEAM

Note: The rule above is the rule with broader criteria.

Rule2: For all incidents on database targets, set owner to DBA-TEAM

Note: The rule above is the rule with more specific criteria.

Using the above rules, if the incident is on a database target, then even if the owner was initially set to SUPPORT-TEAM (per Rule 1), it will be later changed to DBA-TEAM per Rule2 that was executed after Rule1.

Finally, if you have rules containing duration-based criteria in the actions part of the rule (e.g. if incident is of Urgent Priority and has been open for at least 24 hours…), then try to put these rules as last in execution order. It is likely that the conditions specified in these rules are based on incident attributes (e.g. priority, etc.) and since such attributes can be changed by other rules, putting duration-based rules as last ensures the values of the attributes that they rely on are the latest values.

3.   Broader vs. more specific rule criteria

In the rule set defined in Table 4 above, the rule to generate incidents for metric alerts involves specifying individual metrics. One advantage to this approach is that you will know exactly which metric alert (events) generate incidents for your group. However, if you want to monitor additional metrics for any of the targets in the group, and you would like to create incidents for these, you will need to remember to update the rule by adding the additional metrics.   An alternative solution to this is to *not* specify individual metrics in the rule criteria but to expand the rule criteria to all metric alert events of warning or critical severity as shown in Figure 16 below:
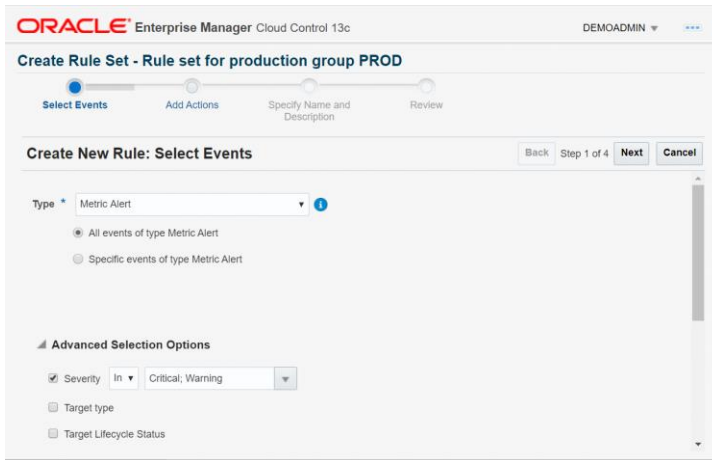
Figure 16. *This rule applies to metric alerts (events) that are in warning or critical severity.*

Again, note that this will apply to all warning and critical metric alerts for the PROD group and not in the entire Enterprise Manager deployment. The advantage to this approach is that it will automatically cover any additional metric alert (events) on the group, i.e., there is no need to update this rule should you decide to monitor any additional metrics for targets in the group. However, this assumes that there is a fair amount of control over setting the appropriate thresholds of the metrics across all targets in the group (hopefully done through the Administration Group – Template Collection feature). Otherwise if you have some spurious metric alert (events) due to incorrect metric thresholds, this rule will end up creating some unnecessary (unwanted) incidents for your group.

4. Email notifications: the 'To' and 'Cc' lines

When specifying the recipients of email notifications, you will notice there are 2 fields on which you can specify the recipients: "E-mail To' and 'E-mail Cc'
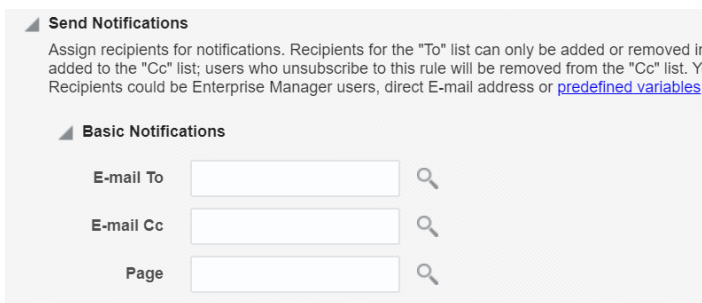

Figure 17. *This shows the options for the e-mail recipients in a rule.*

The 'E-mail To' line should be used to specify the primary recipients of the email notifications, and these are typically the administrators who are expected to take some action in response to the event/incident/problem for which they are being notified. The selection of these recipients are

determined as part of the planning process and not an ad hoc choice for any individual administrator. As such, changing the recipient in the 'To' line can only be done by changing the rule itself (which is restricted to the rule creator and other designated administrators as will be discussed in the next section).   The 'E-mail CC' line should be used for other interested parties (e.g. business owners, etc.) who would like to be notified of the event/incident/problem's occurrence. As such, these users can subscribe or un-subscribe to these rules based on their own interest.

5.   Email recipients need privileges on the targets

When creating rules that send email/page notifications to users, these users will only be able to receive notifications if they have privileges (i.e. at least View privileges) on the targets on which the events/incident have been created.

6.   Determine who creates the rules in Enterprise Manager

When actions on events/incidents (and problems) are specified in rules, these are executed by Enterprise Manager using the privileges of the Enterprise Manager administrator who *created* the rule. Thus, it is important to ensure that the creator of the rule has the necessary privileges on target(s) of the rule set for the actions to be executed successfully; otherwise, the rule action will fail. If the target of the rule set is a privilege propagating group, then the privilege can be granted once on the group and it will automatically be applied to all members of the group. For example, if one of the rule actions is to 'Create an incident' for an event, then the rule creator must have at least *Manage Target Events* privilege on the target on which the event has occurred. The checking of privileges of the rule creator is done at rule execution time so it is important that you proactively check that the rule creator has sufficient privileges (e.g. Manage Target Events) on the targets in the rule, otherwise rule actions will fail. For a complete list of rule actions that use the privileges of the rule creator, refer to the "Incident Management" chapter of the *Oracle Enterprise Manager Cloud Control Administrator's Guide 13c*.

7.   Audit changes to the rules

Since rules are an integral part of your monitoring and incident management processes, consider enabling auditing of your rules to help you can keep track of who and when rule changes were made. You can use EM CLI *update_audit_settings* to enable such operations to be audited. Refer to the section on "Configuring and Managing Audit" of the *Oracle Enterprise Manager Cloud Control Security Guide* for details on enabling audit for incident rule set changes.

8.   Verify rule setup using 'Simulate Rules' feature

You can use the 'Simulate Rules' feature to help you verify your rule set setup. Simulate Rules can be accessed in the main Incident Rules page. The simulate rule feature will list out all the actions that Enterprise Manager would take if a specified event occurs. To use this feature, you will first be asked to choose an event from a list of existing events. Then you can run the simulation. The output of the simulation is a list of rules and actions that the rules would take if the event occurs. The actions will not actually be executed, thus providing a safe way for you to test your rules. If there are missing rules or actions in the simulation output, then that indicates some issue in the rule setup. It could be due to missing privileges on the rule creator (e.g. Manage Target Events privilege is required to create an incident) or perhaps the event conditions in the rules are incomplete.

Implement the Rules in Enterprise Manager

Once the rules have been planned, they can be implemented, i.e. entered in the Enterprise Manager console (Setup → Incidents → Incident Rules). Note that you will need the *Create Enterprise Rule Set*

resource privilege to create these rules. All the rule sets that have been discussed thus far are of type "Enterprise". Enterprise rule sets are used to implement your IT operational processes (send email, open helpdesk tickets, escalate incidents, etc.). Because these rule sets can perform a wide variety of actions, the creation of these rule sets is protected by a resource privilege called *Create Enterprise Rule Set*. It is recommended that specified persons are designated to create these rules on behalf of the team (this is done as part of the planning of job responsibilities discussed in prior sections of the paper). If multiple people are designated as having this responsibility, then once one person has created the rule set, the other users can be made Co-authors of the rule set. This entitles them to also edit the rule set as needed. Finally, it is also possible to enter the rule sets in advance of actual operations, i.e. in advance of having the rules be operational in your environment. This can be done by disabling the rule set and re-enabling it once you are ready to go live.

Benefitting From Economies of Scale: A Fully Automated and Scalable Monitoring Setup

At this point you've setup the groundwork for an automated and scalable monitoring solution. The 'cost' or effort of setting up new target monitoring is now very minimal. Whether you're adding one new target or hundreds of new targets to Enterprise Manager, all you need to do as far as monitoring setup is concerned is one step: set the target properties of the target such that it matches the membership criteria of the Administration Group to which is should belong. Once you've done this, everything else is automatically handled by Enterprise Manager:

- The target is automatically monitored using the appropriate monitoring settings.
    - o This happens because the target is automatically added to the appropriate Administration Group. Once it is in the Administration Group, the monitoring template from the associated Template Collection is automatically applied to the target.

- Administrators can start managing the target because they have the appropriate privileges to manage the newly added target.
    - o This happens because the role that has been granted to them contains privileges on the Administration Group and since Administration Groups are privilege-propagating, it automatically includes privileges on the newly added target.

- Rules for creating incidents, sending notifications, and other operational procedures automatically include the newly added target.
    - o This happens because the Administration Group has been specified as the target for the rule set and newly added targets for the group are automatically included in the rule set.
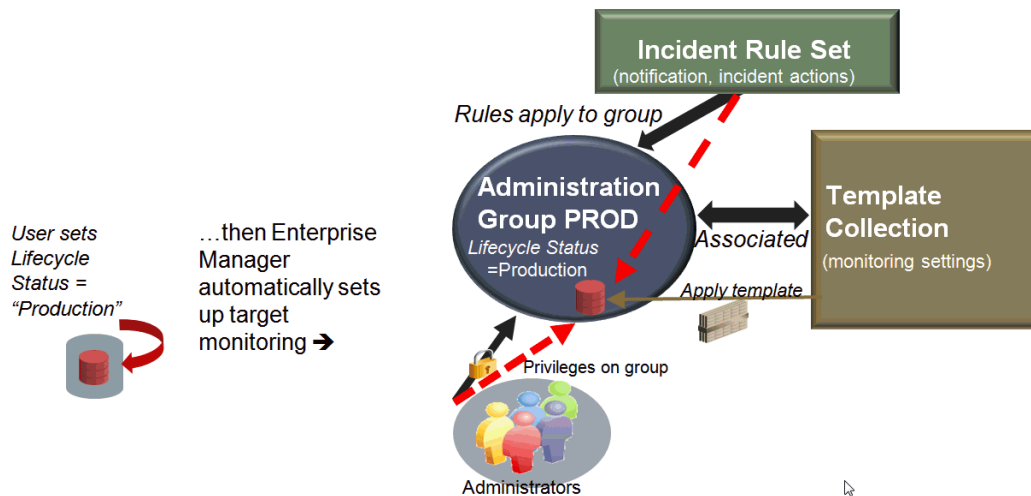
Figure 18. *Monitoring set up for new targets is fully automated.*

MANAGING INCIDENTS

Once monitoring has been setup, events will be raised, and incidents created for these events. Incident Manager (available within the Enterprise Manager console) provides a centralized way to manage all these different incidents. Here are some guidelines on how you can effectively manage incidents:

1. Leverage views in Incident Manager

The views in Incident Manager enable you to focus on the subset of incidents you're interested in. There are several out-of-box views that cover a common set of incidents that a datacenter might be interested in. For example, an administrator can use the out-of-box view 'My open incidents and problems' to view all the incidents and problems assigned to him. This list can be sorted by priority so that he can work on the most important ones first. You can also create additional (custom) views to filter the list of incidents based on your specified search criteria. A good example would be to create a view containing all incidents for an important group of targets (e.g. production group PROD). These custom views can also be shared with other users. Hence you can designate a user in your team to create views on behalf of the team, then share the custom views created.

2. Automate incident assignment

Automate the assignment of incidents to administrators using rules if the set of incidents that should be assigned to administrators is well-defined (e.g. all host incidents are assigned to user A, all database tablespace incidents go to user B, etc.)

3. Manage by priority

The priority field of an incident should be leveraged in order to provide guidance on which incidents are more important and thus should be managed first. If the semantics for determining which incidents are the highest priorities can be determined in advance, you can use rules to automatically set the incident priority. In Incident Manager, when administrators are working from a list of incidents (within a view), it is recommended that they sort the incidents by priority to manage the more important ones first.

4. Working off a queue of incidents

If you operate by having a team work off a common queue of incidents, you can emulate the 'queue of incidents' for the team by creating an Enterprise Manager user to represent the team. For example, you can create a user called 'DBA-TEAM' to represent the DBA team. If there is a well-defined set of incidents that need to be owned by the DBA team, then create a rule to auto-assign these incidents to user DBA-TEAM. This can be done as part of the rule that creates the incident. (For example, create a rule on events and choose specific database metric alerts; in the actions part of the rule, create an incident and assign the incident to DBA-TEAM.) In Incident Manager, create a view to filter the list of incidents to all incidents owned by user DBA-TEAM. The DBA team can then use this view to look at all the incidents assigned to their team. When they actually want to work on an incident, they should take ownership of the incident so that the owner field is set to their own Enterprise Manager user account.

If there is an additional requirement to send email notifications to the DBA team when an incident is assigned to them, then you can specify the DBA-TEAM as the recipient of the email notification in your event/incident rule and specify an email distribution list as the email preference for the DBA-TEAM user. The email distribution list should include the email addresses of the members of the DBA team.

5. Automate incident escalation

Automate the escalation process using rules (e.g. set escalation level and send notifications to appropriate people based on type of incident and duration). In the rule, you can specify a duration condition on the incident that determines when the incident should be escalated.

6. Check for missing incidents

Occasionally review the set of events in the view *Events without incidents.* There may be some important events for which no incidents have been created and thus are left unattended. If you find such events, then you can manually create an incident for them in Incident Manager. Also consider creating a rule for these to automatically create an incident for the event the next time it occurs.

7. Leverage incident resolution status to track the progress of your incidents

The 'status' field of an incident (or problem) allows you to track the resolution of the incident. All newly created incidents are set to 'New'. As administrators work on the incident, they should change the status appropriately (e.g. Work In Progress). This enables you to see how much (or how little) progress has been made to the resolution of an incident. Additional values for the 'status' can also be added to support your operational processes. For example, you can add values such as 'Waiting on Vendor' or 'Waiting on Subject Matter Expert' etc. To add additional values, use the EM CLI verb *create_resolution_state*. You will need to have Enterprise Manager super administrator privileges to execute this verb.

Many incidents in Enterprise Manager cannot be manually cleared. They can only be cleared if their underlying event has been cleared by the agent. If an administrator believes he has fixed the incident but is waiting for Enterprise Manager to clear the event (and incident), he can set the incident status to 'Resolved'. Once the event (and incident) is cleared by the agent, the status will be automatically set to 'Closed'.

8. Use Lifecycle Status target property to identify the most important targets

The Lifecycle Status target property is meant to designate the operational status and importance of a target. Possible values are (in order of descending priority): Mission Critical, Production, Staging, Test, and Development. If your Enterprise Manager site is under heavy load, then events from the higher priority targets (e.g. Mission Critical and Production targets) are processed ahead of the lower priority targets (e.g. Development targets). Target Availability events are always processed ahead of other types of events. Set the appropriate value for the Lifecycle Status target property to your important targets to ensure their events are processed ahead of others in case Enterprise Manager is under heavy

load. You can set the value of target properties using the Target Properties page which is accessible from the target's menu (Target menu → Target Setup → Properties). Or, to do this operation in bulk, use EM CLI *set_target_property_value*. You cannot add or remove values from pre-defined set of values for Lifecycle Status. However, you can change the display name using the EM CLI verb `modify_lifecycle_stage_name.`

ADDITIONAL MONITORING REQUIREMENTS AND RECOMMENDATIONS

Here are some additional common requirements, issues and recommendations related to monitoring:

Auto-fixing Alerts using Corrective Actions

There are certain types of alerts that have standard remediation actions. For example, if a listener is down, the remediation action would be to restart the listener. It is more efficient to automate these remediation actions in Enterprise Manager rather than to have these manually executed by on-call staff. To automate these actions, you can use Corrective Actions. Corrective Actions are jobs that are executed by the agent in response to an alert. Enterprise Manager provides some built-in corrective actions such as: Startup Listener, Startup Database, Add Space to Tablespace. You can also specify custom logic using the OS Command or SQL Script corrective actions.

Corrective actions are specified as part of a target's metric threshold settings. When corrective actions are executed, it must run using the credentials of a specified user. It is recommended that corrective actions are included as part of the monitoring template that is in turn part of the Template Collection that is associated with the Administration Group. If it is setup in this recommended way, then when the corrective action executes, it will use the preferred credentials of the Enterprise Manager user who associated the Template Collection with the Administration Group.

For more details on using corrective actions, refer to the "Utilizing the Job System and Corrective Actions" chapter of the *Oracle Enterprise Manager Cloud Control Administrator's Guide 13c*.

Monitoring When Enterprise Manager is Under Planned Maintenance

Occasionally, administrators need to bring down Enterprise Manager's Oracle Management Service and Repository for planned maintenance. While the duration of this maintenance window may be short, incident management and notification features are not available during this downtime and yet businesses may require 24x7 continuous monitoring of their mission critical targets. In order to meet the continuous monitoring requirements, you can use Always-On Monitoring. Always-On Monitoring is a separate, complimentary application that can be configured to receive alerts from agents and send email notifications for alerts. In Enterprise Manager 13.4, it will send email about the most important alerts: Target Availability and Metric Alerts. You can configure it include additional alerts (e.g. metric errors) if needed. For more details, refer to the "Always-On Monitoring" chapter of the *Oracle Enterprise Manager Cloud Control Administrator's Guide 13c*.

Sharing the Administration Group Hierarchy across Different Teams

In some datacenters, one Enterprise Manager deployment could be shared across different teams that are responsible for managing their own set of targets. One way for different teams to share the same Administration Group hierarchy is to have the first level of the hierarchy be used to divide the targets based on the teams that manage them. As a simple example, say the DBA team is responsible for the databases and listeners, the Middleware team is responsible for the WLS (WebLogic) targets, and the Sysadmin team is responsible for the host targets.   The first level of the Administration Group hierarchy can be based on the Target Type property such that the target types owned by a team are part of the same group:
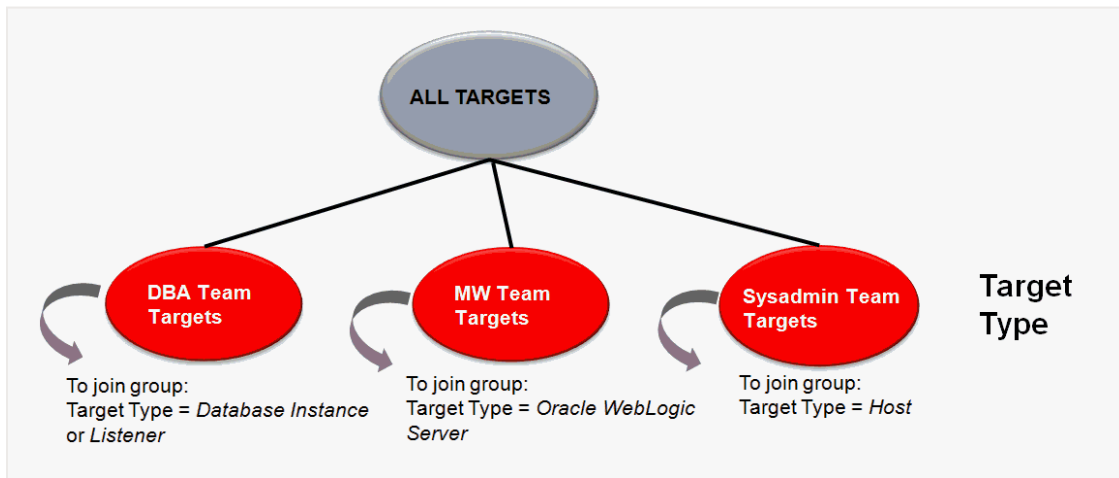
Figure 19. *The same Administration Group hierarchy can be shared by different teams.*

Additional levels, if needed, can be added to the Administration Group, but they would have to be the same across all groups.   For example, another level based on Lifecycle Status can be added to further divide the targets for each team into production and non-production targets if there are different monitoring settings for production and non-production targets.   Since the levels must be the same across all groups, these levels would need to be discussed and agreed-upon by the various teams. However, it should be noted that all the levels need not be used by all teams.   For example, in the Administration Group hierarchy below (Figure 20), a Lifecycle Status level has been added to create Production and Non-Production groups for each team:
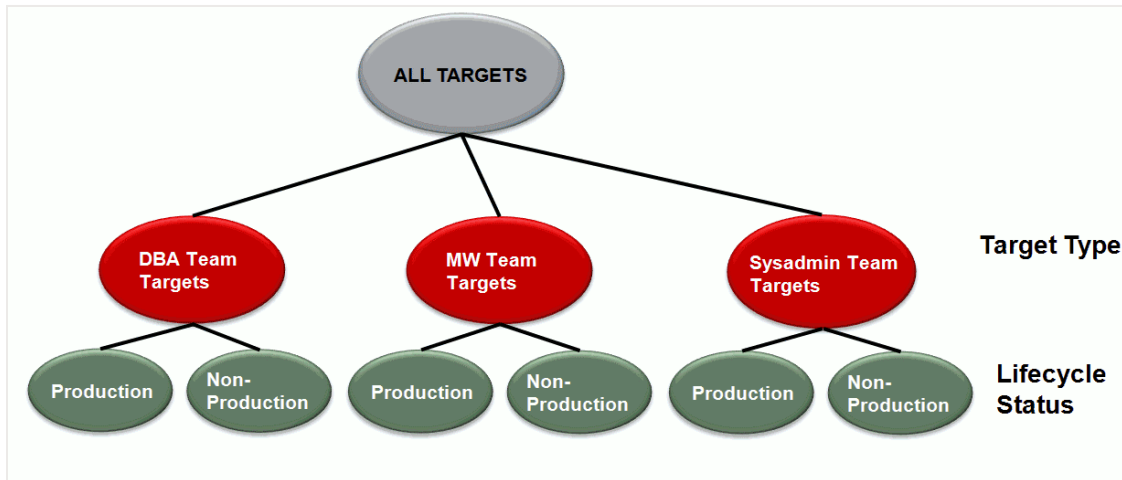
Figure 20. *This Administration Group shared across different teams has the same levels across all groups.*

This enables each team to define different monitoring settings for Production targets and Non-Production targets.   In this example, it is important to note two points:

- If a team, say the Sysadmin Team, has the same monitoring setting for all host targets, then they do not have to use the additional Production and Non-Production groups for template application purposes. They can create a Template Collection containing the monitoring template for host targets and associate the Template Collection with the group 'Sysadmin Team Targets'. This template will then apply to both subgroups Production and Non-Production.

- In order for a target to join the Administration Group hierarchy, it must match all the criteria used in defining the hierarchy. In an Administration Group hierarchy, non-group targets are always found in the lowest level of the hierarchy. So, in our example hierarchy, the target must match both the target type and lifecycle status criteria in order to join the group. For host targets managed by the Sysadmin team, even if there is no difference in monitoring settings between production and non-production targets, the targets' lifecycle status must be set to match the lifecycle criteria of the Administration Group for the host target to join the Administration Group.

Finally, another benefit of using the first level hierarchy to divide targets by teams is ease of use of managing privileges.   If all members of a team have the same level of privileges for all targets they manage and only view privileges (or no privileges) on the other targets that they don't manage, then a role can be created that grants the appropriate level of privilege on the first level groups. For example, the role created for the DBA Team can consist of *Full* privileges on the 'DBA Team Targets' group, *View* privileges on the 'Sysadmin Team Targets' group, and perhaps no privileges on the 'MW Team Targets' group. This is depicted in Figure 21 on the next page.
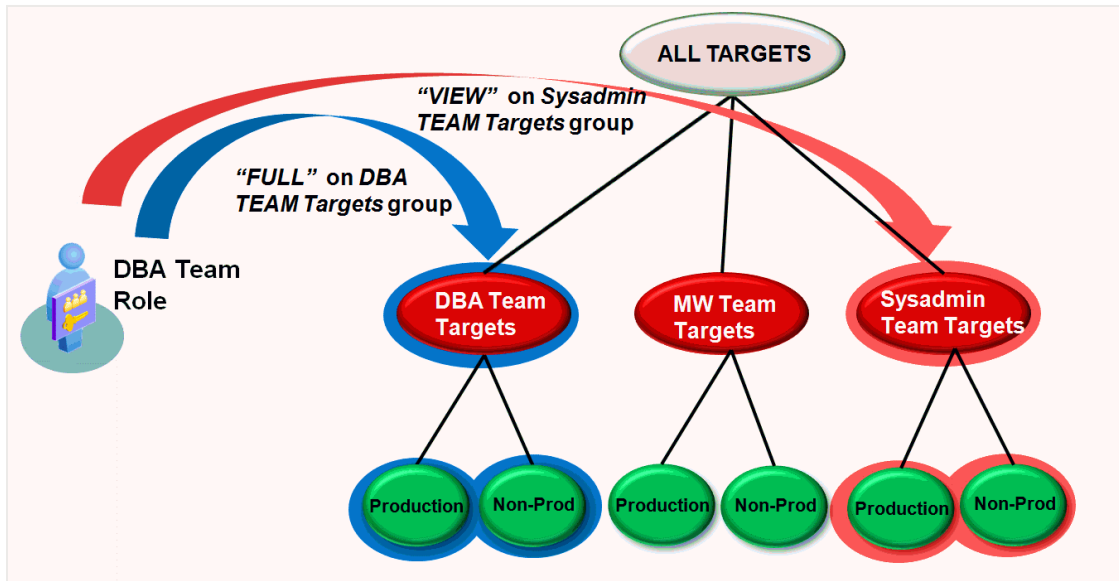
Figure 21. *A role containing privileges on an Administration Group automatically propagates to all its subgroups.*

Note that you need *View Any Target* privilege to see the full Administration Group hierarchy in the Administration Group pages. Hence if you have a setup such as the one depicted in Figure 19 wherein you only have privileges on a subset of groups in the hierarchy, then use the Groups page (accessible from Targets → Groups) to view and perform operations on your groups.

Using Administration Groups for Other Group Operations

Administration Groups can be used in all features that support groups. This includes granting privileges on these groups to your roles, using these as the targets for incident rule sets, blackouts, system dashboard, etc.   However, there can be scenarios where the membership criteria defined for Administration Groups may not exactly match the desired membership criteria for these other group operations. For example, you may have an Administration Group hierarchy based on Lifecycle Status as shown below in Figure 22.
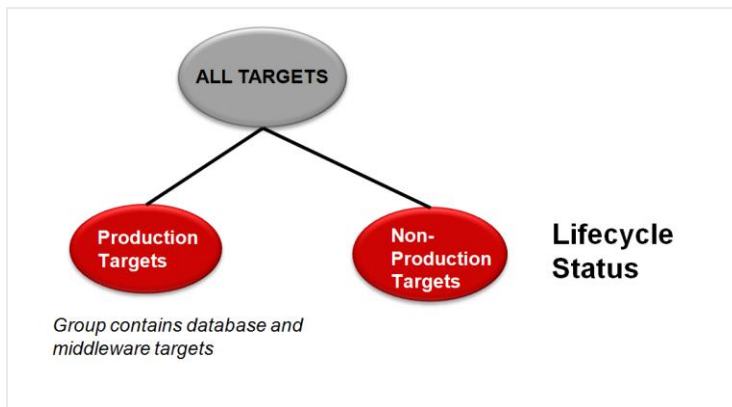
Figure 22. *The Administration Groups in the hierarchy contain database and middleware targets.*

The Administration Groups in the hierarchy contain database and middleware targets. For privilege management purposes you may be unable to grant privileges on these groups to your roles because you only want to grant the privileges on a subset of targets in the group. For example, you may want to grant Full privileges on only the databases and not the middleware targets in the group to the DBA team. One option would be to extend the Administration Group hierarchy by adding an additional property to further breakdown the group (in our example, the new property would be Target Type). This option however could make the hierarchy more complex especially if you already have a 3 or more level hierarchy. The other option is to create *dynamic groups* that better satisfy the grouping requirements. Dynamic groups are similar to Administration Groups in that dynamic groups are defined by membership criteria. Once defined, Enterprise Manager automatically adds targets to the appropriate dynamic group(s) if they match the dynamic group's membership criteria. In our example, you could create a dynamic group with membership criteria: Lifecycle Status = Mission Critical or Production, Target Type = Oracle Database. Also set the Privilege Propagation attribute on the dynamic group so that privileges granted on the group to a role or user automatically applies to all members of the group. So for monitoring template apply purposes, Administration Groups would be used; for privilege management purposes dynamic groups can be used. The two options are summarized in Figure 23 below.
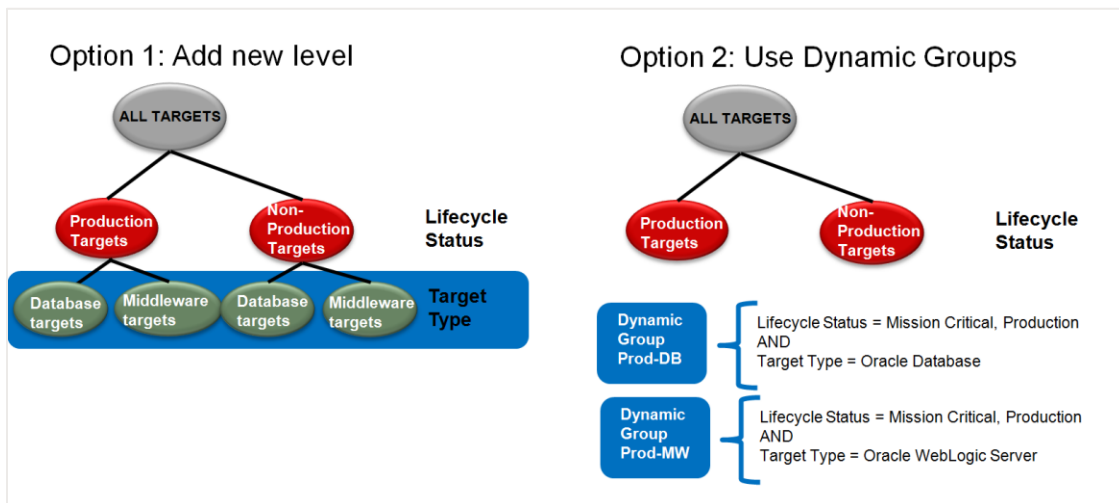
Figure 23. *Options when Administration Groups do not meet other group requirements*

In the end, you need to balance the tradeoffs between the options. Adding more levels to your Administration Group hierarchy to accommodate non-monitoring requirements gives you the benefit of being able to centrally manage all groups from a single group hierarchy structure but it adds more complexity to your Administration Group hierarchy. Since the Administration Group hierarchy is height-balanced, it could result in additional number of groups that will remain un-used. If you choose to go the dynamic group route, then it doesn't add unnecessary complexity to your Administration Group hierarchy but it does require the management of a separate set of groups outside the Administration Group hierarchy. While there is no one correct answer, the general recommendation is to always create the Administration Group hierarchy based on monitoring requirements in order to keep the hierarchy as simple and streamlined as possible. If there are different group membership requirements for other non-monitoring needs, create dynamic groups to meet these other requirements.

Verifying Targets are Part of the Administration Group Hierarchy

It is always good to check for orphaned targets, i.e. that there are no targets that should have been part of an Administration Group but aren't. Use the Unassigned Targets Report for this purpose (accessible from the Associations page of Administration Groups).

Figure 24. *The Unassigned Targets page helps you track targets that have been left out of the Administration Group hierarchy.*

The Unassigned Targets page shows you the targets that have not been added to any Administration Group for any of these reasons:

1. The target property values do not match the Administration Group criteria. The report shows you the target properties that are used for Administration Group criteria (represented by columns in the table) as well as their current values.   The top part of the page shows you the values of the target properties used as criteria for the Administration Group hierarchy.
2. If the target is an aggregate target (i.e. a target that contain other targets as members), it is not privilege propagating. This is indicated by the checkbox under the Non-Privilege Propagating Aggregate column.

3. The target is still scheduled for membership evaluation into an Administration Group. The time in which the target is scheduled for membership evaluation is shown in the 'Pending Membership Evaluation' column.   At that scheduled time, Enterprise Manager checks the target's properties to see if it matches the criteria of any Administration Group. If it does, it adds the target to the appropriate Administration Group.

To address the first issue, set the target property values such that they match the Administration Group criteria. You can set target properties via the Edit Target Properties page or via EM CLI *set_target_property_value*. The target has to match all Administration Group criteria in order to join an Administration Group. To address the second issue, set the aggregate target to be privilege propagating via EM CLI *modify_system* with the *–privilege_propagation=true* option set.

Changing the Administration Group Hierarchy after Initial Creation

Even with careful planning of the Administration Group hierarchy, it is possible that you might have to change it to accommodate new requirements. These changes include adding new target property values (e.g. adding new Line of Business values which add more groups horizontally), merging two or more groups, adding new levels, deleting levels, etc. It is possible to do these changes without re-building the entire hierarchy. If you add a new level, that is the same as adding a new target property to the Administration Group criteria. As such, you will need to 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. When deleting a level, that is the same as deleting a target property criterion. This causes groups at that level to be deleted. 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.

If you want to merge two or more Administration Groups, this is done by merging their corresponding target property criteria in the Administration Group hierarchy definition. The actual merge semantics involve retaining one of the groups to be merged and 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 get to choose the group to be retained. If the group has subgroups, then since the actual targets reside in the leaf level (lowest level) Administration Groups, the actual movement of targets will occur in the leaf level Administration Groups. The upper level Administration Groups will have updated criterion based on the merged criteria.

As an example, consider the Administration Group hierarchy shown in the figure below. It is based on Target Type and Lifecycle Status. The criteria associated with the Administration Groups are also shown.
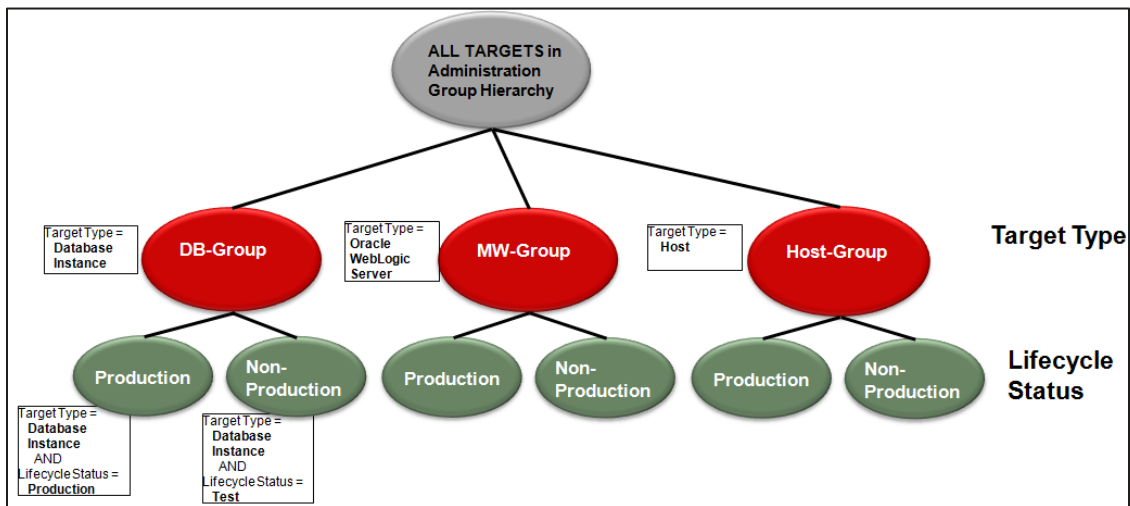

Figure 25. *Administration Group hierarchy before the merge*

Let us assume you want to merge the groups DB-Group and MW-Group. This is done by merging their corresponding target property criteria 'Database Instance' and 'Oracle WebLogic Server'.
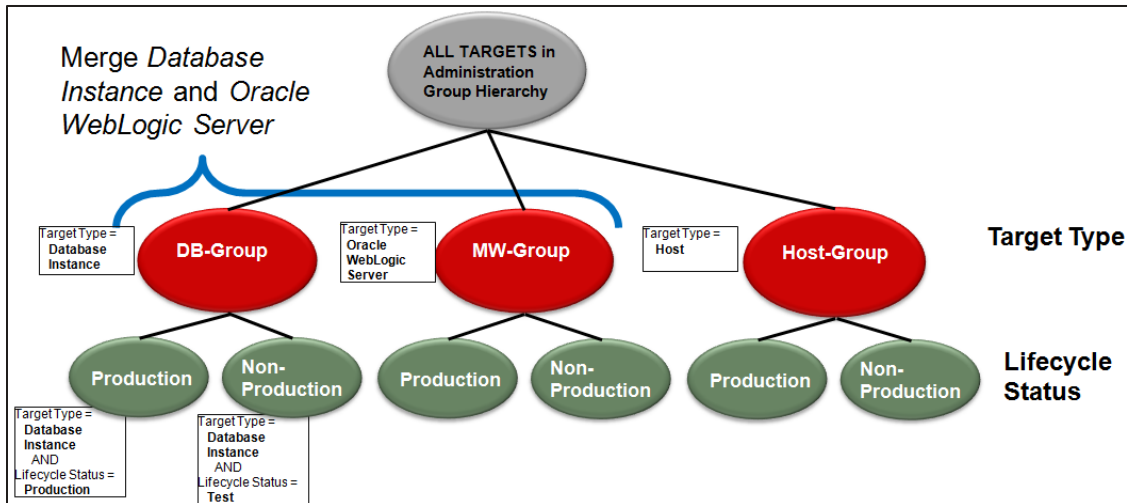
Figure 26. *Merging groups is done by merging their corresponding target property criteria.*

In an Administration Group with multiple levels, each target property criterion corresponds to group(s) containing that criterion. So, merging 'Database Instance' and 'Oracle WebLogic Server' target property criteria means merging of DB-Group and MW-Group and their subgroups. As part of the merge, you will be asked a question on which group to retain. Specifically, if you are merging 'Database Instance' and 'Oracle WebLogic Server', do you want to retain groups with criteria 'Database Instance' or groups with criteria 'Oracle WebLogic Server'? Assuming you choose to retain 'Database Instance', then the following will occur:

- MW-Group will merge into DB-Group. This means the group called DB-Group will have its criteria updated to include Oracle WebLogic Server.

-

- All member targets in Production group under MW-Group will move over to the Production group under DB-Group. The Production group under DB-Group will now include 'Oracle WebLogic Server' as part of its criteria.

-

- All member targets in Non-Production group under MW-Group will move over to the Non-Production group under DB-Group. The Non-Production group under DB-Group will now include 'Oracle WebLogic Server' as part of its criteria.

-

- After the targets have been moved, the MW-Group, Production and Non-Production groups under MW-Group will all be deleted.

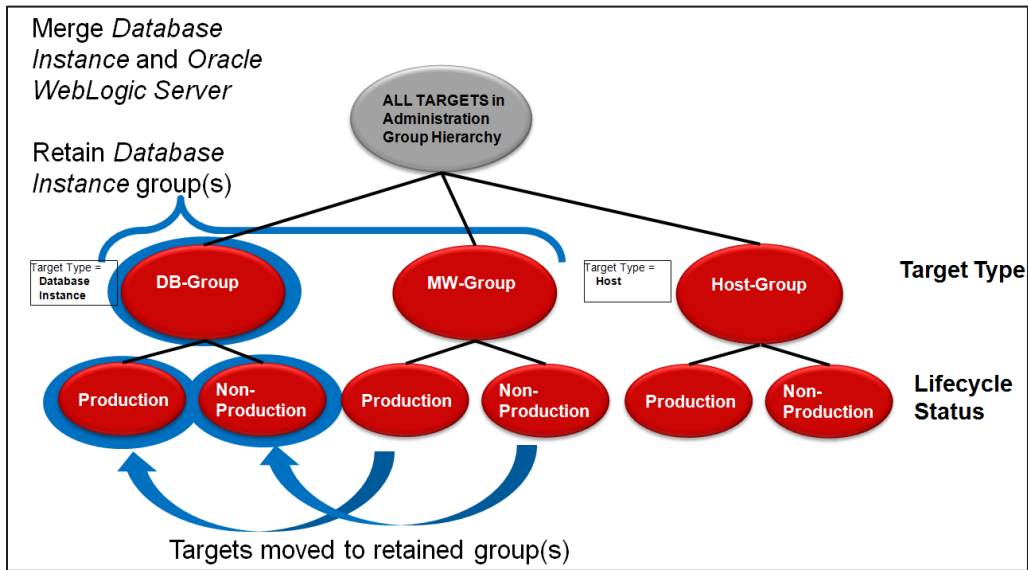The figure below shows a pictorial view of the merge process.

Figure 27. *Merging involves moving over targets to the retained groups and updating its group criteria.*

After the merge has been completed, the Administration Group now looks like this.
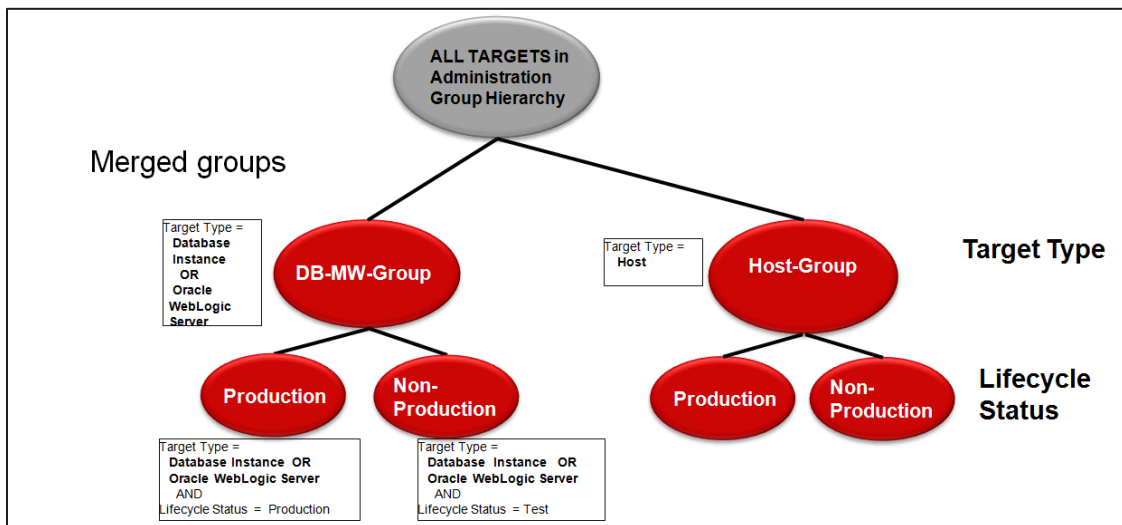

Figure 28. *Updated Administration Group hierarchy after the merge*

When choosing the group to retain, consider keeping the group that is used in most group operations (e.g. incident rule sets, system dashboard, roles) so that the impact of the merge is minimized. The operations on the retained group(s) will continue to remain intact and after the merge is done, the operations will also apply on the targets newly added to the group.

Verifying Targets are in Sync with Your Monitoring Standards

As mentioned, templates are auto-applied to targets that are part of Administration Groups. In order to verify whether or not targets in an Administration Group are in sync with their associated monitoring templates, you can look at the Synchronization Status region in the group's homepage:

◢ **Synchronization Status**

Each target in the Administration Group is synchronized with the items in the Template Collection where applicable. If an error occurred during synchronization, click the link under 'Failed Targets' to see the details in the message column of opened pop-up.

Last Synchronization Apr 1, 2020 12:00:05 AM PDT   **Next Synchronization** Apr 2, 2020 12:00:00 AM PDT

View ▾   **Start Synchronization**   ⬛ **Exclude from Synchronization**   🗉 **Go To Administration Groups and Template Collections**   »

Name	Synchronized Targets	Pending Targets	Running Targets	Failed Targets	Excluded Targets	N/A Targets
Monitoring Templates	9	3	0	0	0	3
Compliance Standards	0	0	0	0	0	15
Cloud Policies	0	0	0	0	0	15

Figure 29. The Synchronization Status region helps you verify if targets are in sync with the associated monitoring templates.

In the figure above, in the row Monitoring Template, the Synchronized Targets column shows the number of targets that are in sync with the associated monitoring templates. This means the targets' monitoring settings are the same as in the monitoring template[5]. For this group, nine targets are in sync with their associated monitoring templates. The next column, Pending Targets, shows the number of targets that have pending template apply operations scheduled on them; hence their current monitoring settings are not necessarily the same as the associated monitoring template. The date for when this synchronization process will happen is indicated by the "Next Synchronization" date located at the top of the table. You can wait for this sync to occur or perform the sync process now by clicking on the 'Start Synchronization" option at the top of the table.

The number under the Failed Targets column shows the number of targets on which the monitoring template apply operation has failed. Finally, the N/A Targets show the number of targets on which there is no associated monitoring template. A non-zero number under N/A Targets column is an indication that you are missing some monitoring templates in your template collection. You can click on the number under that column to find out the target types for which you have no monitoring template included in your template collection.

If an administrator wants to check the sync status of all targets across all Administration Groups, go to the homepage of the topmost Administration Group, and check its Synchronization Status region. It should provide a rollup of the sync status across all targets in the Administration Group hierarchy.

---

[5] The Synchronization Status region also shows whether or not targets are in sync with associated compliance standards and cloud policies. These are, however, outside the scope of this paper.

For more details on the Synchronization Status region, refer to the "Implementing Administration Groups and Template Collections" chapter of the *Oracle Enterprise Manager Cloud Control Administrator's Guide 13c*.

Enabling Events for Jobs

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. By default, events are 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 will need to specify the set of targets for which you would like job-related events to be generated. You can do this using the Job Event Generation Criteria page which is accessible from the Setup → Incidents → Job Events menu.

Integrating with Third Party Event Systems and Service Desks

Many datacenters often require integration with other management systems and/or service desk systems as part of the monitoring solution. Enterprise Manager supports integration with other event management systems to enable sharing of Enterprise Manager events with these other systems. This type of integration can be done using Event Connectors. Once the appropriate event connector is installed and configured, you can reference the event connector within your event rules. For more details, refer to the event connector specific documentation that is available as part of the Enterprise Manager documentation set.

The other type of integration is with service desk systems.   Service Desk (or Helpdesk) connectors are also available to enable the automatic opening of service desk tickets for events /incidents raised in Enterprise Manager. Service desk specific documentation is available to guide you through the installation and configuration of the connector. Once that is setup, then in your event or incident rules, you can specify an action to open a service desk ticket for an event or an incident. The service desk ticket ID and status are tracked and visible within Incident Manager. Thus, you have a central way to monitor and track the status of all incidents, including those that are managed by service desk staff.

Too Many Alerts

Often you might think there are too many alerts (events) generated in your Enterprise Manager site. To address these, think about controlling these events at the source.   This means reviewing and controlling the source of events rather than only focusing on clearing events after they have been generated. Unnecessary events cause unnecessary load on your Enterprise Manager site. For metric alerts specifically, it is important to review the target's metric settings. To implement changes in metric settings across your targets, you can make these metric setting changes in the monitoring templates that you've included in your Template Collections:

- Disable metric collection for metrics you don't care about

  Most metrics for a target are enabled by default. If there are some metrics that you will not use in your Enterprise Manager deployment, consider disabling the metric collection of these metrics to avoid the unnecessary collection and storage of these metric values.

- Set thresholds only on metrics you care about

If your team typically ignores a set of metric alerts, then consider removing thresholds for these metrics to avoid the generation of these metric alerts. You should only set thresholds for metrics whose events (incidents) you will manage. You can still continue to collect metric data values without alerting if the metric historical data is important for you (e.g. for reporting or trending purposes).

- Adjust metric thresholds based on metric trend

Enterprise Manager 13c supports a threshold suggestion region in context of setting thresholds for a metric. This is available as part of the Metric and Collection Settings page and All Metrics page. This feature basically shows the value of the metric over recent periods of time e.g. last 31 days:

Figure 30. *The threshold suggestion region shows you the value of the metric over a recent period of time.*

Use this region to review the historical value of the metric in order to fine-tune thresholds defined for this metric. If, for a specific target, you end up deciding to use a threshold value for the metric that is different from the associated monitoring template, remember to click the 'Template Override' checkbox for the metric in the Metric and Collection Settings page. This ensures that any monitoring template applied to the target will not override your metric's threshold settings.

- Set number of occurrences

When a metric alert (event) should be generated only after a sustained period of time (e.g. metric value is high for at least 30 minutes) use the *number of occurrences* parameter. If this parameter is specified, it is used to determine the consecutive number of metric collection intervals for which the metric must exceed its threshold before a metric alert (event) is generated. If a metric is collected every 5 minutes, then if you specify a *number of occurrences*

parameter of 6, this means the metric needs to exceed its threshold for 30 minutes before a metric alert (event) is generated (i.e. 5-minute collection interval multiplied by 6 *number of occurrences*).

- Use corrective actions to auto-clear metrics alerts

  If there are metric alerts that can be resolved in an automated manner, consider creating Corrective Actions for these. Corrective Actions are job tasks associated with metrics such that when an alert for the metric is generated, the corresponding corrective action (task) is automatically executed. If the Corrective Action resolves the issue (e.g. restart the listener for the listener down event), then the metric alert will be cleared by the agent during the metric's subsequent collection schedule. When setting up rules to send notifications or create incidents for metric alerts, remember to consider any corrective action that might be in place to avoid potential duplication of work between the recipient of the notification and the corrective action.

- Avoid accumulation of resolved events by auto-clearing manually clearable events

  There is a subset of events that require manual clearing by the administrator. Examples of these are log-based events. As part of the workflow of managing the event, the administrator should manually clear the incident that contains the event after the issue has been resolved. (There is an option to clear an incident within Incident Manager). However, as a backup option, you can consider creating a rule that will automatically clear these events after a period of time. To avoid performance issues, this rule should be considered as a last resort to avoid the accumulation of events and not used as the primary way to clear a large number of events. To avoid prematurely clearing these events before someone has had a chance to look into them, the period of time specified should cover the agreed-upon period of time by which these events should have been resolved by the appropriate IT staff. For example, you might have an operational procedure requiring these events be responded to and resolved within 14 days. In that case, you can set up a rule that clears these manually clearable events only after 14 days. Note that any rule you set up will apply to events that occur *after* the rule has been created, i.e. the rule will not retroactively apply to events that already existed prior to the rule's creation. To clear these metric alert (events) that already exist, and as a general procedure to perform bulk clearing of metric alert (events), use the EM CLI verb *clear_stateless_alerts*.
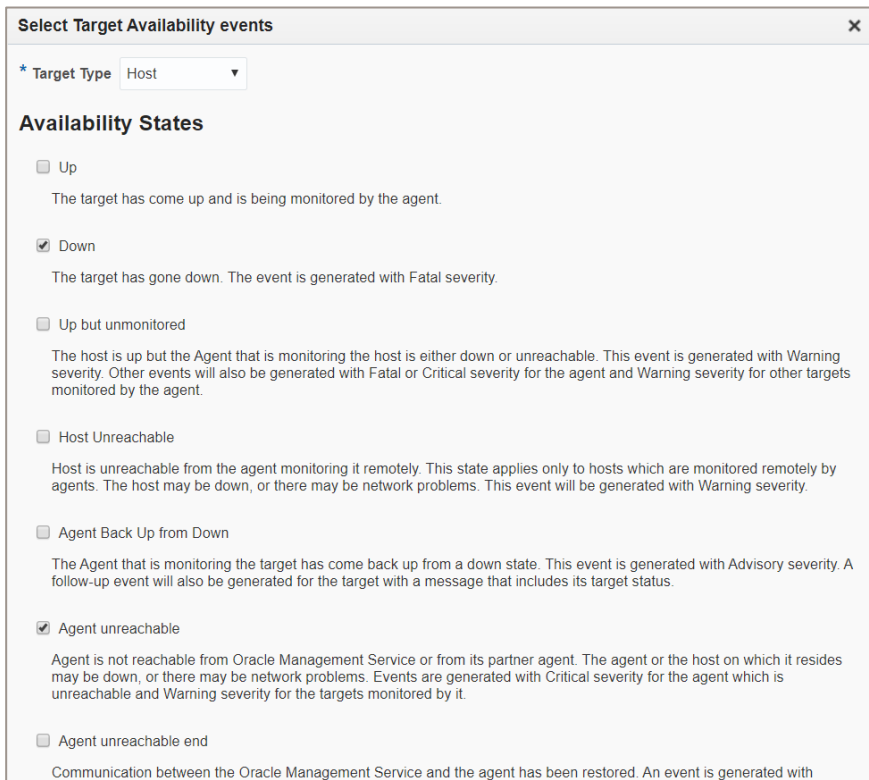
- Leverage the event compression functionality of incident rules to compress or combine related events into a single incident.

Too Many Emails When a Host Goes Down

When a host goes down, all targets on the host go down as well. This will cause either Down or Agent Unreachable events to be generated on the host, agent and all other targets on the host. In situations such as these, it might be desirable to only notify the primary administrator(s) who need to address the issue and avoid a flood of notifications to other administrators who own targets on the host server. The system administrator is most likely the primary administrator who needs to address the problem at

hand by investigating and restarting the host. Any host start up script should also include the startup of the agent to ensure monitoring is restored as soon as the host itself is back up. Do the following to send notifications only to the designated system administrators in this situation.

- For system administrators who would like to know if a host goes down, create an event rule to create an incident for selected 'Target Availability' events for the host target, specifically, the 'Down' and 'Agent Unreachable' availability states. If a host goes down and if its agent has a partner agent, then an event will be raised on the host for the 'down' status. If there is no partner agent, then an event will be raised on the host for the 'Agent Unreachable' status. Thus, choosing both states in the event rule will cover both scenarios.

Select Target Availability events   ✕

* **Target Type** Host ▼

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.

☐ Up but unmonitored

   The host is up but the Agent that is monitoring the host is either down or unreachable. This event is generated with Warning severity. Other events will also be generated with Fatal or Critical severity for the agent and Warning severity for other targets monitored by the agent.

☐ Host Unreachable

   Host is unreachable from the agent monitoring it remotely. This state applies only to hosts which are monitored remotely by agents. The host may be down, or there may be network problems. This event will be generated with Warning severity.

☐ Agent Back Up from Down

   The Agent that is monitoring the target has come back up from a down state. This event is generated with Advisory severity. A follow-up event will also be generated for the target with a message that includes its target status.

☑ Agent unreachable

   Agent is not reachable from Oracle Management Service or from its partner agent. The agent or the host on which it resides may be down, or there may be network problems. Events are generated with Critical severity for the agent which is unreachable and Warning severity for the targets monitored by it.

☐ Agent unreachable end

   Communication between the Oracle Management Service and the agent has been restored. An event is generated with

Figure 31. Choose 'Down' and 'Agent Unreachable' for host targets

For notification purposes, it is also recommended that you create an event rule for the 'Up but unmonitored' state of the host and send notifications for these. This state occurs when the agent that is monitoring the host is down but the host itself is still up. Since the agent is down, this also means there is no monitoring of the host's performance, space and other metrics.

- Create another rule to send notifications to the appropriate administrators for the host incident.

- Owners of the targets on the host will want to be notified if their targets are down (outside of the host down scenario). For these administrators, create a rule to create an incident for the 'Target Availability' event for these targets (database, listener, etc) and choose the 'Down' state. Also create a rule to send notifications for this incident:

Figure 32. *Choose 'Down' availability state for targets on the host.*

With the rules set up in this manner, then if a host goes down, only the system administrators will receive notifications for the 'agent unreachable' or 'host down' event/incident. All other target owners will not be flooded with notifications because their targets are in 'agent unreachable' state and not down state. Only if their targets are detected to be down (with the host up), will they be sent notifications.

Using Root Cause Analysis for Target Down Events

When targets are detected to be down, root cause analysis for target down events will automatically occur. This results in identifying whether a target down event is the root cause of other target down events, or if the target down event is a symptom of other target down events, i.e. it has gone down because the target on which it depends on has gone down. A new attribute called *Causal Analysis Update* has been introduced to identify whether the event is a symptom or root cause. As additional target down events come in, these events are analyzed and could result in updates to prior causal analysis results. For example, a target down event that was neither root cause nor symptom could later be identified as a root cause of another target down event when that second target down event is reported 5 minutes later. Incidents containing these events will be classified accordingly, either as root cause or symptom.

When root cause analysis is performed on target down events, the dependency relationship between targets that are down is used in analyzing whether the target down event is a 'root cause' or 'symptom'. The c*ausal analysis update* attribute is used to show the results of this analysis. If the event is not related to any other target down event, it will not have any value for the *causal analysis update* attribute. The incident containing the target down event will have its *causal analysis update* attribute reflect the same value as the *causal analysis update* attribute of the event.

You can use the root cause analysis feature in a several ways:

- Filter the System Dashboard and/or Incident Manager UI to exclude symptoms

In scenarios wherein a target down could in turn cause other targets to go down, you may want to focus primarily on the 'root cause' event (or incident containing this) and exclude symptom events (or incidents containing these). In the System Dashboard, you can specify an option to 'Exclude Symptoms' from the Incidents and Problems table.
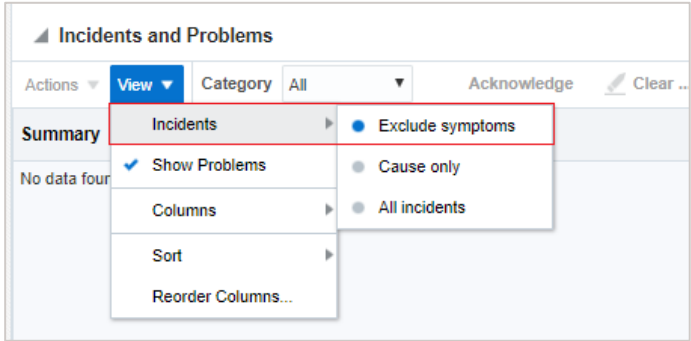

Figure 33. *You can exclude symptoms from the Incidents and Problems table of the System Dashboard.*

Similarly, in Incident Manager you can create a view and use the search field Causal Analysis Update to exclude symptom events/incidents from your custom view.
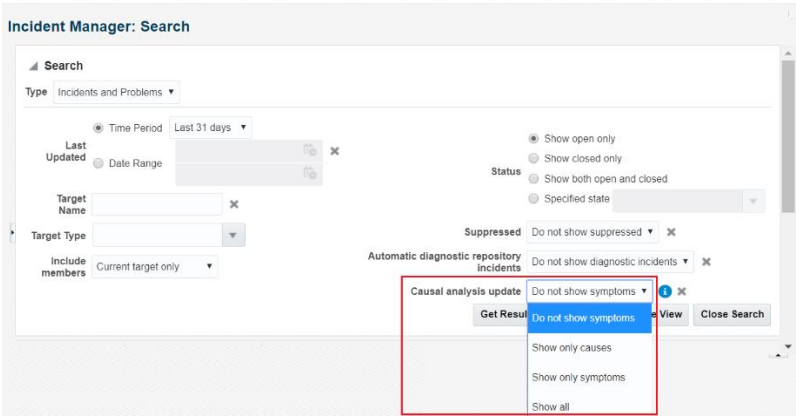
Figure 34. *You can filter out symptom events/incidents in your Incident Manager views.*

- Create rules to create incidents only on non-symptom events

You may also want to leverage the root cause analysis feature in rules that create incidents. Specifically, in scenarios where there could be a multitude of target down (symptom) events and only a few (or single) target down event, you might want to create incidents and send notifications on only the non-symptom events. To do this, create an event rule and select the Target Availability event type and Target Availability State = Down. Also choose the 'Causal analysis update' option and select 'event is marked as a cause' and 'event is not a cause and not a symptom' as shown in the figure below.

Figure 35. *When creating an event rule, you can exclude symptoms as part of the rule criteria.*

With the selections above, events that are marked as symptoms will be not be included in the rule. It is possible that an event is initially reported as neither a root cause nor symptom. However, a few minutes later, when other target down events are reported and analyzed, the original target down event might later be classified as a symptom. Hence, you might consider introducing a small delay (e.g. 5 to 10 minutes) in your rule to allow for causal analysis to have more recent updates. For example, if your rule creates incidents based on non-symptom events, and you did not have a delay in the rule, you could potentially have an incident created for an event that later turned out to be a symptom. If the benefits of getting a potentially updated causal analysis outweigh the need to perform an immediate action on an event (such as creating an incident or sending notification), add a small waiting period in your rule by making the rule action a conditional action based on a duration. This means you want the action in the rule to execute only if it meets a specified condition. In general, the condition for an action can be based on criteria such as event severity or it can be based on a length of time the event has been open. This latter option is what we want to use to introduce a small waiting period in the event action. The figure below gives a specific example of how you could configure the action in your rule with a small waiting period. In the figure below, we've selected the option 'Only execute the actions if specified conditions match' and also 'Event has been open for specified duration' and duration of 5 minutes. Introducing a 5-minute waiting period allows time for other target down events to be reported and an updated causal analysis to occur before acting on the event, i.e. before creating an incident for the event if it is not a symptom.

Figure 36. *You can add a duration-based action with a small delay to allow for more updated causal analysis.*

Managing Diagnostic Incidents and Problems

Most incidents created in Enterprise Manager can be considered operational in nature and are expected to be resolved by a datacenter's IT staff. There is, however, a special type of incident called a 'Diagnostic Incident' or ADR (Automatic Diagnostic Repository) Incident. These are incidents that are automatically raised by Oracle software when it encounters a critical error in its software code. As such, it is expected that IT staff reach out to Oracle (via Oracle Support) to help resolve these incidents. When such incidents occur, an incident (object) and problem (object) are generated automatically in Enterprise Manager.   The incident represents an occurrence of the error and the problem (object) is meant to help resolve these incidents by addressing the root cause of these incidents. The problem object contains the incident's "problem key", i.e. something which uniquely identifies the software error. Each occurrence of the incident will create a (diagnostic) incident that will automatically be associated with the problem. To permanently resolve these incidents, the recommendation is to manage the problem by using Support Workbench to package the appropriate diagnostic logs and open a Service Request (SR) with Oracle Support. A link to Support Workbench is available in the Guided Resolution section of the problem in Incident Manager:
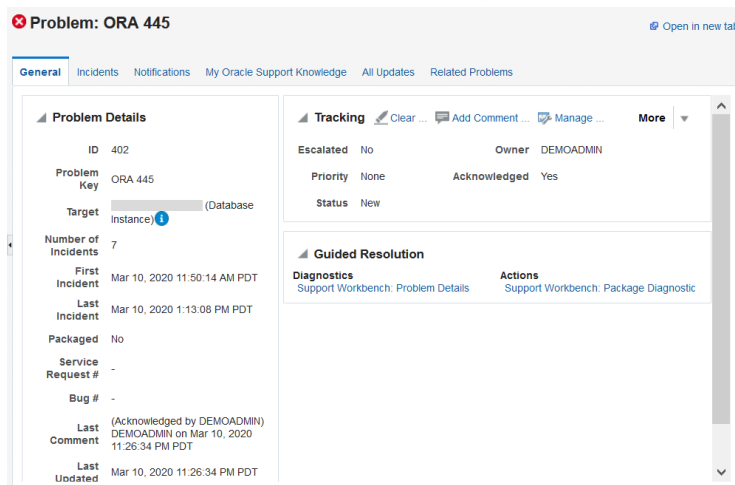
Figure 37. *A Problem (object) in Incident Manager enables management of critical errors in Oracle software.*

The Support Workbench will assist in the creation of a draft SR including the packaging of the appropriate diagnostic log information. Once a draft SR is created, you should complete the creation of the SR within My Oracle Support. The SR will automatically be associated with the problem so that it can be easily tracked with Incident Manager. The Problem object should be assigned an owner, (resolution) status, etc. in the same way that other incidents are owned and managed in Incident Manager. Once a fix to the problem has been provided (e.g. patch available for the software error), after the fix has been applied, you can manually close the Problem in Incident Manager. This will automatically close all its related incidents as well.

CONCLUSION

Comprehensive monitoring of applications and their supporting infrastructure continues to be a critical requirement of today's datacenters. To implement an effective monitoring solution, it is important for IT staff to have the right management tool used in conjunction with best practice standards and processes. It is also equally important to be able to automate IT processes to reduce administrative overhead and enable IT staff to manage more with less while still providing high quality services to their users. Oracle Enterprise Manager 13c provides a rich set of capabilities to meet the enterprise monitoring demands of today's dynamic datacenters and support and automate many of its operational processes.   The strategies outlined in this paper are designed to help IT staff plan and optimize their use of Enterprise Manager 13c to meet their monitoring requirements in an effective and scalable way.

CONNECT WITH US

Call +1.800.ORACLE1 or visit oracle.com.
Outside North America, find your local office at oracle.com/contact.

blogs.oracle.com   facebook.com/oracle   twitter.com/oracle

Strategies for Scalable, Smarter Monitoring Using Oracle Enterprise Manager Cloud Control 13c
May, 2024
Author: Ana McCollum
Contributing Authors: Karilyn Loui, H Siddaramaiah