Oracle® Communications

Policy Management

Configuration Management Platform Cable User's Guide

Release 12.5

E96549-02

October 2019

Policy Management Configuration Management Platform Cable User's Guide, Release 12.5

E96549-02

Contents

2 Configuring the Policy Management Topology

3 Managing Multimedia Policy Engine Devices

4 Configuring Protocol Routing

5 Configuring Advanced Device Settings

6 Configuring Debug Logs

7 Managing Network Elements

8 Managing Record Keeping Servers

9 Managing Event Messaging

10 Managing Management Agent Servers

11 Managing Bandwidth on Demand

12 System-Wide Reports

13 Upgrade Functions

14 Global Configuration

15 System Administration

A CMP Modes

B Generated Statistics

About This Guide

This chapter contains an overview of the manual, describes how to obtain help, where to find related documentation, and provides other general information.

Introduction

This guide describes variables that can be used in policy rules. These variables provide information about the device, subscriber, or quota for which a policy rule is being executed.

How This Guide is Organized

The information in this guide is presented in the following order:

- About This Guide provides general information about the organization of this guide, related documentation, and how to get technical assistance.
- Oracle Communications Policy Management System provides an overview of the Multimedia Policy Engine (MPE), which manages multiple network-based client sessions; the network in which the MPE device operates; policies; and the Configuration Management Platform (CMP) system, which controls MPE devices and associated applications.
- Configuring the Policy Management Topology describes how to set the topology configuration.
- Managing Multimedia Policy Engine Devices describes how to use a CMP system to configure and manage the MPE devices in a network.
- Configuring Protocol Routing describes how to configure protocol routing.
- Configuring Advanced Device Settings describes how to specify advanced settings for MPE devices.
- Configuring Debug Logs describes how to configure device- and system-level settings for troubleshooting system operations using logs.
- Managing Network Elements describes how to manage network elements.
- Managing Record Keeping Servers describes how to configure and manage the record keeping server (RKS) that receives event messages.
- Managing Event Messaging describes how to configure and manage event messaging.
- Managing Management Agent Servers describes how to configure and manage management agent (MA) servers.
- Managing Bandwidth on Demand describes the basic configuration for Bandwidth on Demand (BoD) devices in the CMP system.
- System-Wide Reports describes the reports available on the function of Policy Management systems in your network.

- Upgrade Functions describes the Upgrade Manager pages and the elements found on the pages.
- Global Configuration describes how to configure global settings in the CMP system.
- System Administration describes functions reserved for CMP system administrators.
- The appendix, CMP Modes, lists the functions available in the CMP system, as determined by the operating modes and sub-modes selected when the software is installed.
- The appendix, Generated Statistics, lists statistics generated as part of scheduled tasks.

Scope and Audience

This document is intended for the following trained and qualified service personnel who are responsible for Policy Management devices:

- Application administrators, who install and upgrade Policy Management applications and perform advanced system administration
- Operators, who monitor Policy Management systems daily and perform adjustments
- System administrators, who control access to the CMP system
- System architects, who design carrier network system architectures, including planning for Policy Management systems
- Network administrators, who manage carrier networks

Related Publications

For information about additional publications related to this document, refer to the Oracle Help Center site. See Locate Product Documentation on the Oracle Help Center Site for more information on related product publications.

Policy and Protocol Specifications

The following specifications provide information on protocols:

- PCMM CableLabs specifications:
  - PKT-SP-MM-I05: PacketCable™ Multimedia Specification
  - PKT-SP-DQOS-I12-050812: PacketCable™ Dynamic Quality-of-Service Specification
- Internet Engineering Task Force (IETF) specifications:
  - RADIUS RFCs:
    - RFC 2865: RADIUS
    - RFC 2866: RADIUS Accounting

        - RFC 3576: Dynamic Authorization Extensions to RADIUS
    - Diameter RFCs:
        - RFC 3539: Authentication, Authorization and Accounting (AAA) Transport Profile
        - RFC 3588: Diameter Base Protocol
    - RFC 3164: The BSD syslog Protocol

Locate Product Documentation on the Oracle Help Center Site

Oracle Communications customer documentation is available on the web at the Oracle Help Center (OHC) site, http://docs.oracle.com. You do not have to register to access these documents. Viewing these files requires Adobe Acrobat Reader, which can be downloaded at http://www.adobe.com.

1. Access the Oracle Help Center site at http://docs.oracle.com.
2. Click `Industries`.
3. Under the Oracle Communications subheading, click the `Oracle Communications documentation` link.

   The Communications Documentation page appears. Most products covered by these documentation sets will appear under the headings "Network Session Delivery and Control Infrastructure" or "Platforms."
4. Click on your Product and then the Release Number.

   A list of the entire documentation set for the selected product and release appears.
5. To download a file to your location, right-click the `PDF` link, select `Save target as` (or similar command based on your browser), and save to a local folder.

Customer Training

Oracle University offers training for service providers and enterprises. Visit our web site to view, and register for, Oracle Communications training:

http://education.oracle.com/communication

To obtain contact phone numbers for countries or regions, visit the Oracle University Education web site:

www.oracle.com/education/contacts

My Oracle Support

My Oracle Support is your initial point of contact for all product support and training needs. A representative at Customer Care Center can assist you with My Oracle Support registration.

Call the My Oracle Support main number at 1-800-223-1711 (toll-free in the US), or call the Oracle Support hotline for your local country from the list at http://www.oracle.com/us/support/contact/index.html. When calling, make the selections in the sequence shown below on the Support telephone menu:

1. Select 2 for New Service Request
2. Select 3 for Hardware, Networking and Solaris Operating System Support
3. Select one of the following options:
   - For Technical issues such as creating a new Service Request (SR), Select 1
   - For Non-technical issues such as registration or assistance with MOS, Select 2

You are connected to a live agent who can assist you with My Oracle Support registration and opening a support ticket.

My Oracle Support is available 24 hours a day, 7 days a week, 365 days a year.

Emergency Response

In the event of a critical service situation, emergency response is offered by the Customer Access Support main number at 1-800-223-1711 (toll-free in the US), or by calling the Oracle Support hotline for your local country from the list at http://www.oracle.com/us/support/contact/index.html. The emergency response provides immediate coverage, automatic escalation, and other features to ensure that the critical situation is resolved as rapidly as possible.

A critical situation is defined as a problem with the installed equipment that severely affects service, traffic, or maintenance capabilities, and requires immediate corrective action. Critical situations affect service and/or system operation resulting in one or several of these situations:

- A total system failure that results in loss of all transaction processing capability
- Significant reduction in system capacity or traffic handling capability
- Loss of the system’s ability to perform automatic system reconfiguration
- Inability to restart a processor or the system
- Corruption of system databases that requires service affecting corrective actions
- Loss of access for maintenance or recovery operations
- Loss of the system ability to provide any required critical or major trouble notification

Any other problem severely affecting service, capacity/traffic, billing, and maintenance capabilities may be defined as critical by prior discussion and agreement with Oracle.

List of Figures

List of Tables

1

Oracle Communications Policy Management System

This chapter provides an overview of the Policy Management system and its components. The major components include:

- The Oracle Communications Policy Management Configuration Management Platform (CMP) system controls MPE devices and associated applications.
- The Multimedia Policy Engine (MPE) device manages multiple network-based client sessions.
- The Bandwidth on Demand (BoD) Application Manager (AM) provides an abstract interface for creating dynamic service requests allowing application developers to integrate dynamic Quality of Service (QoS) resources into any application.
- The Management Agent server (MA Server) is a distributed system that collects topology and network information for use with CableLabs' PacketCable Multimedia (PCMM) protocol message routing and policy decisions.

Introduction to Policy Management

Figure 1-1 illustrates a simple cable network. At the subscriber's home network, customer premises equipment (CPE) devices (for example, routers, set-top boxes, and computers) connect to the cable modem's (CM) local area network (LAN). When the CM powers on, the modem connects the subscriber's home network to the operator's core network over hybrid fiber/coaxial (HFC) to a CMTS (cable modem termination system). The CMTS routes packets between the downstream, subscriber-side HFC channels and the operator's upstream back office and core network channels.

Figure 1-1 Policy Management System and its Components

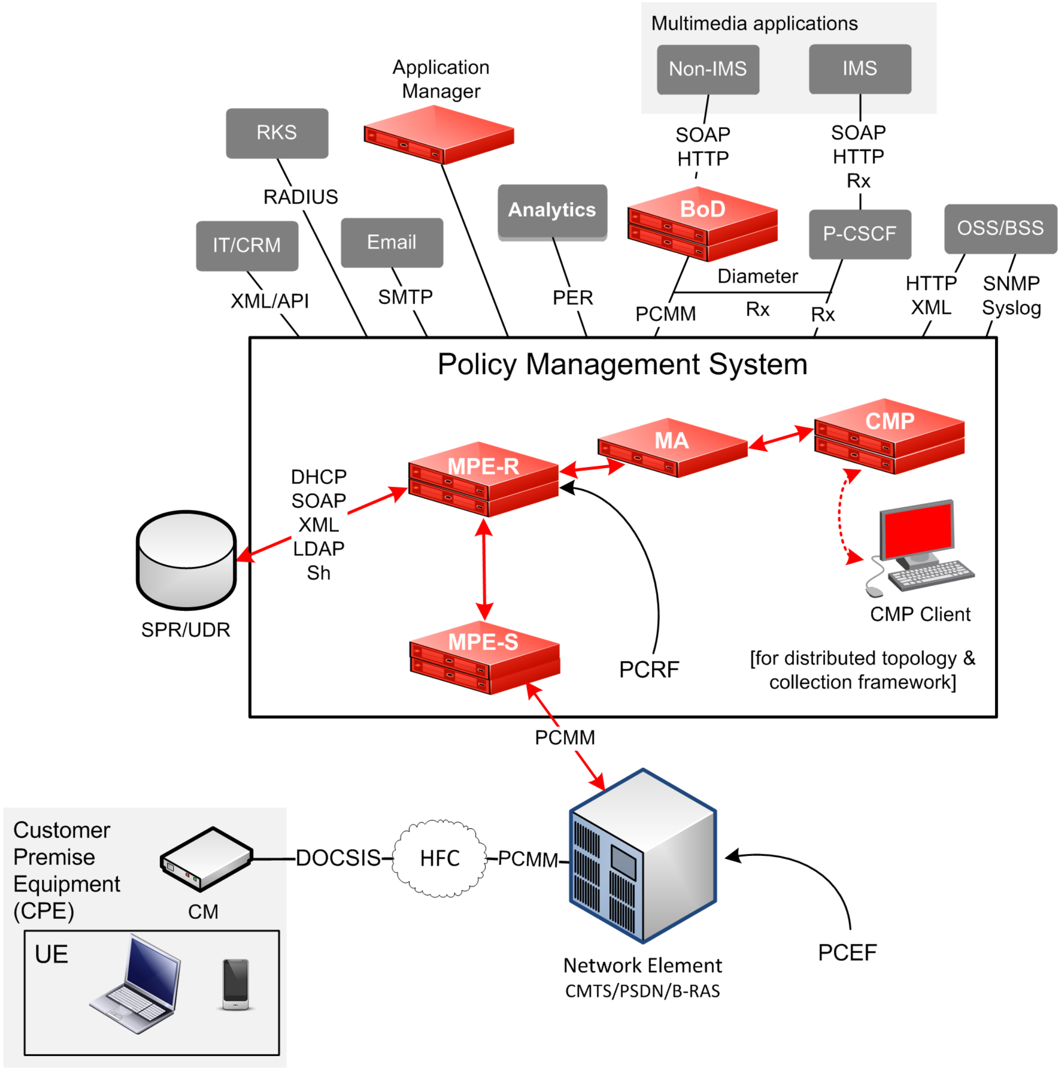

About Policy Management Service Flows

The Oracle Communications Policy Management system provides the mechanism that allows communications service providers (carriers) to control and charge for subscribers' access to network resources. The central component of the Policy Management system is the Multimedia Policy Engine (MPE) device. The MPE device functions as a dynamic policy and charging rules function (PCRF).

The MPE device establishes service flows between subscribers and the application servers that provide multimedia services.

A service flow is activated only after the contents of its QoS request are examined and approved by the MPE device. If approved, the request is forwarded to the intended destination network node.

As illustrated in Figure 1-1, when a subscriber wishes to open an IP-streaming session, the following actions occur:

1. The cable modem (CM) connects through the hybrid fiber/coaxial (HFC) network to the network equipment, for example, a Cable Modem Termination System (CMTS).

2. An application receives the subscriber's request and sends a QoS request to the MPE device for the associated network element, requesting that the requested network resources be provisioned for use for the requesting application.

3. The MPE device examines the QoS request before it gets to the network element and processes the request against the policy rules within its policy repository.

4. The MPE device then makes a decision based on the defined policy rules to accept or reject the request.

5. Depending on the decision made, the MPE device performs one of the following actions:

   - **Accepts** the QoS request and forwards it to the network element, where the required network resources are provisioned, allowing the service flow for IP-streaming to be admitted and activated.

   - **Rejects** the QoS request, in which case an error message is sent back to the application and no service flow is established.

   Note: When provisioned resources are no longer required and deleted, the network resources are recovered for use elsewhere.

Major Components of Policy Management

The Oracle Communications Policy Management system (see Figure 1-1) provides the capability to:

- Create dynamic Quality of Service (QoS) service flows between subscribers and application servers that provide multimedia services
- Manage network resources more efficiently and flexibly
- Manage policies

The major components of the Oracle Communications Policy Management system include:

- #unique_38
- #unique_39
- Bandwidth on Demand Device
- Management Agent Server
- Record Keeping Server
- Notification Server

Configuration Management Platform Server

The CMP server provides the following functionality:

- Provides the policy console for managing the following:
  - Policy objects like MPE and BoD devices
  - Policies

  - Configuration templates
  - Network elements
  - System topology
- Contains a centralized database of policy rules, policy objects, and network objects
- Communicates Policy Management network management information with network management systems (NMSs) using Simple Network Management Protocol (SNMP)

See #unique_45/unique_45_Connect_42_FIGURE_1 for more information.

Multimedia Policy Engine Device

MPE devices perform the following actions:

- Provides policy and charging rules function (PCRF) for controlling policy decisions and flow-based charging

> **Note:** Refer to *Policy Wizard Reference* for information on how to create, organize, and manage policies and the elements they control.

- Obtains subscriber information, evaluates the applicable policies, and directs the PCEF network element to handle the session based on policy rules when the MPE device receives a request for a policy decision for a subscriber session
- Communicates with clients using Diameter application interfaces (for example, Rx)
- Sends Short Message Service (SMS) or Simple Mail Transfer Protocol (SMTP) notifications to subscribers, if enabled
- Operates as two-level devices (MPE-R and MPE-S) for handling large volume loads

> **Note:** You can increase the capacity of the Policy Management network by adding additional MPE devices.

See Multimedia Policy Engine Devices for detailed information.

Bandwidth on Demand Device

BoD Application Manager provides an interface for creating dynamic service requests that allow an application developer to integrate dynamic Quality of Service (QoS) resources into nearly any application. This is achieved by providing Hypertext Transfer Protocol (HTTP) and Simple Object Access Protocol (SOAP) based interfaces that can be integrated into most application development environments.

See Bandwidth on Demand Application Manager for more information. Refer to *Bandwidth on Demand Cable User's Guide* for detailed information.

Management Agent Server

The MA Server collects topology and network information for use with CableLabs' PacketCable Multimedia (PCMM) protocol message routing and policy decisions.

See Management Agent Server for more information.

Record Keeping Server

A Record Keeping Server (RKS) is a repository for PCMM event messages. It gathers billing event messages and passes them on to OSS/BOSS. To use event messaging, you must configure profiles for one or more Record Keeping Servers, and then associate them with MPE devices.

See Managing Record Keeping Servers for more information.

Notification Server

Notifications are generated by a policy action. The destination, content and attributes of the notification are configurable by the operator and allow for flexible notifications within an HTTP request message.

Note: Notification servers are only available when SMPP or XML mode is enabled. See CMP Modes for details.

Refer to *Policy Wizard Reference* for details on managing notification servers.

Hardware System Requirements

The various Policy Management applications (for example, MPE and BoD devices) run on a variety of hardware platforms:

- Oracle Server X5-2 and Oracle X5-2 rack-mount servers (RMS)
- Hewlett-Packard (HP) Enterprise ProLiant DL360/Gen8 RMS

Note: Depending on the hardware and the selected enclosure type, the system also requires the appropriate switches and connectors.

On all the hardware platforms Policy Management applications can execute as virtual machines within a network functions virtualization (NFV) infrastructure.

Maintaining and Managing the System

Oracle Communications Policy Management uses external utilities, applications, and tools for maintaining and managing devices.

TPD and TVOE

Oracle Communications Policy Management servers use the Tekelec Platform (TPD, also known as Tekelec Platform Distribution) operating system. TPD provides tools that configure third-party hardware and platform components that make up Platform 7.5. Configurable hardware components include HP Enterprise ProLiant and Oracle X5-2 rack mount servers (RMS) and Cisco switches, HP c7000 enclosures with HP blade servers, HP and Cisco switches, and HP external storage systems. Platform components include the firmware for various hardware components as well as the Platform Management & Configuration (PMAC) application to provision and manage those components when hosting feature applications.

TPD is the operating system for many Oracle Communications products including:

- Oracle Communications Policy Management
- Oracle Communications Subscriber Database Management
- Oracle Communications User Data Repository
- Oracle Communications Enhanced Subscriber Profile Repository
- Oracle Communications Diameter Signaling Router

In addition to TPD, TVOE (Tekelec Virtual Operating Environment) platform application provides tools to manage virtual machines (VMs).

Refer to Oracle Communications Tekelec Platform documentation: *Configuration Guide* and *TPD Initial Product Manufacture* for more information.

Refer to Oracle Communications Policy Management documentation: *Bare Metal Installation Guide* and *Platform Configuration User's Guide* for information specific to installing and configuring the Policy Management system.

PMAC

A network management server with TPD and TOE installed acts as a Virtual Host environment and hosts the Product Management and Configuration (PMAC, also known as PM&C) application.

The PMAC application, configured on Policy Management devices during initial system installation, provides system-level management functions at specific sites. The PMAC application supports platform-related maintenance, software installation, provisioning, and upgrade activities. PMAC uses an internal control network (IntCtrl) with internal, non-routable addresses. The PMAC application is independent of the Oracle Communications Policy Management system.

See the Oracle Communications Tekelec Platform PMAC documents *PM&C Incremental Upgrade* and *PM&C Disaster Recovery* for additional information.

See the Oracle Communications Policy Management documents *Bare Metal Installation Guide* and *Platform Configuration User's Guide* for additional information specific to Policy Management.

iLO and ILOM

In addition to signaling interfaces and networks, the Policy Management system allows for hardware platform management through out-of-band remote access to individual devices, hardware enclosure onboard administrators (OA), and enclosure switches.

The Hewlett Packard (HP) Enterprise hardware platform management tool is called Integrated Lights-out (iLO) management. This application operates independently of the Policy Management applications running on individual devices. The iLO network allows for access across devices restarts, which is needed for maintenance activities such as installations and upgrades.

Note: For support purposes, the iLO addresses must be remotely accessible.

Oracle Integrated Lights Out Manager (ILOM) manages and monitors Oracle Server X5-2.

Refer to the end-user documentation for your specific hardware platform for detailed information on using these management applications.

About Specialized Communications Protocols

The Oracle Communications Policy Management system supports a variety of specialized communications protocols.

About RADIUS Protocol

The RADIUS (Remote Authentication Dial-In User Service) networking protocol supports Authentication, Authorization, and Accounting (AAA) management for users. One of the Internet Engineering Task Force (IETF) standards, RADIUS is a client/server protocol that manages network access.

The RADIUS protocol uses access credentials (like a username and password) to authenticate and authorize (via a RADIUS server) access to the requested network for a user or an equipment. After the access request is accepted, an `Accounting Start` message is transmitted by the network access server (NAS) to the RADIUS server to indicate the start of the network access for the user or equipment.

When the user or equipment closes the network session, the NAS sends an `Accounting Stop` message to the RADIUS server including information for billing, like usage time and packets/data transferred.

The Policy Management system uses the RADIUS protocol to connect with P-CSCF servers and Record Keeping Servers (RKS). See Policy and Protocol Specifications for IETF RADIUS specifications.

About Diameter Protocol

The Diameter networking protocol supports Authentication, Authorization, and Accounting with Secure Transport (AAAS) management for users and equipment. One of the Internet Engineering Task Force (IETF) standards, Diameter Base protocol implements a peer-to-peer architecture where a host acts either as a client or a server depending on its deployment within the network.

The Diameter protocol defines Diameter messages that send commands or deliver notifications to other Diameter-enabled devices.

The Policy Management system uses Diameter interfaces (for example, Gx and Rx) to connect to many of its components and with components outside of the system. See #unique_45/unique_45_Connect_42_FIGURE_1 for the Diameter-specific interfaces.

See #unique_58 for a list of supported IETF Diameter specifications.

About PCMM Protocol

The Policy Management system uses PCMM (PacketCable Multimedia) to connect CMTS network elements to MPE devices. Additionally, MPE devices use the PCMM protocol to notify Record Keeping Servers (RKSes) of event messages. An MPE device or a CMTS can send event messages. A CMTS sends event messages only when instructed to do so by the MPE device (via signaling that is part of the PCMM protocol).

The MA Server (Management Agent) collects topology and network information for use with PCMM message routing and policy decisions.

See Policy and Protocol Specifications for a list of supported PCMM specifications.

More About Policy Management Components

Oracle Communications Policy Management supports a variety of components that perform a wide range of functions.

Multimedia Policy Engine Devices

The Multimedia Policy Engine (MPE) device provides a policy and charging rules function (PCRF) as defined in the 3rd Generation Partnership Project (3GPP) technical specification “Policy and charging control architecture” (TS 23.203). It fully supports all 3GPP Release 7, 8, 9, and 10 policy and charging control (PCC) interfaces. The MPE device includes a simple, powerful, and flexible policy rules engine. The policy rules engine operates on triggers from any interface or from internal timers; evaluates conditions; and then performs appropriate actions. Through the use of policy rules, you can modify the behavior of an MPE device dynamically as it processes protocol messages.

A policy is a set of operator-created business rules. These business rules control how subscribers, applications, and network resources are used. Policies define the conditions and actions used by a carrier network to determine:

- How network resources are allocated and used
- How applications and subscribers are treated

Refer to *Policy Wizard Reference* for information on how to create, organize, and manage policies and the elements they control.

To accommodate high volume load demands, the MPE device can function in a two-level hierarchical environment. As a Level-2 device (called an MPE-S, or serving, device), it statefully services subscriber flows. As a Level-1 device (called an MPE-R, or routing, device), it statelessly routes subscriber flows to MPE-S devices. Stateless Rx routing is enabled on the MPE-R (see Configuring PCMM Routing) and Rx-to-PCMM translation (see Configuring Rx-to-PCMM Routing) occurs at MPE-S devices. Normally, high volume load is split by one to two MPE-R devices among 10 to 14 MPE-S devices.

See Managing Multimedia Policy Engine Devices for details on managing MPE devices.

Bandwidth on Demand Application Manager

The Bandwidth on Demand (BoD) Application Manager (AM) server provides a simplified and abstract interface for creating dynamic service requests that allow the application developer to integrate dynamic QoS resources into nearly any application. This is achieved by providing HTTP and SOAP based interfaces that can be easily integrated into most application development environments.

Additionally, the BoD AM server maintains and manages all of the state information that is associated with each request, allowing applications to be stateless in their operation.

The BoD AM server presents a SOAP-based remote procedure call (RPC) interface and a pure HTTP request interface. These interfaces provide similar functionality and are designed to let application developers use whichever interface best suits their application.

For example, the HTTP interface allows a parameterized URL to be associated with the `onClick` action of a turbo button or simply allow any application to embed an HTTP `POST` message to dynamically adjust service. Alternatively, the SOAP interface provides easy session control through the RPC mechanism. The decision whether to use HTTP or SOAP largely depends on the personal preferences of the developers of the calling application.

Within the BoD AM server, you can define a number of service names that translate into a particular traffic profile. For example, a generic service name `turboService` could be defined with an associated best-effort upstream flow and a high-priority downstream flow. Additionally, a specific service name such as `uploadService` could simply define a high-priority upstream flow.

Each of the interface bindings allows an application to create a new session, specifying a service name and also supplying a number of specialization parameters such as bandwidth. For example, within a web portal, a number of links or buttons can be defined, all of which use the same `turboService` profile, each specifying a different upstream and downstream bandwidth. This can be used to vary the resulting QoS flows based either on the application context or perhaps a subscriber tier.

The BoD AM server also allows calling of an application to specify the duration of QoS resource allocation. The application may choose to completely manage the life cycle of the resources, in which case it is the responsibility of the application to free the resources at the appropriate time, either after a defined period, or when an application has completed its function. Alternatively, the application may simply tell the BoD AM server to keep the resources active for a specified time, or until there is inactivity for a defined period.

BoD AM servers are configured and managed through the Oracle Communications Policy Management Configuration Management Platform (CMP) system. For information on using the BoD AM server, refer to *Bandwidth on Demand Application Manager User Guide*.

Management Agent Server

The Management Agent (MA) server is designed specifically for network architectures that require a distributed topology and collection framework. An MA Server is not an actively managed device, but rather a distributed system that collects topology and network information for use with PCMM protocol message routing and policy decisions.

The MA Server sits between the CMP system and one or more MPE devices. The number of MA Servers and MPE devices depends on the size of the network. The groupings that define the MPE devices managed by an MA Server and the MA Servers managed by the CMP system depends on the network topology.

See Managing Management Agent Servers for details on managing MA Servers.

Record Keeping Servers

A Record Keeping Server (RKS) is a repository for PCMM event messages. It gathers billing event messages and passes them on to BOSS. To use event messaging, you must configure profiles for one or more RKSes and associate them with MPE devices.

See Managing Event Messaging and Managing Record Keeping Servers for more information.

Notification Servers

> **Note:** Notification servers are only available when SMPP or XML mode is enabled. See #unique_65 for details.

Within the Policy Management system, an MPE device configured for generic notifications connects to a notification server over HTTP. See Configuring MPE Protocol Options.

The notification server processes event notifications in response to policy actions for HTTP messages. These policy actions include the ability to:

- Send notifications using a dynamic URL
- Send notifications using a static URL

Refer to *Policy Wizard Reference* for details on managing notification servers.

The audit log records all notification server actions (create, modify, and delete), policy creation and modification, and associations (both policy servers and configuration templates).

Configuration Management Platform Server

The Configuration Management Platform (CMP) server provides centralized management and administration of policy rules, Policy Management devices, associated applications, and manageable objects, all from a single management console. This browser-based management console supports the following features and functions:

- Configuration and management of MPE devices
- Configuration and management of MRA devices
- Configuration and management of Mediation devices
- Configuration of connections to subscriber profile repository (SPR) servers, including Oracle Communications User Data Repository (UDR) and Oracle Communications Enhanced Subscriber Profile Repository (ESPR) systems
- Definition of network elements
- Management and deployment of policy rules
- Management of objects that can be included in policy rules
- Monitoring of individual product subsystem status
- Administration and management of CMP users
- Upgrading the software on Policy Management devices

Specifications for Using the CMP Server

You interact with the CMP server through a browser-based graphical user interface (GUI). To use the GUI, Oracle recommends the following:

Web Browsers for Wireless and Cable modes

- Mozilla Firefox® release 31.0 or later
- Google Chrome version 40.0 or later

Monitor

- Resolution of 1024 x 768 or greater

Logging in to the CMP System

The CMP system supports either HTTP or HTTPS access. Access is controlled by a standard username and password login scheme.

Note: If you are using the CMP system for the first time, Oracle recommends that you change the default password for your user name. See #unique_70 for details.

Before logging in, you need to know the following:

- The IP address of the CMP system
- Your assigned username
- Your account password

Note: The profile `admin` has full access privileges and is the assumed profile used in all procedures described in this document. You cannot delete this user profile.

To log in to the CMP system:

1. Open a web browser and enter the IP address for the CMP system.

   Note: You can configure the title and text that appear on the login page. For information on changing this page, see Configuring System Settings.

2. Enter your **Username** and **Password**.

   Note: See your System Administrator if you experience problems logging in.

3. Click **Login**.

   The CMP main page opens.

You are logged in.

Logging In to a Standby or Secondary-Site CMP System

Most of the procedures in this document begin with you logged in to the active server of the primary CMP system. A few procedures require you to log in to the active server of a secondary CMP system, and it is also possible to log in to the standby server of a CMP cluster. The functions available on other servers are limited.

- If you log in to the standby server of a primary CMP cluster, the work area displays a message indicating that you are signed into the Primary Standby server.
- If you log in to the active server of a secondary CMP cluster, the work area displays a message indicating that you are signed in is the Secondary Active server.
- If you log in to the standby server of a secondary CMP cluster, the work area displays a message indicating that you are signed in is the Secondary Standby server.

In all cases, you are limited to the **Platform Setting** functions: **Platform Configuration Settings** and **Topology Settings**. Status information for all other servers is not available and is displayed as **out-of-service**.

GUI Overview

You interact with the CMP system through an intuitive and portable graphical user interface (GUI) supporting industry-standard web technologies (the SSL family of secure communication protocols, HTTP, HTTPS, IPv4, and XML). Figure 1-2 shows the layout of the CMP GUI.

Figure 1-2 Layout of the CMP Window—Cable Mode

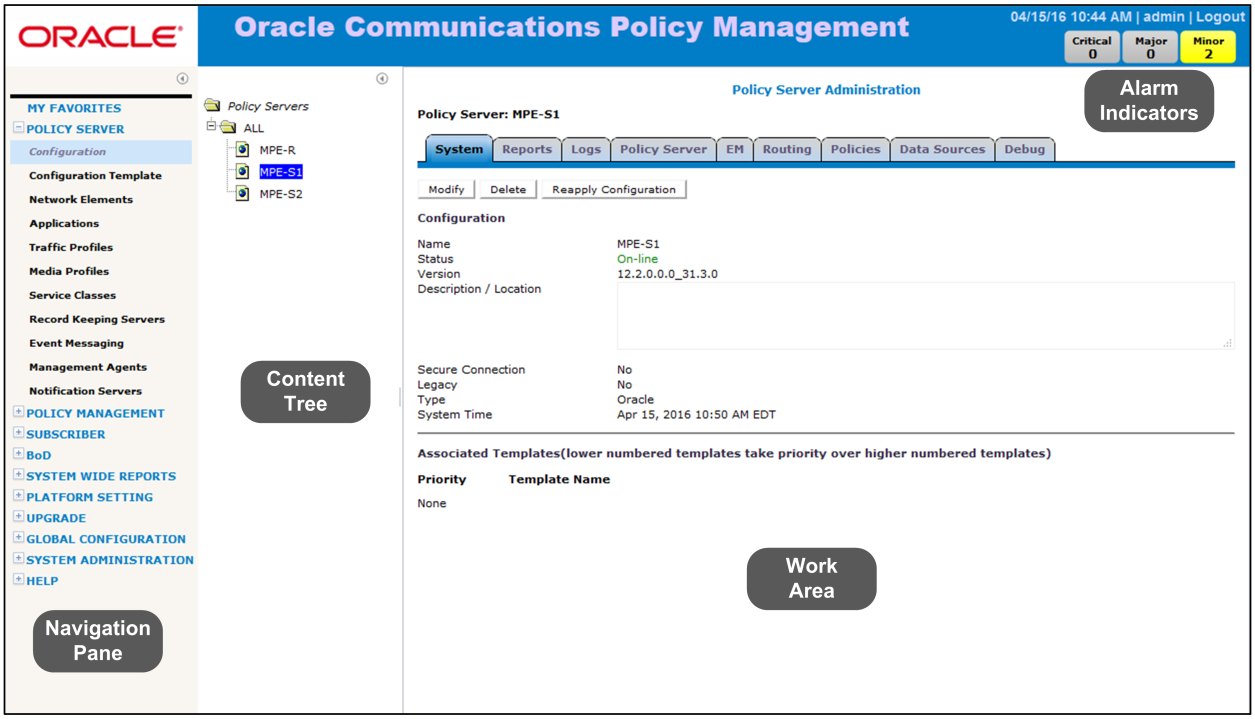

Policy Server: MPE-S1

Configuration

Name	MPE-S1
Status	On-line
Version	12.2.0.0.0_31.3.0
Description / Location	
Secure Connection	No
Legacy	No
Type	Oracle
System Time	Apr 15, 2016 10:50 AM EDT

Associated Templates(lower numbered templates take priority over higher numbered templates)

Priority	Template Name
None	

The CMP system's window is divided into the following sections:

Navigation Pane
Provides access to the various available options configured within the CMP system.

You can bookmark options in the navigation pane by right-clicking the option and selecting **Add to Favorite**. Access the bookmarks by clicking the **My Favorites** folder at the top of the navigation pane. Within the **My Favorites** folder, you can arrange or delete options by right-clicking the option and selecting **Move Up**, **Move Down**, or **Delete from Favorite**.

You can collapse the navigation pane to make more room by clicking the button in the top right corner of the pane (◀). Click the button again to expand the pane.

Content Tree
Contains an expandable/collapsible listing of all the defined items for a given selection. For content trees that contain a group labeled **ALL**, you can create customized groups that display in the tree.

Note: The content tree section is not visible with all navigation selections.

You can collapse the content tree to make more room by clicking the button in the top right corner of the pane (◀). Click the button again to expand the tree. You can also resize the content tree relative to the work area.

Work Area
Contains information that relates to choices in both the navigation pane and the content tree. This is the area where you perform all work.

Alarm Indicators
Provides visual indicators that show the number of active alarms.

CMP Icons

The CMP interface provides the following icons to perform actions or indicate status:

Add
Use this icon to add an item to a list.

Calendar
Use this icon to select a date and, in some cases, a time.

Clone
Use this icon to duplicate a selection in a list.

⊗ **Critical error**
Displays in reports to indicate a critical error during the server replication process.

🗑 **or** ✕ **Delete**
When visible in the work area, selecting the Delete icon deletes an item, removing it from the device.

Note: Deleting an item from the **ALL** folder also deletes the item from any associated group. A delete verification window opens when this icon is selected.

Details
This binoculars icon displays when there is more details for an item.

Edit
Use this icon to modify a selection in a list.

External Connection
When visible in the work area, indicates which server currently has the external connection (the active server).

⚙ **Gear**
Displays when a policy references another policy or policy group.

Hide
When visible in the work area, selecting this icon removes the item from the current view but does not delete the item.

> **Note:** The item is only hidden during the current session. The item will be visible the next time a user logs into the CMP server.

L **Manual**
Displays when a field is configured by the user. Hover over this icon to see the name of the device.

! **Major error**
Displays in reports to indicate a major error during the server replication process.

⚠ **Minor error**
Displays in reports to indicate a minor error during the server replication process.

or ▲▼ **Up/Down**
These arrow icons are displayed when you can change the sequential order of items in a list.

<-- --> **Left/Right**
These arrow icons are displayed when it is possible to move an item from one list to another.

✓ **OK status**
Displays in reports to indicate a that the blade replication process completed without error.

✂ **Remove**
Removes an item from the group. The item is still listed in the **ALL** group and any other group that has an association with the item. For example, if you remove the device PS_1 from group PS_Group2, PS_1 still displays in the **ALL** group.

● **Selection**
This icon occurs in the Policy Wizard. The icon is used to select conditions and actions to add to a policy rule.

Synch broken
When visible in the Upgrade Manager, indicates that the CMP server does not have current information on a server.

Template
Displays when a field is configured by template. Hover over this icon to see the name of the template. Click the icon to view the template.

Virtual Machine
Displays when a Policy Management application is running on a virtual machine (VM).

View Cart
Displays the list of configurable objects selected for the **Export** action.

Overview of Major Tasks

The major tasks involved in using Policy Management system are configuring the system, defining manageable elements and profiles, creating and deploying policy rules, managing subscribers, and administering the system and CMP users.

Configuration Tasks for Setting Up the System

The system configuration tasks are a series of required steps necessary for setting up the Policy Management system. These tasks must be completed in the following order:

1. Configure the Policy Management topology. See #unique_77 for detailed information.
2. Configure MPE devices by creating Policy Server profiles and then configuring protocol options for each device. See Managing Multimedia Policy Engine Devices for detailed information.
3. Configure protocol routing which enables a Policy Management device to forward requests to other Policy Management devices for further processing. See Configuring Protocol Routing for detailed information.
4. Configure BoD devices by creating BoD profiles and then configuring protocol options for each device. Refer to *Bandwidth on Demand Application Manager Cable User 's Guide* for detailed information.

Definition Tasks for Setting Up the Network

The network element and profile definition tasks you need to perform depend on what external systems exist in your network. These tasks can be done in any order at any time. The set of tasks are as follows:

- Create network element profiles, including protocol options, for each network element with which the MPE devices interact and specify which MPE device will interact with which network elements. See Managing Network Elements for detailed information.
- Create record keeping server profiles. See Managing Record Keeping Servers for detailed information.

- Manage events that are sent to record keeping servers. See Managing Event Messaging for detailed information.
- Create management agent profiles. See Managing Management Agent Servers for detailed information.

Administration Tasks for Managing the System

The management and administrative tasks are optional and performed on an as needed basis. These tasks are:

1. View reports on the performance of the Policy Management system and devices in your network. See System-Wide Reports for more information.
2. Manage CMP users, accounts, access, authorization, and operation. See System Administration for detailed information.
3. Upgrade software using the Upgrade Manager. See Upgrade Functions for information.

2

Configuring the Policy Management Topology

This chapter describes how to add Policy Management sites and devices to create a Policy Management system. It also describes how to modify the topology, configure SNMP settings, and configure the global configuration settings.

About Policy Management Topology

Before you can use the Oracle Communications Policy Management system, you must configure the network topology for the Policy Management applications:

- CMP server
- Multimedia Policy Engine (MPE) devices
- Bandwidth on Demandserver (BoD)
- Management Agent servers (MA Servers)

The topology determines the following:

- How clusters are set up
- Which sites are primary and which are secondary
- How Policy Management devices communicate with each other
- How configuration data is replicated
- How incidents (events and alarms) get reported to the CMP system or external network management systems.

Figure 2-1 illustrates a Policy Management topology consisting of a primary (Site 1) and secondary (Site 2) CMP cluster, two georedundant BoD clusters, an MA server cluster, two routing MPE-R clusters, and a series of georedundant serving MPE-S clusters.

Note: These terms are defined in subsequent topics.

Figure 2-1 Example Policy Management Topology

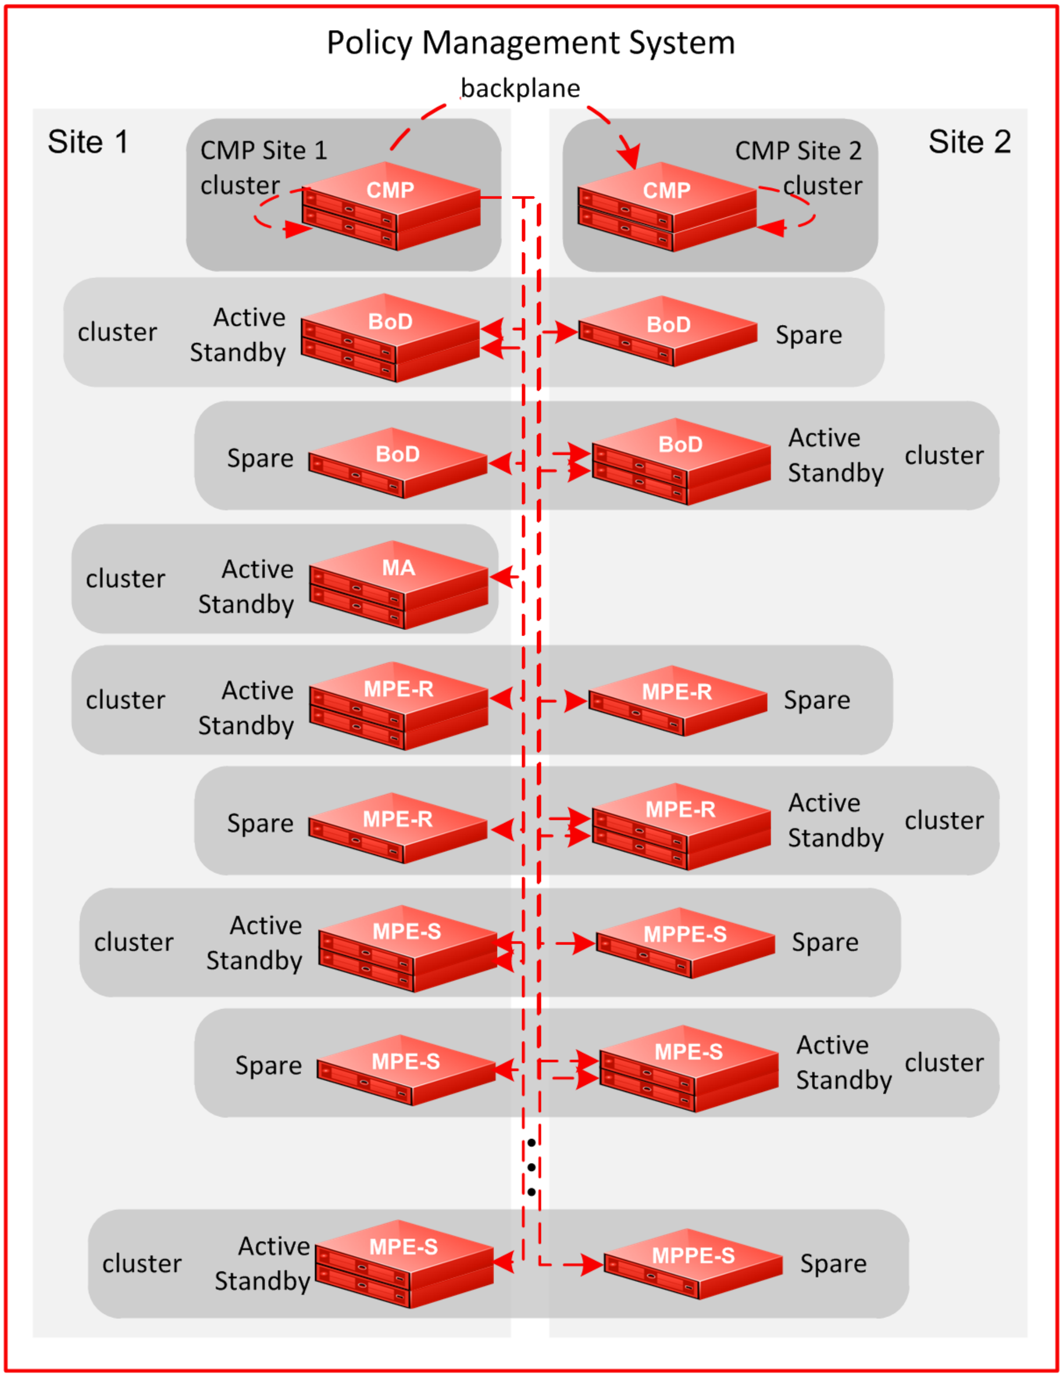

As the figure shows:

- The active CMP Site 1 server replicates its data to its standby CMP server and the active CMP server at CMP Site 2.
- In turn, the active CMP Site 2 server replicates its data to its standby CMP server.
- Additionally, the active Site 1 CMP server replicates data to all servers in any MPE, MA Server and BoD clusters in the topology, regardless of status (active, standby or spare).

- In turn, all servers and clusters merge status, events, alarms, and log data back to the active CMP server at Site 1.

High Availability

Policy Management provides High Availability (HA) to all Policy Management cluster configurations. Policy Management accomplishes HA by using two servers per cluster, an active server and a standby server. For georedundancy, a third, spare server provides additional backup support. Servers are continually monitored by the in-memory database. As shown in Figure 2-2, the active server processes network traffic and is accessible and connected to external devices, clients, gateways, and so forth. Only one server in a cluster can be the active server.

Figure 2-2 High Availability

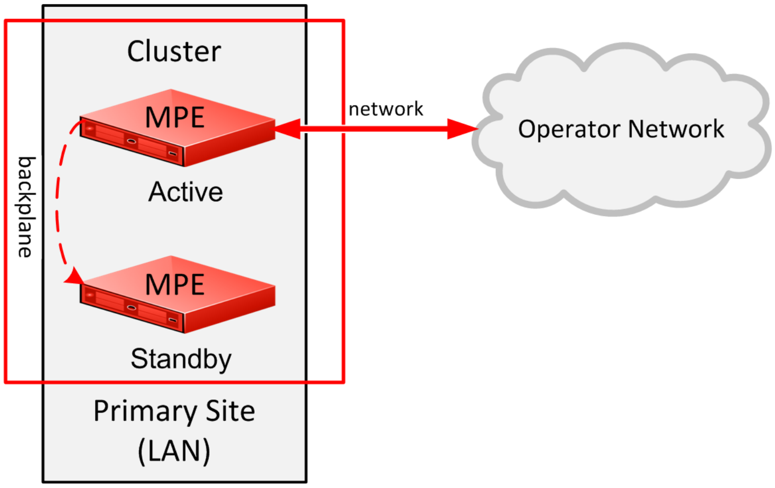

Within the cluster, the servers are connected together and work collaboratively, as follows:

1. The active and standby servers communicate using a TCP connection over the backplane network (direct-link High Availability) to replicate current state data, monitor server heartbeats, and merge alarms.

> **Note:** For georedundancy, a third, spare server is part of the cluster and receives replication data and heartbeats.

2. The servers share a virtual IP (VIP) cluster address to support automatic failover. The active server controls the VIP address.
3. The standby server does not receive any live traffic load, but holds an up-to-date copy of the active session state data at all times, replicated by High Availability. (This is sometimes called a warm standby.)
4. The HA database runtime processes on each server constantly monitor server status using heartbeat signals.
5. If the active server fails, indicated by missing a succession of three heartbeats:
    a. The standby server queries the active server. If the active server fails to respond, the standby server assumes the active state and takes over the VIP address and connections.

b. Because it continually receives session state and data updates through replication, the standby server can assume processing of ongoing sessions, so the failover is automatic and transparent to other components.

The terms active and standby denote roles, or states, that the servers assume, and these roles can change based on decisions made by the underlying HA database, automatically and at any time. If necessary, the standby server assumes control and becomes the active server. (For example, this would occur if the active server became unresponsive as determined by lack of a heartbeat signal.) When this happens, the server that was previously the active server assumes the role of the standby server.

When the failed server recovers, it becomes the standby server, and current state data for the cluster is replicated to the server. This behavior is non-revertive; that is, if an active server fails and then recovers, it becomes the standby server, rather than resuming its role as the active server.

CMP Georedundancy

As shown in Figure 2-3, georedundancy is implemented for CMP clusters by pairing a primary site CMP cluster with a secondary site cluster. The active server from the Site 1 CMP cluster will continuously replicate configuration, provisioning, and policy data, using HA, to the active server of the Site 2 cluster.

Figure 2-3 CMP Georedundancy

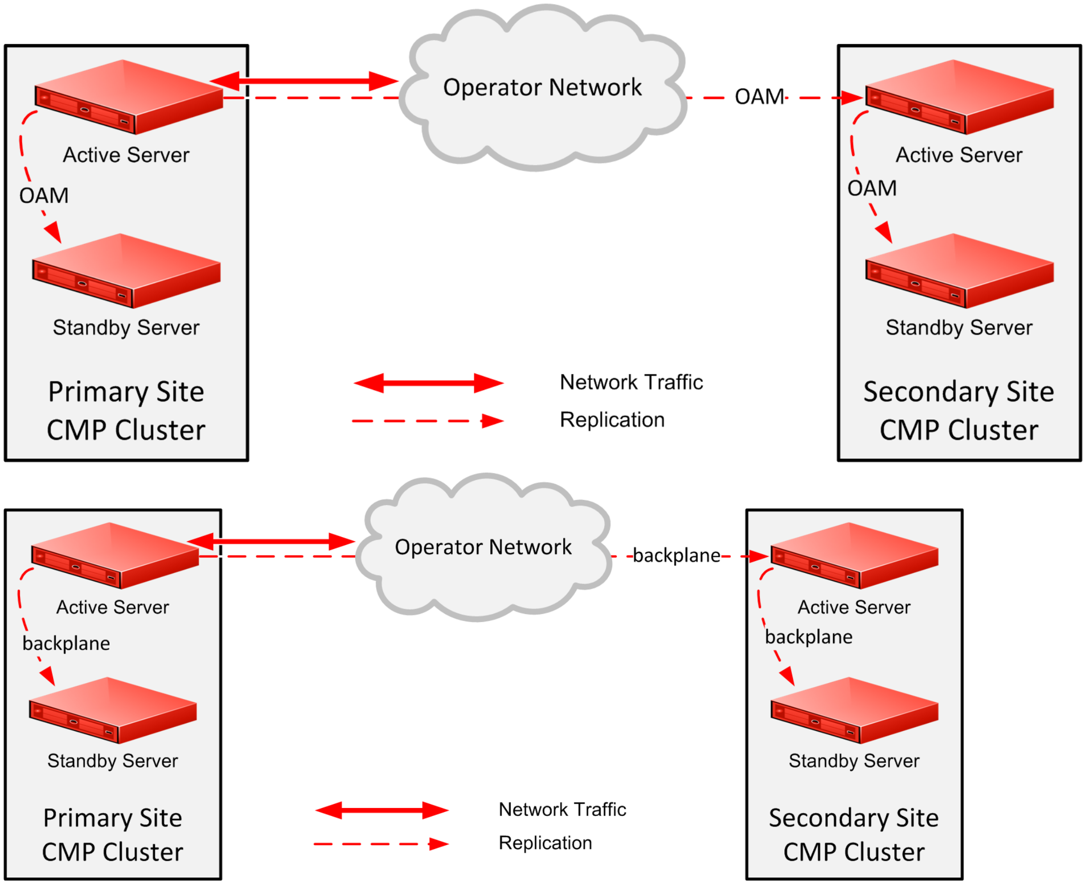

The secondary cluster does not have to be physically close to the primary cluster. The terms primary and secondary denote roles, or states, that the servers or clusters assume, and you can change these roles manually. If the Site 1 CMP cluster goes offline (as in a disaster scenario), you would log in to the active server of the Site 2 CMP cluster and manually promote this cluster to become the primary (Site 1) CMP cluster to manage the Policy Management network.

Promotion of a CMP cluster is always a manual operation (see Promoting a Georedundant CMP Cluster for details). The preferred sequence of operation is to first demote the active CMP server at the primary site and then promote the active CMP server at the secondary site, but this is not required. For example, in a disaster-recovery scenario in which the primary site is inaccessible, you can promote the active CMP server at the secondary site immediately. (This may trigger alarms.) The servers record the timestamp when a role is assigned. Policy Management systems recognize the CMP server with the most recent promotion timestamp as the primary cluster (that is, the recognized authority).

In a georedundant topology, c-Class servers (HP ProLiant BL460c Gen8 servers with a 1x4 mezzanine card) can communicate over a dedicated backup (BKUP) network. This network is set up using the Platform Configuration utility. See the *Platform Configuration User's Guide* for detailed information.

Georedundancy for Non-CMP Servers

Georedundancy is an optional configuration provided for non-CMP clusters in which the spare server can be located in a separate geographical location, as shown in Figure 2-4. The active server replicates state data to the standby and spare servers. If the two servers at one site become unavailable, the third server, located at the other site, automatically becomes the active server and continues to provide service. You can designate sites as primary and secondary.

Georedundancy supports both session-stateful (at the MPE device level) and binding-stateful failover between a pair of geographically separate (or geo-diverse) Policy Management sites. This includes the ability to maintain ongoing sessions and existing bindings that were in progress on the failed site at the time of failure, as well as being able to initiate and handle all new sessions and bindings on the secondary site for the duration of the failure.

Figure 2-4   Non-CMP Georedundant Configuration for Cable

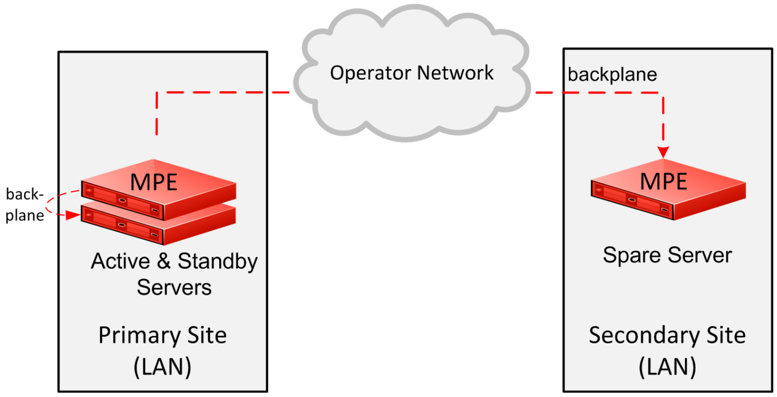

Within a georedundant cluster, the servers are connected through both the backplane and the OAM network. The servers work collaboratively as follows:

- The active and standby servers communicate using the backplane network to perform replication, monitor heartbeats, and merge trace-log and alarm data.
- The active and standby servers share a virtual IP (VIP) cluster address to support failover (georedundancy).

- The spare server has its own virtual IP address.
- The HA database runtime process constantly monitors the status of all servers in the cluster.
- If the active server fails, it instructs the standby server to take over and become the active server.
- If both the active and standby servers fail, it instructs the spare server to take over and become the active server.

Using this configuration, if one site fails, clients retain connectivity to the other site, and established sessions remain active. As servers at the failed site recover, they become standby servers, and current state data for the clusters are replicated to them. After the recovered servers are synchronized with the state data of the active servers, they are automatically returned to active roles. This behavior is called revertive which means that if an active server fails and then recovers, it becomes the active server again.

Within a georedundant cluster, the active and standby servers are connected through a local area network (LAN), that uses a single TCP/IP socket connection or stream. The active and spare servers, located at separate sites, are connected through a wide area network (WAN) Figure 2-5. Since every WAN has distinct bandwidth and packet loss characteristics, the connection can optionally be configured to use up to eight streams to maintain throughput in cases of network congestion or packet loss.

Figure 2-5 Cable Georedundant System with Spare CMP Cluster

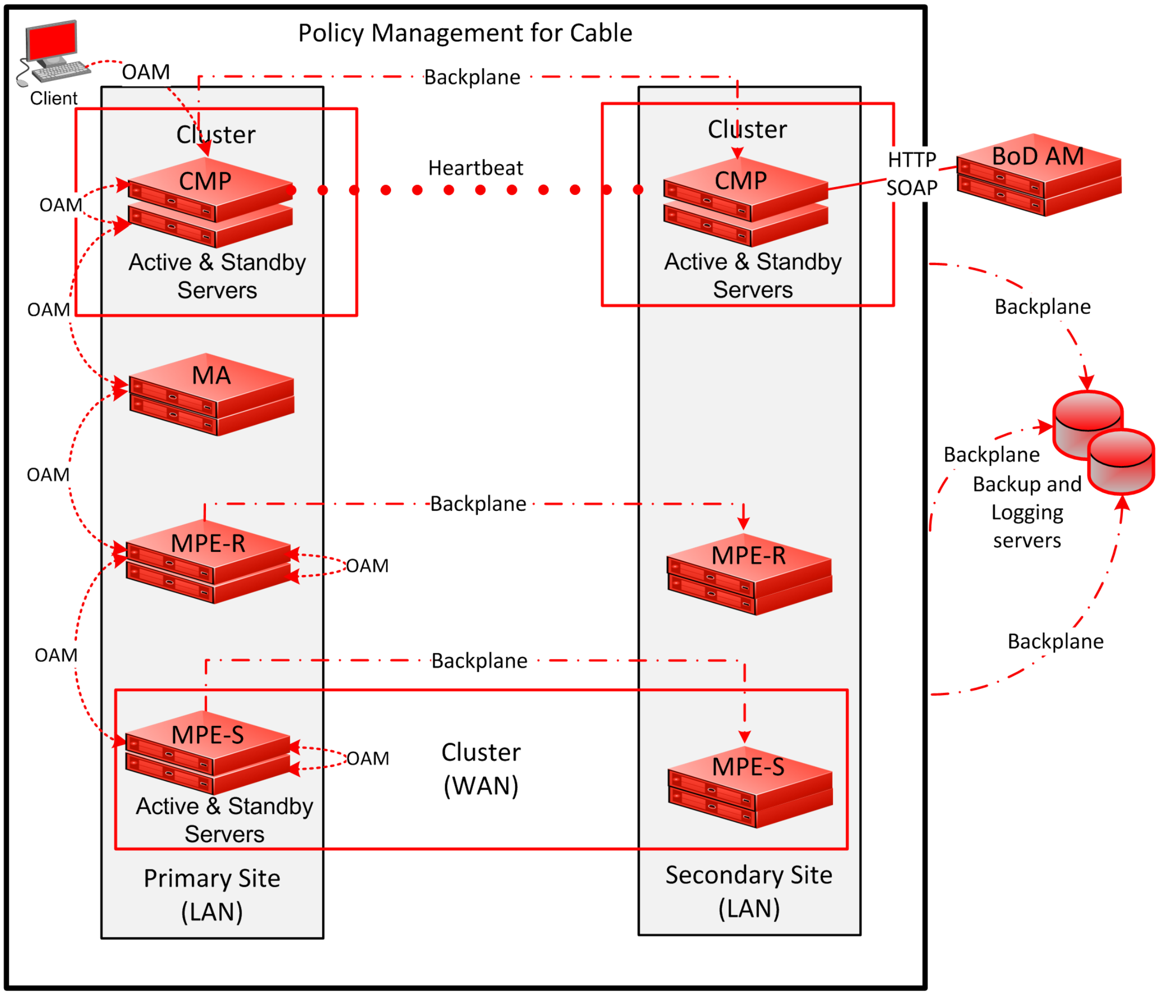

Diameter Signaling

Diameter signaling traffic is carried on a virtual LAN (VLAN) Signaling A (SIG-A) network, a SIG-B, or SIG-C network. Database replication and high-availability (HA) heartbeat traffic within a site (that is, between the active and standby servers) is sent on an Operation, Administration, and Management (OAM) VLAN network. You can configure the Policy Management topology to send replication and HA heartbeat data between sites (that is, between the active and spare servers) using different VLANs. Replication data can be sent between sites on the OAM (default), SIG-A, SIG-B, SIG-C, or a dedicated replication (REP) network. (Replication traffic between CMP servers always uses the OAM network.) For information on configuring a REP network, see Setting Up a Non-CMP Cluster. In a georedundant topology, servers (with a 1x4 mezzanine card) can communicate with logging and backup servers over a dedicated backup (BKUP) network. However, for Policy Management applications, only backup of CMP systems is typical.

Replication (REP) packets can be marked with a symbolic differentiated services code point (DSCP) value to determine per-hop behavior (PHB). The supported code points are class selector (CS), assured forwarding (AF), and expedited forwarding (EF). The available class selectors are CS1 through CS7. The following AF points are available:

Drop Probability	Class 1	Class 2	Class 3	Class 4
Low	AF11	AF21	AF31	AF41
Medium	AF12	AF22	AF32	AF42
High	AF13	AF23	AF33	AF43

A cluster can be configured to use a secondary HA heartbeat path between georedundant sites in case the primary HA heartbeat network fails. The secondary HA heartbeat path can be configured to use the OAM, SIG-A, SIG-B, SIG-C, or REP network. If the primary HA heartbeat network fails, then the secondary HA heartbeat path continues to send heartbeats between the active and spare servers.

The primary HA heartbeat path is the same as the replication path. The default primary HA heartbeat and replication path is the OAM network. If you configure a different network to carry replication traffic, then that network is also used as the primary HA heartbeat network. In this case, the OAM network could be configured as the secondary HA heartbeat network.

Replication traffic, including a threshold of outstanding updates to a standby or spare server (see Configuring the Upsync Log Alarm Threshold), is displayed in an MPE/BoD Replication Stats report (see Viewing the MPE/BoD Replication Statistics Report).

Georedundant Spare Servers

As shown in Figure 2-6, an MPE-S or BoD cluster can contain an additional georedundant server, called a spare server. The active server will replicate its database to the standby server as well as the spare server. In this configuration, the standby server is first in line to take over from the active server and the spare is second in line.

Figure 2-6 Clusters with Active, Standby, and Spare Servers

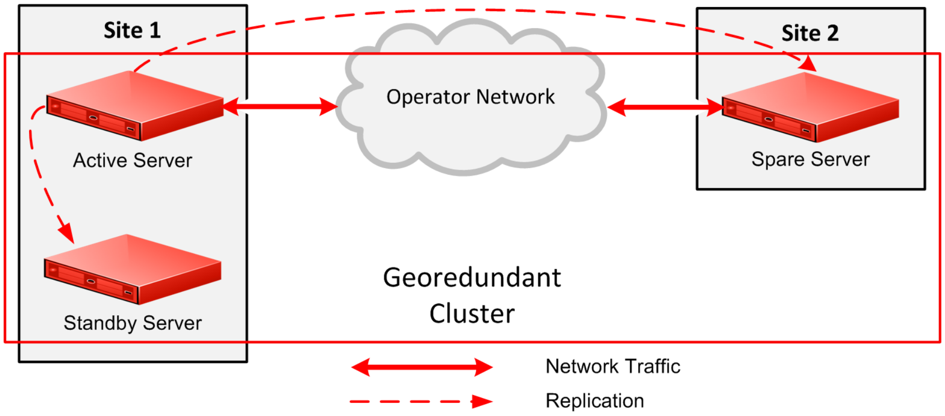

Active, standby, and spare servers interoperate as follows:

1. The servers communicate using WAN TCP streams to perform replication, monitor heartbeats, and merge events.

2. The active and standby servers share a common virtual IP (VIP) cluster address to support automatic failover.

3. The spare server has a unique VIP cluster address.
4. The HA database runtime process constantly monitors the status of all servers.
5. If the standby server does not receive three consecutive heartbeats, it attempts to communicate with the active server. If the standby server receives no response from the active server, it assumes the active role.
6. When HA misses the three heartbeats to the spare server, it instructs the spare server to assume the standby role.

The terms active, standby, and spare denote roles, or states, that the servers assume, and these roles can change automatically and at any time based on decisions made by the underlying HA database. If both the active and standby servers become unavailable, the spare server automatically assumes the active role and continues to provide service.

Primary and Secondary Sites

In the Policy Management topology architecture, primary refers to the preferred option for sites, servers, and connections. Under normal conditions, for any cluster, a server at the primary site is the active server that services traffic or manages the Policy Management network. All clients and gateways are connected to this primary site.

Secondary refers to the georedundant backup site, servers, and connections. MPE-S and BoD clusters can be dispersed between a primary site and a secondary site. This dispersal mates the primary and secondary sites together. (CMP clusters can be paired, but not georedundant. MPE-R and MA clusters are neither paired nor georedundant.) For signaling traffic, the primary and secondary sites use different VIP addresses.

If, for some reason, the active server at a primary site can no longer provide service, the cluster fails over to the standby server at the primary site. The server assuming the service becomes the active server.

If and only if no servers are available at an MPE-S or BoD primary site, the cluster fails over to the secondary site, and a spare server takes over as the active server in the cluster and provides service. When one of the servers at the primary site becomes able to provide service, then the active status transitions back to the server at the primary site. In contrast, CMP server failover is manual.

Note: MPE-R and MA Server clusters do not support failover.

You configure primary and secondary sites as initial states. After MPE-S and BoD clusters are in operation, failover from a primary site to a secondary site is automatic.

It is not meaningful to describe a site as primary except in the context of where the active server of a cluster is located. For example (see Figure 2-7), you could establish a topology with two sites and two MPE-S clusters, with the spare server of each cluster located at the other site. In this topology, the primary site of MPE cluster (Server A and B) is also the secondary site of MPE cluster (Server C), and vice versa.

Figure 2-7 Primary and Secondary Sites

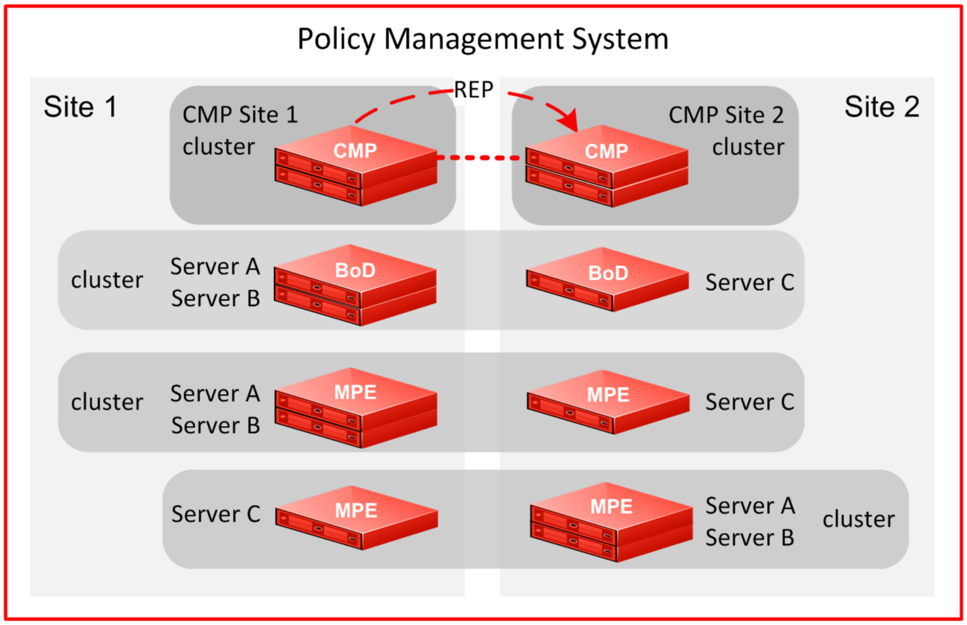

Georedundant Site Preferences

When you configure a georedundant MPE-S or BoD cluster, you initially set the High Availability site preference to **Normal** to designate that the primary site is preferred. This determines which site contains the active server and initially processes traffic.

After the servers are defined, you can reverse this preference, which designates that the secondary site is preferred. Reversing site preference makes the spare server take over as the active server. The former active and standby servers become the standby and spare servers (which server assumes which role is not determined).

Reversing site preference is useful in situations where you need to troubleshoot, service, upgrade, or replace the active server.

The **Cluster Settings** table on the Cluster Configuration page lists information about MPE-S or BoD cluster preferences under the heading **Site Preference**. A cluster preference is one of the following:

- **Normal**
- **Reverse**
- **N/A** (Not Applicable; CMP clusters cannot be reversed)

Server Status

You can view the status of a server in the **Cluster Information Report** (see Cluster Information Report).

Note: The display refreshes every 10 seconds. Click **Pause** to freeze the page.

The status of a server can be thought of as its current role. The status describes what function the server is currently performing, or performing in a cluster. Statuses can change from server to server within a cluster, but no two servers in the same cluster should ever have the same status.

> **Note:** Two servers in the same cluster with the same status is an error condition.

The server status values are as follows:

Active
An active server is externally connected. In a cluster, the active server is the only server that is handling connections and servicing messages and requests. Only an active server writes to the database. An active server at the primary site remains active unless it cannot provide service. An active server at the secondary site will remain active as long as no server at the primary site is available to provide service.

Standby
A standby server is the server in a cluster that is prepared to immediately take over in the event that the current active server is no longer able to provide service. If the standby server takes over, it becomes the active server. When the previously active server has recovered, it reverts to its former status of standby server.

Out of Service
If a server has failed and is unavailable to assume any of the other roles, then its status is out of service. A server is reported as out of service in two scenarios:

- The CMP system can reach the server, but the software service on the server is down.
- The CMP system cannot reach the server.

No Data
The CMP system cannot reach the server. This status value provides backward compatibility with earlier Policy Management releases. It is seen during the upgrade process.

About Virtualization

The Policy Management system can operate as a standard virtualized hardware system with virtual machines loaded onto qualified hardware elements or as an ETSI-defined Network Functions Virtualization (NFV) Management and Organization (MANO) system where orchestrator software (such as, OpenStack) handles the network-wide orchestration and management of the NFV (infrastructure and software) resources and the NFV service topology on the NFV infrastructure. See NF Agent for VNF Management for more information.

The Policy Management system functions on a range of hardware platforms. Alternatively, you can deploy Policy Management applications in a virtualized hardware environment, as virtual machines (VMs). The VMs run on industry-standard high-volume servers (also called Hosts) with switches and storage. VMs are abstracted from Host systems (also called compute nodes) and function as if running alone, although actually multiple VMs (or Guests) can run on a single Host system.

The qualified hardware configurations include:

- HP Enterprise DL360/DL380 Gen8

- Oracle Server X5-2

The qualified VM manager/hypervisor software includes:

Hypervisor	VM Manager
Oracle Virtual Machine Server (OVM-S)	Oracle Virtual Machine Manager (OVM-M)
Kernel-based Virtual Machine (KVM)	With or without OpenStack
VMware ESXi	VMware vSphere
Kernel-based Virtual Machine (KVM)	No VM Manager (for server configuration)

VMs are deployed within a network functions virtualization (NFV) infrastructure that is hardware independent. The NFV infrastructure includes an environment manager known as a VM manager. The VM manager monitors and manages VMs running on a single host system. The VM manager performs the following management functions:

- Dynamically allocates resources, such as CPU, RAM, and storage among VMs to maximize hardware utilization and balance load
- Instantiates new VMs on demand for increased capacity or in the case of VM failure
- Moves VMs from one host system to another before upgrades or in the case of hardware failure

Figure 2-8 illustrates the virtualization architecture.

Figure 2-8 Virtualization Architecture

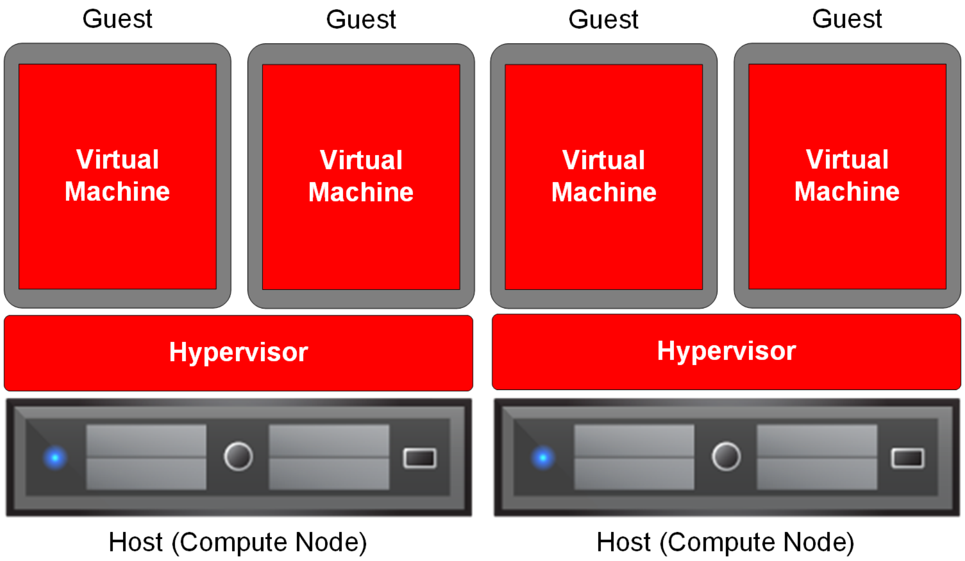

The virtualized environment supports the multiple network interface connections (OAM, SIG A, SIG B, SIG C, and REP, and BKUP) used by Policy Management applications by mapping virtual Ethernet devices as if they were separate physical devices. The virtualized environment supports HA by defining affinity, so that the hypervisor maintains a standby VM on a different host system from the active VM.

The virtualized environment calculates key performance indicators (KPI) such as performance, capacity, and load factor by dynamically obtaining the current resources, and the **KPI Dashboard** indicates if a server is running as a virtual machine (see KPI Dashboard).

Note: A Policy Management topology can combine both virtual and physical (or bare-metal) machines.

An alternate supported configuration is to include all the Policy Management VMs on a pair of Oracle X5-2 RMS host systems or HP Enterprise DL360/DL380 Gen8 blade host-based systems, running Oracle Enterprise Linux (OEL) with either the KVM, OVM-M, or VMware ESXi hypervisor, as a fully functional, high-availability, entry-level, minimal-footprint solution consisting of the following:

- 1 clustered, high-availability (2-server) CMP system
- 1 clustered, high-availability (2-server) MRABoD system
- 2 clustered, high-availability (2-server) MPE systems

Figure 2-9 shows this minimum Policy Management virtualized topology.

Note: The standby VNFC devices reside on separate host systems to ensure proper system function in the event a host system fails.

Figure 2-9 Policy Management Minimum Virtualized Topology

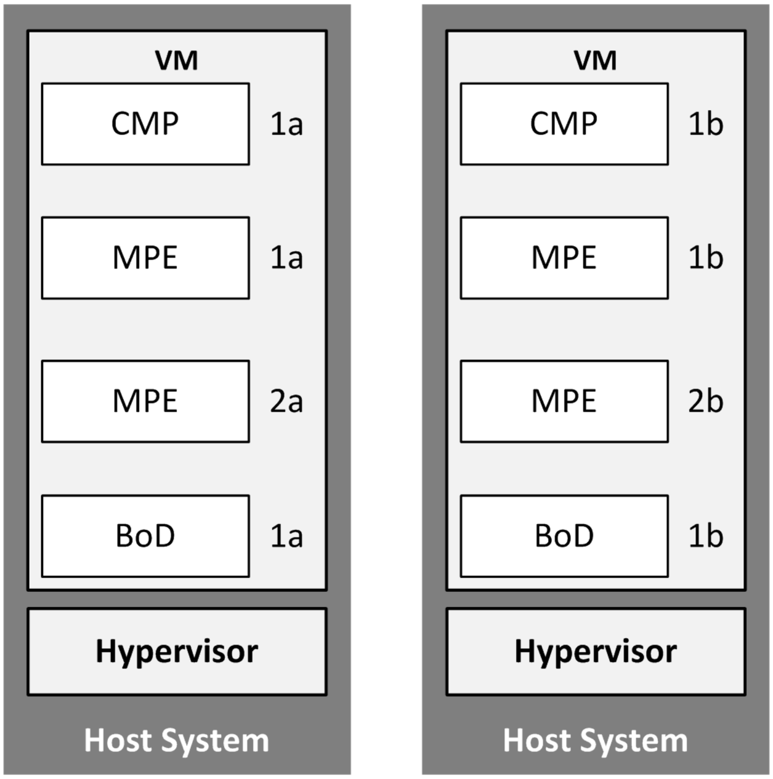

Refer to *Virtual Installation Guide* for detailed information on creating Virtual Network Function Components (VNFCs) (for example, CMP, MPE, and BoD devices) on a host system.

Once a virtualized host system is implemented, adding VNFCs to the Policy Management system is as simple as adding a server. See About Setting Up the Topology for details.

NF Agent for VNF Management

The Network Function Agent (NF Agent) provides VNF (Virtual Network Function) management services as well as acting as a point of integration with Orchestrator software (like OpenStack) and VIMs (Virtual Infrastructure Managers) APIs that manage the VNF/VM lifecycles. The NF Agent provides the logical interface and mappings between virtual deployments and internal Policy objects and logic. The NF Agent has the responsibility of handling specific network functions that run on one or more virtual machines (MPE servers).

VNF management provides a set of services and functionality that allow for a virtual instance of an application (that is, VNF) to be instantiated, managed, and destroyed.

The NF Agent is a web service hosted on the same server that hosts the CMP server. As an independent service, the NF Agent encapsulates virtual operations and VIM and Orchestrator interfaces. The NF Agent keeps mappings between logical MPE devices as well as VNF, VM, VNFD (Virtual Network Function Descriptor) instances. The NF Agent provides support for the following VIM connection types:

- OpenStack API
- OpenStack HEAT API
- VMWare vCloud

The NF Agent functions as a service with a northbound RESTful API and multiple southbound client interfaces for various VIMs. The architecture provides sufficient flexibility for the easy implementation of additional VIM clients. The NF Agent expects the following VM profile and deployment information to the VIM so it can instantiate instances of the described VNF:

- Required vCPUs
- Required vNICs
- Required Networks and IP addressing
- Memory size
- Storage size
- Anti-affinity/Affinity requirements

Orchestration cases are manually implemented through the **Topology Settings** command. The user specifies operations on a new VNF.

Before Setting Up the Topology

Prior to setting up the topology, the server hardware must have been set up and configured as described in *TPD Initial Product Manufacture Software Installation Procedure*. This includes installing TPD, TVOE, and PMAC software, setting up the enclosure and network connections, and installing Policy Management application software (for CMP, MPE and BoD devices).

About Planning the Topology

Before beginning to set up your topology, you will need to gather the information needed to input into the servers' set up forms.

The Platform Configuration (`platcfg`) utility is used to set up and configure the hardware as well as install the proper Policy Management application software (that is, MPE, CMP, BoD).

> **Tip:** This information can be collected at any time before beginning the topology set-up procedure without interrupting service.

For more information, refer to *Platform Configuration User's Guide*.

Topology planning information includes the following:

- Names of existing clusters
- Names for any sites
- The maximum primary site failure threshold, to record site failures (0 is recommended)
- The OAM VIP address of the existing primary site CMP system and, if applicable, the georedundant CMP system
- (Optional) a designated network path, either OAM, SIG-A, SIG-B, or SIG-C for backup (secondary) HA heartbeats between sites
- (Optional) a designated network path, either OAM, SIG-A, SIG-B, or SIG-C for WAN replication traffic between sites
- If DSCP marking for WAN replication traffic is used, the type of DSCP marking
- If multi-stream WAN replication traffic is used, the replication stream count

Information entered using the Platform Configuration utility includes the following:

- Initial provisioning information for servers:
  - A host name
  - For CMP server access, an OAM Real IP address and subnet mask (IPv4 or IPv6)
  - An OAM IPv4/IPv6 default route (default gateway)
  - A list of network time protocol (NTP) server IP addresses
  - A list of domain name system (DNS) server IP addresses
  - Bond interface for the OAM device
  - Backplane bond interface of the OAM device
  - For IPv4-based network elements, an IPv4 VIP address and subnet mask on the SIG-A network

  - For inter-topology communication or any IPv6-based network elements, an IPv6 VIP address and subnet mask on the SIG-A network

- For each existing HA cluster:

  - Verify that firewall rules are correctly provisioned

About Setting Up the Topology

Topology configuration consists of defining Policy Management sites and clusters, including their addresses and hierarchy. You can add MPE, MA Server, and BoD clusters to the topology before configuring the individual servers themselves. You can define all the servers in a cluster in the same operation.

The recommended sequence of creating the Policy Management topology follows:

1. Configure the primary CMP cluster.

   a. You start to build a topology by logging in to the active CMP server at the primary site.

   b. Configure the CMP cluster settings. See Setting Up a CMP Cluster for details.

   c. The settings are replicated (or pushed) to the standby CMP server. Together, the two servers form a primary, or Site 1, CMP cluster. This is the primary CMP site for the whole topology network.

   Note: The primary site cannot be deleted from the topology.

2. (Optional) Use the primary CMP cluster to configure a secondary, or Site 2, CMP cluster.

   Note: A secondary CMP cluster can provide georedundancy.

3. Configure non-CMP (MPE, MA Server, and BoD) clusters. See Setting Up a Non-CMP Cluster for details.

   Enter MPE, MA Server, and BoD cluster settings on the active CMP server on the primary site.

   Note: You can define the topology before defining the servers themselves.

4. After defining the topology, the configuration information is replicated as follows:

   a. The CMP system replicates the topology configuration, including the cluster settings, to active, standby, and (if present) spare servers using the OAM network. These servers form an MPE, MA Server, or BoD cluster based on the topology configuration.

   b. Active servers communicate with standby servers using LAN connections over the OAM network. Active servers communicate with spare servers using WAN connections over the OAM, SIG-A, or SIG-B.

   c. Active and standby servers share a virtual IP (VIP) cluster address to support automatic failover. (If present, the spare server has a unique VIP address.)

d. The HA database runtime process constantly monitors the status of the servers in each cluster. If an active server in a cluster fails, it instructs the standby server to take over and become the active server. In a georedundant topology, if both the active and standby servers in a cluster fail, it instructs the spare server to take over and become the active server.

e. (Optional) For georedundancy, configure additional sites for MPE and BoD clusters.

After you define the topology, use the **System** tab of each server to determine if there are any topology mismatches. See About Reapplying a Configuration for more information.

About Setting Up a CMP Cluster

You must set up at least one CMP cluster before proceeding with setting up the topology. The first CMP cluster you set up is called the CMP Site1 (or Primary) cluster. You can optionally set up a CMP Site2 (or Secondary) cluster.

Before defining the CMP Site1 cluster, ensure the following:

- The CMP application is installed on all servers in the cluster.
- The servers have been configured with network time protocol (NTP), domain name server (DNS), IP Routing, and OAM IP addresses.
- The CMP server IP connection is active.
- The CMP application is running on at least one server.

Setting Up a CMP Cluster

To set up a CMP cluster:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

   If a CMP Site1 cluster is not yet defined, a message appears asking you to add CMP Site 1 cluster.

2. Click **Add CMP Site1 Cluster**.

   The Topology Configuration page opens.

   Note: The **Name** and **Appl Type** fields are fixed.

3. Select the **HW Type** from the list.

   Available options are:

   - **Oracle RMS**—Oracle Server X5-2
   - **HP ProLiant G8 RMS**—HP Enterprise ProLiant DL360 or DL380 Gen 8 (rack-mounted server)
   - **VM**—virtual machine

4. Select the **Degrade on failure of** setting.

This is the signaling network that, if it fails, the server status changes to **Degraded**. Available options are:

- **OAM**
- **SIG-A**
- **SIG-B**
- **Both SIG-A and SIG-B**

5. Click **Add New VIP**.

   The New OAM VIP dialog box appears:

   a. Enter the OAM VIP and the mask.

      This is the IP address the CMP server uses to communicate with a Policy Management cluster.

      Note: Enter the IPv4 address in standard dot format and its subnet mask in CIDR notation from 0 to 32.

   b. Click **Save**.

      The OAM VIP and mask are saved.

   c. Repeat this step for a second OAM VIP, if needed.

6. (Optional) To enter up to four signaling VIPs, click **Add New VIP**.

   The New Signaling VIP dialog box appears:

   a. Enter the signaling VIP and the mask.

      This is the IP address the CMP server uses to communicate with an external signaling network.

      Note: Enter the IPv4 address in standard dot format and its subnet mask in CIDR notation from 0 to 32.

   b. From the **Interface** list, select one of the following:

      - **SIG-A**
      - **SIG-B**

   c. Click **Save**.

      The signaling VIP and mask are saved.

   d. Repeat this step for up to three additional signaling VIPs.

7. Configure the active server by doing the following:

   a. Click **Add New IP**.

      The New IP dialog box appears.

b. Enter the IP address for the server.

   Up to two IP addresses can be entered (one IPv4 and one IPv6). Use the IPv4 standard dot-formatted IP address string and the IPv6 standard 8-part colon-separated hexadecimal string format.

c. Select the preferred IP address format.

   The server will preferentially use the IP address of the selected format.

Note: The following restrictions apply:

- If neither an IPv6 OAM IP nor a static IP address is defined, IPv6 is not available.
- If neither an IPv4 OAM IP nor a static IP address is defined, IPv4 is not available.

d. Click **Save**.

   The IP address for the active server is saved as Server A.

8. (Optional) To enter a second **IP** address, repeat the previous step.

Note: Up to two IP addresses can be entered (one IPv4 and one IPv6).

9. Enter the host name for the server.

   The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_). This name must exactly match the host name for this server (that is, the output of the Linux command `uname -n`).

   - If the server has a configured server IP address, click **Load**, which retrieves the remote server host name. If retrieval fails, you must enter the host name.

10. Select **Forced Standby**, which forces the server into standby mode.

Note: The state is set automatically when a new server is added to a cluster or if a server setting is modified and another server already exists in the cluster.

11. Click **Save**.

    A confirmation message appears.

12. Click **OK**.

    A restart message appears.

13. Click **OK**.

    The active server restarts.

14. Log back in to the CMP server.

15. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

The Cluster Configuration page opens; the initial group is **All Clusters**.

16. From the content tree, select the **CMP Site 1 Cluster**.

    The Topology Configuration page opens.

17. Select **Modify Server-B**, and enter the appropriate information for the secondary server of the cluster.

18. Click **Save**.

The CMP cluster topology is defined.

After you define the topology, use the **System** tab of each server to determine if there are any topology mismatches. See About Reapplying a Configuration for more information.

After you define the primary CMP cluster, you can repeat this procedure to define a georedundant secondary CMP cluster. See Setting Up a Georedundant Site before adding a secondary CMP cluster.

About Setting Up a Non-CMP Cluster

A non-CMP cluster includes one of the following server types:

- BoD
- Management Agent

Note: The list of available server types depends on the CMP modes configured. See CMP Modes for more information.

Note: If you are creating a cluster in a georedundant system, see Setting Up a Georedundant Non-CMP Cluster.

Setting Up a Non-CMP Cluster

Before defining a non-CMP cluster, ensure the following:

- The server software is installed on all servers in the cluster.
- The servers have been configured with network time protocol (NTP), domain name server (DNS), IP Routing, and OAM IP addresses.

Caution: Before you add a non-CMP (MPE, BoD, or MA Server) cluster, you must create a site. See Setting Up a Georedundant Site for details. Failure to set up a site may result in alarms being generated on the CMP server.

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The **Cluster Configuration** page opens; the initial group is **All Clusters**.

2. From the work area, select **Add MPE/BoD/MA Cluster/Mediation Cluster**.

Note: The list of available cluster types to add to the topology depends on the CMP modes configured. See the *CMP Wireless User's Guide* for more information.

The Topology Configuration page opens.

3. In the **Cluster Settings** section of the page:

   a. (Required) Enter the **Name** for the cluster.

   The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_). The maximum length is 250 characters.

   b. Select the **Appl Type** from the list.

   Available options are:

   - MPE (default)
   - BoD
   - MA (Management Agent)
   - Mediation

   Note: The list of available application types depends on the CMP modes configured. See the *CMP Wireless User's Guide* for more information.

   ---

   c. Select the **Degrade on failure of** settings.

   This is the signaling network that, if it fails, the server status changes to Degraded. Available options are:

   - OAM
   - SIG-A
   - SIG-B
   - Both SIG-A and SIG-B

   d. Select the **HW Type** from the list.

   Available options are:

   - **Oracle RMS**—Oracle Server X5-2
   - **HP ProLiant Gen8 RMS**
   - **VM** (virtual machine)
   - **VM(Automated)** (VM managed by NF Agent)

     See Setting Up a VM (Automated) Non-CMP Cluster for details on adding a VM (Automated) cluster.

   e. If needed, repeat the process for the second OAM VIP.

f. (Optional) To enter up to six **Signaling VIPs** addresses (up to two each for each of SIG-A, SIG-B, and SIG-C), click **Add New VIP**.

   The signaling VIP is the IP address a PCEF device uses to communicate with the cluster. A non-CMP cluster supports redundant communication channels, named SIG-A, SIG-B, or SIG-C for carriers who use redundant signaling channels.

   The New Signaling VIP dialog appears.

   i. Enter the **Signaling VIP** address and the **Mask**.

      This is the IP address the CMP server uses to communicate with an external signaling network.

> **Note:** Enter the IPv4 address in standard dot format and its subnet mask in CIDR notation from 0 to 32.

   ii. Select the **Interface** from the list.

      Available options are:

      - SIG-A
      - SIG-B
      - SIG-C

   iii. Click **Save**.

      The **Signaling VIP** address and **Mask** are saved.

g. Repeat the process for any remaining Signaling VIPs.

4. To configure Server-A hardware, in the **Server-A** section of the page:

a. (Required) To enter the **IP** address, click **Add New IP**.

   The Add New IP dialog box appears.

   i. Enter the **IP** address in either IPv4 or IPv6 format.

      The IP address of the server. For an IPv4 address, enter it in the standard IP dot-format. For an IPv6 address, enter it in the standard 8-part colon-separated hexadecimal string format.

   ii. Select the **IP Preference**.

      Either **IPv4** or **IPV6**. If **IPv6** is selected, the server will preferentially use the IPv6 address for communication.

> **Note:** If neither an IPv6 OAM IP nor a static IP address is defined, **IPv6** cannot be selected. If neither an IPv4 OAM IP nor a static IP address is defined, **IPv4** cannot be selected.

b. Enter the **HostName** of the server.

The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_). This must exactly match the host name provisioned for this server (the output of the Linux command `uname -n`).

> **Note:** If the server has a configured server IP, you can click **Load** to retrieve the remote server host name. If the retrieve fails, you must enter the host name.

c. Select **Forced Standby** to put Server-A into forced standby status.

By default, Server-A will be the initial active server of the cluster.

5. (Optional) Click **Add Server-B** and enter the information for the standby server of the cluster.

   Server-B is defined for the cluster.

6. Click **Save**.

   A confirmation message appears.

7. Click **OK**.

Setting Up a VM (Automated) Non-CMP Cluster

Before defining a VM (Automated) non-CMP cluster, ensure the system is configured for virtualization and VIM Connections are defined.

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

2. If you do not see the Cluster Configuration page, click **All Clusters**.

3. From the work area, select **Add MPE/BoD/MA Cluster**.

> **Note:** The list of available cluster types to add to the topology depends on the CMP modes configured. See the *CMP Wireless User's Guide* for more information.

   The Topology Configuration page opens.

4. In the **Cluster Settings** section of the page:

   a. (Required) Enter the **Name** for the cluster.

   Enter up to 250 characters, excluding quotation marks (") and commas (,).

   b. Select the **Appl Type** from the list.

   Available options are:

   - **MPE** (default)
   - **BoD** (Cable Mode)

- **MA** (Management Agent)
- **Mediation**

> **Note:** The list of available application types depends on the CMP modes configured. See the *CMP Wireless User's Guide* for more information.

c. Select the **Degrade on failure of** settings.

   This is the signaling network that, if it fails, the server status changes to Degraded. Available options are:

   - OAM
   - SIG-A
   - SIG-B
   - Both SIG-A and SIG-B

d. Select **VM(Automated)** from the **HW Type** list.

e. If needed, repeat the process for the second OAM VIP.

f. (Optional) Click **Add New VIP**.You can enter up to six **Signaling VIPs** addresses (up to two for each SIG-A, SIG-B, and SIG-C).

   The signaling VIP is the IP address a PCEF device uses to communicate with the cluster. A non-CMP cluster supports redundant communication channels, named SIG-A, SIG-B, or SIG-C for carriers who use redundant signaling channels.

   The New Signaling VIP dialog box appears.

   i. Enter the **Signaling VIP** address and the **Mask**.

      This is the IP address the CMP server uses to communicate with an external signaling network.

> **Note:** Enter the IPv4 address in standard dot format and its subnet mask in CIDR notation from 0 to 32.

   ii. Select the **Interface** from the list.

      Available options are:

      - **SIG-A**
      - **SIG-B**
      - **SIG-C**

   iii. Click **Save**.

      The **Signaling VIP** address and **Mask** are saved.

g. Repeat the process for any remaining Signaling VIPs.

5. Configure Server-A using VM (Automated). In the **Server-A** section of the page:

a. Select the **VIM Connection** from the list.

   If the list is empty or your connection is not listed, it may need to be created. See Creating a VIM Connection for details about creating connections.

b. Verify that the VIM Connection Type is correct.

   You cannot change this field. If the connection is not correct, select another VIM connection, or create a new one. See Creating a VIM Connection for details about creating a new connection.

c. If the server type is VMWare vCloud, configure the following fields:

   i. Select the **Virtual Data Center** from the list of network ports or networks.

   ii. Select the **Catalog** from the list.

   iii. Select a **VM** (Virtual Machine) from the list.

   iv. Enter a **vApp Name** site name is default, or it can be manually entered.

      The vApp Name indicates what vApp to associate with the current Virtual Machine.

   v. Enter the **NTP Server**.

d. If the server type is OpenStack API or OpenStack Heat, configure the following fields:

   i. Select the **Image** from the list.

   ii. Select the**Flavor** from the list.

   iii. Select an **Availability Zone** from the list.

   iv. Verify the Config Drive.

      You cannot change this field.

   v. Enter the **NTP Server**.

   vi. Click **Add New** to add a DNS server.

   vii. Click **Add New** to add a DNS search.

   viii. Click **Manage** to add Security Groups.

e. Click **Add New IP** to add an **IP** address.

f. Select the **IP Preference** as either **IPv4** or **IPv6**.

g. Click **Add New IP** to add an **IP** address.

   This is a fixed IP address for the VM device.

h. Enter the **HostName**.

i. Select to have the server in **Forced Standby**, see Changing Server Status to Forced Standby.

j. Click **Add New** to add a new **Static IP** address.

6. (Optional) Click **Add Server-B** and enter the information for the standby server of the cluster.

   Server-B is defined for the cluster.

7. Click **Save**.

   A confirmation message appears.

8. Click **OK**.

About Setting Up a Georedundant Cluster

Before setting up georedundant clusters, you must create one or more georedundant sites. You can only create georedundant sites when **Manage Geo-Redundant** mode is enabled (see CMP Modes). See Setting Up a Georedundant Site for detailed information.

Note: Before setting up sites, you should plan your Policy Management topology to determine site naming conventions.

Georedundant sites can contain one or more of the following clusters:

- MPE
- BoD
- Mediation
- MA

Note: A site name is required when configuring georedundant non-CMP devices.

Setting Up a Georedundant Site

Note: Sites may only be created when **Manage Geo-Redundant** mode is enabled. See #unique_110 for details.

To set up a georedundant site:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

2. From the content tree, select **All Sites**.

   The Site Configuration page opens.

3. Click **Create Site**.

   The New Site page opens.

4. (Required) Enter the **Name** for the site.

The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_). The maximum length is 35 characters.

5. Enter the number for **Max Primary Site Failure Threshold**.

   If the number of cluster pair failures reaches this threshold, the system generates a trace log entry and a major alarm. A pair failure is recorded when both servers at a primary site are either out of service or in forced standby. The default value is no threshold.

   > **Note:** You can optionally enter a number up to the total number of servers provisioned at this site.

6. Select the **HW Type** from the list.

   The available options are:

   - **Oracle RMS** for a Oracle X5-2 server
   - **HP ProLiant G8 RMS** rack mounted server
   - **VM** for a virtual machine
   - **VM (Automated)** for a VM managed by NF Agent

7. If the hardware type is **C-Class**, **C-Class(Segregated Traffic)**, or **Oracle RMS**, configure the **General Network** settings.

   Virtual LAN (VLAN) IDs are in the range of 1 to 4095.

   a. Enter the **OAM VLAN ID**.

   b. Enter the **SIG-A VLAN ID**.

   c. (Optional) Enter the **SIG-B VLAN ID**.

   d. (Optional) Enter the **SIG-C VLAN ID**.

8. If the hardware type is **C-Class** or **C-Class(Segregated Traffic)**, enter the **VLAN ID** for the **User Defined Network**.

   Virtual LAN (VLAN) IDs are in the range of 1 to 4095.

9. Click **Save**.

The CMP database saves the site configuration.

To define multiple sites, repeat the procedure starting at step 3.

Setting Up a Georedundant Non-CMP Cluster

> **Note:** Georedundancy requires that you configure the CMP system with **Manage Geo-Redundant** enabled. See the *CMP Wireless User's Guide* for more information.

Before defining a cluster, ensure the following conditions are met:

- The server software is installed on all servers in the cluster.
- The servers have been configured with network time protocol (NTP), domain name server (DNS), IP Routing, and OAM IP addresses.

A georedundant non-CMP cluster is one of the following server types:

- BoD
- Management Agent
- Mediation server

Note: The list of available server types depends on the CMP modes configured. See CMP Modes for more information.

Note: If your system is not set up for georedundancy, see Setting Up a Non-CMP Cluster.

To set up a georedundant non-CMP cluster:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

2. From the work area, select **Add MPE/BoD/MA ClusterAdd MPE/MRA/Mediation Cluster**.

   The Topology Configuration page opens.

3. In the **Cluster Settings** section of the page:

   a. (Required) Enter the **Name** for the site.

      The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_). The maximum length is 35 characters.

   b. Select an **Appl Type**.

      Note: The list of available cluster types to add to the topology depends on the CMP modes configured.

      Note: The Management Agent server does not support georedundancy.

   c. Select the **Site Preference**.

      Available options are **Normal** (default) or **Reverse**. See Georedundant Site Preferences for more information.

   d. Select the type of **DSCP Marking** (Differentiated Services Code Point) for replication traffic.

      The valid code points are:

Assured Forwarding
AF11, AF12, AF13, AF21, AF22, AF23, AF31, AF32, AF33, AF41, AF42, AF43

Class Selector
CS1, CS2, CS3, CS4, CS5, CS6, CS7

Expedited Forwarding
EF

Default (for no marking)
PHB (None)

For information on DSCP marking, see

e. Select the **Replication Stream Count**.

This is the number of redundant TCP/IP socket connections (streams) to carry replication traffic between sites. Up to 8 streams can be configured. The default value is `1` stream.

f. Select a **Replication & Heartbeat** network to carry inter-site replication and heartbeat traffic.

This field only is visible if the system supports georedundancy:

- **None** (default)
- **OAM**
- **SIG-A**
- **SIG-B**
- **SIG-C**

> **Note:** When saving a configuration using **SIG-C**, a confirmation appears. Click **OK**. The **RMS** option for **HW Type** is removed until all configured Signaling C VIPs or **SIG-C** interfaces in static IP are removed.

A warning icon (⚠) indicates that you cannot select a network until you define a static IP address on all servers of both sites.

g. Select a **Backup Heartbeat** network to carry inter-site backup heartbeat traffic.

> **Note:** When saving a configuration using **SIG-C**, a confirmation message appears. Click **OK**. The **RMS** option for **HW Type** is removed until all configured Signaling C VIPs or **SIG-C** interfaces in static IP are removed.

> **Note:** This field only is visible if the system supports georedundancy.

Available options are:

- **None** (default)

- **OAM**
- **SIG-A**
- **SIG-B**
- **SIG-C**

A warning icon (⚠) indicates that you cannot select a network until you define a static IP address on all servers of both sites.

4. In the **Primary Site Settings** section of the page:

   a. Select the **Site Name** from the list.

   Select **Unspecified** (default) or the **Name** of a previously defined site. You can assign multiple clusters to the same site.

   > **Note:** If you select **Unspecified**, you create a non-georedundant site and cannot add a secondary site.

   b. Select the **HW Type** from the list.

   Available options are:

   - **Oracle RMS** – Oracle Server X5-2
   - **HP ProLiant G8 RMS**
   - **VM** (virtual machine)
   - **VM(Automated)** (VM managed by NF Agent)

     See Setting Up a VM (Automated) Non-CMP Cluster for details on adding a VM (Automated) cluster.

   c. (Required) To enter up to two **OAM VIP** (one IPv4 and one IPv6) addresses, click **Add New VIP**.

   The New OAM VIP dialog box appears.

   i. Enter the **OAM VIP** address and the **Mask**.

   This is the IP address the CMP server uses to communicate with a Policy Management cluster.

   > **Note:** Enter the IPv4 address in standard dot format and its subnet mask in CIDR notation from 0 to 32.

   ii. Click **Save**

   The **OAM VIP** address and **Mask** are saved. Repeat the process for the second OAM VIP.

   d. (Optional) To enter up to six **Signaling VIPs** addresses (up to two each for each of SIG-A, SIG-B,and SIG-C), click **Add New VIP**.

The signaling VIP is the IP address a PCEF device uses to communicate with the cluster. A non-CMP cluster supports redundant communication channels, named SIG-A and SIG-B, for carriers who use redundant signaling channels.

The New Signaling VIP dialog box appears.

i. Enter the **Signaling VIP** address and the **Mask**.

   This is the IP address the CMP server uses to communicate with an external signaling network.

> **Note:** Enter the IPv4 address in standard dot format and its subnet mask in CIDR notation from 0 to 32.

ii. Select the **Interface** from the list.

   Available options are:

   - SIG-A
   - SIG-B
   - SIG-C

iii. Click **Save**.

   The **Signaling VIP** address and **Mask** are saved.

5. To configure Server-A, in the **Server-A** section of the page:

   a. (Required) To enter the **IP** address, click **Add New IP**.

   The Add New IP dialog box appears.

   i. Enter the **IP** address in either IPv4 or IPv6 format.

   This is the IP address of the server. For an IPv4 address, enter it in the standard IP dot-format. For an IPv6 address, enter it in the standard 8-part colon-separated hexadecimal string format.

   ii. Select the **IP Preference**: **IPv4** or **IPV6**.

   The server will preferentially use the IP address in the specified format for communication.

   - If neither an IPv6 OAM IP nor a static IP address is defined, **IPv6** cannot be selected.
   - If neither an IPv4 OAM IP nor a static IP address is defined, **IPv4** cannot be selected.

   b. Enter the **HostName** of the server.

   The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_). This must exactly match the host name provisioned for this server (the output of the Linux command `uname -n`).

Note: If the server has a configured server IP, you can click **Load** to retrieve the remote server host name. If the retrieve fails, you must enter the host name.

---

c. Select **Forced Standby** to put Server-A into forced standby.

   By default, Server-A will be the initial active server of the cluster.

d. In the **Path Configuration section**, to add a **Static IP**, click **Add New**.

   The New Path dialog box appears.

Note: If an alternate replication path and secondary HA heartbeat path is used, a server **Static IP** address must be entered in this field.

---

   i. Enter a **Static IP** address and **Mask**.

   ii. Select the **Interface**:

      - **SIG-A**
      - **SIG-B**
      - **SIG-C**
      - **BKUP**

6. (Optional) To configure Server-B, in the **Server-B** section of the page:

   a. (Required) To enter the **IP** address, click **Add New IP**.

      The Add New IP dialog box appears.

      i. Enter the **IP** address in either IPv4 or IPv6 format.

         The IP address of the server. For an IPv4 address, enter it in the standard IP dot-format. For an IPv6 address, enter it in the standard 8-part colon-separated hexadecimal string format.

      ii. Select the **IP Preference**: **IPv4** or **IPV6**.

         The server will preferentially use the IP address of the specified format for communication.

         - If neither an IPv6 OAM IP nor a static IP address is defined, **IPv6** cannot be selected.
         - If neither an IPv4 OAM IP nor a static IP address is defined, **IPv4** cannot be selected.

   b. Enter the **HostName** of the server.

      The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_). This must exactly match the host name provisioned for this server (the output of the Linux command `uname -n`).

      If the server has a configured server IP, you can click **Load** to retrieve the remote server host name. If the retrieve fails, you must enter the host name.

c. Select **Forced Standby** to put Server-B into forced standby.

   By default, Server-A will be the initial active server of the cluster.

d. In the **Path Configuration section**, to add a **Static IP**, click **Add New**.

   The New Path dialog box appears.

   > **Note:** If an alternate replication path and secondary HA heartbeat path is used, a server **Static IP** address must be entered in this field.

   i. Enter a **Static IP** address and **Mask**.

   ii. Select the **Interface**:

   - **SIG-A**
   - **SIG-B**
   - **SIG-C**
   - **BKUP**

7. Click **Save**.

   A confirmation message appears.

8. Click **OK**.

9. If you are setting up multiple clusters, repeat this procedure.

The cluster is defined.

Figure 2-10 and Figure 2-11 show the configuration for a georedundant (two-site) BoD cluster, using no replication network and OAM for the backup heartbeat network.

Figure 2-10 Sample BoD Cluster Site 1 Configuration

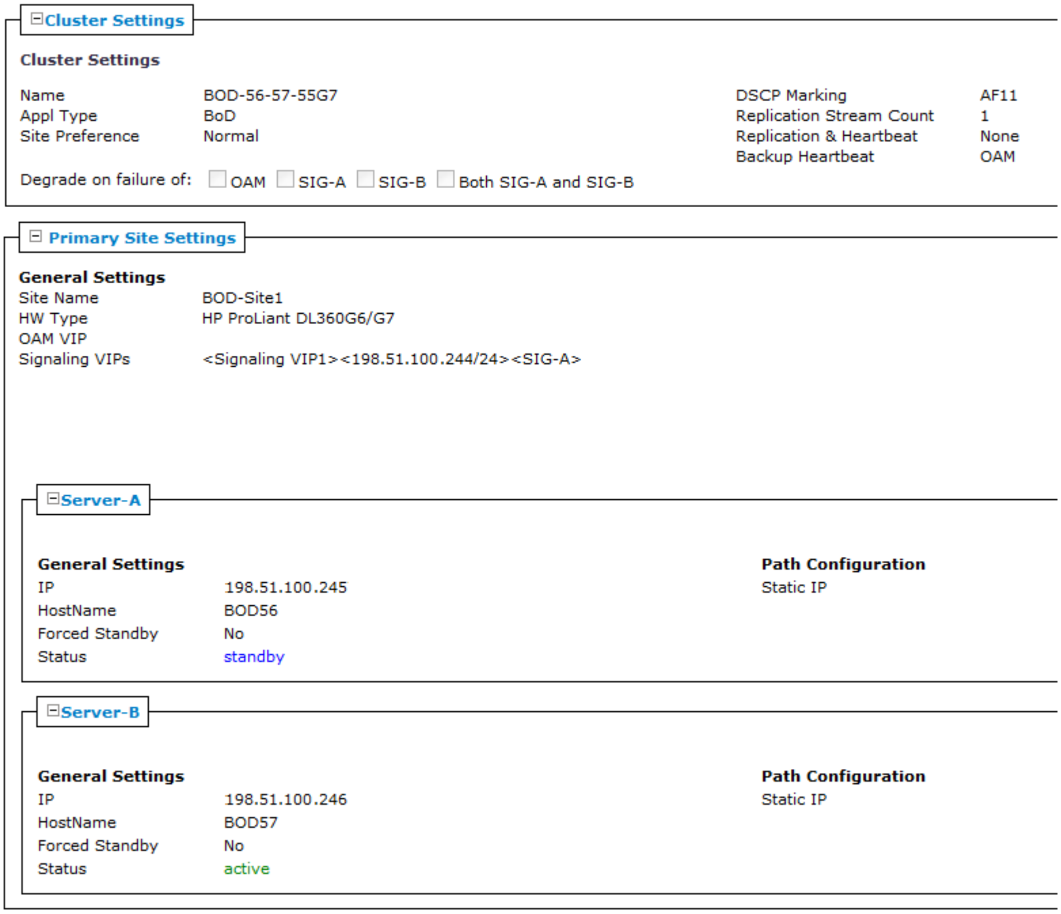

Cluster Settings

Cluster Settings

Name	BOD-56-57-55G7	DSCP Marking	AF11
Appl Type	BoD	Replication Stream Count	1
Site Preference	Normal	Replication & Heartbeat	None
		Backup Heartbeat	OAM

Degrade on failure of: ☐ OAM ☐ SIG-A ☐ SIG-B ☐ Both SIG-A and SIG-B

Primary Site Settings

General Settings

Site Name	BOD-Site1
HW Type	HP ProLiant DL360G6/G7
OAM VIP	
Signaling VIPs	<Signaling VIP1><198.51.100.244/24><SIG-A>

Server-A

General Settings		**Path Configuration**
IP	198.51.100.245	Static IP
HostName	BOD56	
Forced Standby	No	
Status	standby	

Server-B

General Settings		**Path Configuration**
IP	198.51.100.246	Static IP
HostName	BOD57	
Forced Standby	No	
Status	active	

Figure 2-11 Sample BoD Cluster Site 2 Configuration

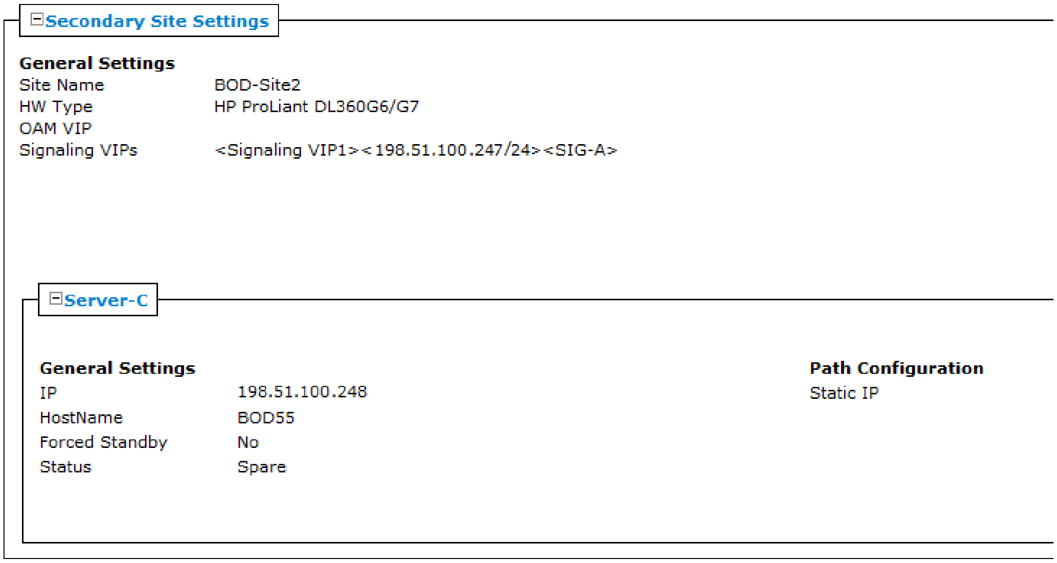

Secondary Site Settings

General Settings

Site Name	BOD-Site2
HW Type	HP ProLiant DL360G6/G7
OAM VIP	
Signaling VIPs	<Signaling VIP1><198.51.100.247/24><SIG-A>

Server-C

General Settings		**Path Configuration**
IP	198.51.100.248	Static IP
HostName	BOD55	
Forced Standby	No	
Status	Spare	

Example: Adding Georedundancy to an Existing Topology

This example describes how to:

- Add a georedundant secondary site, Site-2, to an existing Policy Management topology
- Add a third spare server (Server-C), located at Site-2, to an existing active/standby MPE cluster located at the primary site, Site-1

> **Note:** If the primary site was to fail, the spare server would assume the active role.

The example includes recommended verification steps and refers to tasks described elsewhere.

> **Note:** Before undertaking this procedure, contact My Oracle Support (MOS) for assistance.

Before creating a georedundant cluster, ensure the following:

- All servers in the topology are using the latest Policy Management software.
- The new server (Server-C) is of a supported hardware type and has been delivered with the latest firmware and TPD software pre-installed.

Gather required information. See #unique_113 for details.

After gathering the required information, you can proceed to performing your tasks. See Step 1: Setting Up Server-C for an example.

Step 1: Setting Up Server-C

Before you begin, be sure to have the required information available for reference. See #unique_113.

> **Caution:** This procedure interrupts service.

To prepare Server-C for addition to a georedundant secondary site:

1. Using the Platform Management & Configuration (PMAC) utility, install the MPE application software on Server-C.

   For more information, refer to the relevant chapter of Oracle Communications Policy Management *Bare Metal Installation Guide*. If you have problems or questions, contact My Oracle Support for assistance.

2. Using the Platform Configuration (`platcfg`) utility, provision Server-C with the following configuration information:

   a. Host Name

   b. OAM Real IP Address

   c. OAM Default Route

   d. NTP Server

   e. DNS Server A

   f. DNS Server B (optional)
   g. DNS Search
   h. Device

   For more information, refer to *Platform Configuration User's Guide*.

3. For Oracle X5-2 platforms, configure the following:
   a. OAM VLAN ID
   b. SIG A VLAN ID
   c. SIG B VLAN ID (optional)
   d. SIG C VLAN ID (optional)
4. Using the Platform Configuration utility, export routing configuration information from Server-A or Server-B and import it into Server-C.

Server-C is ready to be added to the secondary site.

Proceed to Step 2: Setting Up the Georedundant Sites.

Step 2: Setting Up the Georedundant Sites

Prior to performing this step you must have competed Step 1: Setting Up Server-C.

> **Caution:** This procedure interrupts service.

To create georedundant primary and secondary sites:

1. Log in to the CMP system, using its OAM VIP address.

2. If this is the first georedundant cluster in your topology, configure the CMP system to enable **Manage Geo-Redundant** mode.

   See #unique_65.

   With georedundancy enabled, the content tree shows the **All Sites** group when you select **Topology Settings** from the **Platform Setting** section of the navigation pane.
3. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens.
4. From the content tree, select the **All Sites** group.

   The Site Configuration page opens.
5. Click **Create Site**.

   The New Site page opens.
6. Enter the **Name** for the site name: `Site-1`.
7. Enter the number of **Max Primary Site Failure Threshold**.

   The default value is zero (0).

8. Select the **HW Type** from the list.

9. Click **Save**.

   See #unique_116 for more information.

   The sites become visible on the Site Configuration page.

Figure 2-12 Successful Site Creation

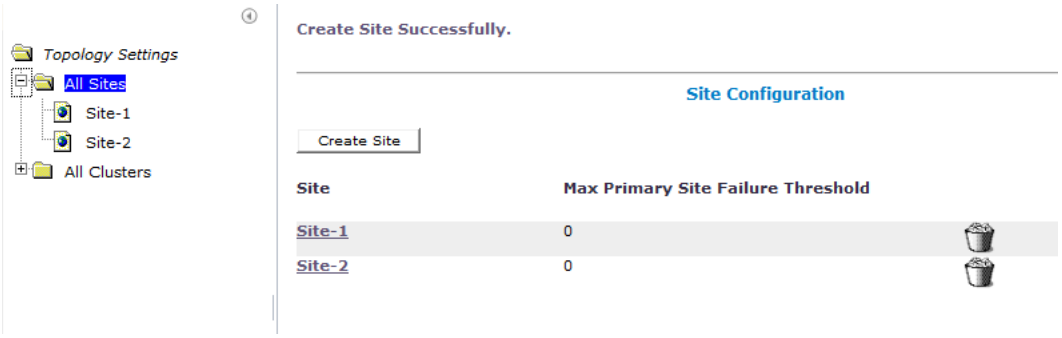

Create Site Successfully.

Site Configuration

Site	Max Primary Site Failure Threshold
Site-1	0
Site-2	0

10. From the content tree, select the **All Clusters** group.

    The Cluster Configuration page opens, displaying the defined clusters.

11. On the Cluster Configuration page, for the MPE cluster you are expanding with Server-C, click the **View** operation.

    The Topology Configuration page opens for the selected MPE cluster.

12. Click **Modify Primary Site**.

    The fields in the **Primary Site Settings** section of the page becomes editable.

13. In the **Primary Site Settings** section of the page:

    a. In the **Site Name** field, select the primary site name (**Site-1** in this example).

    b. Confirm the configuration settings in the **HW Type** field, **Network Configuration** section, and **Signaling VIPs** field.

14. In the **Server-A** section of the page:

    a. Confirm the settings in the **General Settings** section.

    b. In the **Path Configuration** section, click **Add New**.

       The New Path dialog box appears.

       i. Enter the **Static IP** address and subnet **Mask**.

       ii. Select the **Interface** (for this example, the SIG-A network).

       iii. Click **Save**.

15. Repeat step 11 for **Server-B**.

    The Primary Site Settings are defined.

16. Click **Save**.

A restart message appears.

17. Click **OK**.

    Server-A restarts.

The two sites, primary and secondary, and the georedundant primary MPE cluster are configured.

Proceed to Step 3: Setting Up Primary Site Cluster.

Step 3: Setting Up Primary Site Cluster

Before you proceed, you must have completed Step 2: Setting Up the Georedundant Sites.

> **Caution:** This procedure interrupts service.

To configure the primary site cluster:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

2. From the content tree, select **All Clusters**.

   The Cluster Configuration page opens.

3. Click **View** for the MPE cluster you are expanding with Server-C.

   The Topology Configuration page for the selected MPE cluster opens.

4. Click **Modify Primary Site**.

   The fields in the **Primary Site Settings** section of the page becomes editable.

5. In the **Primary Site Settings** section of the page:

   a. Select the primary **Site Name** from the list: **Site-1**.

   b. Confirm the configuration settings: **HW Type**, **Network Configuration**, and **Signaling VIPs**.

6. In the **Server-A** section of the page, confirm the settings in the **General Settings** section.

7. In the **Path Configuration** section, click **Add New**.

   The New Path dialog box appears.

   a. Enter the **Static IP** address and subnet **Mask**.

   b. Select the **Interface** (for this example, the **SIG-A** network).

   c. Click **Save**.

8. Repeat the process for **Server-B**.

9. Click **Save**.

A restart message appears.

10. Click **OK**.

    Server-A restarts.

The Primary Site Settings are defined.

Proceed to Step 4: Setting Up the Secondary Site with Server-C.

Step 4: Setting Up the Secondary Site with Server-C

Before you proceed, you must have completed Step 3: Setting Up Primary Site Cluster.

> **Caution:** This procedure interrupts service.

To configure the secondary site and Server-C settings:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

2. From the content tree, select **All Clusters**.

   The Cluster Configuration page opens.

3. Click **View** for the MPE cluster you are expanding with Server-C.

   The Topology Configuration page for the selected MPE cluster opens.

4. Click **Modify Secondary Site**.

   The fields in the **Secondary Site Settings** section of the page become editable.

5. In the **Secondary Site Settings** section of the page:

   a. Select the secondary **Site Name** from the list: **Site-2**.

   b. Confirm the configuration settings: **HW Type**, **Network Configuration**, and **Signaling VIPs**.

6. In the **Server-C** section of the page, click **Add Server-C**.

   The Server-C section becomes editable.

7. To enter the **OAM IP** address, click **Add New VIP**.

   The New IP dialog box appears.

   a. Enter the **IP** address.

   b. Click **Save**.

8. Select the **IP Preference**: either **IPv4** or **IPv6**.

   The server will preferentially use the selected format's IP address for communication.

Note: If neither an IPv6 OAM IP nor a static IP address is set, IPv6 cannot be selected. If neither an IPv4 OAM IP nor a static IP address is set, IPv4 cannot be selected.

9. Enter the **HostName**.
10. Select **Forced Standby**.
11. In the **Path Configuration** section, click **Add New**.

    The New Path dialog box appears.

    a. Enter the **Static IP** address and subnet **Mask**.

    b. Select the **Interface** from the list (for this example, the **SIG-A** network).
12. If the **REP** network is used, in the **User Defined Network** section, enter the **VLAN ID** for the **REP** network.
13. Click **Save** (at the bottom of the page).

    A restart message appears.
14. Click **OK**.

    Server-C restarts.

Note: The status of Server-C is `Out of Service` and critical alarm 31283 is raised; this is expected.

15. Click the **Status** of Server-C.

    The status changes to **Spare**.
16. Click **Save**.

    The configuration is saved.

Site-2 and Server-C are defined, and Server-C is placed in Force Standby status.

Proceed to Step 5: Updating the Georedundant Cluster.

Step 5: Updating the Georedundant Cluster

Before you proceed, you must have completed Step 4: Setting Up the Secondary Site with Server-C

This procedure updates the MPE cluster settings, if needed, and verifies the server is functioning properly.

Caution: This procedure interrupts service.

To update the georedundant MPE cluster:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

2. From the content tree, select **All Clusters**.

   The Cluster Configuration page opens.

3. Click **View** for the MPE cluster you are expanding with Server-C.

   The Topology Configuration page for the selected MPE cluster opens.

4. Click **Modify Cluster Settings**.

   The fields in the **Cluster Settings** section of the page become editable.

5. In the **Cluster Settings** section of the page:

   a. If DSCP marking is used, select the type of **DSCP Marking** from the list.

   b. If replication streams are used, select the number of stream from the **Replication Stream Count** list.

   c. Select the network used for the **Replication & Heartbeat** (or **None** to use the system default).

   d. If used, select the network used for the **Backup Heartbeat** (or **None** to disable the feature).

6. Click **Save**.

   The **Cluster Settings** are saved.

7. Verify the **Status** of Server-C by viewing the cluster.

   Server-C is shown as part of the cluster in **Force Standby** with replication on.

8. Use the **Alarm History Report** and filter in all alarms on the cluster name to verify that no new alarms have been raised.

   For more information, see Viewing the Alarm History Report.

   Alarm 31102 (DB Replication from a master DB failed) is in the report, but with severity of Clear.

9. Using the Platform Configuration utility on Server-C, exchange SSH keys with the other servers of the cluster.

   Refer to *Platform Configuration User's Guide* for detailed information.

10. Using the Platform Configuration utility on the CMP system, exchange SSH keys with all other CMP systems in the topology.

11. Using the CMP system, modify the cluster configuration to cancel the **Force Standby** status of Server-C.

    See #unique_121 for details.

    The status of Server-C changes to **Spare**.

12. Use the **KPI Dashboard** to verify that Server-C is reporting its status as part of the cluster.

    For more information, see KPI Dashboard.

13. (Optional) Use the **Policy Checkpoint** function to create a policy checkpoint.

    Tip: If the function is not available, ensure that the system settings allow policy checkpoints. See Configuring System Settings.

    For more information on policy checkpoints, refer to *Policy Wizard Reference*.

14. To configure the connections on Server-C to existing data sources, use the **Data Sources** tab.

    For more information, see Connecting to Data Sources.

Proceed to #unique_123

Modifying the Topology

You can modify the topology to:

- Correct errors
- Add a server to a cluster
- Define new clusters
- Add clusters to an existing site
- Define new sites
- Change which cluster is primary and which secondary
- Put an active server into standby status

Note: You can modify a cluster even if the standby server is offline. However, you cannot modify or delete the active server of a cluster.

About Managing Georedundant Sites

Note: Sites are only available for a system with **Manage Geo-Redundant** enabled. See CMP Modes for detailed information.

Modifying and configuring sites, hardware, and IP addresses requires Platform Configuration (`platcfg`) and PMAC (also known as PM&C) software. For detailed information, refer to the relevant Policy Management documentation:

- *Platform Configuration User's Guide*
- *Bare Metal Installation Guide*

Note: The *Bare Metal Installation Guide* includes PMAC procedures for Policy Management.

For detailed information, refer to the following Tekelec Platform PMAC documentation:

- *PM&C Incremental Upgrade*
- *PM&C Disaster Recovery*

You can manage a site using the following procedures:

- Setting Up a Georedundant Site
- Modifying a Georedundant Site
- Updating Site Information
- Removing a Site from the Topology

Modifying a Georedundant Site

Note: You must enable **Manage Geo-Redundant** to modify sites within the Policy Management topology. See CMP Modes for more information.

To modify a georedundant site:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

2. From the content tree, select **All Sites**.

   The Site Configuration page opens.

3. Select the site you want to modify.

   The Site Configuration page displays information about the site.

4. Click **Modify**.

   The Modify Site page opens.

5. Modify site information.

   For a description of the fields contained on this page, see Setting Up a Georedundant Site.

6. Click **Save**.

Your changes to the site are saved.

Removing a Site from the Topology

You can remove a site only if the site is not referenced by a Server C-level cluster. If you try to delete a site that is in use by a cluster, you will receive an information message indicating that the cluster cannot be removed because it is being referred to by at least one other cluster. The message also lists the referring clusters.

To remove a site from the topology:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

2. From the content tree, select **All Sites**.

   The Site Configuration page opens.

3. Delete the site using one of the following methods:

   - From the work area, click 🗑 (Delete icon), located to the right of the site.
   - From the content tree, select the site and click **Delete**.

   A confirmation message appears.

4. Click **OK**.

The site is removed from the topology.

Updating Site Information

You can update site information (for example, IP address, host name, or VLAN IDs) for a secondary site in a cluster. For more information on modifying a cluster, see Modifying a non-CMP Cluster.

To update the secondary site in a cluster:

1. On the CMP server, remove the secondary site from the topology.

2. Using the Platform Configuration (`platcfg`) utility, re-create the configuration for the secondary site to change the IP address, host name, or VLAN ID settings.

   Refer to *Platform Configuration User's Guide* for information on configuring a site.

3. On the CMP server, from the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

4. In the CMP system, add the secondary site.

   A confirmation message appears.

5. Click **OK**.

6. After the site has been updated, reapply the configuration for the changes to take effect.

   See About Reapplying a Configuration for more information.

The site has been updated in the cluster.

About Managing Clusters

Managing clusters, both CMP and non-CMP servers, entails the following:

- Setting Up a CMP Cluster
- Setting Up a Non-CMP Cluster
- Setting Up a Georedundant Non-CMP Cluster
- Modifying a non-CMP Cluster

- Modifying a CMP Cluster
- Removing a Cluster from the Topology
- Reversing Georedundant Cluster Preference
- Demoting a Georedundant CMP Cluster
- Promoting a Georedundant CMP Cluster
- Changing Server Status to Forced Standby
- Changing Server Status from Forced Standby

Modifying a non-CMP Cluster

To modify an non-CMP cluster:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

2. From the content tree, select **All Clusters**.

   The Cluster Configuration page opens.

3. Click **View** for the cluster you want to modify.

   The Topology Configuration page opens, displaying information about the cluster.

4. Click the button for the changes you want to make:

   - To modify cluster settings, click **Modify Cluster Settings**.
   - To modify the primary server, click **Modify Server-A**.
   - To modify the secondary server, click **Modify Server-B**.
   - To delete a server configuration, click the appropriate **Modify** button and then click **Delete**.

   The appropriate section on the Topology Configuration page becomes editable.

5. Make changes as required.

   You must make changes to each section individually.

   - You can remove all servers from a cluster.
   - You can select **Forced Standby** on one or more servers in the cluster.

   Caution: If you force all servers in a cluster into the Standby state, then no server can be active, which effectively removes the cluster from service.

   Note: If you add, remove, or modify a server, the active server restarts.

6. Click **Save**.

   A warning message appears.

7. Click **OK**.

The cluster is modified. You can determine if there is a topology mismatch by viewing the **System** tab for the specific server.

Modifying a CMP Cluster

To modify a CMP cluster:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

2. From the content tree, select **All Clusters**.

   The Cluster Configuration page opens.

3. Click **View** for the CMP cluster you want to modify.

   The Topology Configuration page opens, displaying information about the CMP cluster.

4. Click the button for the changes you want to make:

   - To modify cluster settings, click **Modify Cluster Settings**.
   - To modify the primary server, click **Modify Server-A**.
   - To modify the secondary server, click **Modify Server-B**.

   The appropriate section on the Topology Configuration page becomes editable. For information on configurable settings, see Setting Up a CMP Cluster.

5. Make the changes as required.

   You must make changes to each section individually.

   - You can remove either server from the cluster, but not both.
   - You can select **Forced Standby** on either server of the cluster, but not both, and not at all if the cluster has only one server.

   Note: If you add, remove, or modify a server, the active server restarts.

6. Click **Save**.

   A restart message appears.

7. Click **OK**.

The changes to the CMP cluster are saved. You can determine if there is a topology mismatch by viewing the **System** tab for each policy server profile.

Removing a Cluster from the Topology

You can remove a non-CMP or secondary CMP cluster from the topology.

Note: You cannot remove the primary CMP cluster from the topology.

To remove a cluster from the topology:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

2. From the content tree, select **All Clusters**.

   The Cluster Configuration page opens.

3. In the **Cluster Settings** table, in the row listing the cluster you want to remove, click **Delete**.

   A confirmation message appears.

4. Click **OK**.

   An instructional message with further instructions appears.

The cluster is removed from the topology.

After the cluster is removed, use the Platform Configuration (PlatCfg) utility to remove cluster information. For more information, see the *Platform Configuration User's Guide*.

Reversing Georedundant Cluster Preference

If your system has been configured for georedundancy (**Manage Geo-Redundant** mode is enabled), there can be situations when you need to change the preference of the servers in a cluster to be active or spare. See Georedundant Site Preferences for more information.

To reverse a georedundant cluster preference:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

2. From the content tree, select **All Clusters**.

   The Cluster Configuration page opens.

3. Click **View** for the cluster you want to modify.

   The Topology Configuration page opens, displaying information about the cluster.

4. Click **Modify Cluster Settings**.

5. In the **Cluster Settings** section of the page:

   - To set the preference to reverse (where the active Site 1 becomes the inactive site and Site 2 becomes the active site), toggle to **Reverse**.

   - To set the preference to normal (where the active Site 2 becomes the inactive site and Site 1 becomes the active site), toggle to **Normal**.

6. Click **Save**.

The cluster preferences are reversed.

Demoting a Georedundant CMP Cluster

In a two-cluster CMP topology, you can demote the primary cluster (which is typically the Site 1 cluster) to secondary status. You would do this, for example, prior to performing site-wide maintenance that affects service (such as replacing a server) or if the primary cluster has failed completely and is unreachable.

Note: This is a manual process.

When you demote a CMP cluster, the secondary site (which is typically the Site 2 cluster) can be promoted to the primary site (see Promoting a Georedundant CMP Cluster for details). This promoted status will persist until you manually demote the new primary site or the primary site fails over for some reason.

Caution: Demote the primary CMP cluster before promoting another CMP cluster to avoid having both georedundant clusters active at the same time. Continuous and rapid failovers (flopping back and forth) between georedundant clusters is not recommended and should be avoided. Improper cluster failover can result in loss of data or interruption of network services on the CMP cluster.

To demote a georedundant CMP cluster:

1. Log in to the currently active georedundant CMP cluster:

   a. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

      The Cluster Configuration page opens; the initial group is **All Clusters**.

   b. From the content tree, select **All Clusters**.

      The Cluster Configuration page opens.

      Note: The name of the primary CMP cluster is marked with **(P)**, and the name of the secondary cluster is marked with **(S)**.

      You should see the operations **View** and **Demote**.

2. Open a second browser window and log in to the secondary CMP cluster.

   The page displays the a message indicating that you are signed into Secondary Active server.

   Note: The state of the servers of the primary cluster is not available to the secondary active server and appears as **Out-of-Service**.

3. Verify the status of the secondary cluster by doing the following on the secondary CMP cluster:

   a. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

The Cluster Configuration page opens; the initial group is **All Clusters**.

b. From the content tree, select **All Clusters**.

The Cluster Configuration page opens.

> **Caution:** If you do not see the same information in this step as you did in 2, stop this procedure and do not try to change the current active georedundant cluster. Contact My Oracle Support before proceeding.

4. Return to the browser window showing the primary CMP cluster.

   You should still be on the Cluster Configuration page.

5. In the **Cluster Settings** table, in the row listing the primary CMP cluster, click **Demote**.

   A confirmation message displays.

6. Click **OK**.

7. Log out of the primary CMP system for the cluster you have just demoted.

The primary CMP cluster is demoted to secondary status.

After demoting a primary cluster, you must promote the secondary cluster for it to become active. See Promoting a Georedundant CMP Cluster for detailed information.

Promoting a Georedundant CMP Cluster

Prior to performing this procedure, you must demote the primary active cluster. See Demoting a Georedundant CMP Cluster for detailed information.

In a two-cluster CMP topology and after demoting the primary cluster, you can promote the secondary cluster (which is typically the Site 2 cluster) to primary active status. You would do this, for example, prior to performing site-wide maintenance that affects service (such as replacing a server) or if the primary cluster has failed completely and is unreachable.

> **Note:** This is a manual process.

When you promote a CMP cluster, the secondary site (which is typically the Site 2 cluster) becomes the primary site. This status will persist until you manually demote the new primary site or the primary site fails over for some reason.

> **Caution:** Demote the primary CMP cluster before promoting another CMP cluster to avoid having both georedundant clusters active at the same time. Continuous and rapid failovers (flopping back and forth) between georedundant clusters is not recommended and should be avoided. Improper cluster failover can result in loss of data or interruption of network services on the CMP cluster.

To promote a georedundant CMP cluster:

1. Log in to the secondary CMP cluster:

   The page displays a message indicating that you are signed into Secondary Active server.

   > **Note:** The state of the servers of the primary cluster is not available to the secondary active server and appears as **Out-of-Service**.

   a. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

   b. From the content tree, select **All Clusters**.

   The Cluster Configuration page opens.

   > **Note:** The name of the primary CMP cluster is marked with **(P)**, and the name of the secondary cluster is marked with **(S)**.

   For the secondary cluster, you should see the operations **View** and **Promote**.

   > **Caution:** If you do not see the same information in this step as you did in 2, stop this procedure and do not try to change the current active georedundant cluster. Contact My Oracle Support before proceeding.

2. If you have just demoted a primary cluster, wait 2 minutes.

3. In the **Cluster Settings** table, in the row listing the secondary CMP cluster, click **Promote**.

   A confirmation message appears.

4. Click **OK**.

5. Log out of the CMP system for the cluster you have just promoted.

6. Log in to the CMP system for the cluster you have just promoted.

7. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

8. From the content tree, select **All Clusters**.

   The Cluster Configuration page opens.

   The newly promoted primary cluster is marked with **(P)**, and the name of the demoted secondary cluster is marked with **(S)**. The old primary cluster may briefly display as off-line.

   > **Note:** For the new primary cluster, you should see the operations **View** and **Demote**. All functions available for the primary CMP cluster should now appear and be accessible.

9. Wait 10 minutes.

10. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

    The Cluster Configuration page opens; the initial group is **All Clusters**.

11. From the content tree, select **All Clusters**.

    The Cluster Configuration page opens.

12. Verify that both the primary and secondary CMP clusters are available and have the correct status.

The secondary CMP cluster is promoted to primary status.

Changing Server Status to Forced Standby

You can change the status of a server in a cluster to forced standby. A server placed into forced standby status cannot become active. You would do this, for example, to an active server prior to performing maintenance on it. Oracle recommends this method to switch over from an active server or to resolve issues where more than one server in a cluster is active.

When you place a server into forced standby status, the following actions occur:

- If the server is active, the server is demoted.
- The server will not assume the active role, regardless of its status or the roles of the other servers in the cluster.
- The server continues as part of its cluster and reports its status as Forced Standby.
- The server coordinates with the other servers in the cluster to take the role Standby or Spare.

> **Caution:** If you set all servers in a cluster into forced standby status, you can trigger an outage.

To change a server to forced standby status:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

2. From the content tree, select **All Clusters**.

   The Cluster Configuration page opens.

3. Click **View** for the cluster you want to change.

   The Topology Configuration page opens, displaying information about the cluster.

4. Click **Modify Server-A** or **Modify Server-B** (whichever server needs the status change).

5. Select **Forced Standby**.

6. Click **Save**.

The server status is changed to forced standby.

Changing Server Status from Forced Standby

A server placed into forced standby status is not active. You can change the status of a server in a cluster from forced standby. You would do this, for example, to return a server to active status after performing maintenance on it.

When you take a server from forced standby, the server coordinates with the other servers in the cluster to take either the role **Standby** or **Spare**.

To take a server from a forced standby status:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

   The Cluster Configuration page opens; the initial group is **All Clusters**.

2. From the content tree, select **All Clusters**.

   The Cluster Configuration page opens.

3. Click **View** for the cluster you want to change.

   The Topology Configuration page opens, displaying information about the cluster.

4. Click **Modify Server-A** or **Modify Server-B** (whichever server needs the status change).

5. Deselect **Forced Standby**.

6. Click **Save**.

The server status is changed from forced standby.

About NF Management

The NF Management function allows the creation of VIM connections for use with **VM (Automated)** devices. VIM connections work with the OpenStack API and OpenStack Heat API.

The NF Agent provides the following VM profile and deployment information to the VIM so it can instantiate instances of the described VNF:

- Required vCPUs
- Required vNICs
- Required Networks and IP addressing
- Memory size
- Storage size
- Anti-affinity/Affinity requirements

Creating a VIM Connection

To create a VIM connection:

1. From the **NF Management** section of the navigation pane, select **VIM Connections**.

The VIM Connections page opens; the initial group is **VIM Connections**.

2. From the content tree, select the **VIM Connections** group.

   The VIM Connections page opens.

3. Click **Create VIM Connection**.

   The Create VIM Connections page opens.

4. Enter a **Name** for the VIM connection.

   The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_).

   Note: Once saved, the VIM Connection Name cannot be changed.

5. (Optional) Enter a **Description** for the VIM connection.

6. Select the **VIM Type** from the list.

   Available options include:

   - **OpenStack API—**Indicates the connection will use the OpenStack API
   - **OpenStack Heat** —Indicates the connection will use the OpenStack Heat API
   - **VMWare vCloud—**Indicates the connection will use the vCloud API

7. Enter the **Host** name.

   Enter an IP address or the FQDN of the VIM host.

8. Enter the **Port** number.

   This is the port to connect to the VIM host. Enter a number from 1 to 65535. A typical port number is 5000.

9. (For OpenStack VIM types) Select to use a **Secure Connection**.

   If enabled, the connection will use an HTTPS connection to encrypt the connection.

10. Enter the **Username**.

11. (For OpenStack VIM types) Enter the **Tenant** name.

12. Enter the **Password**.

    Select **Show Password** to view the password in clear text.

13. Click **Save**.

The CMP server saves the VIM connection to the database.

Modifying a VIM Connection

To modify a VIM connection:

1. From the **NF Management** section of the navigation pane, select **VIM Connections**.

The VIM Connections page opens.

2. From the content tree, select the VIM connection to modify.
3. Click **Modify**.
4. Change the values of the configuration settings as needed.
5. Click **Save**.

Deleting a VIM Connection

To delete a VIM connection:

1. From the **NF Management** section of the navigation pane, select **VIM Connections**.

   The NF Management page opens.

2. From the content tree, select the VIM connection to be deleted.
3. Click **Delete**.
4. Click **OK** to confirm the delete.

Configuring SNMP Settings

You can configure SNMP settings for the CMP system and all Policy Management servers in the topology network. You can configure the Policy Management network such that the CMP system collects and forwards all traps to up to five external systems (SNMP managers) or such that each server generates and delivers its own traps.

> **Note:** SNMP settings configuration must be done on the active CMP server in the primary cluster. A warning displays if the login is not on the active primary CMP system.

To configure SNMP settings:

1. From the **Platform Setting** section of the navigation pane, select **SNMP Settings**.

   The SNMP Settings page opens, displaying the current settings.

2. Click **Modify**.

   The Edit SNMP Settings page opens.

3. For each SNMP manager, enter a valid host name or an IPv4/IPv6 address.

   The **Hostname/IP Address** field is required for an SNMP Manager to receive traps and send SNMP requests. The field has the following restrictions:

   - The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_).
   - The maximum length is 20 characters.
   - The name is case insensitive (uppercase and lowercase are treated as the same).

   By default, these fields are blank.

4. (Optional) You can configure a port for each SNMP manager by entering a port value between 1 and 65535 in the **Port** field. If left blank, the default value is 162.

5. From the **Enabled Versions** list, select one of the following versions:

   - **SNMPv2c**
   - **SNMPv3**
   - **SNMPv2c and SNMPv3** (default)

6. If you selected **SNMPv2c** or **SNMPv2c and SNMPv3** from the **Enabled Versions** list, configure the following:

   a. **Traps Enabled—**Specifies whether sending SNMPv2 traps is enabled. The default is enabled.

   Note: To use the **SNMP Trap Forwarding** feature, enable this option.

   b. **Traps from individual Servers—**Specifies whether sending SNMPv2 traps from individual servers is enabled. If disabled, SNMPv2 traps are only sent from the active CMP system only. The default is disabled.

   Note: To use the **SNMP Trap Forwarding** feature, disable this option.

   c. **SNMPv2c Community Name—**Enter the SNMP read-write community string. This field has the following restrictions:

   - The field is required if SNMPv2c is enabled.
   - The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_).
   - The name cannot exceed 31 characters in length.
   - The name cannot be either `private` or `public`.

   The default value is `snmppublic`.

7. If you selected **SNMPv3** or **SNMPv2c and SNMPv3** from the **Enabled Versions** list, configure the following:

   a. **SNMPv3 Engine ID—**Enter an Engine ID for SNMPv3. The Engine ID can be 10 to 64 digits long and must use only hexadecimal digits (0-9 and a-f). The default is no value (null).

   b. **SNMPv3 Security Level—**Select the level of SNMPv3 authentication and privacy from the list:

   - **No Auth No Priv—**Authenticate using the **Username**. No Privacy.
   - **Auth No Priv—**Authenticate using MD5 or SHA1 protocol.
   - **Auth Priv** (default)—Authenticate using MD5 or SHA1 protocol. Encrypt using the AES or DES protocol.

   c. **SNMPv3 Authentication Type—**Select an SNMPv3 authentication protocol from the list:

- **SHA-1**—Use Secure Hash Algorithm authentication.
- **MD5** (default)—Use Message Digest authentication.

d. **SNMPv3 Privacy Type**—Select an SNMPv3 privacy protocol from the list:

- **AES** (default)—Use Advanced Encryption Standard privacy.
- **DES**—Use Data Encryption Standard privacy.

e. **SNMPv3 Username**—Enter a user name. The user name can contain 0 to 32 characters and must only contain alphanumeric characters. The default is `TekSNMPUser`.

f. **SNMPv3 Password**—Enter an authentication password. The password must contain between 8 and 64 characters and can include any character.

Note: The SNMPv3 password is also used for `msgPrivacyParameters`.

8. Click **Save**.

The SNMP settings for the network are configured.

Platform Configuration Settings

The Platform Configuration Settings page sets global options that are applicable for all other components that have the same georedundant arrangement.

Configuring the Upsync Log Alarm Threshold

You can configure the threshold of outstanding updates to a secondary server that triggers an alarm. When the outstanding updates reaches a configured percent of the upsync log capacity, an event is issued and the current condition of the connection (volume of outstanding data, current throughput, time of the event, and so forth) is logged.

The events are tracked in the MPE/BoD replication report. See Viewing the MPE/BoD Replication Statistics Report for more information.

To configure the upsync log alarm threshold:

1. From the **Platform Setting** section of the navigation pane, select **Platform Configuration Setting**.

   The Platform Configuration page opens.

2. Click **Modify**.
3. In the **Upsync Log Alarm Threshold** field, enter the percent of upsync log capacity at which the upsync log alarm will be triggered.

   Valid values are 50 through 90.

4. Click **Save**.

3

Managing Multimedia Policy Engine Devices

This chapter describes how to use the Configuration Management Platform (CMP) system to configure and manage Multimedia Policy Engine (MPE) devices in a network.

> **Note:** The MPE device is also called the Policy Server.

About Managing an MPE Device

To manage an MPE device:

1. You must first create its profile (see Managing Policy Server Profiles).
2. After creating the profile, you must configure the device:
   - Manually (see Configuring a Policy Server Profile)
   - By applying a configuration or virtual template (see Managing Configuration and Virtual Templates)

After you have configured a policy server profile for an MPE device in your Policy Management network, you can associate network elements with it (see Managing Network Elements section).

Managing Policy Server Profiles

A policy server profile contains the configuration information for an MPE device (which can be a single server, a two-server cluster, or a three-server cluster). The CMP system stores policy server profiles in a configuration database. After you create and configure policy profiles, you deploy them to MPE devices across the network.

The following sections describe how to manage policy server profiles:

- Creating a Policy Server Profile
- Deleting a Policy Server Profile

For information on deploying defined policies to an MPE device, see *Policy Wizard Reference*.

Creating a Policy Server Profile

> **Note:** You must establish the Policy Management network topology before you can create policy server profiles.

To create a policy server profile:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of server groups; the initial group is **ALL**.

2. From the content tree, select the **ALL** group.

   The Policy Server Administration page opens in the work area.

3. Click **Create Policy Server**.

   The New Policy Server page opens.

4. (Required) Select the **Associated Cluster** with which to associate this MPE device.

   See Configuring the Policy Management Topology for details on adding clusters to the topology.

5. (Required) Enter the **Name** for this device.

   The default is the associated cluster name. A name is subject to the following rules:

   - The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_).
   - The name is case insensitive (uppercase and lowercase are treated as the same).
   - The maximum length is 255 characters.

6. (Optional) Enter the **Description / Location**.

   Information that defines the function or location of this MPE device.

7. (Optional) Select to enable **Secure Connection**.

   This setting determines whether or not to use the HTTPS protocol for secure communication between Policy Management devices. If selected, devices communicate over port 8443. See the *Platform Configuration User's Guide* for information on creating and exchanging security certificates within and between Policy Management clusters to support secure communication.

8. Select the **Type** from the list.

   This setting defines the policy server type:

   - **Oracle** (default)

     The policy server is an MPE device and can be fully managed by the CMP system.

   - **Unmanaged**

     The policy server is not an MPE device and therefore cannot be actively managed by the CMP system. This selection is useful when an MPE device is routing traffic to a third-party policy server.

9. (Optional) Select **Associate Templates**.

   See Managing Configuration and Virtual Templates for information about configuration and virtual templates.

10. Click **Save**.

The server profile appears in the list of policy servers. You have defined the policy server profile.

Proceed with configuring the policy server. See Configuring a Policy Server Profile.

About Policy Server Administration

After creating and saving a policy server profile, you can proceed with configuring the device using the Policy Server Administration page.

The Policy Server Administration page contains the following tabs:

- **System**

  Defines the system information associated with this policy server, including the name, host name or IP address in IPv4 or IPv6 format, information about the policy server, and whether or not the policy server uses a secure connection to any management system (such as the CMP system). See Creating a Policy Server Profile for details.

- **Reports** (read only)

  Displays various statistics and counters related to the physical hardware of the cluster, policy execution, and network protocol operation. Reports cannot be modified. See #unique_148 for details.

- **Logs**

  Displays the Trace Log and Syslog configurations. See #unique_149 for details.

- **Policy Server**

  Lets you associate applications and network elements with an MPE device and configure protocol information. See Configuring MPE Protocol Options for details.

- **EM**

  Lets you view and configure event messages. See #unique_150 for details.

- **Diameter Routing**

  Lets you configure the Diameter peer and route tables. See #unique_151 for details.

- **Routing**

  Lets you organize large networks of policy servers into a hierarchical configuration, applicable for network designs with either centralized application architectures, or distributed application architectures. See #unique_151 for details.

- **Policies**

  Lets you manage policies that are deployed on the policy server. Refer to *Policy Wizard Reference* for details.

- **Data Sources**

  Lets you configure interfaces to DHCP (Dynamic Host Configuration Protocol) systems. See #unique_152 for details.

- **Debug**

Lets you troubleshoot and modify component-specific level logging. The files logback-tomcat.log, logback-rc.xml, and logback-bod.xml can be modified using the CMP interface for each target Policy Management device. See Configuring Debug Logs.

Configuring a Policy Server Profile

To configure a policy server profile:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of server groups; the initial group is **ALL**.

2. From the content tree, select the policy server.

   The Policy Server Administration page opens in the work area.

3. Select the tab that contains the information you want to configure or modify and click **Modify**.

4. Edit the information:

   - **Logs**

     See #unique_149 for details.

   - **Policy Server**

     See Configuring MPE Protocol Options for details.

     > **Note:** You must configure attribute information on the **Policy Server** tab for most protocols to function correctly.

   - **EM**

     See #unique_153 for details.

   - **Routing**

     See #unique_151 for details.

   - **Policies**

     Refer to *Policy Wizard Reference* for details.

   - **Data Sources**

     See #unique_154 for details.

   - **Debug**

     See Configuring Debug Logs for details.

5. Click **Save**.

After you have configured a policy server profile for an MPE device in your Policy Management network, you can associate network elements with it (see #unique_155).

Modifying a Policy Server Profile

To modify a policy server profile:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of server groups; the initial group is **ALL**.

2. From the content tree, select the policy server.

   The Policy Server Administration page opens in the work area.

3. Select the tab that contains the information you want to configure or modify and click **Modify**.

4. Edit the information:

   - **System**

     See Creating a Policy Server Profile for details.

   - **Logs**

     See #unique_149 for details.

   - **Policy Server**

     See Configuring MPE Protocol Options for details.

     > **Note:** You must configure attribute information on the **Policy Server** tab for most protocols to function correctly.

   - **EM**

     See #unique_153 for details.

   - **Routing**

     See #unique_151 for details.

   - **Policies**

     Refer to *Policy Wizard Reference* for details.

   - **Data Sources**

     See #unique_154 for details.

   - **Debug**

     See Configuring Debug Logs for details.

5. Click **Save**.

Deleting a Policy Server Profile

Deleting a policy server profile for an MPE device from the ALL group also deletes it from any associated group.

> **Note:** You cannot delete a policy server profile if the profile is configured in an MPE pool. Refer to *Policy Front End Wireless User's Guide* for more information.

To delete a policy server profile:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of server groups; the initial group is **ALL**.

2. From the content tree, select the **ALL** group.

   The Policy Server Administration page opens in the work area.

3. Use one of the following methods to select the MPE device profile to delete:

   - From the work area, click 🗑 (trash can) located next to the MPE device profile you want to delete.
   - From the policy server group tree:

     a. Select the MPE device.

        The Policy Server Administration page opens.

     b. Select the **System** tab and click **Delete**.

   A confirmation message appears.

4. Click **OK** to delete the MPE device profile.

   The profile is removed from the list.

The policy server profile is deleted.

Managing Configuration and Virtual Templates

Configuration and Virtual Templates provide a more efficient means of normalizing common configurations between multiple MPE instances. Any given device can be associated with no template, one, or many templates. In addition, users can add, remove, clone, and prioritize templates.

Virtual Templates are similar to symbolic links in Linux. Virtual Templates are particularly efficient when users want to replace a template that has been associated to multiple MPE or MRA devices with another template.

About Configuring Templates

You can create both virtual and configuration templates for an MPE device in the **Policy Server** section of the navigation pane of the CMP interface.

After templates are created and associated with a device, the templates can be viewed and managed from the **System** tab of the MPE device.

Creating a Configuration Template

Note: This procedure applies to both MPE devices.

Note: You must create a configuration template before creating a virtual template because a virtual template references, and is dependent on, a configuration template.

Use this procedure if you want to make a template that you will use many times.

To create a configuration template:

1. From the **MRA** or **Policy Server** section of the navigation pane, select **Configuration Template**.

   The content tree displays a list of **All Templates** including **Virtual Templates** and **Configuration Templates.**

2. From the content tree, select **Configuration Templates**.

   The Configuration Template Administration page opens.

3. Click **Create Template**.

   The New Configuration Template page opens.

4. Enter the **Name** of the template.

   Note: This is an alphanumeric field that is limited to 255 characters. Single quotes, double quotes, spaces, commas, and backslash characters are not valid.

5. (Optional) To use an existing template as a base for the new template, select an existing template from the **Copy From** list.

6. (Optional) Enter a **Description / Location**.

   The text box is limited to 255 characters.

7. Click **Save**.

The new template appears in the list in the content pane.

After creating the template, proceed with configuring the template.

Changing the Template Priority

You would reorder templates in a list to prioritize templates according to configuration values applied to a given MPE instance. For example, different configurations will provide different prioritizations depending on the order (the lower the number the higher the prioritization) as it is listed in the Associated Templates section of the Modify System Settings screen.

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of **All Policy Servers** or **MRA** devices.

2. From the content tree, select the device.

   The Administration page opens with the device configuration.

3. Select the **System** tab.

The device's system configuration settings display on the page.

4. Click **Modify**.

   The administration page becomes enabled for editing.

5. In the **Associated Templates** section, edit the **Priority** value to change the number to a higher or lower value.

6. Click **Update Order**.

   The priority order of the **Associated Templates** is changed.

Creating a Virtual Template

Because an MPE device can exist independently of one another, you can create both virtual and configuration templates in two locations in the CMP interface. Depending on your needs, the CMP interface enables you to create templates in the the **Policy Server** section of the navigation pane.

Because virtual templates are based on configuration templates, modifying a configuration template associated with a virtual template automatically modifies the virtual template. After the template is created, the template has the functionality that is specific to that instance (that is, MPE). After templates are created and associated, the templates can be viewed and managed from the **System** tab of the MPE device.

Note: You must create a configuration template before creating a virtual template because a virtual template references, and is dependent on, a configuration template.

Note: This procedure applies to both MPE devices.

Use this procedure if you have virtual template capability.

To create a virtual template:

1. From the **Policy Server** section of the navigation pane, select **Configuration Template**.

   The content tree displays a list of **All Templates** including **Virtual Templates** and **Configuration Templates.**

2. From the content tree, select **Virtual Templates**.

   The Virtual Template Administration page opens.

3. Click **Create Virtual Template**.

   The New Virtual Template page opens.

4. Enter the **Name** of the template.

Note: This is an alphanumeric field that is limited to 255 characters. Single quotes, double quotes, spaces, commas, and backslash characters are not valid.

5. Select a template from the **Associated Configuration Template** list.
6. (Optional) Enter a **Description**.
7. Click **Save**.

The settings are saved for the template, and applied to all associated MPE devices.

About Overlaps

Overlaps occur when both a template and an MPE server are assigned an identical value for the same attribute or field. For example, the index of a user name is true in template A, and the index of a user name is also true in an MPE server. The result is that when the template and MPE server are associated, the index of the user name becomes an overlapped field. When an overlap occurs, a prompt appears stating, The server configuration has overlaps with the associated template(s). You can take one of two actions:

- Remove the overlaps and use the settings from the template.
- Keep the overlaps and use the settings from the server.

Associating Templates with a Device

Note: This procedure applies to both MPE devices.

You would use this procedure if you had a number of devices that required the same instance.

To associate templates with an MPE device:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of server groups; the initial group is **ALL**.
2. From the content tree, select the device.

   The administration page opens in the work area.
3. Select the **System** tab.
4. Click **Modify**.

   The administration page becomes editable.
5. In the Associated Templates section, click **Add**.

   The Add Associated Templates dialog appears.
6. Select one or more templates from the list and click **Add**.

   The Associated Templates list updates to include the selected templates.
7. To order the **Priority** of the associated templates, change the values for each listed template.

> **Note:** Lower-numbered templates have higher priority than higher-numbered templates. This means that settings configured with a lower-value priority template can override the settings of a higher-value priority template.

8. Click **Save**.

The specified templates' configurations are applied to the specified device.

Configuring MPE Protocol Options

To configure protocol options on an MPE device:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of server groups; the initial group is **ALL**.

2. From the content tree, select the MPE device.

   The Policy Server Administration page opens.

3. Select the **Policy Server** tab.

   The current configuration options are displayed.

4. Click **Modify** and define options as necessary.

   Selecting or choosing **undefined** signifies that the value comes from a configuration profile. If there is not a configuration profile, then the default value is used.

   The following sections define the available options. (The options you see vary depending on the mode configuration of your system.)

   - Associations Configuration Options
   - #unique_164
   - General Configuration Options
   - #unique_166
   - #unique_167
   - Diameter AF Default Profiles Configuration Options
   - #unique_169
   - SMTP Configuration Options

5. Click **Save**.

You have defined the protocol options for this MPE device.

Associations Configuration Options

Attribute **Description**

Applications

The application profiles associated with this MPE device. To modify this list, click **Manage**. For more information on application profiles, see *Policy Wizard Reference*.

Network Elements
The network elements associated with this MPE device. To modify this list, click **Manage**. For more information on network elements, see #unique_155.

Network Element Groups
The network element groups associated with this MPE device. To modify this list, select or deselect groups. For more information on network element groups, see #unique_155.

Notification Servers
The notification servers associated with this MPE device. To modify this list, click **Manage**. For more information on notification servers, see Notification Servers.

General Configuration Options

Attribute **Description**

Management Agent
Visible if your network contains Management Agent servers. For more information, see #unique_171.

PCMM Configuration Options

Attribute **Description**

Validate the application
When enabled, all PCMM requests are checked to ensure that there is an application defined that can be associated with the request (typically by matching the application manager ID, or AMID, in the request). If there is no such application, the MPE device rejects the request.

Validate the service class
When enabled, any PCMM requests that refer to a Service Class Name in a traffic profile are checked to ensure that the service class is known to be valid for the destination CMTS.

Validate the gate ID
When enabled, all PCMM requests that refer to an existing gate are checked against the MPE device's database of existing gates. If the request refers to a gate ID that does not exist, then it is rejected without forwarding to the CMTS.

Validate traffic profile envelopes
When enabled, all PCMM requests that include traffic profiles are checked to ensure that the parameters for the Authorized, Reserved, and Committed envelopes are valid, as defined in the PCMM Specification.

Enable MGPI
Enable Multiple Grants Per Interval (MGPI) for all Rx applications. By default, not selected (that is, MGPI is disabled). For more information, see Configuring Protocol Routing .

Note: If MGPI is enabled, flow aggregation begins with the next call that creates or modifies an application flow.

Upstream Flow Limit for Triggering MGPI
The number of upstream service flows above which MGPI is triggered. A value from 1 through 99; the default is 8 flows.

Maximum Number of Grants per Interval
The maximum number of grants per interval allowed on one gate (that is, the maximum number of sub-flows aggregated on one service flow). A value from 2 through 99; the default is 8 grants.

Default Local Time Mode
Select the time used within the session for a user from the list: **System Local Time** to use the local time of the MPE device (default) or **User Local Time** to use the user's local time.

Note: If the time zone was never provided for the user equipment, system local time is applied.

Diameter AF Default Profiles Configuration Options

Note: To select a profile of any of the attributes, you must first create a Diameter profile in the general profile configuration.

Attribute **Description**

Default
Define the bandwidth parameters that are used when a request from an Application Function (AF) does not contain sufficient information for the MPE device to derive QoS parameters. These profiles are defined per media type: The **Default** profile is used when a profile for a media type is not defined.

Audio
The profile for the audio.

Video
The profile for the video.

Data
The profile for data.

Application
The profile for application.

Control
The profile for control.

Note: To select a profile, first create a Diameter profile in the general profile configuration.

Text
The profile for text.

Message
The profile for messages.

Other
The profile for all other media types.

SMTP Configuration Options

Attribute **Description**

SMTP Enabled
Select **true** to enable Simple Mail Transport Protocol (SMTP) messaging (email) to subscribers. SMTP notifications are triggered from policy action and sent through an SMS Relay (SMSR) function to an external mail transfer agent (MTA).

Note: There is no delivery receipt for the SMTP messages sent from the SMSR, only confirmation that it reached the configured MTA.

MTA Host
Enter the FQDN or IP address of the Mail Transfer Agent server, which accepts SMTP messages from the SMSR function.

MTA Port
Enter the port number on which the MTA server is listening for SMTP messages. The default port is 25.

MTA Username
Enter the system ID of the SMSR function. Sending the ID and password values authenticates the SMSR function as a trusted source.

Note: This value must be configured on the MTA.

MTA Password
Enter the password of the SMSR function. Sending the ID and password values authenticates the SMSR function as a trusted source.

Note: This value must be configured on the MTA.

Confirm MTA Password
Re-enter the password for verification.

Note: This is a new configuration setting for the SMTP connection.

Default From Address(es)
Enter the source address for an SMTP message. Enter up to five comma-separated static values, or up to five comma-separated references to custom fields in the subscriber profile. The default is none.

Note: The total number of To, CC, and BCC addresses is limited to five.

SMTP Connections
The number of SMTP connections. Enter a number from 1 through 10.

Note: SMTP connections can be increased to support a higher throughput. Contact My Oracle Support for more information.

Default Reply-To Address(es)
Enter the email address automatically inserted into the To field when a user replies to an email message. For most email messages, the From and Reply-To fields are the same, but this is not necessarily so. If no Default Reply-To is specified here, the From address is used. Optionally, enter a static email address to use for Reply-To.
The default is none.

Default CC Address(es)
Enter the copy address for an SMTP message. Enter up to five comma-separated static values, or up to five comma-separated references to custom fields in the subscriber profile. The default is none.

Note: The total number of To, CC, and BCC addresses is limited to five.

Default BCC Address(es)
Enter the blind copy recipient address for an SMTP message. Enter up to five comma-separated static values, or up to five comma-separated references to custom fields in the subscriber profile. The default is none.

Note: The total number of To, CC, and BCC addresses is limited to five.

Default Signature
Enter the text that should appear as the signature in an SMTP message.
The default is none.

Generic Notification Configuration Options

Attribute	Description

Notification Enabled
If SMPP, XML, or CMPP mode is enabled, select **true** to enable notifications using notification servers. For more information about notification servers, see the *Policy Wizard Reference*.

Connecting to Data Sources

MPE devices establish connections with data sources to retrieve information about subscribers from a database. An MPE device queries a data source using a key attribute that uniquely identifies a subscriber and stores the results in its cache. A data source uses this key attribute (for example, the phone or account number of a subscriber) to index the information contained in the database.

The CMP system supports Dynamic Host Configuration Protocol (DHCP) data sources.

You can use a single data source that holds all subscriber information or you can have the subscriber data distributed across multiple data sources. When using a single data source, you configure the server connection and the query format.

About DHCP Data Sources

In the home network of a subscriber, several CPE devices (for example, routers, set-top devices, and computers) can connect to the local area network (LAN) of the cable modem. When the cable modem powers on, it connects the home network to the core network for the operator (for example, hybrid fiber/coaxial) through a CMTS (cable modem termination system). The CMTS routes packets between the downstream HFC channels and the upstream back office and core network channels.

The network for the operator uses a provisioning system that includes DHCP servers that provide the cable modem with the initial configuration information (for example, IP address) when the cable modem powers on. The modem provides the provisioning system information that can be used to make device and service configuration decisions during the provisioning process.

After the MPE device receives the information from the customer equipment, it uses this data to drive policy decisions for the subscriber.

The Bandwidth on Demand Application Manager (BoD AM) device which interfaces with MPE devices for requesting QoS-based service.

Adding a DHCP Data Source

To add a data source:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of server groups; the initial group is **ALL**.

2. From the content tree, select the policy server.

   The Policy Server Administration page opens in the work area.

3. Select the **Data Sources** tab.

   The page lists the existing data sources including the following information:

   - Administrative state
   - Name
   - Role
   - Type
   - Primary host
   - Secondary host
   - Tertiary host

4. Click **Modify**.

5. Add a data source.

   a. Click **Add**.

   b. Select the data source type from the list.

   c. Configure the values. For DHCP data sources, see Configuring a DHCP Data Source.

   d. Click **Save**.

6. (Optional) Add, modify, delete, or reorder data sources.

   - Cloning an entry in the table

     a. Select an entry in the table.

     b. Click **Clone**. The Clone window opens with the information for the entry.

     c. Make changes as required.

     d. Click **Save**. The entry is added to the table.

   - Editing an entry in the table

     a. Select the entry in the table.

     b. Click **Edit**. The Edit Response window opens, displaying the information for the entry.

     c. Make changes as required.

     d. Click **Save**. The entry is updated in the table.

   - Deleting a value from the table

     a. Select the entry in the table.

     b. Click ✕ **Delete**. A confirmation message displays.

     c. Click **Delete** to remove the entry. The entry is removed from the table.

   - Ordering the list.

     If you define multiple entries, they are searched in the order displayed in this list. To change the order:

     a. Select an entry.

     b. Click ↑ **Up** or ↓**Down**. The search order is changed.

7. Click **Save**.

8. The following general settings are available:

   - **Merge Search Results** — If you define multiple data sources and a search returns results from more than one source, the results are displayed in source order. To display one sorted list instead, select this option.

- **Subscription Enabled Via Policy Only** — For detailed information, see the SPR documentation.

9. Click **Save**.

Configuring a DHCP Data Source

For DHCP, you can configure connections one or two DHCP servers.

To configure a DHCP data source:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the policy server.

   The Policy Server Administration page opens.

3. Select the **Data Sources** tab.

   The current data sources are displayed.

4. Click **Modify**.

5. Click **Add** and select **DHCP**.

   The Add Data Source window opens.

6. **Admin State** — Select to enable this data source.

   Selected by default.

7. **Primary** — FQDN or IP address in IPv4 or IPv6 format of primary DHCP server.

8. **Secondary** — FQDN or IP address in IPv4 or IPv6 format of secondary DHCP server.

9. **Timeout (ms)** — Length of time to wait before a DCHP request times out.

   The default timeout is 1000 ms (one second).

10. **Fail on Unassigned Lease** — Action to take if the DHCP server returns an unassigned lease.

    By default, the action fails.

11. **4388 Compliant Mode** — Compliant with *RFC4388 - Dynamic Host Configuration Protocol (DHCP) Leasequery*.

12. Click **Save**.

The DHCP data source is configured.

Working with Policy Server Groups

For organizational purposes, you can aggregate MPE devices in your network into groups. For example, you can use groups to define authorization scopes. The following subsections describe how to manage policy server groups.

Creating a Policy Server Group

To create a policy server group:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the **ALL** group.

   The Policy Server Administration page opens in the work area.

3. Click **Create Group**.

   The Create Group page opens.

4. Enter the name of the new policy server group.

   The name cannot contain quotation marks (") or commas (,).

5. Click **Save**.

You have created a policy server group and the new group appears in the content tree.

Adding a Policy Server to a Policy Server Group

To add a policy server to a policy server group:

1. From the Policy Server section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the policy server group.

   The Policy Server Administration page opens in the work area displaying the contents of the selected policy server group.

3. Click **Add Policy Server**.

   The Add Policy Server page opens, displaying the policy servers not already part of the group.

4. Click the policy server you want to add; press Ctrl or Shift-Ctrl to select multiple policy servers.

5. Click **Save**.

The policy server is added to the selected group.

Creating a Policy Server Sub-group

You can create sub-groups to further organize your policy server network. To add a policy server sub-group to an existing policy server group:

1. From the Policy Server section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the policy server group.

   The Policy Server Administration page opens in the work area, displaying the contents of the selected policy server group.

3. Click **Create Sub-Group**.

   The Create Group page opens.

4. Enter the name of the new sub-group.

   The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_).

5. Click **Save**.

The sub-group is added to the selected group.

Removing a Policy Server Profile from a Policy Server Group

Removing a policy server profile from a policy server group or sub-group does not delete the profile. To delete a policy server profile, see #unique_183.

To remove a policy server profile from a policy server group or sub-group:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the policy server group or sub-group.

   The Policy Server Administration page opens in the work area, displaying the contents of the selected policy server group or sub-group.

3. Remove the policy server profile using one of the following methods:

Note: The profile is removed immediately, without a confirmation message.

- Click the **Remove** (🗑) icon located next to the policy server you want to remove.
- From the content tree, select the policy server. The Policy Server Administration page opens. Select the **System** tab and click **Delete**.

The policy server is removed from the group or sub-group.

Deleting a Policy Server Group

Deleting a policy server group also deletes any associated sub-groups. However, any policy server profiles associated with the deleted group or sub-groups remain in the ALL group. You cannot delete the ALL group.

To delete a policy server group or subgroup:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the policy server group or sub-group.

   The Policy Server Administration page opens in the work area, displaying the contents of the selected policy server group or sub-group.

3. On the Policy Server Administration page, click **Delete**.

   A confirmation message displays.

4. Click **OK** to delete the group.

The policy group is deleted.

About Reapplying a Configuration

You can reapply the configuration to an individual Policy Management device (server), or to all Policy Management devices in a group. When you reapply the configuration, the CMP system completely reconfigures the servers with topology information, ensuring that the configuration matches the data in the CMP system. This action is not needed during normal operation but is useful in the following situations:

- When the servers of a cluster are replaced, the new servers come up initially with default values. Reapplying the configuration lets you redeploy the entire configuration rather than reconfiguring the server field by field. You should also use the Rediscover Cluster operation to clear the failed status in the Cluster Information Report. See Rediscovering a Cluster for more information.
- After upgrading the software on a server, it is recommended that you reapply the configuration from the CMP system to ensure that the upgraded server and the CMP system are synchronized.
- The server configuration may go out of synchronization with the CMP system (for example, when a break in the network causes communication to fail between the CMP system and the server). If such a condition occurs, the CMP system displays the server status on the **System** with a message that there is a configuration

mismatch. You can click the notice to display a report comparing the server configuration with the CMP database information. Reapplying the configuration brings the server back into synchronization with the CMP database.

CMP provides the following methods for reapplying a configuration:

- Reapplying the Configuration to a Single Device
- Reapplying the Configuration to a Group of Devices

Reapplying the Configuration to a Single Device

To reapply the configuration associated with an MPE or BoD device:

1. From the appropriate section of the navigation pane (for example, **Policy Server** or **BoD**), select **Configuration**.

   The content tree displays a list of Policy Management device groups; the initial group is **ALL**.

2. From the content tree, select the device.

   The Policy Server Administration page opens to the **System** tab, displaying information for that device.

3. Click **Reapply Configuration**.

The device is synchronized with the CMP system.

Reapplying the Configuration to a Group of Devices

To reapply the configuration associated with a group of MPE or BoD devices:

1. From the appropriate section of the navigation pane (for example, **Policy Server** or **BoD**), select **Configuration**.

   The content tree displays a list of Policy Management device groups; the initial group is **ALL**.

2. From the content tree, select the group.

   The appropriate Administration page opens in the work area.

3. From the **Operations** menu, select **Reapply Config**.

   The Bulk Reapply Config dialog displays stating the number of agents affected.

4. Specify the delay time (in seconds) for applying the operation to each server in the group.

   The number of seconds is 0 to 60. 0 is the default.

5. Click **Reapply Config** to reapply the configuration.

   An in-progress message appears. After the action completes, a message stating the reapply was successful with a list of the affected devices appears.

All of the servers in a group are synchronized with the CMP server.

Rediscovering a Cluster

After reapplying a configuration or deleting a failed server, use the Rediscover Cluster operation to refresh the Cluster Information Report. The Rediscover Cluster operation rediscovers the cluster, deleting any failed servers that have been removed from service or refreshing the status of any failed servers after reapplying the configuration.

To rediscover a cluster:

1. From the appropriate section of the navigation pane (for example, **Policy Server** or **BoD**), select **Configuration**.

   The content tree displays a list of Policy Management device groups; the initial group is **ALL**.

2. From the content tree, select the server or cluster.

   The corresponding administration page opens in the work area.

3. Click the **Reports** tab.

   The **Reports** tab opens.

4. Click **Rediscover Cluster**.

   The Cluster Information Report is updated.

The cluster is rediscovered.

Resetting Counters

The **Reset Counters** option is included in the **Operations** menu when the **Stats Reset Configuration** option is set to **Interval**. The **Reset All Counters** option is included in the **Operations** menu when the **Stats Reset Configuration option** is set to **Manual**. See Configuring Stats Settings for more information.

To reset the counters associated with a group of MPE servers:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the group that contains the servers of interest.

   The Policy Server Administration page opens in the work area.

3. From the **Operations** menu, select **Reset Counters** or **Reset All Counters**.

   The Bulk Reset All Counters or Bulk Reset Counters dialog displays showing the number of servers affected.

4. Specify the delay time for applying the operation to each server. The number of seconds is 0 to 60. 0 is the default.

The counters are reset.

Checking the Status of an MPE Server

The CMP lets you view the status of MPE servers, either collectively (all servers within the topology) or individually.

Group View
Select **ALL** from the policy server content tree to view all the defined MPE servers, or select a specific policy server group or sub-group to view just the servers associated with that group. The display in the work area includes a status column that indicates the following states:

- **On-Line**

  All servers in the cluster are operational.

- **Degraded**

  One server is not functioning properly (for example, an interface is down) or has failed, but the cluster continues to function with the standby or spare server. This state sets alarm ID 70005 with severity Major.

> **Note:** If a cluster status is **Degraded**, but the server details do not show any failures or disconnections, then the cluster is performing a database synchronization operation. Until the synchronization process has completed, the server cannot perform as the active server.

- **Failed**

  All servers in the cluster are no longer functioning.

- **Off-line**

  Communication to the cluster has been lost.

- **Config Mismatch**

  The MPE device configuration does not match the CMP database.

Policy Server Profile View
Select a server from the content tree, then click the **System** tab to view the current operating status of the device (**On-line** or **Off-line**) and profile configuration.

Policy Server Group View
Select a group from the content tree to view the current operating status of the servers in the group.
Figure 3-1 shows an example of a Group View in which one of the servers is degraded.

Figure 3-1 Group View

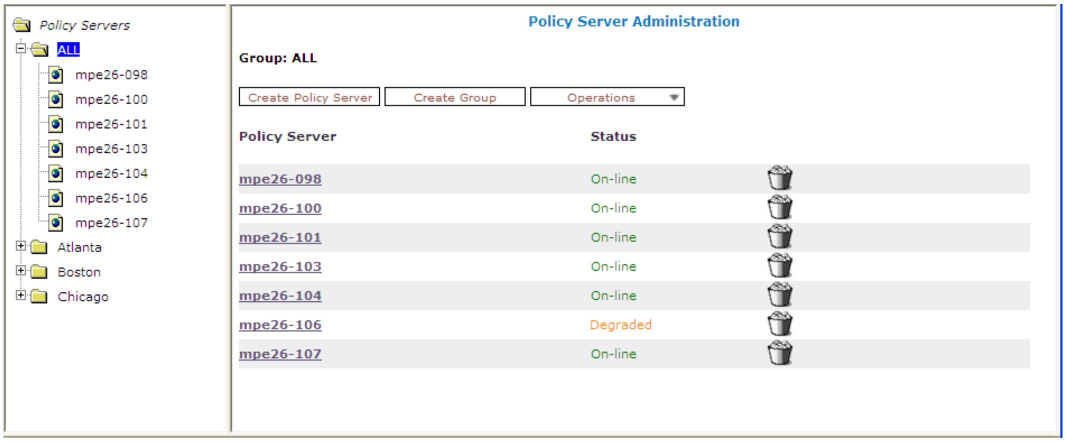

Group: ALL

Policy Server	Status	
mpe26-098	On-line	🗑
mpe26-100	On-line	🗑
mpe26-101	On-line	🗑
mpe26-103	On-line	🗑
mpe26-104	On-line	🗑
mpe26-106	Degraded	🗑
mpe26-107	On-line	🗑

Trash can icon
Click 🗑 (trash can icon) to delete an MPE server.

Policy Server Reports

The Reports tab lets you view a hierarchical set of reports that you can use to monitor both the status and the activity of a specific policy server.

Each report page provides the following information:

- **Mode** — Shows whether data collection is currently **Active** or **Paused**.
- **Buttons** — The buttons let you navigate between reports, or control the information displayed within the report. The following list describes the buttons; which buttons are available depend on your configuration and differ from one report page to the next:
  - **Reset Counters** — Resets display counters, such as those that indicate blade failures.
  - **Rediscover Cluster** — Rediscovers the cluster, deleting any failed servers that have been removed from service.
  - **Pause/Resume** — Stops or restarts automatic refreshing of displayed information. The refresh period is 10 seconds.

The report also displays various statistics and counters related to the following:

- **Cluster Information** — Information about the cluster.
- **Blades** — Information about the individual physical components in the cluster.
- **Policy Statistics** — Information about the execution of policy rules.
- **Protocol Statistics** — Information about the active network protocols.
- **Latency Statistics** — Information about protocol latency.
- **Error Statistics** — Information about any errors, arranged by protocol.
- **Data Source Statistics** — Information about activity with configurable data sources.

- **Database Statistics** — Information about LDAP activity.
- **KPI Interval Statistics** — Information about the configured reporting interval for key performance indicator (KPI) statistics.

> **Note:** The Cluster Information Report is also available as a selection on the navigation pane.

Cluster Information Report

The fields that are displayed in the Cluster Information Report section include the following:

- **Cluster Status** — The status of the cluster:
  - **On-line**: If one server, it is active; if two servers, one is active and one is standby; if three servers, one is active, one is standby, one is spare.
  - **Degraded**: One server is active, but at least one other server is not available.
  - **Out-Of-Service**: No server is active.
  - **No Data**: The CMP system cannot reach the server.
- **Site Preference** — The preference of the cluster (Normal or Reversed). Default status is Normal.

Also within the Cluster Information Report is a listing of all the servers (blades) contained within the cluster. A symbol ( ) indicates which server currently has the external connection (the active server). The report also lists the following server-specific information:

- **Overall** — Displays the current topology state (Active, Standby, or Forced-Standby), number of server (blade) failures, and total uptime (time providing active or standby policy or GUI service). For the definitions of these states, see Server Status.
- **Utilization** — Displays the percentage utilization of disk (of the `/var/camiant` file system), CPU, and memory.

The **Actions** buttons let you restart the Policy Management software on the server or restart the server itself.

Policy Statistics

The Policy Statistics section summarizes policy rule activity within the MPE device. This is presented as a table of statistics for each policy rule that is configured for the MPE device.

The following statistics are included:

Name
Name of the policy being polled.

Evaluated
Number of times the conditions in the policy were evaluated.

Executed
Number of times policy actions were executed. This implies that the conditions in the policy evaluated to be true.

Ignored
Number of times the policy was ignored. This can happen because the policy conditions refer to data which was not applicable given the context in which it was evaluated.

To see statistics per policy, click **(details...)**. All existing policies are displayed in a statistics table, with Evaluated, Executed, and Ignored counter values listed for each.

To see details for a specific policy with the distribution of execution time, click the policy name. In addition to Evaluated, Executed, and Ignored, the following details are displayed:

Total Execution Time (ms)
The summary of all execution durations, where execution duration is measured starting at the beginning of the policy conditions evaluation until the execution completion.

Maximum Execution time (ms)
The longest execution duration of the policy.

Average Execution time (ms)
The average of all execution durations of the policy.

Processing Time Statistics
The number of policies processed per time range, in milliseconds. Ranges include:

- 0 to 20
- 20 to 40
- 40 to 60
- 60 to 80
- 80 to 100
- 100 to 150
- 150 to 200
- 200 to 250
- greater than 250

Session Cleanup Statistics

The Session Cleanup Statistics section summarizes the activity of removing stale or stranded PCMM sessions within the MPE device.

For information on configuring session cleanup, see Configuring Advanced Device Settings.

The following statistics are included:

- **Ready for Cleanup**

Number of sessions that are stale (created at least 24 hours ago).

- **Removed on unknown session id**

  Number of sessions removed because the session ID is no longer valid.

- **Reauthorized**

  Number of sessions reauthorized.

- **Reauthorization Timeout**

  Number of sessions for which the reauthorization request timed out.

- **Removed for Expiration**

  Number of sessions removed.

Protocol Statistics

The Protocol Statistics section summarizes the protocol activity within the MPE device. This information is presented as a table of summary statistics for each protocol. Some protocols are broken down into sub-entries to distinguish between the different types of protocol activity.

The summary protocol statistics are the following:

Connections
If the protocol is connection oriented, this value represents the current number of established connections using each protocol.

Total client messages in / out
The total number of incoming and outgoing messages received and sent using each protocol.

Total messages timeout
The total number of incoming and outgoing messages that timed out using each protocol.

Figure 3-2 shows a sample.

Figure 3-2 Sample Protocol Statistics

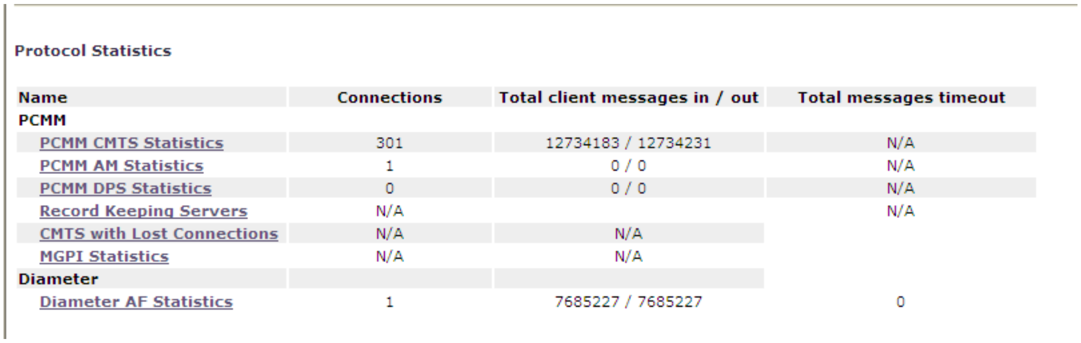

Protocol Statistics

Name	Connections	Total client messages in / out	Total messages timeout
PCMM			
PCMM CMTS Statistics	301	12734183 / 12734231	N/A
PCMM AM Statistics	1	0 / 0	N/A
PCMM DPS Statistics	0	0 / 0	N/A
Record Keeping Servers	N/A		N/A
CMTS with Lost Connections	N/A	N/A	
MGPI Statistics	N/A	N/A	
Diameter			
Diameter AF Statistics	1	7685227 / 7685227	0

You can click the name of each entry in the Protocol Statistics table to display a detailed report page. For most protocols, this report page displays a set of counters that break down the protocol activity by message type, message response type, errors, and so on.

Many of the protocol report pages also include a table that summarizes the activity for each client or server with which the MPE device is communicating through that protocol. These tables let you select a specific entry to further examine detailed protocol statistics that are specific to that client or server.

Since many of these statistics contain detailed protocol-specific summaries of information, the specific definitions of the information that is displayed are not included here. For more specific information, see the appropriate technical specification that describes the protocol in which you are interested (see #unique_195).

Note:

1. Statistical information is returned from the MPE server as a series of running peg counts. To arrive at interval rate information, such as session success and failure counts, two intervals are needed to perform the difference calculation. Also, statistical information, such as session activation counts, is kept in memory and is therefore not persisted across the cluster. After a failover, non-persistent metrics must be repopulated based on a sampling from the newly active primary server. Therefore, when an MPE server is brought online, or after a failover, one or more sample periods will display no statistical information.

2. Historical network element statistical data is inaccurate if configuration values (such as capacity) were changed in the interim. If the network element was renamed in the interim, no historical data is returned.

Latency Statistics

The Latency Statistics section summarizes latency information, for Diameter and PCMM protocols, within the MPE device. This is presented as a table of statistics for each configured protocol. Each protocol lists the number of connections.

To see details for a specific protocol, click the protocol name. Statistics are displayed for the maximum and average transaction time for messages sent and received, as well as the distribution of execution times.

You can control the information displayed within the detailed report using the following buttons:

Show Absolute/Show Deltas
Switches between absolute mode (statistics between last reset) and delta mode (statistics since last display).

Pause/Resume
Stops or restarts automatic refreshing of displayed information. The refresh period is ten seconds.

Cancel
Returns to the previous page.

Error Statistics

The Error Statistics section summarizes any protocol-related errors reported by the MPE device. This is presented as a table of overall statistics for each protocol that is configured for the MPE device. Figure 3-3 shows a sample.

Figure 3-3 Sample Error Statistics

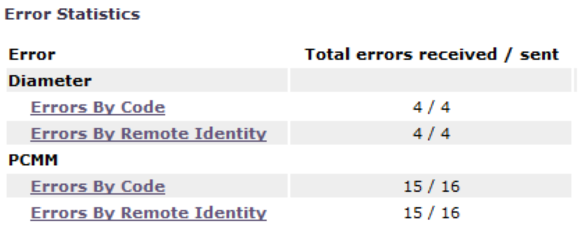

Error Statistics

Error	Total errors received / sent
Diameter	
Errors By Code	4 / 4
Errors By Remote Identity	4 / 4
PCMM	
Errors By Code	15 / 16
Errors By Remote Identity	15 / 16

The following summary statistics are displayed:

Error
List of protocols configured on this MPE device.

Total errors received/sent
Total number of errors received or sent in this protocol.

You can click the name of each entry in the Error Statistics table to display a detailed report page. For most protocols, this report page displays a set of counters that break down the errors by error code and the remote identity of each client or server with which the MPE device is communicating through that protocol.

Data Source Statistics

The Data Source Statistics section summarizes the data source activity within the MPE device. Information is available for each data source. You can click the name of each entry in the Data Source Statistics table to display a detailed report page.

Dynamic Host Configuration Protocol Statistics

For a Dynamic Host Configuration Protocol (DHCP) data source, the DHCP Data Source Statistics page displays the following statistics:

- **Number of successful searches**
- **Number of unsuccessful searches**
- **Number of searches that failed because of errors**
- **Number of search errors that triggered retry**
- **Max Time spent on successful search (ms)**
- **Max Time spent on unsuccessful search (ms)**
- **Average time spent on successful searches (ms)**
- **Average time spent on unsuccessful searches (ms)**

Database Statistics

The Database Statistics section summarizes the read/write activity for the MPE device database. Click **Database Status Statistics** to display the last reset time (that is, the last time that you clicked **Reset All Counters**), the last collection time, and cumulative read/write activity. Data is collected every 10 seconds.

KPI Interval Statistics

The KPI Interval Statistics section summarizes the maximum key performance indicator (KPI) values recorded by the Policy Management cluster during the previous recording interval. Intervals are recorded on the quarter hour.

The following interval statistics are displayed:

Interval StartTime
Timestamp of when the current interval started.

Configured Length (Seconds)
Configured interval length. The value of 900 seconds (15 minutes) is fixed.

Actual Length (Seconds)
Actual interval length. When data is collected over a full interval, this value matches the Configured Length value.

Is Complete
Displays 0 or 1, where 1 indicates that data was collected for a full interval.

Interval MaxSessionCount
The highest value of the counter MaxSessionCount during the previous interval.

You can control the information displayed within the detailed report using the following buttons:

Pause/Resume
Stops or restarts automatic refreshing of displayed information.

Cancel
Returns to the previous page.

> **Note:** If a cluster has just started up and no data is available, the Interval StartTime is displayed as Undefined and the maximum values are displayed as 0. If a cluster has started up and a recording interval has completed but it is less than 15 minutes, the value of Actual Length will not match Configured Length, and the maximum values are displayed as 0.

Mapping Reports Displays to KPIs

The Reports page displays a variety of statistics and measurements for configured protocols. The following tables map these statistics to the statistics returned from OSSI XML queries.

For more information on OSSI XML statistics, see the *OSSI XML Interface Definitions Reference Guide*.

- Table 3-1 shows information for these protocols:
  - PCMM CMTS (Cable Modem Termination System)
  - PCMM AM (Application Manager)
  - PCMM DPS

- Table 3-2 shows information for Record Keeping Servers (RKSs).
- Table 3-3 shows information for individual CMTS systems with lost connections.
- Table 3-4 shows information for the MGPI protocol.
- Table 3-5 shows information for the Diameter AF protocol.
- Table 3-6 shows information for these statistics:
  - Diameter AF
  - PCMM AM
  - PCMM CMTS
  - PCMM DPS
- Table 3-7 shows information for these statistics:
  - Diameter
  - PCMM
- Table 3-8 shows information for these statistics:
  - Diameter
  - PCMM
- Table 3-9 shows information for the KPI collection interval.
- Table 3-10 shows information for policy execution.

Table 3-1 PCMM (PacketCable MultiMedia) Protocol Statistics

Reports Display Name	OSSI XML Name
Connections	Conn Count
Total messages in / out	Msg In Count\Msg Out Count
Gate set messages	
Gate set ack / error messages processed	
Gate info messages	
Gate info ack / error messages processed	
Gate delete ack / error messages processed	
Gate report messages	
Messages dropped	
Currently active gates	
Highest number of active gates seen so far	
Last stats reset time	

Table 3-2 Record Keeping Servers Protocol Statistics

Reports Display Name	OSSI XML Name
Connections	Conn Count
Total messages in / out	Msg In Count\Msg Out Count
Event messages attempted	
Undeliverable event messages	
Policy request messages sent	
Policy update messages sent	
Policy delete messages sent	
Policy change messages sent	
Record Keeping Servers Stats (in Record Keeping Servers window)	
IP Address : Port	
Event messages attempted	
Ack messages received	
Undeliverable event messages	
Policy request messages sent	
Policy update messages sent	
Policy delete messages sent	
Time change messages sent	
Messages sent to primary	
Ack messages received from the primary	
Messages sent to secondary	
Ack messages received from the secondary	

Table 3-3 CMTS with Lost Connections Statistics

Reports Display Name	OSSI XML Name
CMTS Name	
CMTS IP Address	
Last Connection Time	
Last Disconnection Time	

Table 3-4 MGPI Statistics

Reports Display Name	OSSI XML Name
Total flows	

Table 3-4 (Cont.) MGPI Statistics

Reports Display Name	OSSI XML Name
Actual gates	
Multi-flow gates	
Effective gates	

Table 3-5 Diameter AF (Application Function) Statistics

Reports Display Name	OSSI XML Name
Connections	Conn Count
Total messages in / out	Msg In Count\Msg Out Count
AAR messages received / sent	AAR Recv Count\AAR Send Count
AAR Initial messages received / sent	AAR Initial Recv Count\AAR Initial Send Count
AAR Modification messages received / sent	AAR Modification Recv Count\AAR Modification Send Count
AAA success messages received / sent	AAA Recv Success Count\AAA Send Success Count
AAA failure messages received / sent	AAA Recv Failure Count\AAA Send Failure Count
AAR messages timeout	AAR Timeout Count
ASR messages received / sent	ASR Recv Count\ASR Sent Count
ASR messages timeout	ASR Timeout Count
ASA success messages received / sent	ASA Recv Success Count\ASA Send Success Count
ASA failure messages received / sent	ASA Recv Failure Count\ASA Send Failure Count
RAR messages received / sent	RAR Recv Count\RAR Send Count
RAR messages timeout	RAR Timeout Count
RAA success messages received / sent	RAA Recv Success Count\RAA Send Success Count
RAA failure messages received / sent	RAA Recv Failure Count\RAA Send Failure Count
STR messages received / sent	STR Recv Count\STR Send Count
STR messages timeout	STR Timeout Count
STA success messages received / sent	STA Recv Success Count\STA Send Success Count
STA failure messages received / sent	STA Recv Failure Count\STA Send Failure Count

Table 3-5 (Cont.) Diameter AF (Application Function) Statistics

Reports Display Name	OSSI XML Name
Rx-Pcmm Messages Timeout	
Last stats reset time	
Currently active sessions	Active Session Count
Max active sessions	Max Active Session Count
Diameter AF Peer Stats (in Diameter AF Stats window)	
Connect Time	Connect Time
Disconnect Time	Disconnect Time
Connection Type	
IP Address: Port	
Total messages in / out	Msg In Count\Msg Out Count
Total error messages in / out	
AAR messages received / sent	AAR Recv Count\AAR Send Count
AAR Initial messages received / sent	AAR Initial Recv Count\AAR Initial Send Count
AAR Modification messages received / sent	AAR Modification Recv Count\AAR Modification Send Count
AAA success messages received / sent	AAA Recv Success Count\AAA Send Success Count
AAA failure messages received / sent	AAA Recv Failure Count\AAA Send Failure Count
AAR messages timeout	AAR Timeout Count
ASR messages received / sent	ASR Recv Count\ASR Sent Count
ASR messages timeout	ASR Timeout Count
ASA success messages received / sent	ASA Recv Success Count\ASA Send Success Count
ASA failure messages received / sent	ASA Recv Failure Count\ASA Send Failure Count
RAR messages received / sent	RAR Recv Count\RAR Send Count
RAR messages timeout	RAR Timeout Count
RAA success messages received / sent	RAA Recv Success Count\RAA Send Success Count
RAA failure messages received / sent	RAA Recv Failure Count\RAA Send Failure Count
STR messages received / sent	STR Recv Count\STR Send Count

Table 3-5 (Cont.) Diameter AF (Application Function) Statistics

Reports Display Name	OSSI XML Name
STR messages timeout	STR Timeout Count
STA success messages received / sent	STA Recv Success Count\STA Send Success Count
STA failure messages received / sent	STA Recv Failure Count\STA Send Failure Count
Rx-Pcmm Messages Timeout	
Last stats reset time	
Currently active sessions	Active Session Count
Max active sessions	Max Active Session Count

Table 3-6 Latency Statistics

Reports Display Name	OSSI XML Name
Connections	Active Connection Count
Maximum Processing Time received / sent (ms)	Max Trans In Time\ Max Trans Out Time
Average Processing Time received / sent (ms)	Avg Trans In Time\ Avg Trans Out Time
Transactions Processed received / sent [*timeframe*] (ms)	Processing Time [0-20] ms Processing Time [20-40] ms Processing Time [40-60] ms Processing Time [60-80] ms Processing Time [80-100] ms Processing Time [100-120] ms Processing Time [120-140] ms Processing Time [140-160] ms Processing Time [160-180] ms Processing Time [180-200] ms Processing Time [>200] ms

Table 3-7 Protocol Error Statistics

Reports Display Name	OSSI XML Name
Total errors received	In Error Count
Total errors sent	Out Error Count
Last time for total error received	Last Error In Time
Last time for total error sent	Last Error Out Time
Last stats reset time	

Table 3-7 (Cont.) Protocol Error Statistics

Reports Display Name	OSSI XML Name
Diameter Protocol Errors on each error codes	(see specific errors listed in GUI)

Table 3-8 Connection Error Statistics

Reports Display Name	OSSI XML Name
Total errors received	In Error Count
Total errors sent	Out Error Count
Last time for total error received	Last Error In Time
Last time for total error sent	Last Error Out Time
Last stats reset time	
Protocol Errors on each error codes	(see specific errors listed in GUI)

Table 3-9 KPI Interval Statistics

Reports Display Name	OSSI XML Name
Interval StartTime	Interval Start Time
Configured Length (Seconds)	Configured Length (Seconds)
Actual Length (Seconds)	Actual Length (Seconds)
Is Complete	Is Complete
Interval MaxSessionCount	Interval Max Session Count
Interval PCMM MaxTransactionsPerSecond	Interval Maximum PCMM Transactions per Second
Interval Rx MaxTransactionsPerSecond	Interval Maximum Rx Transactions per Second

Table 3-10 Policy Statistics

Reports Display Name	OSSI XML Name
Peg Count	
Evaluated	
Executed	
Ignored	
Policy Details Stats:	
Name	
Evaluated	Eval Count
Executed	Trigger Count

Table 3-10 (Cont.) Policy Statistics

Reports Display Name	OSSI XML Name
Ignored	
Policy write state on session create	
Name	
Evaluated	
Executed	
Ignored	
Total Execution Time (ms)	
Max Execution Time (ms)	
Avg Execution Time (ms)	
Processing Time Stats	
Policy write state on session termination	
Name	
Evaluated	
Executed	
Ignored	
Total Execution Time (ms)	
Max Execution Time (ms)	
Avg Execution Time (ms)	
Processing Time Stats	

Viewing Policy Server Logs

The log files trace the activity of a Policy Management device. You can view and configure the logs for an individual cluster.

To view the log:

1. From the Policy Server section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy server groups.

2. From the content tree, select the Policy Management device.

   The Policy Server Administration page opens in the work area.

3. Select the **Logs** tab.

Depending on your mode and release, you can configure the following logs:

- **Trace log** — Records application-level notifications.
- **Trace Log Forwarding** — Forwards cluster-level notifications.
- **Policy Log Settings** — Records the policy-level messages.
- **Policy Syslog Forwarding** — Records policy-processing activity. Supports the standard UNIX logging system, in conformance with RFC 3164.
- **SMS log** — Contains all Short Messaging Service messages sent by the MPE device as well as any ACK messages received from an SMS Center (SMSC) server or its equivalent

Viewing the Trace Log

The trace log records Policy Management application notifications, such as protocol messages, policy messages, and custom messages generated by policy actions, for individual servers. Trace logs are not replicated between servers in a cluster, but they persist after failovers. You can use the log to debug problems by tracing through application-level messages. You can configure the severity of messages that are recorded in the trace log. For more information, see Configuring Log Settings.

Note: Prior to release 7.5, the trace log was called the event log, which also contained platform events. Platform and connectivity events are now displayed as alarms. Additionally, prior to release 7.5, a policy log file recorded the activity of the Policy Rules Engine, at seven levels: Alert, Critical, Error, Warning, Notice, Info, and Debug. This information is now recorded in the trace log at eight levels: Emergency (ID 4560), Alert (ID 4561), Critical (ID 4562), Error (ID 4563), Warning (ID 4564), Notice (ID 4565) Info (ID 4566), and Debug (ID 4567).

To view log information using the Trace Log Viewer:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of groups; the initial group is **ALL**.

2. From the content tree, select the device.

   The Policy Server Administration page opens in the work area.

3. Select the **Logs** tab.

   Log information for the selected device is displayed.

4. Click **View Trace Log**.

   The Trace Log Viewer window opens. While data is being retrieved, the in-progress message Scanning Trace Logs displays.

   All events contain the following information:

   - **Date/Time** — Event timestamp. This time is relative to the server time.
   - **Code** — The event code. For information about event codes and messages, see the *Policy Management Troubleshooting Reference*.

- **Severity** — Severity level of the event. Application-level trace log entries are not logged at a higher level than Error.
- **Message** — The message associated with the event. If additional information is available, the event entry shows as a link. Click on the link to see additional detail in the frame below.

5. You can filter the events displayed using the following:

   - **Trace Log Viewer for Server** — Select the individual server within the cluster.
   - **Start Date/Time** — Click 📅 (calendar icon), select the starting date and time, then click **Enter**.
   - **End Date/Time** — Click 📅 (calendar icon), select the ending date and time, then click **Enter**.
   - **Trace Codes** — Enter one or a comma-separated list of trace code IDs. Trace code IDs are integer strings up to 10 digits long.
   - **Use timezone of remote server for Start Date/Time** — Select to use the time of a remote server (if it is in a different time zone) instead of the time of the CMP server.
   - **Severity** — Filter by severity level. Events with the selected severity and higher are displayed. For example, if the severity level selected is **Warning**, the trace log displays events with the severity level Warning.
   - **Contains** — Enter a text string to search for. For example, if you enter `connection`, all events containing the word connection appear.

   > **Note:** The **Start Date/Time** setting overrides the **Contains** setting. For example, if you search for events happening this month, and search for a string in events last month and this month, only results from this month are listed.

6. After entering the filtering information, click **Search**.

   The selected events are displayed.

By default, the window displays 25 events per page. You can change this to 50, 75, or 100 events per page by selecting a value from the **Display results per page** pulldown list.

Events that occur after the Trace Log Viewer starts are not visible until you refresh the display. To refresh the display, click one of the following buttons:

- **Show Most Recent** — Applies filter settings and refreshes the display. This displays the most recent log entries that fit the filtering criteria.
- **Next/Prev** — When the number of trace log entries exceeds the page limit, pagination is applied. Use the **Prev** or **Next** buttons to navigate through the trace log entries. When the **Next** button is not visible, you have reached the most recent log entries; when the **Prev** button is not visible, you have reached the oldest log entries.
- **First/Last** — When the number of trace log entries exceeds the page limit, pagination is applied. Use the **First** and **Last** buttons to navigate to the beginning

or end of the trace log. When the **Last** button is not visible, you have reached the end; when the **First** button is not visible, you have reached the beginning.

Click **Close** to close the view.

Syslog Support

Notifications generated by policy actions are sent to the standard UNIX syslog. No other notifications are forwarded to the syslog. For information on policy actions, see the *Policy Wizard Reference*.

Note: These logs are separate from the TPD syslogs.

You can define multiple destinations for notifications and filter notifications by severity level. For more information, see Configuring Log Settings.

The SMS Log

The SMS log, `/var/Camiant/log/smsr.log`, contains all Short Message Service (SMS) messages sent by the MPE device as well as any ACK messages received from an SMS Center (SMSC) server or its equivalent. You can configure the severity level as well as the destination IP addresses of messages that are written to the SMS log. The default severity level is WARN. See Configuring Log Settings for more information.

The SMTP Log

The SMTP log contains all Simple Mail Transfer Protocol (SMTP) messages sent by the MPE device, as well as any ACK messages received from a Mail Transfer Agent (MTA). In SMPP or XML mode, the SMTP log information appears on the **Logs** tab of the Policy Server Administration page. You can modify the severity level of messages that are written to the SMTP log on the MPE configuration page. The default severity is WARN. See Configuring Log Settings to modify the settings.

The HTTP Log

The HTTP log contains all Hypertext Transfer Protocol (HTTP) messages sent by the MPE device. In SMPP or XML mode, the HTTP log information appears on the **Logs** tab of the Policy Server Administration page. You can modify the severity level of messages that are written to the HTTP log on the server configuration page. The default severity is WARN. See Configuring Log Settings for more information.

Configuring Log Settings

To configure the log settings for the servers in a cluster:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of server groups; the initial group is **ALL**.

2. From the content tree, select the **ALL** group.

   The Policy Server Administration page opens in the work area.

3. Select an MPE device from the list.

The Policy Server Administration page opens in the work area and details the configuration settings of the selected device.

4. Select the **Logs** tab.

   The Policy Server Administration page opens and details the logs configuration settings for the specified device.

5. To edit the logs configuration settings, click **Modify**.

   The editable fields open in the work area.

6. In the **Modify Trace Log Settings** section of the page, select the **Trace Log Level** from the list.

   This setting indicates the minimum severity of messages that are recorded in the trace log. These severity levels correspond to the syslog message severities from RFC 3164 *The BSD syslog Protocol*. Adjusting this setting allows new notifications, at or above the configured severity, to be recorded in the trace log. The levels are:

   - **Emergency**

     Provides the least amount of logging, recording only notification of events causing the system to be unusable.

   - **Alert**

     Action must be taken immediately in order to prevent an unusable system.

   - **Critical**

     Events causing service impact to operations.

   - **Error**

     Designates error events which may or may not be fatal to the application.

   - **Warning** (default)

     Designates potentially harmful situations.

   - **Notice**

     Provides messages that may be of significant interest that occur during normal operation.

   - **Info**

     Designates informational messages highlighting overall progress of the application.

   - **Debug**

     Designates information events of lower importance.

   Caution: Before changing the default logging level, consider the implications. Lowering the log level setting from its default value (for example, from **Warning** to **Info**) causes more notifications to be recorded in the log and can adversely affect performance. Similarly, raising the log level setting (for example, from **Warning** to **Alert**) causes fewer notifications to be recorded in the log and may cause you to miss important notifications.

7. You can enable and configure **Trace Log Forwarding Settings** for individual clusters.

   > **Note:** The CMP system provides log forwarding configuration for all products that have trace logs: MPE, MA, BoD, and the CMP itself.

   For each cluster, enter the following:

   a. Select to enable **Enable Trace Log Forwarding** in the **Modify Trace Log Forwarding Settings** section of the page.

      The Trace Log Forwarding settings become editable.

   b. Enter a valid **Hostname/IP Address** for each device receiving the trace logs.

      > **Note:** The system validates the IP address is unique based on the literal value. It does not resolve the host name or check the short pattern IPv6 to the full pattern IPv6 address.

   c. Select the appropriate **Severity** level for the trace logs being forwarded for each cluster. See Step 6 for a description of each level.

8. In the **Modify Policy Log Settings** section of the page, configure the **Policy Log Level**.

   This setting indicates the minimum severity of messages that are recorded in the policy log for all policies. The levels are:

   - **OFF**

     No messages are recorded.

   - **DEBUG**

     All messages are recorded.

   - **INFO**

     Only informational messages are recorded.

   - **WARN** (default)

     Only messages designating potentially harmful situations are recorded.

9. In the **Modify SMTP Log Settings** section of the page, configure the **SMTP Log Level**.

   This setting indicates the minimum severity of messages that are recorded in the SMTP log. These severity levels correspond to the syslog message severities from RFC 3164 *The BSD syslog Protocol*. Adjusting this setting allows new notifications, at or above the configured severity, to be recorded in the SMTP log. The levels are:

   - **OFF**

     Turns off logging.

   - **ERROR**

     Designates error events which may or may not be fatal.

- **WARN** (default)

  Designates potentially harmful situations.
- **INFO**

  Designates informational messages highlighting overall progress.
- **DEBUG**

  Designates information events of lower importance.
- **TRACE**

  Designates informational events of very low importance.
- **ALL**

  Records all logging levels.

10. In the **Modify HTTP Log Settings** section of the page, configure the **HTTP Log Level**.

    This setting indicates the minimum severity of messages that are recorded in the HTTP log. Adjusting this setting allows new notifications, at or above the configured severity, to be recorded in the HTTP log. The levels are:

    - **OFF**

      Turns off logging.
    - **ERROR**

      Designates error events which may or may not be fatal.
    - **WARN** (default)

      Designates potentially harmful situations.
    - **INFO**

      Designates informational messages highlighting overall progress.
    - **DEBUG**

      Designates information events of lower importance.
    - **TRACE**

      Designates informational events of very low importance.
    - **ALL**

      Records all logging levels.
11. Click **Save**.

The log settings are configured.

4

Configuring Protocol Routing

Routing enables a Policy Management device to forward requests to other Policy Management devices for further processing. The following routing messages and protocols are supported:

- PacketCable MultiMedia (PCMM) messages
- Diameter Rx messages

PCMM Routing Architectures

There are two architectures you can employ with PCMM routing: Hierarchical and Mesh.

- **Hierarchical** — In a hierarchical architecure, there is a top-level MPE cluster (an MPE-R cluster) and one or more bottom-level MPE clusters (MPE-S clusters). A PCMM message is directed to the top-level MPE cluster, which then routes the message to the appropriate MPE cluster below based on the subscriber IP address in the message.
- **Mesh** — In a mesh architecture, there is a set of two or more MPE clusters, but there is no top-level cluster. If you imagine three MPE clusters arranged in a triangle, a PCMM message coming into any one of these clusters can be forwarded out to any of the other two MPE clusters. Each cluster points to the other clusters.

In either architecture, a PCMM message is handled by the MPE cluster to which it is sent, and does not have to be forwarded. For example, in a hierarchical architecture, if a PCMM message comes into the top-level MPE cluster, and the appropriate CMTS is associated with that cluster, then the cluster handles the message itself.

Configuring PCMM Routing

Configuring PCMM routine establishes a hierarchical network of MPE-R (routing) and MPE-S systems.

To configure PCMM routing:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the **ALL** group.

   The Policy Server Administration page opens in the work area.

3. Select the MPE device.

   The Policy Server Administration page opens to the **System** tab displaying information for that device.

4. Select the **Routing** tab.

   The routing configuration settings are displayed.

5. Click **Modify**.

   The Modify Routing Configuration page opens. (Figure 4-1 shows an example.)

6. Set the following values:

   a. **Execute Policies for Routed Traffic** — If this checkbox is enabled, the MPE device applies its locally configured policies to any requests before forwarding them to another policy server.

      Typically, this feature is disabled, as the MPE device that is receiving the request is also applying policies. However, this feature is useful in a hierarchical network. Enabling this feature typically causes a reduction in the performance of the routing function.

      Note: MPE devices do not support policy execution on Diameter traffic on the basis of routing, either by normal Diameter routing or by IP address.

   b. **Route to Downstream Policy Servers using IP subnets** — If this checkbox is selected, Rx traffic is routed statelessly (without translation) to other MPE devices.

   c. **Downstream Policy Servers** — A list of MPE-S devices where this MPE-R device can forward requests.

      You can change this setting by clicking on the MPE devices in the list. Highlighted MPE devices are included; others are not.

      Note: If you wish to configure both MGPI and downstream policy servers, you must select either **Execute Policies for Routed Traffic** or **Route to Downstream Policy Servers using IP subnets** here.

7. Click **Save**.

PCMM routing is configured.

Figure 4-1 Modify Routing Configuration Page

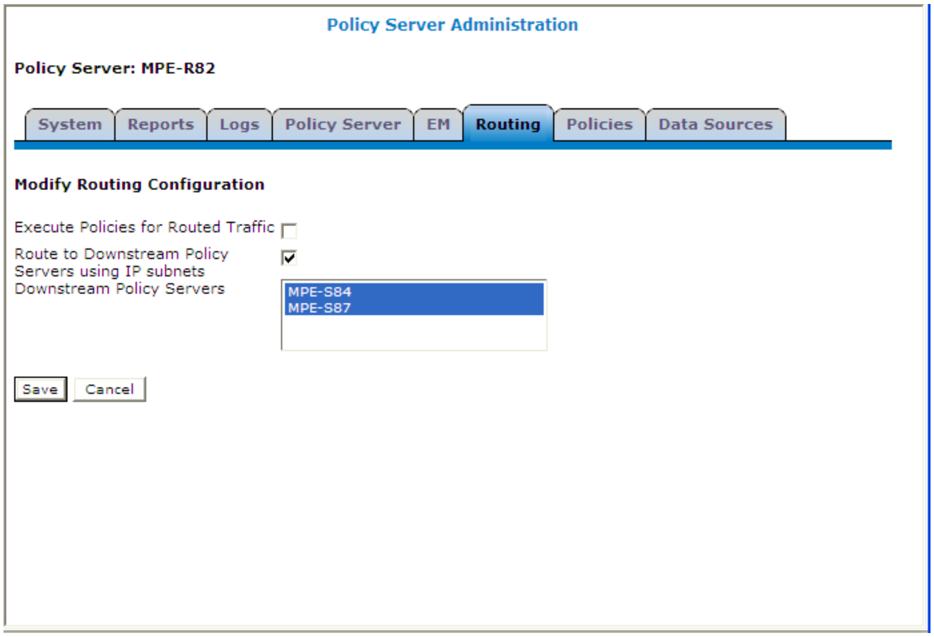

Policy Server Administration

Policy Server: MPE-R82

Modify Routing Configuration

Execute Policies for Routed Traffic □

Route to Downstream Policy Servers using IP subnets ☑

Downstream Policy Servers: MPE-S84, MPE-S87

Configuring Rx-to-PCMM Routing

An MPE device can translate Rx requests to PCMM requests or, in a hierarchical network, route them elsewhere to be translated. For Rx-to-PCMM routing, configure the top-level MPE device for stateless PCMM routing. To do this:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the **ALL** group.

   The Policy Server Administration page opens in the work area.

3. From the **ALL** group, select the MPE device.

   The Policy Server Administration page opens to the **System** tab, displaying information for that device.

4. Select the **Routing** tab.

   The routing configuration settings display.

5. Click **Modify**.

   The Modify Routing Configuration page opens.

6. Select **Route to Downstream Policy Servers using IP subnets**.

7. Deselect **Execute policies for Routed Traffic**.

8. Click **Save**.

Rx-to-PCMM routing is configured.

5

Configuring Advanced Device Settings

This chapter describes how to configure and manage expert settings, service overrides, and load shedding options.

Configuring Expert Settings for an MPE

Expert settings control global settings that are not used regularly. For example, session cleanup options and timers. These setting are set for a specific MPE.

To configure Expert Settings:

1. To view the device list, from the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the device and and select the **Policy Server** tab.

   The configuration settings for the device display.

3. Click **Advanced**.

   The advanced settings for the device display:

   - **Expert Settings**
   - **Service Overrides**
   - **MPE Load Shedding Configuration** (Level 1 to Level 4)

4. Click **Modify**.

   The advanced configuration settings can be edited.

5. Select a configuration key in the **Expert Settings** table and click **Edit**.

   The Edit Expert Setting Value dialog opens.

6. Modify the settings and click **OK** to save.

   See Expert Settings for MPE for information about the configurations keys.

7. Click **Save**.

The settings are applied to the selected device.

Expert Settings for MPE

Table 5-1 Expert Settings for MPE

Category	Configuration Key	Description	Default
Admission	ADMISSION.DIAMETER. RequestProcessingLimit	The maximum amount of time a request can be processed before being dropped, if no answer has been sent. Specified in milliseconds.	5000
Database	DB.User.EnableBillingStartDate	Enables the use of the billing date effective name from the profile for the subscriber. This is used as a start date for the plan and is also used to calculate the next reset time.	true
Diameter	DIAMETER.AF.AuditForAuthLifetime	Enables the configuration of a minimum and maximum lifetime for an AF session.	False
Diameter	DIAMETER.AF.AuthLifetime	The maximum lifetime of an AF session. Otherwise the corresponding AF session would be purged subject to the configured grace period. Specified in seconds. Valid range is 300 to 58060800.	86400 (1 day)
Diameter	DIAMETER.AF.MinAuthLifetime	The minimum lifetime of an AF session. Otherwise the corresponding AF session would be purged subject to the configured grace period. Specified in seconds.	300
Diameter	DIAMETER.AF. EnableGracePeriodForSubscriptionExpiry	Enables the configuration of a grace period for an AF session.	False
Diameter	DIAMETER.AF. GracePeriodForSubscriptionExpiry	Indicates the maximum configured grace period for an AF session, which is added to the negotiated AuthLifeTime to determine if a given AF session can be considered stale and purged. Specified in seconds. Valid range is 0 to 86400.	86400 (1 day)
Diameter	DIAMETER.AF. SignallingSessionAuthLifetime	Indicates the maximum configured period for an Rx sessions containing signaling flow past which the session is considered stale. Specified in seconds. Valid range is 300 to 58060800.	259200 (3 days)

Table 5-1 (Cont.) Expert Settings for MPE

Category	Configuration Key	Description	Default
Diameter	DIAMETER.AppsToEvaluateOnTermination	Indicates applications for which policy evaluation is triggered when a terminate request is received. Specified as a comma-separated string containing the names of the applications for which policy engine should evaluate termination requests. DIAMETER. Valid values are: • Gx • Rx • Sd • S9 PolicyExecutionOnSessionTermination must be set to false to use this configuration.	null
Diameter	DIAMETER.Cleanup.AuditRxSessions	If enabled, an RAR message is sent for auditing. **Note:** This is for future releases and has not been implemented yet.	False
Diameter	DIAMETER.Cleanup.AuditSySessions	If enabled, an SLR (INTERMEDIATE) message is sent for auditing. If disabled, the Sy session is checked for an association with an IP-CAN session. If there are no IP-CAN associations, the Sy session is considered active; otherwise, the session is deleted. **Note:** This is for future releases and has not been implemented yet.	False

Table 5-1 (Cont.) Expert Settings for MPE

Category	Configuration Key	Description	Default
Diameter	DIAMETER.Cleanup.CleanupStaleRxSessions	Determines if the MPE device should consider AF sessions in the regular cleanup cycles. If enabled, AF sessions are considered expired if they have lived longer than the specified AFSessionValidityTime. **Note:** At that point, in future releases, if AuditAFSessions is set to True, an RAR will be sent for auditing the session.	True
Diameter	DIAMETER.Cleanup.OverrideCleanupAudit	This specifies if the regular audit processing for cleaning up a stale session is overridden. When enabled, the cleanup task bypasses the audit process and deletes all sessions that are stale for the session validity time.	False
Diameter	DIAMETER.Cleanup.RXSessionValidityTime	The amount of time in seconds after which the session is expired and is purged if EnabledAFSessionCleanup is enabled.	86400 (1 day)
Diameter	DIAMETER.Cleanup.SessionCleanupInterval	The amount of time in seconds after which the cleanup task will run to look for stale sessions.	21600 (6 hours)
Diameter	DIAMETER.Cleanup.SessionValidityTime	The amount of time in seconds after which a session in a binding is declared stale. Valid range is 1 to 8640000.	432000 (5 days)
Diameter	DIAMETER.Cleanup.SySessionValidityTime	The amount of time in seconds after which the session is declared stale and deemed a candidate for cleanup.	36000 (10 hours)
Diameter	DIAMETER.Cleanup. AuditSySendEmptyPolicyCounterList	If enabled, the Policy Counter Identifier subscription list is not sent as part of an SLR (INTERMEDIATE) message to audit stale Sy sessions. If disabled, the current Policy Counter Identifier subscription list is sent as part of an SLR (INTERMEDIATE) message to audit stale Sy sessions.	True
Diameter	DIAMETER.Cleanup. MaxDurationForSessionIteration	The maximum duration in seconds to iterate through the sessions. Valid range is 1 through 86400.	7200 (2 hours)

Table 5-1 (Cont.) Expert Settings for MPE

Category	Configuration Key	Description	Default
Diameter	DIAMETER.Cleanup.MaxSessionCleanupRate	The rate (in sessions/sec) at which the cleanup task attempts to clean stale sessions. Valid range is 1 through 5000.	50
Diameter	DIAMETER.Cleanup.MaxSessionIterationRate	The rate (in sessions/sec) at which the cleanup task iterates through the sessions database. Valid range is 1 through 100000.	1000
Diameter	DIAMETER.Cleanup.MaxSessionValidityTime	The maximum amount of time in seconds after which the session is cleaned up on any error. Valid range is 1 through 8640000.	172800 (2 days)
Diameter	DIAMETER.Cleanup.MaxSySessionValidityTime	The maximum amount of time in seconds after which the Sy session is cleaned up on any error. Valid range is 1 through 8640000.	172800 (2 days)
Diameter	DIAMETER.Cleanup.SessionCleanupStartTime	Schedules the cleanup task once a day at a specified time. If the start time is specified, then it is scheduled to run once a day at the given time. The value can be specified in either a 24-hr format (*HH*:*mm*) or an exact date and time (*YYYY-MM-dd*T*hh*:*mm*:*ss*) of when it will first run and then repeat at the interval specified.	undefined
Diameter	DIAMETER.ENF.MaxTimeForAnGwFailure	The maximum time allowed after getting indication for SGW failure in which the MPE device does not send any new or updated policies to the P-GW except rules to be removed. Specified in seconds.	3600
Diameter	DIAMETER.ENF.ReevaluateGeneratedDefaultRule	If set to True, on receipt of a session update request the MPE device evaluates the flow which contains the generated default rule even when there is no new default-EPS-Bearer-QoS information in the request.	False
Diameter	DIAMETER.ENF.RegisterForAnGwChangeWithSGWRest	If the SGW-Rest supported feature was negotiated and the value of this parameter is false, the MPE device checks if AN_GW_CHANGED event trigger is one of the armed event triggers to be installed and removes it. As a result, the MPE device will not register for AN_GW_CHANGE.	False

Table 5-1 (Cont.) Expert Settings for MPE

Category	Configuration Key	Description	Default
Diameter	DIAMETER.EnableSessionCleanUp	Enables the DiameterSessionCleanUp Task.	True
Diameter	DIAMETER.PCEF.NetLocSupportedAccesses	A combination of values used for network location (NetLoc) support. The format is *IPCanType*:*RATType*:*AN-Trusted*. If not set, any types or values are applicable. For example: `THREEGPP_GPRS` means to support NetLoc as if IPCanType=THREEGPP_GPRS. `NON_THREEGPP_EPS:WLAN:TRUSTED` means to support NetLoc when IPCanType=NON_THREEGPP_EPS, RATType=WLAN, and ANTrusted=TRUSTED. `NON_THREEGPP_EPS::TRUSTED` means to support NetLoc when IPCanType=NON_THREEGPP_EPS and ANTrusted=TRUSTED. `::` means all accesses support NetLoc.	THREEGPP_GPRS, THREEGPP_EPS, NONTHREEGPP_EPS: WLAN: TRUSTED
Diameter	DIAMETER.PolicyExecutionOnSessionTermination	If enabled (True), policy evaluation will be triggered for all applications when a terminate request is received. This configuration must be disabled (False) to use DIAMETER.AppsToEvaluateOnTermination to selectively trigger policy evaluation by application.	True
Diameter	DIAMETER.SessionUniquenessControl	If enabled (True), the MPE cluster will maintain session uniqueness and avoid stale session processing.	False
Diameter	DIAMETER. SessionUniquenessControlWaitTime	If enabled (True), the Max-Wait-Time AVP is used in conjunction with the message origination time stamp to determine staleness of message.	False
PCMM	PCMM.Cleanup.CleanupStalePcmmSessions	Enables the inclusion of PCMM sessions in regular cleanup cycles. If enabled (True), PCMM sessions are considered expired if they have lived longer than the specified PcmmSessionValidityTime or license timeout duration configured in the application.	

Table 5-1 (Cont.) Expert Settings for MPE

Category	Configuration Key	Description	Default
PCMM	PCMM.Cleanup.PcmmSessionValidityTime	The amount of time in seconds after which the session is deemed expired and is purged if the setting EnabledAFSessionCleanup is enabled.	86400 (1 day)
SH	SH.Retry.Enabled	If enabled, indicates that the MPE device retries the Sh Requests for UDR, PUR, and SNR to the backup server when the primary server returns one of the defined error codes. PNA messages are not retried. If the backup server returns an error, then retry is not performed on the primary server.	False
SH	SH.Retry.EnabledOnTimeout	If enabled with the SH.Retry.Enabled setting, allows the MPE to retry an Sh Request to a backup datasource server when there is a response timeout. If the request to backup server times out, then a retry is not performed on the primary server. This setting depends on the configuration of the SH.Retry.Enabled and SH.ResponseTimeout settings. You can configure the SH.ResponseTimeout in the **Service Overrides** section of the Advanced Settings page.	False
SMPP	SMPP.SendSMSNowWhenDeliveryDateInPast	In SMS:SMPP mode, determines if SMS notifications scheduled to be delivered in the past will be dropped or delivered. If false, SMS notifications scheduled to be delivered on a date in the past will be dropped. If true, SMS notifications scheduled to be delivered on a date in the past will be delivered immediately.	False
SMSXML	SMSXML.SendSMSNowWhenDeliveryDateInPast	In SMS:XML mode, determines if SMS notifications scheduled to be delivered in the past will be dropped or delivered. If false, SMS notifications scheduled to be delivered on a date in the past will be dropped. If true, SMS notifications scheduled to be delivered on a date in the past will be delivered immediately.	False

Table 5-1 (Cont.) Expert Settings for MPE

Category	Configuration Key	Description	Default
SY	SY.Reconciliation.Enabled	Determines whether the Sy Reconciliation is activated and an audit of Sy sessions will be executed on a recovery from a split-brain scenario.	False
SY	SY.Reconciliation.HoldTimer	The time in seconds after receipt of a notification of recovery from a split-brain scenario the Sy Reconciliation task will wait before starting.	180
SY	SY.Reconciliation.MaxSessionReconcileRate	The rate (in sessions/sec) at which the tasks will attempt to send Sy SLR Messages to reconcile Sy sessions.	50
SY	SY.SendOriginationTimestamp	If enabled (True), the message origination timestamp will be included as part of the initial SLR Sy message.	False
Diameter	DIAMETER.Cleanup.MaxSySessionValidity Time		172800
Diameter	DIAMETER.ConnectionTimeOut		3
KPI	KPI.Capacity.Session		1
KPI	KPI.Capacity.TPS		1

Configuring Service Overrides

Caution: Do not attempt to add or change a service override without first consulting with My Oracle Support.

Configuration key changes are made using the Service Overrides section of the Advanced configuration page.

Make service override changes as follows:

1. View the device list.

   - For an MPE device, go to the **Policy Server** section of the navigation pane and select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the device.

   - For an MPE device, select the **Policy Server** tab.

   The configuration settings for the device display.

3. Click **Advanced**.

The advanced settings for the device display.

4. Click **Modify**.

   The advanced configuration settings can be edited.

5. Select a configuration key in the Service Overrides table and click **Edit**.

   - Adding a key to the table:

     a. Click **Add**.

        The Add Configuration Key Value window opens.

     > **Caution:** There is no input validation on values. Also, if you overwrite a setting that is configurable using the CMP GUI, the value adopted by the server is undetermined.

     b. Enter the following values:

        – **Configuration Key —** The attribute to set

        – **Value —** The attribute value (up to 255 characters)

        – **Comments —** Information about the key.

     c. Click **OK**.

        The key is displayed in the table with its defined and default values.

   - Cloning a key in the table:

     a. Select an existing key in the table.

     b. Click **Clone**.

        The Clone Configuration Key Value window opens with the information for the key.

     c. Make changes as required.

     d. Click **Save**.

   - Editing a key in the table:

     a. Select an existing key in the table.

     b. Click **Edit**.

        The Edit Configuration Key Value window opens with the information for the key.

     c. Make changes as required.

     d. Click **Save**.

   - Deleting a key from the table:

     a. Select an existing key in the table.

     b. Click **Delete**. A confirmation message displays.

c. Click **Delete** to remove the key.

6. Click **Save**.

The settings are applied to the selected device.

About Overload Controls

Load Shedding occurs when a Diameter node (an MPE device) has insufficient resources to successfully process all of the Diameter requests that it receives. You can access **Load Shedding Configuration** controls from the MPE Advanced Configuration page where you can configure rules for handling messages during overload conditions. Multiple congestion levels can be configured to accept, reject or drop selected messages at each level.

MPE devices have configurable levels of congestion (busyness) for handling message overload. An MPE device has four congestion levels (Levels 1–4). At each level you can define a default action for the level and create rules to handle specific message types. A level action is an action that is taken if none of the rules configured for the level match a message type. For example, for MPE Level 1, the default level action is **Accept**, which means to bypass load shedding rules instead of rejecting messages.

Configuring MPE and MRA Load Shedding Rules

Use the **Load Shedding Configuration** section of the Advanced Configuration page to edit, reorder, or add new rules at each level of busyness for a device based on the amount of backlog. To reach a configured level of busyness:

- The backlog of outstanding messages in a node crosses a predefined threshold for the level.
- The backlog has been above the busyness level threshold for a minimum amount of time.

At each level, the device can be configured to take one of the following actions (referred to as rules) until the busyness level clears:

- Accept the message.
- Drop the message.
- Reject new messages with a specific result code (the default is DIAMETER_TOO_BUSY).

Refer to MPE Default Load Shedding Rules for more information on default rules.

Note: Configuration keys must also be used in configuring load shedding options. Contact My Oracle Support for assistance.

Configure the load shedding rules as follows:

1. View the device list.
   - For an MPE device, go to the **Policy Server** section of the navigation pane and select **Configuration**.

- For an MRA device, go to the **MRA** section of the navigation pane and select **Configuration**.

The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the device.

   - For an MPE device, select the **Policy Server** tab.
   - For an MRA device, select the **MRA** tab.

   The configuration settings for the device display.

3. Click **Advanced**.

   The advanced settings for the device display.

4. Click **Modify**.

   The advanced configuration settings can be edited.

5. In the **Load Shedding Configuration** section of the page, select the enabled state.

   - **true** (default)

     Enables load shedding.

   - **false**

     Disables load shedding.

   - **undefined**

     The value for this field is taken from the associated Configuration Template. If there is not a configuration template associated, then the default value is used.

6. Set the Level Action for a busyness level.

   This step is optional. The default Level Action of Accept applies to all levels except MPE Level 4 and MRA Level 2, which has a default Level Action of Drop.

   a. Click ▶ (right arrow) next to the level to expand the level.

   b. Select one of the following default Level Actions:

      - **Accept** all messages.
      - **Drop** all messages.
      - Reject all messages and **Answer with** (select a code from the drop-down list).
      - Reject all messages and **Answer With Code** (enter a code) and **Vendor ID** (enter a vendor ID).

7. Configure the rules for the busyness levels:

   a. Click ▶ (right arrow) next to the level to expand the level.

   b. Click **Add** and select the category.

      The Add Load Shedding Rule dialog appears.

c. Enter the values for the load shedding rule:

- **Name**

  Name of the rule. The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_). It cannot begin or end with a hyphen or period, and any labels must be separated by periods.

- **Application**

  Select the application the rule applies to. You can select **Drma**, **Gx**, **Gxx**, **S9**, **Rx**, **Sh**, or **Sy**.

- **Message**

  Type of message the rule applies to (which depends on the application chosen).

- **Request Types** (available only when the CCR message type is selected)

  Select the Request-Type attribute-value pairs (AVPs) that the message must contain. You can select **Initial**, **Update**, and/or **Terminate**.

- **APNs**

  Enter a CSV list of one or more access point names that the message must contain.

- **DRMP** (availability dependent on selected application and message type)

  Enter a CSV list of one or more diameter routing message priority codes that the message must contain. The valid range of values is 0 to 15.

- **Both MPS ID and Reservation Priority Exist** (available only when the AAR message type is selected)

  Determines whether or not to check for the existence of the MPS-Identifier and Reservation-Priority AVPs.

d. Click **OK**.

   The rule is displayed in the table.

8. After a rule is defined, you can clone, edit, or delete it by selecting the rule and clicking the appropriate button.

The settings are applied to the selected device.

MPE Default Load Shedding Rules

You can configure load shedding rules to determine how a device reacts to a processing backlog. This state is called busyness. Levels of busyness can be configured to accept, reject, or drop select messages at each level. An MPE has four busyness levels. With each successive level, the device becomes more aggressive in rejecting or discarding messages in an attempt to prevent the main queue from becoming full. At any level of busyness, requests that have been queued longer than a configurable time are discarded without further processing, since the originator would have abandoned that request.

On the MPE Advanced Configuration page, there is a default action for each Busyness Level. The default level action for levels 1, 2, and 3 is Accept, which means to process the message by bypassing load shedding rules. Level actions are configurable.

The following tables show the default load-shedding rules for an MPE. For configuration information, see the task Configuring MPE and MRA Load Shedding Rules.

> **Note:** The default rules shown in your system may differ than those listed here depending on how your system is configured.

Table 5-2 MPE Busyness Level 2

Rule Name	Actions
DefaultRule7	Reject Rx AAR messages with DIAMETER_TOO_BUSY
DefaultRule15	Accept Gx CCR messages with a DRMP value of 0
DefaultRule17	Accept Rx AAR messages with both the MPS-Identifier and Reservation-Priority AVPs present

Table 5-3 MPE Busyness Level 3

Rule Name	Actions
DefaultRule8	Reject Rx AAR messages with DIAMETER_TOO_BUSY

Table 5-4 MPE Busyness Level 4

Rule Name	Actions
DefaultRule11	Accept Drma LNR with ACCEPT
DefaultRule12	Accept Drma LSR with ACCEPT
DefaultRule13	Accept Drma RUR with ACCEPT

Resetting Configuration Keys to Defaults

All the configuration keys in the Expert Settings table can be reset to the defaults. The configuration keys in the Service Overrides table cannot be reset.

To reset the configuration keys in the Expert Settings table:

1. View the device list.

   - For an MPE device, go to the **Policy Server** section of the navigation pane and select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the device.

   - For an MPE device, select the **Policy Server** tab.

   The configuration settings for the device display.

3. Click **Advanced**.

   The advanced settings for the device display.

4. Click **Modify**.

   The advanced configuration settings can be edited.

5. Click **Set to Default**.

   A confirmation message displays.

6. Click **OK**.

All the configuration keys for Expert Settings are set to default values.

Filtering the Configuration Keys

To limit the number of configuration keys in the Expert Settings or Service Overrides tables, use the filter option.

To filter the configuration key table:

1. View the device list.

   - For an MPE device, go to the **Policy Server** section of the navigation pane and select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the device.

   - For an MPE device, select the **Policy Server** tab.

   The configuration settings for the device display.

3. Click **Advanced**.

   The advanced settings for the device display.

4. (Optional) Click **Modify**.

   The advanced configuration settings can be edited.

5. Click **Filters** to open the filtering popup.

   The filtering popup opens.

6. Specify the filtering parameters using any of the following fields.

Option	Description
Change Status	The change status of the configuration key. • **All** (default)—All keys are listed. • **Changed**—Lists the configuration keys that have been modified from the default setting.

Option	Description
	• **Unchanged—**Lists the configuration keys that have not been modified from the default setting.
Category	The category for the configuration key.
Configuration Key	Enter all or part of a configuration key name.

7. Click **Filter Result**.

   The filtered list of configuration keys displays.

Exporting the Configuration Keys

The Expert Settings or Service Overrides configuration keys can be exported to a comma separated values (CSV) file or to a printable format in a new browser window.

To export the configuration key table:

1. View the device list.

   - For an MPE device, go to the **Policy Server** section of the navigation pane and select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the device.

   - For an MPE device, select the **Policy Server** tab.

   The configuration settings for the device display.

3. Click **Advanced**.

   The advanced settings for the device display.

4. Click **Export**.

   The export list opens.

5. Select the export type.

Option	Description
Save as CSV	A comma-separated value (CSV) file named `CSV_report.csv` is generated, suitable for a spreadsheet application, and a standard File Download window opens, so you can save or open the file.
Printable Format	The configuration key list displays in a separate window for printing.

6

Configuring Debug Logs

This chapter describes how to configure class-level logging options. Debug logs can be configured at the System Administration level or at the Policy Management device (MPE or BoD) level.

About Debug Logs

Use the CMP server to modify the configuration for each target component on a Policy Management (MPE or BoD) device.

To aid in troubleshooting, you can enable component-specific level logging for the following components:

Component	File Name
Tomcat	`logback-tomcat.log`
RC	`logback-rc.xml`
BoD	`logback-bod.xml`
SMSR	`logback-tomcat-rc.xml`

> **Note:** See #unique_221for details about configuring the default values for debug logs.

About MPE Debug Logs

Debugs logs for the MPE include log configuration for the following components:

- Tomcat Log
- RC Log
- SMSR (SMS Relay)
- DC Log

Figure 6-1 MPE Debug Logs

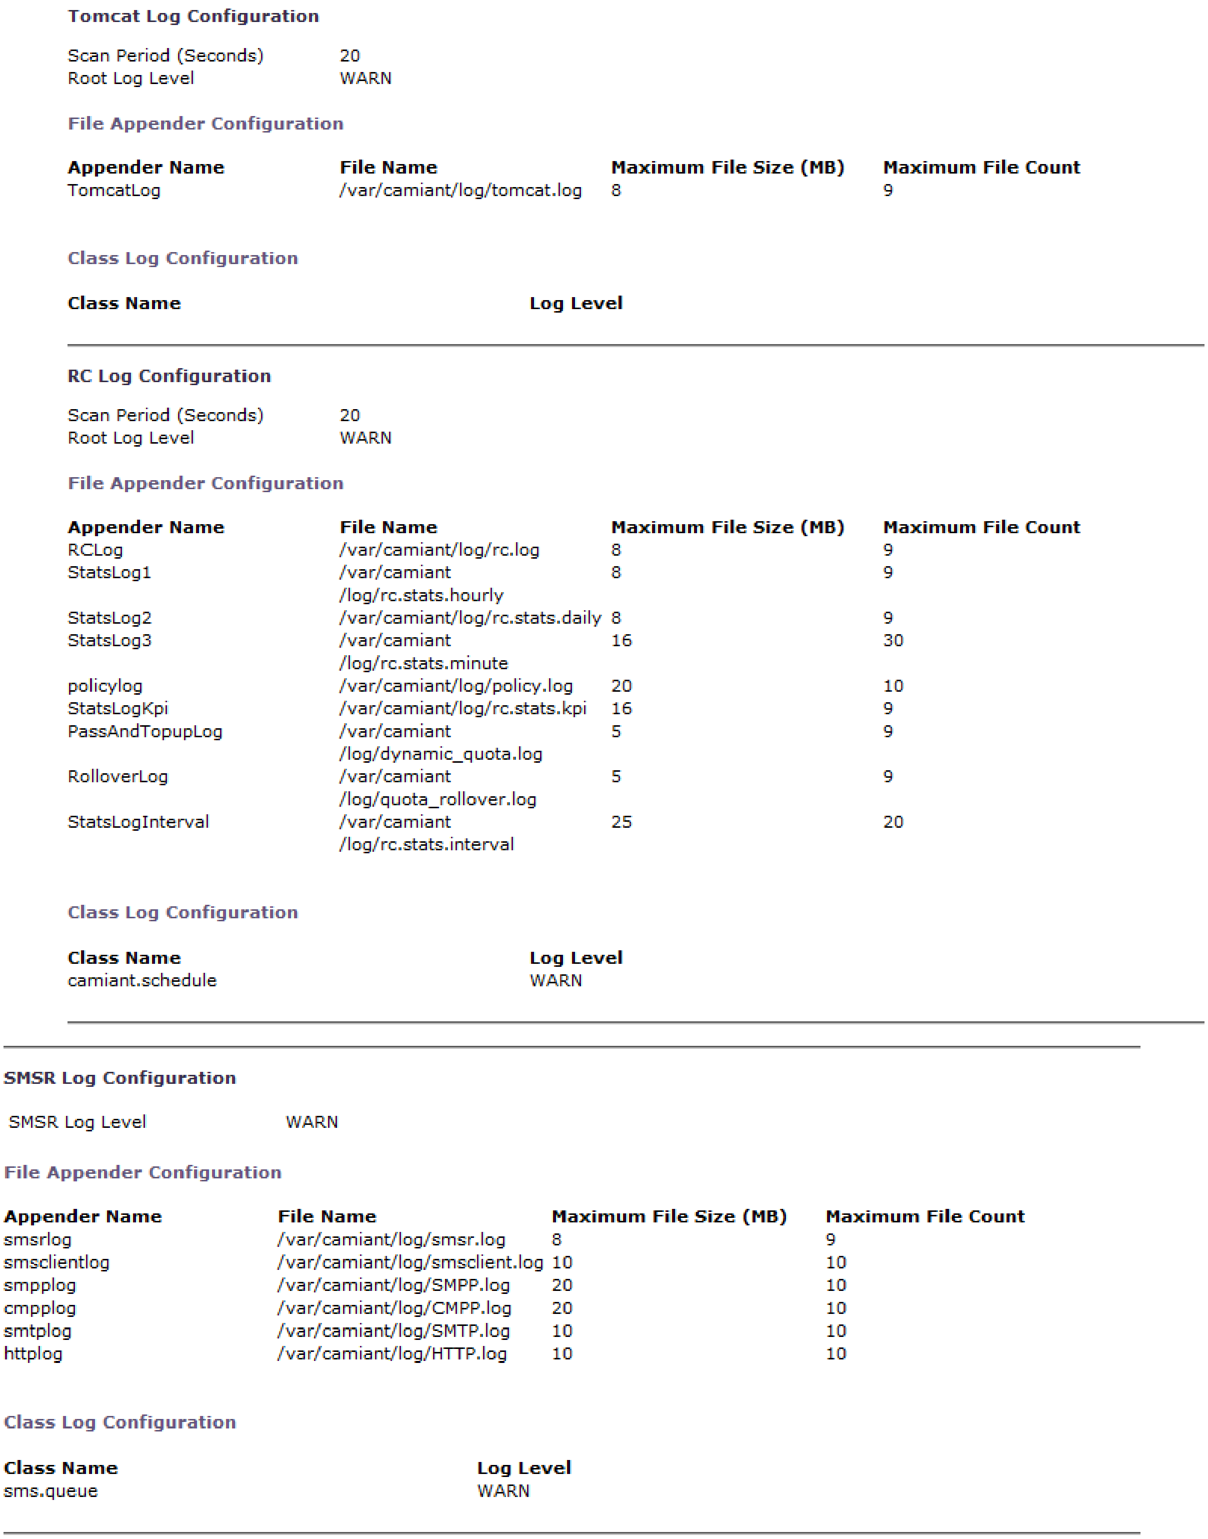

Tomcat Log Configuration

Scan Period (Seconds)	20
Root Log Level	WARN

File Appender Configuration

Appender Name	File Name	Maximum File Size (MB)	Maximum File Count
TomcatLog	/var/camiant/log/tomcat.log	8	9

Class Log Configuration

Class Name	Log Level

RC Log Configuration

Scan Period (Seconds)	20
Root Log Level	WARN

File Appender Configuration

Appender Name	File Name	Maximum File Size (MB)	Maximum File Count
RCLog	/var/camiant/log/rc.log	8	9
StatsLog1	/var/camiant /log/rc.stats.hourly	8	9
StatsLog2	/var/camiant/log/rc.stats.daily	8	9
StatsLog3	/var/camiant /log/rc.stats.minute	16	30
policylog	/var/camiant/log/policy.log	20	10
StatsLogKpi	/var/camiant/log/rc.stats.kpi	16	9
PassAndTopupLog	/var/camiant /log/dynamic_quota.log	5	9
RolloverLog	/var/camiant /log/quota_rollover.log	5	9
StatsLogInterval	/var/camiant /log/rc.stats.interval	25	20

Class Log Configuration

Class Name	Log Level
camiant.schedule	WARN

SMSR Log Configuration

SMSR Log Level	WARN

File Appender Configuration

Appender Name	File Name	Maximum File Size (MB)	Maximum File Count
smsrlog	/var/camiant/log/smsr.log	8	9
smsclientlog	/var/camiant/log/smsclient.log	10	10
smpplog	/var/camiant/log/SMPP.log	20	10
cmpplog	/var/camiant/log/CMPP.log	20	10
smtplog	/var/camiant/log/SMTP.log	10	10
httplog	/var/camiant/log/HTTP.log	10	10

Class Log Configuration

Class Name	Log Level
sms.queue	WARN

About BoD Debug Logs

Debugs logs for the BoD include log configuration for the following:

- Tomcat Log
- RC Log

Figure 6-2 BoD Debug Logs

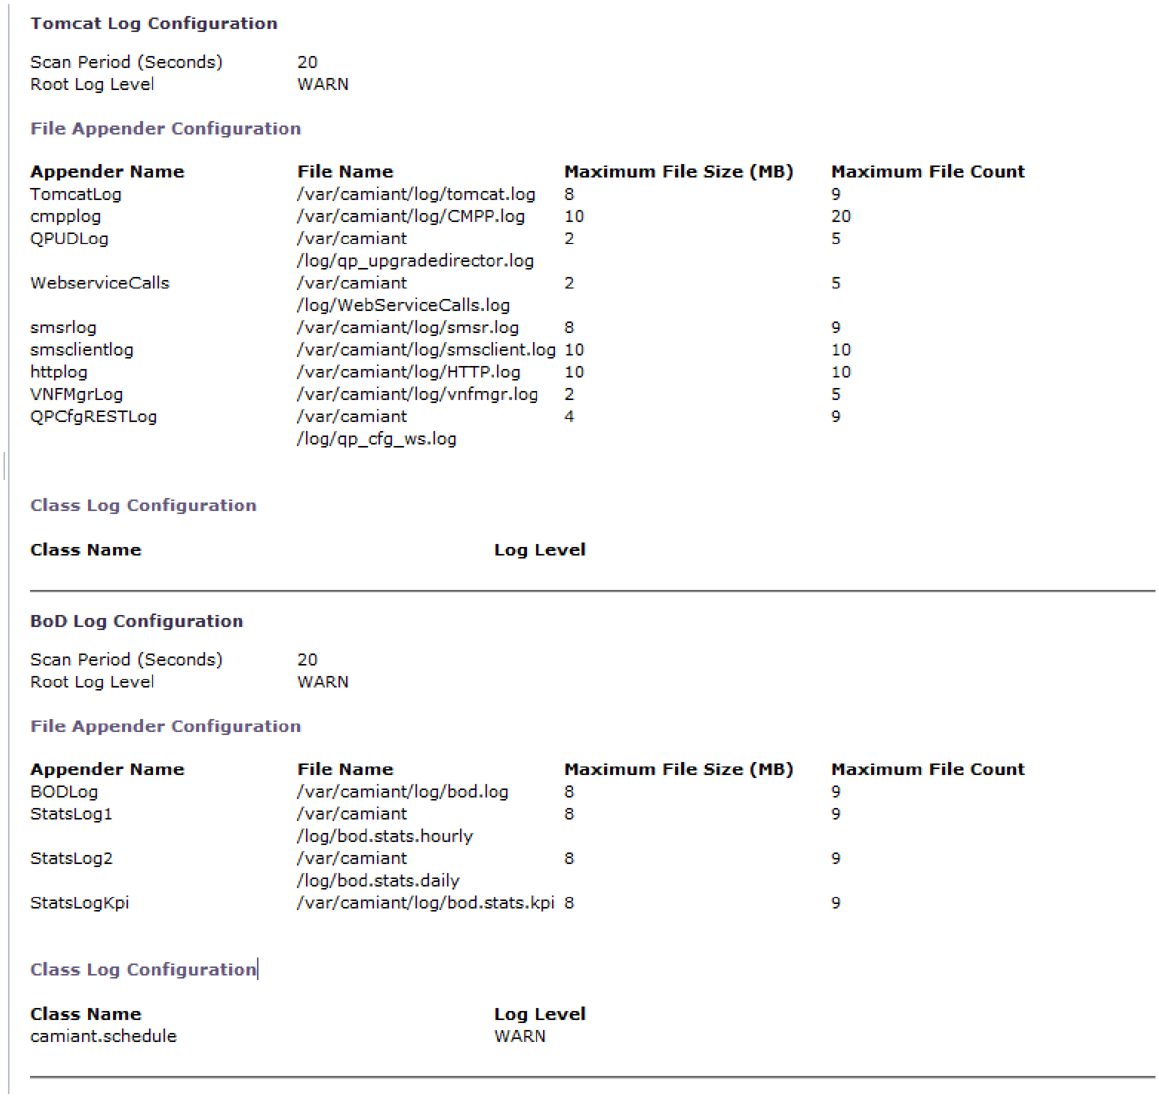

Tomcat Log Configuration

Scan Period (Seconds)	20
Root Log Level	WARN

File Appender Configuration

Appender Name	File Name	Maximum File Size (MB)	Maximum File Count
TomcatLog	/var/camiant/log/tomcat.log	8	9
cmpplog	/var/camiant/log/CMPP.log	10	20
QPUDLog	/var/camiant /log/qp_upgradedirector.log	2	5
WebserviceCalls	/var/camiant /log/WebServiceCalls.log	2	5
smsrlog	/var/camiant/log/smsr.log	8	9
smsclientlog	/var/camiant/log/smsclient.log	10	10
httplog	/var/camiant/log/HTTP.log	10	10
VNFMgrLog	/var/camiant/log/vnfmgr.log	2	5
QPCfgRESTLog	/var/camiant /log/qp_cfg_ws.log	4	9

Class Log Configuration

Class Name	Log Level

BoD Log Configuration

Scan Period (Seconds)	20
Root Log Level	WARN

File Appender Configuration

Appender Name	File Name	Maximum File Size (MB)	Maximum File Count
BODLog	/var/camiant/log/bod.log	8	9
StatsLog1	/var/camiant /log/bod.stats.hourly	8	9
StatsLog2	/var/camiant /log/bod.stats.daily	8	9
StatsLogKpi	/var/camiant/log/bod.stats.kpi	8	9

Class Log Configuration

Class Name	Log Level
camiant.schedule	WARN

Note: The MPE device does not support SMSR (SMS Relay) Log configuration.

Configuring Debug Logs for a Device

Note: While all CMP users can view debug log settings, only users with the **Administrator** user role can configure debug logs.

Note: The MPE device does not support SMSR (SMS Relay) Log configuration.

To configure debug logs for a server:

1. From the appropriate (**Policy Server** or BoD section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy device groups; the initial group is **ALL**.

2. From the content tree, select the policy device.

   The Administration page for the device opens in the work area.

3. Select the **Debug** tab.

   The device Administration page shows the configuration settings for supported components (for example, Tomcat Log, RC Log, SMSR Log, or BoD log).

4. Click **Modify**.

   The device's Administration page becomes editable.

5. In the specific **Log Configuration** section (for example, Tomcat Log, RC Log, SMSR Log, or BoD log) enter a value if you want to override the default**Scan Period (seconds)**.

   The default value is 20 seconds.

   Note: The BoD device does not support SMSR (SMS Relay) Log configuration.

6. To override the default value, select the **Root Log Level** from the list.

   Available options are:

   - **Off**—Turns off logging.
   - **Error**—Designates error events which may or may not be fatal to the application.
   - **Warn** (default)—Designates potentially harmful situations.
   - **Info—**Designates informational messages highlighting overall progress of the application.
   - **Debug—**Designates information events of lower importance.
   - **Trace**—Designates informational events of very low importance.
   - **All**—Records all logging levels.

7. To add a **Class Name**, in the **Class Log Configuration** section:

   a. Click **Add Row**.

      A confirmation message appears.

   b. Click **OK**.

   c. Enter the **Class Name**.

      For example, `example.schedule`.

d. Select the **Log Level** from list.

> **Note:** The CMP server performs no input validation on the entered **Class Name** or whether the **Class Name** belongs to the particular component.

8. Repeat for each debug component you wish to configure a **Class Name** level log.
9. Click **Save**.

The settings are applied to the selected device.

Deleting a Class Name from a Log File

> **Note:** While all CMP users can view debug log settings, only users with the **Administrator** user role can configure debug logs.

To delete a per class log for a component:

1. From the appropriate (**Policy Server** or BoD section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy device groups; the initial group is **ALL**.

2. From the content tree, select the policy device.

   The Administration page for the device opens in the work area.

3. Select the **Debug** tab.

   The device Administration page shows the configuration settings.

4. Click **Modify**.

   The Modify page becomes editable.

5. To delete a **Class Name**:

   a. In the Modify page, locate the **Class Name** you want to delete.

   b. Click **Delete** next to the **Log Level**.

   > **Note:** The Class Names `camiant.schedule` and `SMS.queue` are default class names and cannot be deleted.

6. Repeat for each debug component from which you wish to delete a **Class Name**.
7. Click **Save**.

The **Class Name** and **Log Level** are deleted from the selected device.

7

Managing Network Elements

This chapter describes how to define network elements within the CMP system.

Network elements are the devices, servers, or functions within your network with which Policy Management systems interact.

About Network Elements

A network element is a high-level device, server, or other entity within your network for which you would like to use an MPE device to manage Quality of Service (QoS). Examples include the following:

- Cable modem termination system (CMTS)
- Packet-switched data network (PSDN)
- gateway GPRS support node (GGSN)
- Broadband remote access server (B-RAS)
- Router
- Server
- Zone

After you have defined a network element in the CMP database, you associate it with the MPE device that you will use to manage that element.

There are also lower-level entities within the network that the MPE device manages that are not considered network elements. These are sub-elements, such as a channel within a CMTS or an interface on a router, or devices that are connected directly to network elements, such as a cable modem connected to a CMTS. Typically, there is no need to define these lower-level entities, because when a network element is associated with an MPE device the lower-level devices related to that network element are discovered and associated automatically.

Create a network element profile for each device you are associating with an MPE device. After defining a network element in the CMP database, configure its protocol options. The options available depend on the network element type.

For ease of management, after you define network elements, you can combine them into network element groups.

Creating a Network Element

You must create a network element for each device associated with any of the MPE devices within the network. To create a network element:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.

   The content tree displays a list of network element groups; the initial group is **ALL**.

2. Click **Create Network Element**.

   The New Network Element page opens.

3. Enter information for the network element:

   a. (Required) **Name** — The name you assign to the network element.

      The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_). The maximum length is 250 characters.

   b. (Required) **Host Name/IP Address** — Registered domain name, or IP address in IPv4 or IPv6 format, assigned to the network element.

   c. **Backup Host Name** — Alternate address that is used if communication between the MPE device and the primary address for the network element fails.

   d. **Description/Location** — Free-form text.

      Enter up to 250 characters.

   e. (Required) **Type** — Select the type of network element.

      The supported types are:

      - **CMTS** (default)

        Cable Modem Termination System

      - **AF**

        Application Function

      - **DRA**

        Diameter Routing Agent

      Note: For more information on managing network elements, see the *Configuration Management Platform Wireless User's Guide*.

   f. **SNMP Read Community String** — A password-like field that allows read-only access to the MIBs for the network element that are used for SNMP polling.

      If a value is not entered, SNMP data is not collected from this network element.

   g. **Capacity** — The bandwidth allocated to this network element.

4. In **Policy Servers associated with this Network Element**, select one or more policy servers (MPE devices) to associate with this network element.

5. In **Network Element Groups which contain this Network Element**, select one or more groups (see Adding a Network Element to a Network Element Group).

6. Click **Save**.

You have created the definition for a network element and the network element is listed on the Network Element Administration page.

Note: After saving the new network element, the CMP server automatically discovers all associated subnets. The CMP supports automatically provisioning to the associated MPE-R and MPE-S devices.

Configuring Options for Network Elements

The following sections describe how to configure options for a given network element type. The network element types available depend on the operating mode in which your CMP system is configured, and may differ from the list given here.

Note: Configuration changes made in the CMP system could potentially be reverted on an MPE device if the scheduled run time of the OSSI Distributor task on the Management Agent is before the scheduled rule time for the CMP system. The discrepancy is resolved when the OSSI Distributor Task runs on the CMP system. See Managing Scheduled Tasks for more information.

Configuring a CMTS Network Element

To configure options for a CMTS network element:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.

   The content tree displays a list of network element groups; the initial group is **ALL**.

2. Select a CMTS network element from the content tree.

   The Network Element Administration page opens in the work area.

3. Select the **CMTS** tab and then click **Modify**.

   The Modify Network Element page opens.

4. Configure the following information:

   a. **Configuration**

      - **PCMM Enabled**— Indicates whether the CMTS supports PCMM. If this feature is enabled, the MPE device establishes a PCMM connection to the CMTS.

        Disabling this feature invokes a special MPE feature called Camiant Admission Control (CAC) for this CMTS. When CAC mode is turned on for a CMTS, if the MPE device receives any PCMM messages that should be sent to the CMTS, the MPE device generates simulated responses for those messages rather than rejecting them.

   b. **Subnets**

      - **Subnets Configured Manually** — Within this field you can add or delete subnets.

      - **Subnets Discovered via SNMP** — This read-only field displays subnets that were discovered using SNMP. If additional subnets need to be added, you can add them using the **Subnets Configured Manually** field.

        Click **Rediscover** to update the list.

- **Subnets Obtained from the OSS —** This read-only field displays subnets that were imported via the OSS interface to the CMP server.

c. **Service Classes**

- **Service Classes Discovered via SNMP —** This read-only field displays service classes that were discovered using SNMP.

  Click **Rediscover** to update the list.

5. Click **Save**.

The CMTS device is configured.

Finding a Network Element

The **Search** function lets you find a specific network element within a large configuration. You can also use the function to locate all of the Cable Modem Termination Systems (CMTS) and MPE devices associated with a specified subscriber IP address or subnets. To use the network element search function:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.

   The content tree displays a list of network element groups; the initial group is **ALL**.

2. From the content tree, select **ALL**.

   The Network Element Administration page opens in the work area.

3. Click **Search**.

   The Network Element Search Criteria window opens.

4. Enter the search criteria:

   - **Name**

     The name assigned to the network element.

   - **Host Name/IP Address**

     The domain name or IP address in IPv4 or IPv6 format of the network element.

   - **Description**

     The information pertaining to the network element that helps identify it within the network. Enter up to 250 characters.

     Note: Searches are not case sensitive. You can use the wildcard characters * and ?.

   - **Subnets**

     The subnet and mask of the network element.

   If a subscriber IP address is entered with a mask code (up to 32 for IPv4, or up to 128 for IPv6), then the associated CMTS and MPE device is displayed. If the mask is left blank, then the input IP subnet is treated as an IP address, and the mask code is set automatically to 32 for IPv4 or 128 for IPv6.

5. After entering search criteria, click **Search**.

The Search Results page opens in the work area, displaying the results of the search. The last search results are held in a **Search Results** folder in the content tree until you close the Search Results page.

Modifying a Network Element

To modify a network element:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.

   The content tree displays a list of network element groups; the initial group is **ALL**.

2. From the content tree, select the network element.

   The Network Element Administration page opens in the work area.

3. Click **Modify**.

   The Modify Network Element page opens.

4. Modify the network element information.

   For a description of the fields contained on this page, see Creating a Network Element.

5. Click **Save**.

The network element definition is modified.

Deleting Network Elements

Deleting a network element definition removes it from the list of items that a Policy Management device can support. To delete a network element definition, delete it from the **ALL** group. Deleting a network element from the **ALL** group also deletes it from every group with which it is associated.

To delete a network element:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.

   The content tree displays a list of network element groups; the initial group is **ALL**.

2. From the content tree, select the **ALL** group.

   The Network Element Administration page opens in the work area, displaying all defined network elements.

3. From the work area, click 🗑 (trash can icon) located to the right of the network element.

   A confirmation message displays.

4. Click **OK**.

You have deleted the network element definition.

Deleting Multiple Network Elements

A large network can contain a great many network elements. To perform a bulk delete of network element definitions:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.

   The content tree displays a list of network element groups; the initial group is **ALL**.

2. From the content tree, select **ALL**.

   The Network Element Administration page opens in the work area.

3. Click **Bulk Delete**.

   The Bulk Delete Network Elements page opens.

4. Select the network elements or network element groups to delete.

5. (Optional) Filter the search by entering a search pattern (for example, `cm*`) and click **Filter**.

   By default, the **Search Pattern** entry box contains an asterisk (*) to match all network elements.

6. Click **Bulk Delete**.

   A confirmation message displays.

7. Click **OK**.

The selected network element or group definitions are deleted from the CMP database and all associated MPE devices.

Associating a Network Element with an MPE Device

To associate a network element with an MPE device:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy server groups; the initial group is **ALL**.

2. From the content tree, select the MPE device.

   The Policy Server Administration page opens in the work area.

3. Select the **Policy Server** tab.

   The Associations section lists the network elements associated with the MPE device.

4. Click **Modify**.

   The Modify Policy Server page opens.

5. To the right of the list of network elements in the Associations section, click **Manage**.

The Select Network Elements window opens.

For example:

Figure 7-1 Select Network Elements

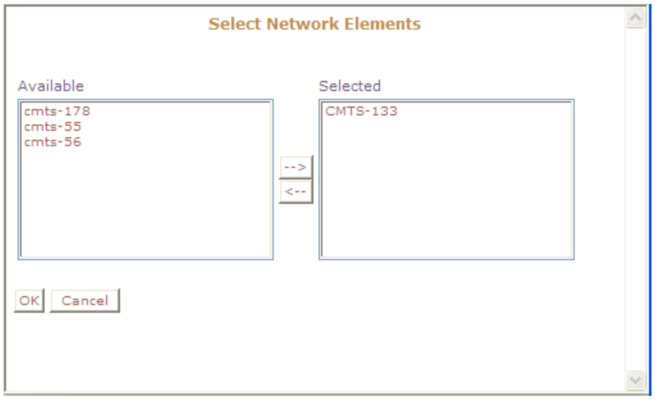

6. Select the network elements in the **Available** list and click **-->**.
7. Click **OK**.

   The selected network elements are added to the list of network elements managed by this MPE device.
8. To associate a network element group with the MPE device, select the group from the list of network element groups located under **Associations**.
9. Click **Save**.

The network element is associated with this MPE device.

Working with Network Element Groups

For organizational purposes, you can aggregate the network elements in your network into groups. For example, you can use groups to define authorization scopes or geographic areas. You can then perform operations on all the network elements in a group with a single action.

Creating a Network Element Group

Network element groups exist in a distributed network to perform specific duties.

Use this procedure if you are creating a network element group to perform specific functions in your distributed network. After you create a network group, you can then create network elements to associate with devices such as an MPE.

To create a network element group:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.

   The content tree displays a list of network element groups; the initial group is **ALL**.
2. From the content tree, select the **ALL** group.

   The Network Element Administration page opens in the work area.
3. Click **Create Group**.

The **Create Group** page opens.

4. Enter the name of the new network element group.

   The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_). The maximum length is 250 characters.

5. Enter a text description and location of the network group.

6. Click **Save**.

You have created a network element group.

Adding a Network Element to a Network Element Group

After a network element group is created, you can add individual network elements to the group.

To add a network element to a network element group:

1. From the **Network** section of the navigation pane, select **Network Elements**.

   The content tree displays a list of network element groups; the initial group is **ALL**.

2. From the content tree, select the network element group.

   The Network Element Administration page opens in the work area, displaying the contents of the selected network element group.

3. Click **Add Network Element**.

   The Add Network Elements page opens. The page supports both small and large networks, as follows:

   - If there are 25 or fewer network elements defined, the page displays the network elements not already part of the group.

   - If there are more than 25 network elements defined, the page does not display any elements. Instead, use the **Search Pattern** field to filter the list. Enter an asterisk (*) to generate a global search, or a search pattern to locate only those network elements whose name matches the pattern (for example, `star*`, `*pGw`, or `*-*`). When you have defined a search string, click **Filter**; the page displays the filtered list.

4. Select the network element you want to add. Use the Ctrl or Shift keys to select multiple network elements.

   You can also add previously defined groups of network elements by selecting those groups.

5. Click **Save**.

The network element is added to the network element group.

Creating a Network Element Sub-group

You can create sub-groups to further organize your network element network. To add a network element sub-group to an existing network element group:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.

   The content tree displays a list of network element groups; the initial group is **ALL**.

2. From the content tree, select the network element group.

   The Network Element Administration page opens in the work area, displaying the contents of the selected network element group.

3. Click **Create Sub-Group**.

   The Create Group page opens.

4. Enter the name of the new sub-group.

   The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_).

5. Enter a text description of the sub-group.

6. Click **Save**.

The sub-group is added to the selected group, and now appears in the listing.

Deleting a Network Element from a Network Element Group

Removing a network element from a network element group or sub-group does not delete the network element from the **ALL** group, so it can be used again if needed. Removing a network element from the **ALL** group removes it from all other groups and sub-groups.

To remove a network element from a network element group or sub-group:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.

   The content tree displays a list of network element groups; the initial group is **ALL**.

2. From the content tree, select the network element group or sub-group.

   The Network Element Administration page opens in the work area, displaying the contents of the selected network element group or sub-group.

3. Remove the network element using one of the following methods:

   - On the Network Element Administration page, click the 🗑 (trash can) icon, located to the right to the network element.

   - From the content tree, select the network element; the Network Element Administration page opens. Click the **System** tab and then click **Remove**.

   A confirmation message appears.

4. Click **OK**.

The network element is removed from the group or sub-group.

Modifying a Network Element Group or Sub-Group

To modify a network element group or sub-group:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.

   The content tree displays a list of network element groups; the initial group is **ALL**.

2. From the content tree, select the network element group or sub-group.

   The Network Element Administration page opens in the work area.

3. Click **Modify**.

   The Modify Group page opens.

4. Modify the name, description, or both.

5. Click **Save**.

The group or sub-group is modified.

Deleting a Network Element Group or Sub-group

Deleting a network element group also deletes any associated sub-groups. However, any network elements associated with the deleted groups or sub-groups remain in the **ALL** group, from which they can be used again if needed. You cannot delete the **ALL** group.

To delete a network element group or sub-group:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.

   The content tree displays a list of network element groups.

2. From the content tree, select the network element group or sub-group.

   The Network Element Administration page opens in the work area, displaying the contents of the selected network element group or sub-group.

3. Click **Delete**.

   A confirmation message displays.

4. Click **OK** to delete the group.

The network element group or sub-group is deleted.

8

Managing Record Keeping Servers

This chapter defines how to use the CMP system to configure and manage record keeping servers (RKSs) that receive event messages.

About Record Keeping Servers

A Record Keeping Server (RKS) is a repository for PacketCable Multimedia (PCMM) event messages. It gathers billing event messages and passes them on to back-office support systems (BOSS). To use event messaging, you must configure profiles for one or more RKSs, and then associate them with MPE devices, either by adding them to the Record Keeping Server List for the MPE, or by defining one as the default Record Keeping Server.

When configuring a RKS, note that a single RKS can correspond to a single external server, but it can also correspond to a pair of external servers. This depends on how the RKS handles failover situations.

A Record Keeping Server is uniquely identified by the following:

- Primary IP Address
- Primary Port
- Secondary IP Address
- Secondary Port

If you have a single server that provides both a primary and secondary address, you can configure it as a single RKS. If you have two servers, each of which only provides a single IP address/port, then you could either configure both of them as a single RKS (that acts as a backup pair) or you could configure them as two separate RKSs, each with a primary address/port and no secondary address/port. However, if an RKS does not have a secondary address/port, then that RKS will not be able to participate in the RKS failover mechanism as defined in the PCMM specification.

Creating a Record Keeping Server Profile

To configure an Record Keeping Server profile, complete the following:

1. From the **Policy Server** section of the navigation pane, select **Record Keeping Servers**.

   The content tree displays the **Record Keeping Servers** group.

2. Select the **Record Keeping Servers** group.

   The Record Keeping Server Administration page opens in the work area.

3. Click **Create Record Keeping Server**.

   The New Record Keeping Server page opens.

4. Enter the following information:

   a. **Name** — The name assigned to the Record Keeping Server profile.

   b. **Description/Location** (optional) — Information about the Record Keeping Server that helps identify it within the network or location.

   c. **Primary Address** — IP address, in IPv4 or IPv6 format, of the primary Record Keeping Server.

   d. **Primary Port** — IP port number of the primary Record Keeping Server. (The port number is typically 1813.)

   e. **Secondary Address** (optional) — IP address of the secondary Record Keeping Server.

   f. **Secondary Port** — IP port number of the secondary Record Keeping Server. (The port number is typically 1813.)

5. Click **Save**.

The Record Keeping Server profile is created.

Modifying a Record Keeping Server Profile

To modify a Record Keeping Server profile:

1. From the **Policy Server** section of the navigation pane, select **Record Keeping Servers**.

   The content tree displays the **Record Keeping Servers** group.

2. Select the **Record Keeping Servers** group.

   The Record Keeping Server Administration page opens in the work area, displaying the list of defined record keeping servers.

3. Select the Record Keeping Server.

   Configuration information for the Record Keeping Server displays.

4. Click **Modify**.

   The Modify Record Keeping Server page opens.

5. Modify configuration information.

6. Click **Save**.

The Record Keeping Server profile is modified.

Deleting a Record Keeping Server Profile

To delete a Record Keeping Server profile:

1. From the **Policy Server** section of the navigation pane, select **Record Keeping Servers**.

   The content tree displays the **Record Keeping Servers** group.

2. Select the **Record Keeping Servers** group.

   The Record Keeping Server Administration page opens in the work area, displaying the list of defined record keeping servers.

3. Delete the Record Keeping Server profile using one of the following methods:

   a. Select the **ALL** folder. A list of BoD servers appears in the work area. Click the 🗑 (trash can) icon next the server that you want to delete.

   b. Select a BoD server from the content tree. Click **Delete** on the BoD Administration page.

   - Click the 🗑 (trash can) icon next the server that you want to delete.

   - Select the profile from the content tree. Click **Delete** on the Record Keeping Server Administration page.

   A confirmation message appears.

4. Click **OK** to delete the Record Keeping Server profile.

The Record Keeping Server profile is deleted.

9

Managing Event Messaging

This chapter defines how to use the CMP system to configure and manage event messaging.

About Event Messaging

Event messaging is the standard mechanism by which an external server can be notified when certain PCMM events occur. The external server is referred to as a record keeping server (RKS). The RKS correlates event messages to derive call detail records (CDRs), service billing information, network resource usage patterns, capacity planning, and so on.

Note: Most of the behaviors described in this chapter are standard behaviors defined in PCMM specification PKT-SP-MM-I03. For more specific details on the algorithms or protocols involved in event messaging, refer to the PCMM specification.

In the PCMM architecture, event messages can be sent from a policy server or a CMTS. A CMTS sends event messages only when instructed to do so by the MPE device (using signaling that is part of the PCMM protocol). This is determined on a per-gate basis. The MPE device only instructs the CMTS to send event messages for gates for which it is also sending event messages.

An application manager does not send any event messages, but it can request the MPE device to send them for any gates that it creates. This is accomplished by including a special object (called an Event Generation Info object) with the gate creation request.

The MPE device uses an algorithm to determine if it should send event messages. As mentioned previously, this algorithm also determines whether the MPE device will instruct the CMTS to send event messages. The algorithm is as follows:

1. If event messaging support is disabled, then no messages are sent.
2. If the required event messaging attributes are not configured, then no messages are sent. The required attributes are the Financial Entity ID (FEID) Domain and the Element ID.
3. If the application manager has included an Event Generation Info object with a gate creation request, the contents of that object are examined:
   - If the object refers to an RKS that is configured on the MPE device, the event messages are sent to that RKS for all operations performed on that gate.
   - If the object refers to an RKS that is not configured on the MPE device, then it is ignored.

4. If a default RKS is configured on the MPE device, then event messages are sent to the default RKS for all operations on that gate. If not, no event messages are sent.

If you want to ensure that event messages are sent for every operation that is performed, then configure a default RKS. However, there is one important limitation to this type of configuration.

When an application manager requests event messages to be sent as part of that request, it includes a piece of information called the Billing Correlation ID (BCID). The purpose of the BCID is to make it easier for the RKS to correlate events that are associated with the same application session. Since this is initiated from the application manager, it can use the same BCID to associate events for multiple gates together. Since most applications use multiple gates for a single application session, this is a very desirable feature.

When event messages are generated by the MPE device using a default RKS, there is no BCID that is available from the application manager. In this situation, the MPE device generates a unique BCID for each gate. Consequently, it is not possible to correlate multiple gates together when using this type of event messaging configuration.

MPE device support of event messaging is configured in the CMP by a set of attributes. Each of these attributes is set either globally (shared by all MPE devices) or per MPE device. You can configure an attribute globally and then override it for a specific MPE device.

Configuring Global Settings for Event Messaging

Before you can configure global event messaging settings, you need to define record keeping servers (RKSs). For more information, see Managing Record Keeping Servers.

To configure global event messaging settings:

1. From the **Policy Server** section of the navigation pane, select **Event Messaging**.

   The Event Messaging Administration page opens, displaying the current global settings.

2. Click **Modify**.

   The Modify Event Messaging page opens. Figure 9-1 shows an example.

3. Configure the attributes as follows:

   a. **Enable** — If selected, event messages can be sent from the MPE device (depending on the algorithm described earlier). If not selected, event messages are not sent.

   b. **FEID Prefix (hex)** — The 8-byte hexadecimal prefix used in the FEID in event messages.

      As defined in the PCMM specification, the first 8 bytes of the FEID constitute operator-defined data. If this value is not defined, these bytes are zero-filled.

   c. **FEID Domain** — The domain name used in the FEID in event messages.

      As defined in the PCMM specification, this is the domain name for the multiple-service operator (MSO), which uniquely identifies the operator for billing and settlement purposes. This domain name is limited to 239 characters.

d. **Record Keeping Server List** — The list of configured RKSs.

   If you are configuring event messaging in your network so that the AM devices request event messages, then configure the same RKSs in both the AM devices and the MPE devices.

e. **Default Record Keeping Server** — Defines the default RKS for event messaging.

4. Click **Save**.

The global event message settings are defined.

Figure 9-1 Modify Event Messaging Page

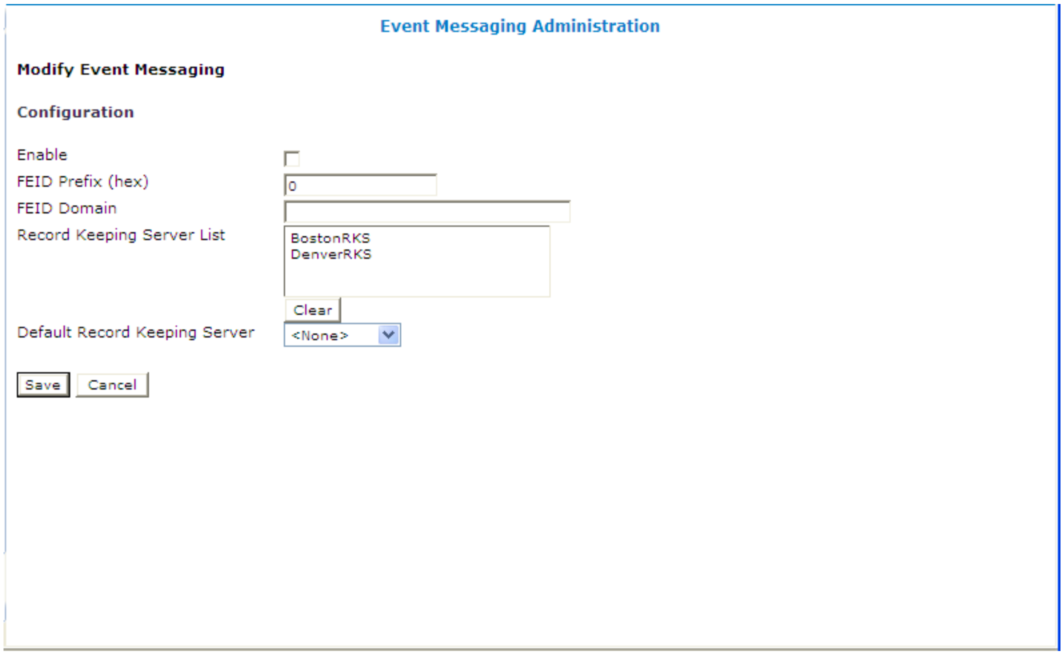

Event Messaging Administration

Modify Event Messaging

Configuration

Enable ☐

FEID Prefix (hex): 0

FEID Domain:

Record Keeping Server List: BostonRKS, DenverRKS

Default Record Keeping Server: <None>

Configuring Event Messaging for an MPE

The CMP system lets you configure how event messages are handled for a specific MPE device.

Note: MPE device event messaging settings override global CMP settings.

To configure event message settings for an MPE device:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.

   The content tree displays a list of policy server groups; the default group is **ALL**.

2. From the content tree, select the MPE device.

   The Policy Server Administration page opens in the work area.

3. Select the **EM** tab.

   The current event messaging settings for the MPE device display.

4. Click **Modify**.

   The Modify Event Messaging page opens.

5. Enter the **Element ID**.

   This attribute is set for each MPE device. The Element ID identifies event messages sent from this MPE device. This is a 5-digit value (between 0 and 99999) that must be unique within the network among all elements that send event messages. Therefore, this value must be unique among all MPE and CMTS devices within your network.

6. Select to **Enable** event messaging.

   Indicates whether event messaging is enabled.

   - If this value is set to **true**, event messages can be sent from the MPE device.
   - If this value is set to **false**, event messages are not sent.

7. Enter the **FEID Prefix (hex)**

   This is the 8-byte prefix used in the FEID in event messages. As defined in the PCMM specification, the first 8 bytes of the FEID constitute operator-defined data. If this value is not defined, these bytes are filled with zeros.

8. Enter the **FEID Domain**.

   This is the domain name used in the FEID in event messages. As defined in the PCMM specification, this is the domain name for the MSO, which uniquely identifies the operator for billing and settlement purposes. This domain name is limited to 239 characters.

9. Select the **Record Keeping Server List** from the list.

   The list of configured Record Keeping Servers. If you are configuring event messaging in your network so that the application managers request event messages, then configure the same Record Keeping Servers in the application managers and the MPE device.

10. Select the **Default Record Keeping Server** from the list.

    This defines the default Record Keeping Server for event messaging.

11. Click **Save**.

Local settings are defined for this MPE device.

10

Managing Management Agent Servers

This chapter describes how to use the CMP system to configure and manage a Management Agent (MA) server.

About Management Agent Servers

The Management Agent server is designed specifically for network architectures that require a distributed topology and collection framework. The Management Agent server is not an actively managed device, but rather a distributed system that collects topology and network information for use with PCMM message routing and policy decisions.

The Management Agent server sits between the CMP system and one or more MPE devices. The number of Management Agent servers and MPE devices depends on the size of the network. The groupings that define the MPE devices are managed by a Management Agent server and the Management Agent servers are managed by the CMP system based on the network topology.

Using the Management Agent server provides the following primary benefits:

- A distributed framework, allowing the complete system to segment and process data in a parallel fashion.
- A reduction in the management traffic across the backbone network.

All communication between the CMP system and the Management Agent server is initiated by the CMP system, and optionally is performed over a secured interface.

Creating a Management Agent Profile

To create an Management Agent profile:

1. From the **Policy Server** section of the navigation pane, select **Management Agents**.

   The content tree displays the **Management Agents** group.

2. Select the **Management Agents** group.

   The Management Agent Administration page opens in the work area.

3. Click **Create Management Agent**.

   The New Management Agent page opens.

4. Enter the following profile information:

   a. **Associated Cluster** — Select the cluster from the list.

   b. **Name** — The name assigned to the Management Agent.

   c. **Description/Location** — Free-form text that defines the function or location of the Management Agent.

   d. **Secure Connection** — Designates whether or not to use SSL as a secure connection for this Management Agent.

5. Click **Save**.

The Management Agent profile is created and added to the list of available profiles.

Modifying a Management Agent Profile

To modify a management agent profile:

1. From the **Policy Server** section of the navigation pane, select **Management Agents**.

   The Management Agent Administration page opens in the work area.

2. From the content tree, select the management agent.

   The management agent is displayed in the Management Agent Administration page.

3. Click **Modify**.

   The Modify System Settings page opens.

4. Edit the profile information. See Creating a Management Agent Profile for descriptions of these fields.

5. Click **Save**.

The management agent profile is modified.

Deleting a Management Agent Profile

To delete a management agent profile:

1. From the **Policy Server** section of the navigation pane, select **Management Agents**.

   The Management Agent Administration page opens in the work area.

2. Use one of the following methods to select the management agent profile to delete:

   - From the work area, click 🗑 (trash can icon) located to the right of the policy.

   - From the policy group tree, select the policy. The management agent displays in the Management Agent Administration page. Click **Delete**.

   A confirmation message displays.

3. Click **OK** to delete the management agent.

The management agent profile is deleted.

Reapplying a Management Agent Profile Configuration

To reapply a configuration to a management agent server:

1. From the **Policy Server** section of the navigation pane, select **Management Agents**.

   The Management Agent Administration page opens in the work area.

2. Select the management agent you want to reconfigure.

   The management agent is displayed in the Management Agent Administration page.

3. Click **Reapply Configuration**.

   The management agent profile information is pushed to the management agent server.

> **Note:** The Reapply Configuration process can take up to 30 minutes. However, this process runs in the background and allows you to continue to use the CMP system, with the exception of the Management Agent feature.

Management Agent Tasks

A set of configurable management agent tasks collect and distribute data:

- **Subnet SNMP Collector —** Collects all subnet information residing on the CMTS devices by polling, using SNMP, all CMTS devices for all subnets and then updates the MA with these subnets.

- **Service Class SNMP Collector —** Collects all service class information residing on the CMTS devices by polling, using SNMP, all CMTS devices for all service class information and then updates the MA with this information.

- **Subscriber SNMP Collector** — Uses SNMP to poll the CMTS devices for their subscriber topology data (such as CPE IPs, CM MACs, and channel data) and then updates the MA with this information.

- **CMTS Distributor** — Distributes CMTS, Subnet, and Service Class data to the MPE devices.

- **Subscriber Distributor —** Reads the subscriber topology data from the MA database and distributes it to the appropriate MPE devices.

Managing Management Agent Tasks

To view the current Management Agent task status and the current scheduled data processing:

1. From the **Policy Server** section of the navigation pane, select **Management Agents**.

   The content tree displays the **Management Agents** group.

2. From the content tree, select the management agent.

   The management agent is displayed in the Management Agent Administration page.

3. Select the **Tasks** tab.

   The configurable tasks are displayed.

In the Status column of the display, Success* means that the task last ran successfully and is scheduled to run again. A value of Success means that the task last ran successfully, but is not currently scheduled to run again. For more information about configuring scheduled tasks, see Managing Scheduled Tasks.

Viewing Details of a Management Agent Task

To view the details of a Management Agent task:

1. From the **Policy Server** section of the navigation pane, select **Management Agents**.

   The content tree displays the **Management Agents** group.

2. From the content tree, select the management agent.

   The management agent is displayed in the Management Agent Administration page.

3. Select the **Tasks** tab.

   The configurable tasks are displayed.

4. Click the task name.

The details of the Management Agent task displays.

Rescheduling a Task on a Management Agent

To reschedule a task on a Management Agent:

1. From the **Policy Server** section of the navigation pane, select **Management Agents**.

   The content tree displays the **Management Agents** group.

2. From the content tree, select the Management Agent.

   The management agent is displayed in the Management Agent Administration page.

3. Select the **Tasks** tab.

   The configurable tasks are displayed.

4. Click the task name.

   Detailed information is displayed.

5. Click **Reschedule**.

   The Scheduled Task Administration page opens.

6. Change the run time, interval, or when the task runs or the task that this task run after.

   - To change the run time or interval, click 📅 (calendar icon), select the date and time, and then click **Enter**.

   - To change the interval, select the hours and minutes. Valid values are from 0 to 24 hours and 0 to 55 minutes (in 5-minute increments).

- To change the task that this task follows, select the task from the list.

7. Click **Save**.

The task is rescheduled

Running a Task on a Management Agent Immediately

To run a task on a Management Agent immediately:

1. From the **Policy Server** section of the navigation pane, select **Management Agents**.

   The content tree displays the **Management Agents** group.

2. From the content tree, select the management agent.

   The management agent is displayed in the Management Agent Administration page.

3. Select the **Tasks** tab.

   The configurable tasks are displayed.

4. Click the task name.

   Detailed information is displayed.

5. Click **Run Now**.

   A confirmation message displays.

6. Click **OK**.

The task runs immediately.

Disabling a Task on a Management Agent

Note: A task must be enabled before it can be disabled. To enable the task see Enabling a Task on a Management Agent.

To disable a task on a Management Agent:

1. From the **Policy Server** section of the navigation pane, select **Management Agents**.

   The content tree displays the **Management Agents** group.

2. From the content tree, select the Management Agent.

   The management agent is displayed in the Management Agent Administration page.

3. Select the **Tasks** tab.

   The configurable tasks are displayed.

4. Click the task name.

   Detailed information is displayed.

5. Click **Enable**.

   A confirmation message displays.

> **Note:** If the **Disable** button is not available, then the task is already enabled. To enable the task see Enabling a Task on a Management Agent.

6. Click **OK**.

The task is rescheduled

Enabling a Task on a Management Agent

> **Note:** A task must be disabled before it can be enabled. To disable the task see Disabling a Task on a Management Agent.

To enable a task on a Management Agent:

1. From the **Policy Server** section of the navigation pane, select **Management Agents**.

   The content tree displays the **Management Agents** group.

2. From the content tree, select the Management Agent.

   The management agent is displayed in the Management Agent Administration page.

3. Select the **Tasks** tab.

   The configurable tasks are displayed.

4. Click the task name.

   Detailed information is displayed.

5. Click **Enable**.

   A confirmation message displays.

> **Note:** If the **Enable** button is not available, then the task is already enabled. To disable the task see Disabling a Task on a Management Agent.

6. Click **OK**.

The task is enabled

Refreshing Task Details on a Management Agent

Refreshing a task updates the information you see in the detail of the task.

To refresh task details on a Management Agent:

1. From the **Policy Server** section of the navigation pane, select **Management Agents**.

   The content tree displays the **Management Agents** group.

2. From the content tree, select the management agent.

   The management agent is displayed in the Management Agent Administration page.

3. Select the **Tasks** tab.

   The configurable tasks are displayed.

4. Click the task name.

   Detailed information is displayed.

5. Click **Refresh**.

   The page updates with the current task details.

The task runs immediately.

11

Managing Bandwidth on Demand

This chapter describes the basic configuration for Bandwidth on Demand (BoD) devices in the CMP system. See the *Bandwidth on Demand Cable User's Guide* for more information on how to define and manage BoD devices.

Bandwidth on Demand Application Manager Overview

The Bandwidth on Demand Application Manager (BoD AM) allows applications to request the setup and tear down of dynamic Quality of Service (QoS) resources within a broadband network, providing the necessary bandwidth and priority to enhance the subscriber's experience.

The BoD AM is designed to provide a simplified and abstract interface that converts 3GPP non-compliant commands into dynamic service requests that are used by an MPE device. This interface uses HyperText Transfer Protocol (HTTP) and Simple Object Access Protocol (SOAP) that enables the application developer to integrate dynamic QoS resources into nearly any application development environment.

Additionally, the BoD AM maintains and manages all of the state information that is associated with each request, allowing applications to be stateless in their operation.

The BoD AM presents a SOAP based Remote Procedure Call (RPC) interface and a pure HTTP request interface. These interfaces provide similar functionality and are designed to allow application developers to use the interface that best suits their application.

For example, the HTTP interface allows a parameterized URL to be associated with the onclick action of a turbo-button, or allows any application to embed an HTTP POST message to dynamically adjust service. Alternatively, the SOAP interface provides easy session control through an RPC mechanism. The capability to use HTTP or SOAP provides the flexibility for to use either interface depending on the developer's preference.

Within the BoD AM, you can define a number of service names that translate into a particular traffic profile. For example, a generic service name, such as turboService, could be defined with an associated best effort upstream flow and a high-priority downstream flow. Additionally, a specific service name can be defined, such as uploadService, that defines a high priority upstream flow.

Each of the interface bindings allows an application to create a new session specifying a service name and also supplying a number of specialization parameters, such as bandwidth. For example, within a web portal, a number of links or buttons can be defined that use the same turboService profile. Each link or button can specify a different upstream and downstream bandwidth. This capability can be used to vary the resulting QoS flows based either on the application context or a subscriber tier.

The BoD AM also allows calling an application to specify the duration of QoS resource allocation. The application can choose to completely manage the lifecycle of the resources, in which case it is the responsibility of the application to perform one of three actions:

- Free the resources at the appropriate time:
  - After a defined period
  - After an application has completed its function
- Tell the BoD AM to keep the resources active for a specified time
- Free up the resource if there is inactivity for a defined period

12

System-Wide Reports

This chapter describes the reports available on the function of Policy Management systems in your network. Reports can display platform alarms, network protocol events, and Policy Management application errors.

KPI Dashboard

The KPI Dashboard provides a multi-site, system-level summary of performance and operational health indicators. The display includes indicators for:

- Offered load (transaction rate)
- System capacity (counters for active sessions)
- Inter-system connectivity
- Physical resource utilization (memory, CPU)
- System status
- Alarms
- Protocol errors

The KPI Dashboard displays the indicators for all MPE KPIs in one table. Each row in the table represents a single MPE device. The table cells are rendered using a color scheme to highlight areas of concern that is well adopted by the telecommunication industry. The table contents are refreshed every 10 seconds; this time period is not configurable. The color changing thresholds are user configurable.

Figure 12-1 illustrates the dashboard’s contents.

Figure 12-1 Example of KPI Dashboard

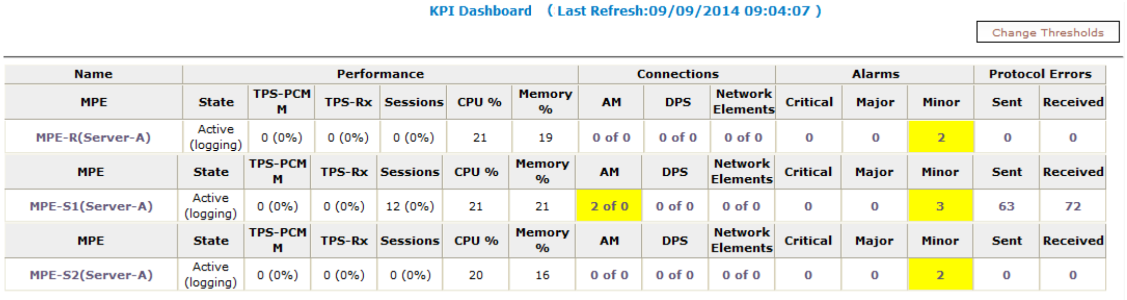

KPI Dashboard (Last Refresh:09/09/2014 09:04:07)

Name	Performance						Connections			Alarms			Protocol Errors	
MPE	State	TPS-PCM M	TPS-Rx	Sessions	CPU %	Memory %	AM	DPS	Network Elements	Critical	Major	Minor	Sent	Received
MPE-R(Server-A)	Active (logging)	0 (0%)	0 (0%)	0 (0%)	21	19	0 of 0	0 of 0	0 of 0	0	0	2	0	0
MPE	State	TPS-PCM M	TPS-Rx	Sessions	CPU %	Memory %	AM	DPS	Network Elements	Critical	Major	Minor	Sent	Received
MPE-S1(Server-A)	Active (logging)	0 (0%)	0 (0%)	12 (0%)	21	21	2 of 0	0 of 0	0 of 0	0	0	3	63	72
MPE	State	TPS-PCM M	TPS-Rx	Sessions	CPU %	Memory %	AM	DPS	Network Elements	Critical	Major	Minor	Sent	Received
MPE-S2(Server-A)	Active (logging)	0 (0%)	0 (0%)	0 (0%)	20	16	0 of 0	0 of 0	0 of 0	0	0	2	0	0

The displayed headings are:

- Name of MPE
- Performance:

  - State
  - PCMM Transactions per second (TPS-PCMM)
  - Rx Transactions per second (TPS-RX)
  - Active Sessions
  - CPU utilization percentage (%)
  - Memory utilization percentage (%)
- Connections
  - Application Managers (AMs)
  - Downstream policy servers (DPS)
  - Network Elements
- Alarms
  - Critical
  - Major
  - Minor
- Protocol Errors
  - In messages sent
  - In messages received

The **Change Thresholds** button allows you to change threshold settings used to determine cell coloring. See Configuring the KPI Threshold and Resetting the KPI Thresholds for detailed information.

Each MPE cluster has one row in the table for each server configured in the cluster:

- The first row displays information for the first server (Server-A).
- The second row displays information for the second server (Server-B), if present.
- The third row displays information for the third server (Server-C), if present.

Several of the KPI columns are not populated for the standby or spare server (because those servers are not active). The only columns that contain data are: Status, CPU %, and Memory %. For Connections, Alarms, and Protocol Errors, the column data is a hyperlink that opens a more detailed report.

If a monitored system is unreachable, or if the data is unavailable for some reason, then the state is set to “Off-line” and the values in all the associated columns are cleared. In this situation, the entire row is displayed with the error color (red). If a monitored system does not support KPI retrieval then the status is set to “N/A” and the values in all the associated columns are cleared. No coloring is applied.

The columns that display information in the form of X (Y%) (that is, “TPS” and “Sessions”) correspond to the following: X represents the actual numeric value and Y represents the % of rated system capacity that is being consumed.

The columns that display connection counts are displayed in the form "X of Y" where X is the current number of connections and Y is the configured number of connections. When X and Y are not the same, the column uses the warning color to indicate a connectivity issue, unless X is 0, in which case the error color is displayed.

The Alarm and Protocol Errors columns display the number of current events. If there are any Critical or Major alarms, then these cells will be colored red or yellow, respectively.

Note: To learn more about an alarm and how to resolve it, refer to the *Troubleshooting Reference* for this release.

Mapping Display to KPIs

#unique_267/unique_267_Connect_42_V5837369 explains how each of the columns in the KPI dashboard are mapped to a specific statistic in the KPI statistics. On the initial KPI Dashboard window, KPIs for each MPE device are shown. Since the tables contain row entries for the active, standby, and spare servers, the mapping is described for all servers.

Table 12-1 KPI Definitions for MPE Devices

KPI Dashboard Column	Mapping to Statistics	
	Active Server	**Standby or Spare Server**
Name	Not derived from statistics.	Not derived from statistics.
State	Label representation of the PrimaryServerStatus	Label representation of the SecondaryServerStatus
TPS-PCMM	CurrentPcmmTransactionsPerSecond and CurrentPcmmTPSPercentageOfCapacity	None
TPS-Rx	CurrentRxTransactionsPerSecond and CurrentRxTPSPercentageOfCapacity	None
Sessions	CurrentSessionCount and CurrentSessionPercentageOfCapacity	None
CPU %	PrimaryCPUUtilizationPercentage	SecondaryCPUUtilizationPercentage
Memory %	PrimaryMemoryUtilizationPercentage	SecondaryMemoryUtilizationPercentage
AM Connections	A value in the form "X of Y", where: X is CurrentAmConnectionCount Y is ConfiguredAMConnectionCount	None
DPS Connections	A value in the form "X of Y", where: X is CurrentDpsConnectionCount Y is ConfiguredDpsConnectionCount	None

Table 12-1 (Cont.) KPI Definitions for MPE Devices

KPI Dashboard Column	Mapping to Statistics	
Network Element Connections	A value in the form "X of Y", where: X is CurrentConnectedNECount Y is ConfiguredConnectedNECount	None
Critical Alarms	Not derived from statistics	Not derived from statistics
Major Alarms	Not derived from statistics	Not derived from statistics
Minor Alarms	Not derived from statistics	Not derived from statistics
Protocol Errors Sent	CurrentProtocolErrorSentCount	None
Protocol Errors Received	CurrentProtocolErrorReceivedCount	None

Configuring the KPI Threshold

You can customize the warning and error thresholds for KPIs.

To customize the KPI thresholds:

1. From the **System Wide Reports** section of the navigation pane, select **KPI Dashboard**.

   The KPI Dashboard page opens in the work area.

2. Click **Change Thresholds**.

   The KPI Dashboard Configuration displays.

3. Change the threshold percentages.

   You can configure warning and error thresholds for:

   - TPS
   - Sessions
   - CPU
   - Memory

4. Click **Save**.

The KPI thresholds are configured.

Resetting the KPI Thresholds

You can reset custom warning and error thresholds for KPIs to the defaults.

To reset the KPI thresholds:

1. From the **System Wide Reports** section of the navigation pane, select **KPI Dashboard**.

   The KPI Dashboard page opens in the work area.

2. Click **Change Thresholds**.

   The KPI Dashboard Configuration window displays.

3. Click **Reset**.

   The percentages change to the default values.

4. Click **Save**.

The KPI thresholds are reset.

Viewing the Trending Reports

To view the trending reports, from the **System Wide Reports** section of the navigation pane, select **Trending Reports**.

1. From the **System Wide Reports** section of the navigation pane, select **Trending Reports**.

   The Trending Reports are listed in the content tree.

2. Select the report to view.

   - **PCMM Transaction Per Second** — The number of PCMM requests and answer pairs processed in a second. For more information about the PCMM Transaction Per Second report, see Viewing PCMM Transaction Per Second.
   - **Rx Transaction Per Second** — The number of Rx requests and answer pairs processed in a second. For more information about the Rx Transaction Per Second report, see Viewing Rx Transaction Per Second.
   - **Session Count** — The maximum number of sessions per interval which were maintained over a period of time in selected or all MPE devices. For more information about the Session Count report, see Viewing Session Count.

The report opens in the work area. To create a custom report, see Custom Trending Reports.

Viewing PCMM Transaction Per Second

PCMM transactions per second is defined as the number of PCMM transactions processed in a second, graphed over time periods equal to the KPI interval length (by default 15 minutes). Transactions are recorded by the counter IntervalMaxPcmmTransactionsPerSecond.

To view the PCMM Transaction Per Second trending report:

1. From the **System Wide Reports** section of the navigation pane, select **Trending Reports**.

   The content tree displays a list of trending reports.

2. From the content tree, select **PCMM Transaction Per Second**.

   The report page displays the selected graph.

The following report options are available:

- **Refresh** — You are provided with the most recently updated graph.
- **Search Filter** — You can specify which Policy Management devices are graphed (all or specific devices) and which counters to graph (all or TPS for each class of Policy Management device). You can also specify the graph parameters:
  - **Start Date & Time** — The start date and time for the graph. Use the calendar window to select or enter the year, month, day, and time. The graph uses after the set duration.
  - **Duration** — Displays the time duration of the data. A pulldown list provides the following options:
    - **24 hours** (the default)
    - **2 days**
    - **3 days**
    - **4 days**
    - **5 days**
    - **6 days**
    - **7 days**

    Note: The durations available depend on the settings of the OM Statistics scheduled task.
  - **Show Aggregation** — If you check this box, the aggregated data for all selected devices is displayed in the graph.
- **Settings** — The table parameters are displayed; click **Run** to generate the graph.
- **Printable Format** — The most recently updated graph is displayed in a separate window.
- **View Raw Data** — The interval data statistics are displayed in a separate window.
- **Export CSV** — A comma-separated value (CSV) file named `Export_PCMM Transaction Per Second.csv` is generated, suitable for a spreadsheet application, and a standard File Download window opens, so you can save or open the file.
- **View Summary** — The distribution of data (average, minimum, and maximum) of the interval statistics for each device are displayed in a separate window.

Viewing Rx Transaction Per Second

Rx transactions per second is defined as the number of Rx transactions processed in a second, graphed over time periods equal to the KPI interval length (by default 15 minutes). Transactions are recorded by the counter CurrentRxTransactionsPerSecond.

To view the Rx Transaction Per Second trending report:

1. From the **System Wide Reports** section of the navigation pane, select **Trending Reports**.

   The content tree displays a list of trending reports.

2. From the content tree, select **Rx Transaction Per Second**.

   The report page displays the selected graph.

The following report options are available:

- **Refresh** — You are provided with the most recently updated graph.
- **Search Filter** — You can specify which Policy Management devices are graphed (all or specific devices) and which counters to graph (all or TPS for each class of Policy Management device). You can also specify the graph parameters:
  - **Start Date & Time** — The start date and time for the graph. Use the calendar window to select or enter the year, month, day, and time. The graph uses after the set duration.
  - **Duration** — Displays the time duration of the data. A pulldown list provides the following options:
    - **24 hours** (the default)
    - **2 days**
    - **3 days**
    - **4 days**
    - **5 days**
    - **6 days**
    - **7 days**

    Note: The durations available depend on the settings of the OM Statistics scheduled task.

  - **Show Aggregation** — If you check this box, the aggregated data for all selected devices is displayed in the graph.
- **Settings** — The table parameters are displayed; click **Run** to generate the graph.
- **Printable Format** — The most recently updated graph is displayed in a separate window.

- **View Raw Data** — The interval data statistics are displayed in a separate window.
- **Export CSV** — A comma-separated value (CSV) file named `Export_Rx Transaction Per Second.csv` is generated, suitable for a spreadsheet application, and a standard File Download window opens, so you can save or open the file.
- **View Summary** — The distribution of data (average, minimum, and maximum) of the interval statistics for each device are displayed in a separate window.

Viewing Session Count

The session counts determine the number of Rx or PCMM sessions maintained in the MPE device, graphed over time periods equal to the KPI interval length (by default 15 minutes). The session count is recorded by the counter MaxSessionCount.

To view the Session Count trending report:

1. From the **System Wide Reports** section of the navigation pane, select **Trending Reports**.

   The content tree displays a list of trending reports.

2. From the content tree, select **Session Count**.

   The Session Count page displays the Session Count for policy server (MPE) device graph.

The following report options are available:

- **Refresh**

  You are provided with the most recently updated graph.

- **Search Filter**

  You can specify which MPE devices are graphed (all or specific devices) and which counters to graph (all or session counters for MPE devices, which for this report is the same thing). You can also specify the graph parameters:

  – **Start Date & Time**

    The start date and time for the graph. Click 📅 (calendar icon) to select or enter the year, month, day, and time. The graph uses after the set duration.

  – **Duration**

    Displays the time duration of the data. A list provides the following options:

    - **24 hours** (default)
    - **2 days**
    - **3 days**
    - **4 days**
    - **5 days**
    - **6 days**

- **7 days**

Note: The durations available depend on the settings of the OM Statistics scheduled task.

- **Show Aggregation** — If you check this box, the aggregated data of all selected MPE content is displayed in the graph.

- **Settings**

  The table parameters are displayed; click **Run** to generate the graph.

- **Printable Format**

  The most recently updated graph is displayed in a separate window.

- **View Raw Data**

  The interval data statistics are displayed in a separate window.

- **Export CSV**

  A comma-separated value (CSV) file named `Export_Session Count.csv` is generated, suitable for a spreadsheet application, and a standard File Download window opens, so you can save or open the file.

- **View Summary**

  The distribution of data (average, minimum, and maximum) of the interval statistics for each device are displayed in a separate window.

Custom Trending Reports

Along with the pre-configured trending reports, you can create custom trending reports based on one or more counters.

The following groups of MPE statistics are available for graphing:

- DiameterAfStats
- GateStats

Within each group, a set of counters is available.

After creation, customized trending reports appear in the **Trending Reports** list following the pre-configured Trending Reports in alphabetical order.

Creating a Custom Trending Report

To create a custom trending report:

1. From the **System Wide Reports** section of the navigation pane, select **Trending Reports**.

   The Trending Report Definition Administration page opens.

2. Click **Create Trending Report Definition**.

   A new Trending Report Definition Administration page opens, containing fields for configuring a customized trending report.

   See Figure 12-2 shows a sample.

Figure 12-2 Trending Report Definition Configuration Page

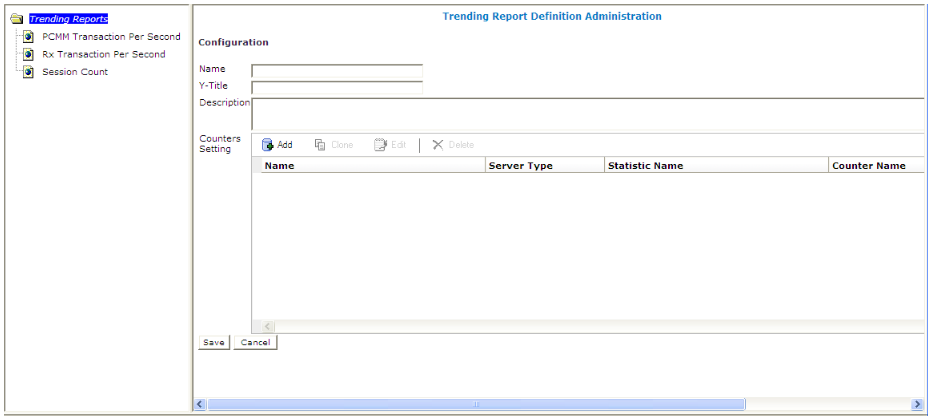

3. Enter the following information for the new trending report:

   a. **Name**—The name of the trending report.

      The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_).

   b. **Y-title**—The title of the Y series.

      The title can contain up to 40 characters and cannot begin or end with a space.

   c. **Description**—The description of the trending report.

      The description can contain up to 250 characters and cannot begin or end with a space.

4. Add counters to the report:

   a. Click **Add** next to the **Counters Setting** field.

      The Add Stats Definition popup opens.

   b. Enter a name for the counter in the **Name** field.

      The name can contain up to 40 characters, cannot contain double quotes (") or commas (,), and cannot begin or end with a space.

   c. Select the server type from the **Server Type** list.

   d. Select a statistic from the **Statistic Name** list.

      After selecting a statistic, all counters supported by that statistic populate the **Counter Name** list.

   e. Select a counter from the **Counter Name** list.

   f. Click **Save** to add the counter to the **Counters Setting** list.

      You have added a single counter to the trending report. You can continue to add individual counters to the report, using this step. You can also add counters by cloning an existing counter (described in the following step).

5. (Optional) Add, edit, or delete reports.

- Cloning an entry in the table

  a. Select an entry in the table.

  b. Click **Clone**. The Clone window opens with the information for the entry.

  c. Make changes as required.

  d. Click **Save**. The entry is added to the table.

- Editing an entry in the table

  a. Select the entry in the table.

  b. Click **Edit**. The Edit Response window opens, displaying the information for the entry.

  c. Make changes as required.

  d. Click **Save**. The entry is updated in the table.

- Deleting a value from the table

  a. Select the entry in the table.

  b. Click **Delete**. A confirmation message displays.

  c. Click **Delete** to remove the entry. The entry is removed from the table.

6. Click **Save**.

You have defined and saved a custom trending report. The custom trending report appears, in alphabetical order by name, in the list of custom trending reports.

Editing a Custom Trending Report

You can edit any of the configured information for an existing custom trending report. You can also add, edit, or delete the counters associated with the report.

To edit a custom trending report:

1. From the **System Wide Reports** section of the navigation pane, select **Trending Reports**.

   The Trending Report Definition Administration page opens.

2. Select the custom trending report.

   The report opens.

3. Click **Settings.**

   The Trending Report Definition Administration page displays for the report.

4. Click **Modify**.

   You can edit the **Name**, **Y-Title**, or **Description** of the report. You can also add, edit, or delete the counters associated with the report. See Creating a Custom Trending Report for additional information.

Deleting a Custom Trending Report

You can delete any of the existing custom trending reports. You cannot delete the pre-configured trending reports.

To delete a custom trending report:

1. From the **System Wide Reports** section of the navigation pane, select **Trending Reports**.

   The Trending Report Definition Administration page opens.

2. Select the custom trending report.

   The report opens.

3. Click **Settings.**

   The Trending Report Definition Administration page displays for the report.

4. Click **Delete**.

   A confirmation message displays.

5. Click **OK**.

You have deleted the report.

Alarm Reports

There are two types of alarm reports in the CMP system.

Active Alarms
The Active Alarms report provides an aggregate view of time stamped alarm notifications for Policy Management systems. The display is refreshed every ten seconds and appears in the upper right corner of all CMP pages. Alarms remain active until they are reset.

Alarm History
The Alarm History Report displays historical alarm information.

About Active Alarm Reports

The Active Alarms report provides an aggregate view of time stamped alarm notifications for Policy Management systems. The display is refreshed every ten seconds and appears in the upper right corner of all CMP pages. Alarms remain active until they are reset.

The alarm levels are as follows:

Critical
Service is being interrupted. (Critical alarms are displayed in red.)

Major
Service may be interrupted if the issue is not corrected. (Major alarms are displayed in orange.)

Minor
Non-service affecting fault. (Minor alarms are displayed in yellow.)

Notifications, which have a severity of Info, are not displayed in the Active Alarms report, but are written to the trace log. For more information, see About the Trace Log.

Figure 12-3 shows a sample active alarms report.

The following options are available:

- To sort the report on any column, click the column title.
- To display online help for an alarm, click on its ID.
- To pause the display of alarms, click **Pause**. To resume the display, click **Refresh**.
- To select what information is displayed, click **Columns** and select from the pulldown list.
- To control what alarms and alarm classes are displayed on the page, click **Filters** and select from the pulldown list:
  - The **Server** control lets you display alarms from all servers (the default) or a specific server.
  - The **Server Type** control lets you display alarms from all Policy Management products (the default) or just **CMP** or **MPE** systems.
  - The **Severity** control lets you display alarms of all severities (the default), critical and major alarms, critical alarms, major alarms, or minor alarms.
- To save your formatting changes to the report page, click **Save Layout**.
- **Printable Format** — The current alarms are displayed in a separate window.
- **Save as CSV** — A comma-separated value (CSV) file named `report.csv` is generated, suitable for a spreadsheet application, and a standard File Download window opens, so you can save or open the file.
- **Export PDF** — A Portable Document Format (PDF) file named `report.pdf` is generated, suitable for a spreadsheet application, and a standard File Download window opens, so you can save or open the file.

Figure 12-3   Sample Active Alarms Report

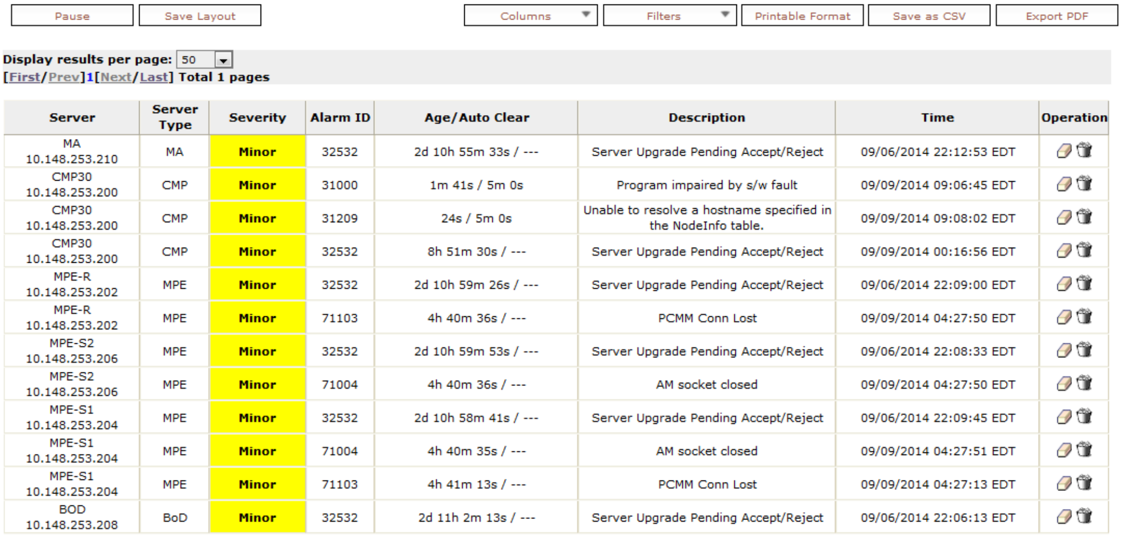

Display results per page: 50
[First/Prev]1[Next/Last] **Total 1 pages**

Server	Server Type	Severity	Alarm ID	Age/Auto Clear	Description	Time	Operation
MA 10.148.253.210	MA	Minor	32532	2d 10h 55m 33s / ---	Server Upgrade Pending Accept/Reject	09/06/2014 22:12:53 EDT	
CMP30 10.148.253.200	CMP	Minor	31000	1m 41s / 5m 0s	Program impaired by s/w fault	09/09/2014 09:06:45 EDT	
CMP30 10.148.253.200	CMP	Minor	31209	24s / 5m 0s	Unable to resolve a hostname specified in the NodeInfo table.	09/09/2014 09:08:02 EDT	
CMP30 10.148.253.200	CMP	Minor	32532	8h 51m 30s / ---	Server Upgrade Pending Accept/Reject	09/09/2014 00:16:56 EDT	
MPE-R 10.148.253.202	MPE	Minor	32532	2d 10h 59m 26s / ---	Server Upgrade Pending Accept/Reject	09/06/2014 22:09:00 EDT	
MPE-R 10.148.253.202	MPE	Minor	71103	4h 40m 36s / ---	PCMM Conn Lost	09/09/2014 04:27:50 EDT	
MPE-S2 10.148.253.206	MPE	Minor	32532	2d 10h 59m 53s / ---	Server Upgrade Pending Accept/Reject	09/06/2014 22:08:33 EDT	
MPE-S2 10.148.253.206	MPE	Minor	71004	4h 40m 36s / ---	AM socket closed	09/09/2014 04:27:50 EDT	
MPE-S1 10.148.253.204	MPE	Minor	32532	2d 10h 58m 41s / ---	Server Upgrade Pending Accept/Reject	09/06/2014 22:09:45 EDT	
MPE-S1 10.148.253.204	MPE	Minor	71004	4h 40m 35s / ---	AM socket closed	09/09/2014 04:27:51 EDT	
MPE-S1 10.148.253.204	MPE	Minor	71103	4h 41m 13s / ---	PCMM Conn Lost	09/09/2014 04:27:13 EDT	
BOD 10.148.253.208	BoD	Minor	32532	2d 11h 2m 13s / ---	Server Upgrade Pending Accept/Reject	09/06/2014 22:06:13 EDT	

Viewing Alarms

To view alarms or the alarms history:

1. From the **System Wide Reports** section of the navigation pane, select **Alarms**.
2. Select the report to view.

The navigation pane displays the available alarms reports.

Viewing Active Alarms

The Active Alarms report provides details about active alarms. To view the Active Alarms report:

1. From the **System Wide Reports** section of the navigation pane, select **Alarms**.

   The **Alarms** section expands to show the available alarm reports.

2. Select **Active Alarms**.

   The Active Alarms report opens in the work area.

Working with the Active Alarms Report

To work with the Active Alarms report:

1. From the **System Wide Reports** section of the navigation pane, select **Alarms**.

   The **Alarms** section expands to show the available alarm reports.

2. Select **Active Alarms**.

   The Active Alarms report opens in the work area.

3. Perform one or more of the following actions:

- To sort the report on any column, click the column title.
- To display online help for an alarm, click on its ID.
- To pause the display of alarms, click **Pause**. To resume the display, click **Refresh**.
- To select what information displays, click **Columns** and select from the list.
- To control what alarms and alarm classes display on the page, click **Filters** and select from the list:
  - The **Server** control lets you display alarms from all servers (default) or a specific server.
  - The **Server Type** control lets you display alarms from all Policy Management products (the default) or just **CMP** or **MPE** systems.
  - The **Severity** control lets you display alarms of all severities (default), critical and major alarms, critical alarms, major alarms, or minor alarms.
- To save your formatting changes to the report page, click **Save Layout**.
- To print the report, click **Printable Format** to display the current alarms in a separate window for printing.
- To save the report at a comma-separated value (CSV) file, click **Save as CSV**. A file with the name `report.csv` is generated and a standard File Download window opens so you can save or open the file.
- To save the file as a PDF file, click **Export PDF**. A file named `report.pdf` is generated and a standard File Download window opens, so you can save or open the file.

Viewing the Alarm History Report

The Alarm History Report displays historical alarm information.

To view the alarm history report:

1. From the **System Wide Reports** section of the navigation pane, select **Alarms**.

   The **Alarms** section expands to show the available alarm reports.

2. Select **Alarm History Report**.

   The Alarm History report opens.

   Note: If you are using Internet Explorer, the window appears behind the main window.

   The window displays up to 50,000 alarms, sorted by age.

   Note: If you wish to view the most recent alarms, and there are more than 50,000 alarms in the database, specify a start date/time that includes the present.

3. To view older alarms, reduce the number of alarms displayed, or locate a specific alarm or group of alarms, you can define filtering criteria using the following fields:

   - **Start Date**

     Filter out alerts before a specific date/time. Click the calendar icon to specify a date/time.

   - **End Date**

     Filter out alerts after a specific date/time. Click the calendar icon to specify a date/time.

   - **Severity**

     Filter alerts by severity level. Select a level from the list. The default is **All**.

   - **Cluster or Server**

     Select the cluster or server within the cluster to view the alarms.

   - **Active Alarms**

     Select to view only active alarms; the default is to display both active and cleared alarms.

   - **Aggregate**

     Select to aggregate alarms that have the same IP address, alarm ID, and severity. (This function is limited to 50,000 alarms.)

4. After entering filtering information, click **Filter** to refresh the display with the filtering applied.

   The alarm list is filtered.

5. Click **Close**.

Alarms contain the following information:

- **Occurrence**

  The most recent time this alert was triggered.

- **Severity**

  The severity of the alert:

  - **Critical—**Service is being interrupted (displays in red).
  - **Major—**Service may be interrupted if the issue is not corrected (displays in orange).
  - **Minor—**Non-service affecting fault (displays in yellow).
  - **Info—**Informational message only.
  - **Clear—**Alarm has been cleared.

> **Note:** Alarms generated by Policy Management systems running software lower than release 7.5 are mapped to these levels as follows: Emergency or Critical map to Critical; Alert or Error map to Major; Warning or Notice map to Minor.

- **Alarm ID**

  When clicked, the alarm ID provides online help information.

- **Text**

  User-readable text of the alert.

- **OAM VIP**

  OAM IP address in IPv4 address.

- **Server**

  Name and IP address, in IPv4 or IPv6 format, or FQDN of the device from which this alarm was generated.

To view alert details, click (binoculars icon), located to the right of the alert. A window displays additional information.

For example:

Figure 12-4 Alert Details

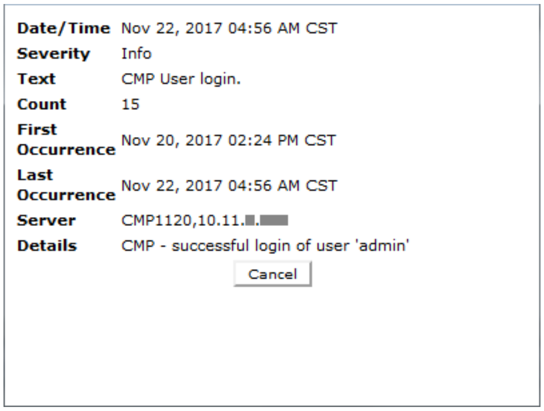

Date/Time	Nov 22, 2017 04:56 AM CST
Severity	Info
Text	CMP User login.
Count	15
First Occurrence	Nov 20, 2017 02:24 PM CST
Last Occurrence	Nov 22, 2017 04:56 AM CST
Server	CMP1120,10.11.[redacted].[redacted]
Details	CMP - successful login of user 'admin'

Viewing Other Reports

To view the miscellaneous reports:

1. From the **System Wide Reports** section of the navigation pane, select **Others**.
2. Select the report to view.

The navigation pane displays the available reports.

Viewing the Connection Status Report

The connection status report provides an aggregate view of connections maintained by managed Policy Management systems. The display is refreshed every ten seconds.

To view the connection status report:

1. From the **System Wide Reports** section of the navigation pane, select **Others**.

   The list of available reports displays in the navigation pane.

2. Select **Connection Status**.

   The Connection Status report opens.

The report columns display the following data:

- **Server** — Name of the associated system.
- **Server Type** — MPE (Multimedia Policy Engine).
- **Remote Identity** — The ID (if known) or IP address of the remote system.
- **Type** — The type of connection (for example, PCMM CMTS, PCMM AM, PCMM DPS, or Diameter AF).
- **Status** — The status of the connection (the possible values are protocol-specific).
- **Up/Down Since** — The timestamp when the connection reached its current state (**N/A** if the connection has never been established).
- **# Total Connect** — The number of times that the connection has been re-established

  Note: This counter is reset if the cluster is restarted.

- **# Active Connect** — The number of active connections.
- **Msgs Sent** — The number of protocol messages that have been sent to the remote system.
- **Msgs Received** — The number of protocol messages that have been received from the remote system.
- **Errors Sent** — The number of protocol error messages that have been sent to the remote system.
- **Errors Received** — The number of protocol error messages that have been received from the remote system.

If a connection is in a non-functional state, the row is displayed in red; if a connection is in a transitional state between functional and non-functional (including when a connection is being established), the row is displayed in yellow.

The following options are available:

- To pause the display, click **Pause**. To resume the display, click **Refresh**.
- To sort the display rows, click on a column heading.
- To select what information is displayed, click **Columns** and select from the pulldown list.
- To control what rows are displayed on the page, click **Filters** and select from the pulldown list:

- The **Server** control lets you display information from all servers (the default) or a specific server.
- The **Server Type** control lets you display information from all Policy Management products (the default) or just **MPE** systems.
- The **Remote Identity** control lets you display information from all remote devices (the default) or a specific remote device selected by its ID or IP address.
- The **Type** control lets you display information for all protocols (the default) or a specific protocol.
- The **Status** control lets you display information for all status values (the default) or a specific status.

- To save the current layout, click **Save Layout**.
- **Printable Format —** The current alarms are displayed in a separate window.
- **Save as CSV —** A comma-separated value (CSV) file named `report.csv` is generated, suitable for a spreadsheet application, and a standard File Download window opens, so you can save or open the file.
- **Export PDF —** A Portable Document Format (PDF) file named `report.pdf` is generated, suitable for a spreadsheet application, and a standard File Download window opens, so you can save or open the file.

Viewing the Protocol Errors Report

The protocol errors report provides an aggregate view of connection errors, with one row for each distinct error code or sub-code. The display is refreshed every ten seconds.

To view the protocol errors report:

1. From the **System Wide Reports** section of the navigation pane, select **Others**.

   The list of available reports displays in the navigation pane.

2. Select **Protocol Errors**.

   The Protocol Errors report opens.

The report columns display the following data:

- **Server —** name of the associated system
- **Server Type — MPE** (Multimedia Policy Engine)
- **Remote Identity** — the ID (if known) or IP address of the remote system
- **Error** — the protocol error
- **# Received** — the number of protocol errors received from the remote system
- **# Sent** — the number of protocol errors sent to the remote system

The following options are available:

- To pause the display, click **Pause**. To resume the display, click **Refresh**.
- To sort the display rows, click on a column heading.
- To select what information is displayed, click **Columns** and select from the pulldown list.
- To control what rows are displayed on the page, click **Filters** and select from the pulldown list:
  - The **Server** control lets you display information from all servers (the default) or a specific server.
  - The **Server Type** control lets you display information from all Policy Management products (the default) or just **MPE** systems.
  - The **Remote Identity** control lets you display information from all remote devices (the default) or a specific remote device selected by its ID or IP address.
  - The **Type** control lets you display information for all protocols (the default) or a specific protocol.
  - The **Status** control lets you display information for all status values (the default) or a specific status.
- To save the current layout, click **Save Layout**.
- **Printable Format** — The current alarms are displayed in a separate window.
- **Save as CSV** — A comma-separated value (CSV) file named `report.csv` is generated, suitable for a spreadsheet application, and a standard File Download window opens, so you can save or open the file.
- **Export PDF** — A Portable Document Format (PDF) file named `report.pdf` is generated, suitable for a spreadsheet application, and a standard File Download window opens, so you can save or open the file.

Viewing the Policy Statistics Report

The policy statistics report provides an aggregate view of policy statistics, with one row for each policy, letting you gauge the performance of individual policies. The display is refreshed every ten seconds.

To view the policy statistics report:

1. From the **System Wide Reports** section of the navigation pane, select **Others**.

   The list of available reports displays in the navigation pane.

2. Select **Policy Statistics Report**.

   The Policy Statistics report opens.

From the report page you can do the following:

- To sort the report on any column, click the column title.
- To pause the display, click **Pause**. To resume the display, click **Refresh**.

- To display another page of the report, click the page number.

You can customize what information is displayed by controlling which table columns appear, using the **Columns** list. The following columns are available:

- **Server Name**

  Name of the associated system

- **Server Type**

  Either **MPE**

- **Policy Name**

  The name of each policy defined and active on the displayed server

- **Evaluated**

  The number of times the displayed policy was evaluated for the displayed server

- **Executed**

  The number of times the displayed policy was executed for the displayed server

- **Ignored**

  The number of times the displayed policy was ignored by the displayed server

- **Total Execution Time (ms)**

  The total execution time for each policy, in milliseconds

- **Average Execution Time (ms)**

  The average amount of time it takes a policy to execute, in milliseconds

- **Maximum Execution Time (ms)**

  The maximum execution time for each policy, in milliseconds

You can filter results by controlling which table rows appear, using the **Filters** list. You can define filtering criteria using the following fields:

- **Server Name**

  Filter in all servers (default) or one specific server.

- **Policy Name**

  Filter in all policies (default) or one specific policy.

You can save formatting changes to the report page. Click **Save Layout**.

You can display the report in a format suitable for printing. Click **Printable Format**; a Policy Statistics Report window opens.

You can save the report in comma-separated value (CSV) format, suitable for importing into a spreadsheet application. Click **Save as CSV**. A file named `report.csv` is generated, and a standard File Download window opens, so you can save or open the file.

You can save the report as a Portable Document Format (PDF) file, suitable for storage or online display. Click **Export PDF**. A file named `report.pdf` is generated, and a standard File Download window opens, so you can save or open the file.

Viewing the MPE/BoD Replication Statistics Report

The MPE/MRABoD replication statistics report provides a view of database replication statistics, with one row for each replication path in an MPE or BoD cluster. The display is refreshed every ten seconds.

To view the replication statistics report:

1. From the **System Wide Reports** section of the navigation pane, select **Others**.
2. Select **MPE/BoD Rep Stats**.

From the report page you can do the following:

- To sort the report on any column, click the column title.
- To pause the display, click **Pause**. To resume the display, click **Refresh**.
- To save the any formatting changes in the page, click **Save Layout**.
- To display another page of the report, click the page number.

You can customize what information is displayed by controlling which table columns appear, using the **Columns** list. The following columns are available:

- **Cluster Name**

  The name of the cluster and the blades participating in replication as well as their high availability (HA) states.

- **Server Type**

  The type of cluster being utilized (MPE or BoD).

- **Blade State**

  Displays the state of the blade replicating with the current active blade.

  Table 12-2 Blade State Values in MPE/BoD Replication Stats Report

Blade Ha State	Value Displayed in the Report	Icon Used in the User Interface
Standby	OK	✓ Green check mark
Spare	OK	✓ Green check mark
Forcestandby	Minor	⚠ Warning Sign
Out of Service	Critical	✗ Red X
Unknown	Critical	✗ Red X

- **Sync State**

  Displays the values reported from COMCOL.

Table 12-3 Sync State Values in MPE/BoDReplication Stats Report

Sync Status	Description	Value Displayed on the CMP	Icon Used in the User Interface
Down	The link is down and there is no current attempt to restore it.	Critical	⊗ Red X
DownListening	The incoming link is down awaiting the other side to initiate the connect attempt.	Critical	⊗ Red X
Down Connecting	The link is down by this side is trying to connect.	Critical	⊗ Red X
DownRejected	The link is down because a connect attempt was rejected in the handshake phase.	Critical	⊗ Red X
Down Handshake	The link is connected but not ready for application use (so it is down logically). The links is being validated in a handshake as legitimate.	Critical	⊗ Red X
Connected	Connected and ready for use.	Critical	⊗ Red X
Connected Reinit	Connected and ready for use, but after an application error where the recovery is start over without either a link drop or a complete application restart.	Critical	⊗ Red X
Connected Incompat	Connected but the schema are incompatible and replication cannot run until (1) the schema has the needed upgrade information or (2) problematic tables are excluded from replication.	Critical	⊗ Red X
RegisterSent	RegisterSent means the link is exchanging application level credentials and information (such as data dictionary information). In this state, registration has been sent from one side and it is being awaited from the other side.	Critical	⊗ Red X
RegisterAcked	In this state, registration has been sent acknowledged from the other side. In most configurations, it is a transitory state, but the end application can hold the link in this state before permitting an audit.	Critical	⊗ Red X
Standby	Standby means the high-availability state is standby, but the applications have exchanged registration messages.	Critical	⊗ Red X
Inhibited	Inhibited means the link administrative state is inhibited (or disabled), but the applications have exchanged registration messages.	Major	❗ Red Exclamation Mark
AuditWait	The audit is awaiting a message that is is OK tor proceed from the remote side.	Critical	⊗ Red X
AuditQueue	The audit is queued because a limit on the number of simultaneous audits.	Critical	⊗ Red X

Table 12-3 (Cont.) Sync State Values in MPE/BoDReplication Stats Report

Sync Status	Description	Value Displayed on the CMP	Icon Used in the User Interface
Audit	Audit means the application is bringing the databases into agreement. It does so by comparing each table one-by-one, and then applying database updates since the audit began.	Major	❗ Red Exclamation Mark
Active	Active means the link is in the normal active steady-state conditions where updates are being transferred to the slave databases with a normal and acceptable delay.	OK	✓ Green Check Mark
ActiveBehind	ActiveBehind is the same as Active but the slave database is unacceptably behind for whatever reasons. After an audit, it would be typical to be in the ActiveBehind state until any queued updates are applied to the slave database.	Major	❗ Red Exclamation Mark
ActiveSwitch	A switchover is being attempted without an audit if the states of the databases allow it.	Major	❗ Red Exclamation Mark
ActivePost Audit	The database is coherent but has not caught back up to current after the preceding audit.	Major	❗ Red Exclamation Mark

- **Cluster State**

  Represents the overall state of the cluster. The Cluster State Column is an aggregation of the Blade State and Sync State columns. The value for the Cluster State is selected based on the maximum severity.

Table 12-4 Priority Table in MPE/BoD Replication Stats Report

Priority	Value	Icon Used in the User Interface
1	Critical	✖ Red X
2	Major	❗ Red Exclamation Mark
3	Minor	⚠ Warning Sign
4	OK	✓ Green Check Mark

You can filter results by controlling which table rows appear, using the **Filters** list. You can define filtering criteria using the following fields:

- **App Type**

  Filter in all applications (default) or filter by MPE or BoD.

- **Server Name**

  Filter in all servers (default) or one specific server.

- **Cluster Name**

  Filter in all clusters (default) or one specific cluster.

You can display the report in a format suitable for printing. Click **Printable Format**. The MPE/BoD Rep Status Report window opens.

You can save the report in comma-separated value (CSV) format, suitable for importing into a spreadsheet application. Click **Save as CSV**. A file named `report.csv` is generated, and a standard File Download window opens, so you can save or open the file.

You can save the report as a Portable Document Format (PDF) file, suitable for storage or online display. Click **Export PDF**. A file named `report.pdf` is generated, and a standard File Download window opens, so you can save or open the file.

13

Upgrade Functions

This chapter describes the functions under the **Upgrade** menu in the CMP system.

For detailed instructions on upgrading your system, download the *Upgrade Guide* for your release from the Oracle Help Center (see #unique_286 for more information).

Overview of Upgrade Functions

> **Note:** Access to the Upgrade functions may be restricted by user role. See About User Roles for more information.

The ISO Maintenance page lets you manage ISO files for the Policy Management servers. Using the **Operations** menu, you can push scripts, upload ISO files to selected servers or clusters, and delete uploaded ISO files from selected servers or clusters. See Viewing the ISO Maintenance Page for more information.

The Upgrade Manager page lets you upgrade software on clusters in the Policy Management network, or roll back an upgrade. An upgrade or rollback automatically processes a multi-server cluster or a georedundant site in the proper order to minimize data loss and downtime. During the process, the Upgrade Manager page displays progress information. The Upgrade Manager page also lets you view an upgrade log to review operations. See Viewing the Upgrade Manager Page for more information.

Viewing the ISO Maintenance Page

To view the ISO Maintenance page:

1. From the **Upgrade** section of the navigation pane, select **ISO Maintenance**.

   The ISO Maintenance page opens.

The ISO Maintenance page lists the clusters and servers in the Policy Management system. Also listed are application type, site, IP address, current release, and any ISO files available for upgrading the server or cluster.

Figure 13-1 ISO Maintenance Page

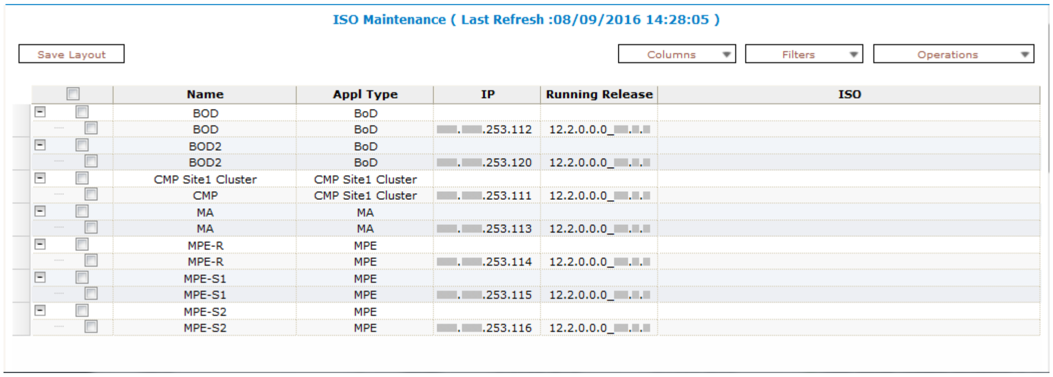

ISO Maintenance (Last Refresh :08/09/2016 14:28:05)

Name	Appl Type	IP	Running Release	ISO
BOD	BoD			
BOD	BoD	[redacted].253.112	12.2.0.0.0_[redacted]	
BOD2	BoD			
BOD2	BoD	[redacted].253.120	12.2.0.0.0_[redacted]	
CMP Site1 Cluster	CMP Site1 Cluster			
CMP	CMP Site1 Cluster	[redacted].253.111	12.2.0.0.0_[redacted]	
MA	MA			
MA	MA	[redacted].253.113	12.2.0.0.0_[redacted]	
MPE-R	MPE			
MPE-R	MPE	[redacted].253.114	12.2.0.0.0_[redacted]	
MPE-S1	MPE			
MPE-S1	MPE	[redacted].253.115	12.2.0.0.0_[redacted]	
MPE-S2	MPE			
MPE-S2	MPE	[redacted].253.116	12.2.0.0.0_[redacted]	

Viewing the Upgrade Manager Page

Note: For detailed instructions on upgrading your system, download the *Upgrade Guide* for your release from the Oracle Help Center (see #unique_286 for detailed information).

To view the Upgrade Manager page:

1. From the **Upgrade** section of the navigation pane, select **Upgrade Manager**.

   The Upgrade Manager page opens.

The Upgrade Manager page lists the clusters and servers in the Policy Management system. Also listed are the current server alarm state, server role, previous release, current release, and the date, time, and result of the last upgrade operation performed on the server. You can display an upgrade log, which records timestamped upgrade events.

Figure 13-2 Upgrade Manager Page

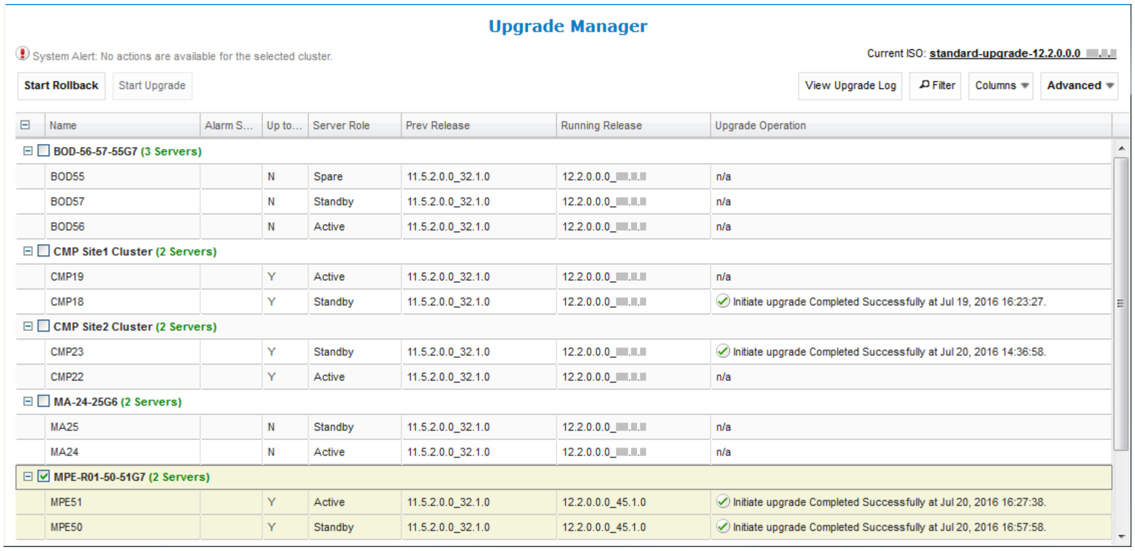

Upgrade Manager

System Alert: No actions are available for the selected cluster.

Current ISO: standard-upgrade-12.2.0.0.0_[redacted]

Name	Alarm S...	Up to...	Server Role	Prev Release	Running Release	Upgrade Operation
□ BOD-56-57-55G7 (3 Servers)						
BOD55		N	Spare	11.5.2.0.0_32.1.0	12.2.0.0.0_[redacted]	n/a
BOD57		N	Standby	11.5.2.0.0_32.1.0	12.2.0.0.0_[redacted]	n/a
BOD56		N	Active	11.5.2.0.0_32.1.0	12.2.0.0.0_[redacted]	n/a
□ CMP Site1 Cluster (2 Servers)						
CMP19		Y	Active	11.5.2.0.0_32.1.0	12.2.0.0.0_[redacted]	n/a
CMP18		Y	Standby	11.5.2.0.0_32.1.0	12.2.0.0.0_[redacted]	✓ Initiate upgrade Completed Successfully at Jul 19, 2016 16:23:27.
□ CMP Site2 Cluster (2 Servers)						
CMP23		Y	Standby	11.5.2.0.0_32.1.0	12.2.0.0.0_[redacted]	✓ Initiate upgrade Completed Successfully at Jul 20, 2016 14:36:58.
CMP22		Y	Active	11.5.2.0.0_32.1.0	12.2.0.0.0_[redacted]	n/a
□ MA-24-25G6 (2 Servers)						
MA25		N	Standby	11.5.2.0.0_32.1.0	12.2.0.0.0_[redacted]	n/a
MA24		N	Active	11.5.2.0.0_32.1.0	12.2.0.0.0_[redacted]	n/a
☑ MPE-R01-50-51G7 (2 Servers)						
MPE51		Y	Active	11.5.2.0.0_32.1.0	12.2.0.0.0_45.1.0	✓ Initiate upgrade Completed Successfully at Jul 20, 2016 16:27:38.
MPE50		Y	Standby	11.5.2.0.0_32.1.0	12.2.0.0.0_45.1.0	✓ Initiate upgrade Completed Successfully at Jul 20, 2016 16:57:58.

14

Global Configuration

This chapter describes how to configure global CMP settings.

Configuring and Enabling Subnet Settings

The Subnet Settings page allows you to enable or disable filtering and/or aggregation. Aggregation and filtering help reduce the amount of data to be processed, which improves server performance.

You can enable filtering for both IPv4 and IPv6 prefixes. Filtering allows IPv4 and IPv6 prefixes that match the configured criteria to be discarded before data is routed to the CMP system and MPE devices.

You can enable aggregation for IPv6 prefixes. Aggregation allows multiple IPv6 prefixes to be aggregated into a single entry.

You can view IP prefix filters and aggregation configuration changes in the audit log (see Viewing the Audit Log). To view prefix filtering and aggregation functionality for each CMTS device, and the net result of the functionality for all CMTS devices, refer to the trace log (see Viewing the Trace Log).

To configure and enable subnet settings:

1. From the **Global Configuration** section of the navigation pane, select **Global Configuration Settings**.
2. From the content tree, select the **Subnet Settings** folder.

   The Subnet Settings page opens in the work area.
3. Click **Modify**.
4. Select **IPv6 Subnet Aggregation Enable** to enable the IPv6 aggregation functionality.
5. Select **IPv4 Subnet Filtering Enable** and **IPv6 Subnet Filtering Enable** to enable filtering functionality.

   Note: If you enable the filtering functionality, additional fields appear, allowing you to configure filtering rules. Up to 1000 filtering rules are supported. If configuring more than 100 rules, validation is recommended to assess the time impact of filtering on the subnet collection task. If no rules are configured, filtering does not occur.

   ---

   a. Enter an IP address and prefix length in the **Subnets** fields (enter the prefix length after the slash). The IP string must be a valid IPv4 or IPv6 address and is case insensitive.

- To filter all IPv6 prefixes, enter * for the IP address and leave the prefix length field blank.
- To filter all IPv6 prefixes with a specific prefix length, enter * for the IP address and the appropriate value for the prefix length.

Note: Wildcards (*) are not supported in IPv4 IP addresses.

b. Click **Add** to add the IP address and prefix length.

   The IPv6 address and prefix length are added to the list of addresses.

c. To remove an IP address/prefix length from the list, select the IP address/ prefix length, and click **Delete**. Click **Delete All** to remove all addresses and prefix lengths from the list.

6. Click **Save**.

   Clicking **Save** deploys the configuration to the Management Agent (MA) if an MA is managed by the CMP system. A message appears, indicating the result of the deployment. If the deployment fails for an MA, reapply the configuration for the corresponding MA (see Reapplying a Management Agent Profile Configuration).

The subnet settings are configured and the subnet is enabled.

About Routing by CMTS IP

The Cable Modem Termination System (CMTS) routing parameter allows the Cable Policy Server to pass destination and CMTS information, in addition to the standard routing parameters (subscriber ID and IP address), to a CMTS within a Carrier Grade Network Address Translation (CG-NAT) network. Standard subnet-based routing is not compatible in CG-NAT networks. The alternate routing function is triggered by the presence of the CMTS parameter.

The CMTS IP routing parameter contains a CMTS management IP address (CMTSIP), which is sent by the Bandwidth on Demand Application Manager (BoD-AM) to a Multimedia Policy Engine Routing Layer (MPE-R) device as a custom extended object. The MPE-R device uses CMTSIP instead of SubscriberID to find the routing path in PCMM requests that contain the CMTS IP parameter. The CMTS IP parameter is then passed to an MPE-Serving Layer (MPE-S) device, which initiates a standard PCMM request to the CMTS.

Note: The CMTS IP parameter is only supported in PCMM messages. For Rx requests, the Cable Policy Server uses the standard routing architecture.

The CMTS IP parameter is not enabled by default and only accepts one valid IPv4 or IPv6 address. The format of the value is as follows:

- IPv4 address: dotted decimal notation format
- IPv6 address: colon hexadecimal notation format

If an MPE-R receives a PCMM request without CMTSIP, or if the **Route by CMTS IP** setting is disabled, the MPE-R routes PCMM requests based on SubscriberID. If an MPE-R device routes a request to an MPE-S device based on CMTSIP and cannot find the corresponding MPE-S based on the CMTSIP, it responds with an error.

The CMP supports the capture of CMTS IP information in a trace log and the BoD session viewer. See the document *Policy Management Troubleshooting Reference* for more information about logs.

See also Configuring and Enabling Subnet Settings.

Configuring and Enabling CMTS IP Routing

You can use Global Configuration Settings to enable and define a Cable Modem Termination System (CMTS) routing parameter, **Route by CMTS IP**, for CG-NAT network compatibility.

Note: Though not required, IPv4 subnet filtering is recommended because it removes unnecessary subnets which might disrupt the CMTS IP routing function. See the topic Configuring and Enabling Subnet Settings to configure IPv4 subnet filtering.

To configure and enable CMTS IP routing:

1. From the **Global Configuration** section of the navigation pane, select **Global Configuration Settings**.
2. From the content tree, select the **Route by CMTS IP** folder.

   The Routing by CMTS IP page opens.
3. Click **Modify**.

   The **Route by CMTS IP** checkbox is visible.
4. Select **Route by CMTS IP** to enable CMTS routing.
5. Click **Save**.

   The Route by CMTS IP value is true.

   Note: If there are existing CMTS devices, enabling **Route by CMTS IP** will change the original hostname (IPv4) to the CMTS management IP. If the original hostname is a fully qualified domain name (FQDN), the system generates a warning on the Routing by CMTS IP page for you to configure the CMTS IP manually (from the Network Elements Administration page).

   If there is an existing CMTS device without a CMTS IP address, the system provides a warning with a list up to 10 conflicting CMTS devices.
6. To create the CMTS IP network element, go to the **Policy Server** section of the navigation pane and select **Network Elements**.

   The Network Element Administration page opens.
7. Click **Create Network Element**.

   The New Network Element page opens.
8. Enter the required fields for the CMTS parameter.

a. **Name** — Enter the name you assign the CMTS IP parameter

b. **IP Address** — Enter the IP address for the CMTS IP parameter.

> **Note:** The CMTS IP parameter does not support wildcards (*) in IPv4 addresses.

c. Click **Save**.

The CMTS IP parameter is added to the Network Elements tree.

Configuring Stats Settings

You can define when and how measurement statistic values are reset.

> **Caution:** Saving changes to the statistics settings causes the historical stats data to be lost.

To configure stats settings:

1. From the **Global Configuration** section of the navigation pane, select **Global Configuration Settings**.

   The content tree displays a list of global configuration settings.

2. From the content tree, select the **Stats Settings** folder.

   The Stats Settings page opens in the work group area.

3. Click **Modify**.

   The fields on the Stats Settings page become editable.

4. Enter values for the **Stats Collection Period** configuration attributes. Specify the time interval to use from the list. Options are in minutes:

   - **5**
   - **10**
   - **15** (default)
   - **20**
   - **30**
   - **60**

5. Click **Save**.

The Stats Settings are configured.

15

System Administration

This chapter describes functions reserved for CMP system administrators.

Note: Some options are visible only when you are logged in with administrative rights to the CMP system. However, the **Change Password** option is available to all users.

Configuring System Settings

Within the CMP system you can define the settings that control system behavior.

To define system settings:

1. From the **System Administration** section of the navigation pane, select **System Settings**.

   The System Settings page opens in the work area, displaying the current system settings.

2. Click **Modify**.

   The System Settings page opens.

3. In the Configuration section, define the following:

   a. **Idle Timeout (minutes; 0=never)**—The interval of time, in minutes, that a user login session is kept alive.

      The default value is 30 minutes; a value of zero indicates the session remains active indefinitely.

   b. **Account Inactivity Lockout (days; 0=never)**—The maximum number of days since the last successful login after which a user is locked out.

      If the user fails to log in for the defined number of days, the user is locked out and cannot gain access to the system until an administrator resets the account. The default value is 21 days; a value of zero indicates no limit (the user is never locked out for inactivity).

   c. **Maximum Concurrent Sessions Per User Account (0=unlimited)—**The maximum number of times a defined user can be logged in simultaneously. A value of zero indicates no limit.

      If more than the configured number of concurrent users try to log in (for example, a second user if this value is set to 1), they are blocked at the login page with the a message indicating that the maximum number of concurrent sessions has been reached.

d. **Password Expiration Period (days; 0=never)**—The number of days a password can be used before it expires. Enter a value from 7 to 365, or 0 to indicate that the password never expires.

e. **Password Expiration Warning Period (days; default=3)**—The number of days before a password expires to begin displaying a window to users after login warning that their password is expiring.

f. **Admin User Password Expiration**—By default, the password for the `admin` user never expires.

   If you select this option, the `admin` user is subject to the same password expiration policies as other users.

g. **Block users when password expires**—By default, after a password expires, the user must immediately change it at the next login.

   If you select this option, if their password expires, users cannot log in at all. (If you select **Admin User Password Expiration** and the `admin` user's password expires, the user can still log in but must immediately enter a new password.)

h. **Alert Destination**—The host name or IP address, in IPv4 or IPv6 format, of the target where all alerts are sent from for the various servers in the network. Normally, this is the address of the CMP system.

Caution: If you defined the address of this server using IPv4 format, then define the alert destination using IPv4 format. If you defined the address of this server using IPv6 format, then define the alert destination using IPv6 format. Otherwise, alerts can be lost.

i. **Minimum Password Length**—The minimum allowable length in characters for a password, from 6 to 64 characters.

   The default is six characters.

j. **Login Banner Title**—The banner that displays on the login page. The default is "Welcome." You can enter up to 30 characters.

k. **Login Banner Text**—The text that displays on the login page. You can enter up to 10,000 characters.

l. **Top Banner Text**—The text that displays in the banner at the top of the GUI page. You can enter up to 50 characters. You can select the font, size, and color of the text.

m. **Allow policy checkpoint and restore (copies; 0=disallow)**—The number of checkpoints allowed in the system. Valid value range is 0 to 10. If set to 0, the Policy Checkpoint/Restore option is turned off and is no longer visible under the Policy Management heading on the navigation panel. The default value is 0.

4. In the Invalid Login Threshold Settings section, define the following:

   a. **Enable**—Enables login threshold control.

      By default, this feature is enabled; deselect the check box to disable this feature.

   b. **Invalid Login Threshold Value**—Defines the maximum number of consecutive failed logins after which action is taken.

Enter a value from 1 through 500; the default is 3 attempts.

c. **Action(s) upon Crossing Threshold—**The system action to take if a user reaches the invalid login threshold:

  - **Lock user account—**Prevents users from logging in if they reach the invalid login threshold.
  - **Send trace log message—**If a user account reaches the threshold, an incident is written to the trace log, including the username and the IP address from which the login attempts was made. The default level is **Warning**; to change the event level, select a different level from the list.

5. The Password Strength Settings section lists four character categories: lowercase letters, uppercase letters, numerals, and non-alphabetic characters. You can specify a password strength policy that requires users to create passwords by drawing from these categories:

   - **Require at least categories below—**By default, this setting is 0 (disabled). Select it to require users to include password characters from between one to four of the categories.
   - **Require at least lower-case letter(s) (1-64)—**By default, this setting is 0 (disabled). Select it to require users to include from 1 to 64 lowercase letters in their passwords.
   - **Require at least upper-case letter(s) (1-64) —**By default, this setting is 0 (disabled). Select it to require users to include from 1 to 64 uppercase letters in their passwords.
   - **Require at least numeral(s) (1-64)—**By default, this setting is 0 (disabled). Select it to require users to include from 1 to 64 numerals in their passwords.
   - **Require at least non-alphanumeric character(s) (1-64)—**By default, this setting is 0 (disabled). Select it to require users to include from 1 to 64 non-alphabetic characters in their passwords.
   - **Force users with weak password to change password at their next login—**By default, this setting is disabled. Select it to require users to conform to a new password policy effective the next time they log in.

6. Click **Save**.
7. To refresh the CMP system, log out and log back in.

The system settings are configured.

Figure 15-1 shows an example of settings that establish a password strength policy requiring user passwords to contain at least one uppercase letter, four numerals, and one non-alphabetic character. (A password that would satisfy this policy is `P@ssword1357`.) Users whose passwords do not meet these requirements will be forced to change their passwords the next time they log in.

Figure 15-1 Sample Password Strength Policy

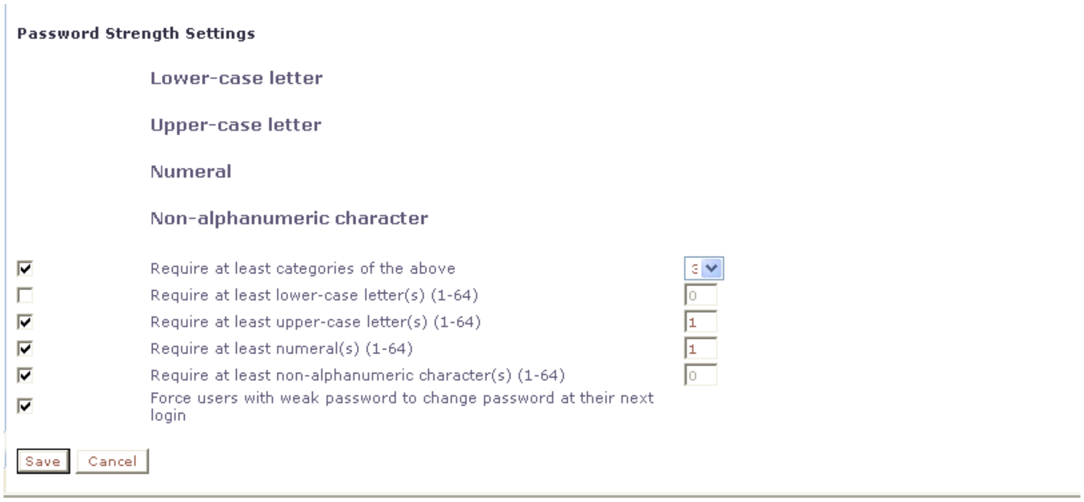

Password Strength Settings

Lower-case letter

Upper-case letter

Numeral

Non-alphanumeric character

☑	Require at least categories of the above	3
☐	Require at least lower-case letter(s) (1-64)	0
☑	Require at least upper-case letter(s) (1-64)	1
☑	Require at least numeral(s) (1-64)	1
☑	Require at least non-alphanumeric character(s) (1-64)	0
☑	Force users with weak password to change password at their next login	

Importing and Exporting Configurable Objects

This section describes how to perform a simple or a bulk export of configurable objects and how to import object configurations into the CMP system.

You can export configuration information as a single ZIP file. This ZIP file contains an inner ZIP file of XML files, an MD5 checksum file for data verification, and an `exportResults.txt` file containing export result messages. You can use these exported XML or ZIP files to update configurable objects in a CMP system or to configure a new system.

User Privileges and Import/Export

To use the **Import/Export** action, you must have a role that includes the **System Administrator Privileges Import / Export** privilege.

> **Note:** If the you have a role with the **Import / Export** privilege, then top-level objects are available for importing and exporting even though you may not have specific privileges to view the objects (that is, the specific object is hidden to your role).

See Creating a Role for details on roles and privileges.

About Importing Configuration Objects

The CMP system exports configurable objects to a ZIP file. This ZIP file consists of an inner ZIP file, `export.zip`, containing XML files of configuration data (by object type), an MD5 checksum file (`md5.txt`), and an `exportResults.txt` message file. By default, during an import, the CMP system uses the MD5 file to verify the integrity of the import file.

The Import function provides a method for importing either XML or ZIP configuration data files. In this way, the CMP configurable objects in a system can be restored to a prior configuration or a new CMP system created and configured.

See About Adding Objects for Export for details on exporting configuration objects.

Checkpoint and Importing Objects

NOT_SUPPORTED: To use the checkpoint option, the **Allow policy checkpoint and restore** setting must be configured in **System Settings**.

During an import action, the CMP system includes the **Perform checkpoint before importing** option. This is a protective measure that uses the Checkpoint function to save a snapshot of all the configuration objects that exist in the system. See the *Policy Wizard Reference* for more checkpoint details. Included in the checkpoint output file are:

- Configurable objects such as policies, policy groups, policy templates, policy tables, traffic profiles, match lists, retry profiles, applications, policy counter IDs, MPE configuration templates, and traffic profile groups, RADIUS CoA templates
- All MPE templates
- Associations between virtual MPE templates and real MPE templates
- Associations between MPEs and virtual MPE templates
- Associations between virtual MPE templates and other configuration objects (for example, policies)

Enabling the checkpoint option ensures that if a newly reconfigured system fails to function as expected, a checkpoint XML file exists that returns the system to the state prior to the import action. An additional benefit is that the checkpoint output file is saved in the CMP database and is accessible using **Policy Checkpoint/Restore** under the **Policy Management** section. Refer to the *Policy Wizard Reference* for details on managing checkpoint files.

In a networked CMP system, the checkpoint function is limited to the NW-CMP server. The **Perform checkpoint before importing** option is not available for S-CMP servers. After a restoration from a checkpoint onto a NW-CMP server, the NW-CMP server pushes the restoration data to all S-CMP servers. The MPE and virtual MPE templates are not updated. Real MPE templates and related associations are restored on the S-CMP servers, so that the policies can be restored to MPE through the virtual template.

The **Checkpoint** function differs from the **Export All** function in the following ways:

- The output file is an XML file that is saved in the CMP database.
- MPE template associations are saved.

See Exporting All Configuration Objects for details on exporting all configurable objects.

Importing Configuration Object Files

To import XML or ZIP files, use the following file naming conventions:

- If you are importing a ZIP file, the filename for the inner ZIP file must be `export.zip`.

- If you are importing a series of XML files, Oracle recommends the file naming convention `obj_type + Export[sequence_number] + *.xml`. For example:
  - `networkelementExport1_20150501.xml`
  - `networkelementExport2_20150501.xml`
  - `networkelementExport3_20150501.xml`

Note: If there are multiple files of one type, the sequence number is necessary. Any object groups for import should be included in the last file of the sequence.

1. From the **System Administration** section of the navigation pane, expand **Import/ Export**.
2. Select **Import**.

   The import file upload form and options appear.
3. Click **Browse** to locate the file to be imported.
4. Select a processing option to use to **Handle collisions between imported items and existing items**:
   - **Delete all before importing**

     The CMP system deletes all objects for each object type matching the import file before importing the object type XML file.

     NOT_SUPPORTED: This import strategy can result in object inconsistency. For example, if you import a ZIP file that only contains traffic profiles, all the traffic profiles are deleted first. However, if existing policies depend on the existing traffic profiles, and the import file does not contain them, the policies can become invalid.
   - **Overwrite with imported version**

     For each object in the import file, if the object exists in the CMP system, the import updates the object with the configuration contained in the import file. If an object does not exist, the CMP system adds the object to the system.
   - **Reject any that already exist**

     For objects that already exist in CMP system, the import action does nothing. For objects that do not exist in the CMP system, the import adds the objects to the system.
   - **Any collisions prevent all importing**

     The CMP system compares all objects in the import file with objects in the system. If any object exists, the entire import is canceled.
   - **Validate without importing**

     If an MD5 file (part of the ZIP file) is present, the CMP system performs an MD5 checksum on the ZIP file and compares the hash value with that in the MD5 file. If the hash values match, the system validates each XML file contained in the inner ZIP file. The CMP system then performs a collision check between the system and the XML files and indicates if any exists.

5. Select one or more **Options**:

   - **Skip checksum**

     If unselected and an MD5 file (part of the ZIP file) is present, the CMP system performs a checksum on the inner ZIP file and compares the hash value with that of the MD5 file from the ZIP file. If the values do not match, the import action is canceled.

     If selected, the CMP system does not perform a checksum and proceeds with the import.

   - **Perform checkpoint before importing**

     If selected, the CMP system uses the **Checkpoint** function to save a file with all the configuration objects that exist in the system. The checkpoint file is saved in the CMP database.

6. Click **Import**.

The configuration objects and their configuration settings are imported into the CMP database. After the import is complete, the window reports the results for each XML file contained in the ZIP file. Results include whether a file was successfully imported or the number of objects successfully updated. Any problems during the import action appear in red text.

> **Note:** As the result of memory limitations, the CMP system only displays the first 2000 characters for each import type file.

The checkpoint file and the checkpoint management functions are accessed using **Policy Checkpoint/Restore** under the **Policy Management** section. Refer to the *Policy Wizard Reference* for details on managing policy checkpoint files.

About Exporting Configuration Objects

The Export action provides a method for exporting CMP configuration information as a ZIP file. Using this file, the configurable objects can be restored to a prior configuration or a new system can be created and configured.

As part of the export action, the CMP system performs the following actions:

1. Exports the configuration objects as individual XML files by object type.
2. Compresses the XML files into a ZIP file named `export.zip`.
3. Creates the `exportResults.txt` file containing export result messages.
4. Calculates MD5 checksum number for `export.zip` and creates an `md5.txt` file containing the MD5 number.
5. Compresses the `md5.txt`, `exportResults.txt`, and `export.zip` files into another ZIP file named `CMP_export__`*`yyyyMMddhhmmss`*`.zip`.
6. Downloads the `CMP_export`ZIP file to the local computer.

To restore or configure a CMP system, see Importing Configuration Object Files.

Export Window Overview

The Bulk Export window is divided into three areas.

- Object List

  The left pane lists all of the configurable object types in CMP that are available for export. The top pane contains the **View Cart** link and a counter for the total number of objects in the cart.

- Work Area

  The work area lists all of the configured objects for the selected object type. Objects are listed by Name in ascending alphabetical order. A maximum of 50 objects per page are shown.

- Navigation Area

  Page navigation links appear at the top left of the work area. A **Search** text box appears at the top right of the work area. You can search the entire list of objects for specific strings and numbers. The **Search** text box accepts the wildcard (%) and underscore (_) characters.

Figure 15-2 Bulk Export Window

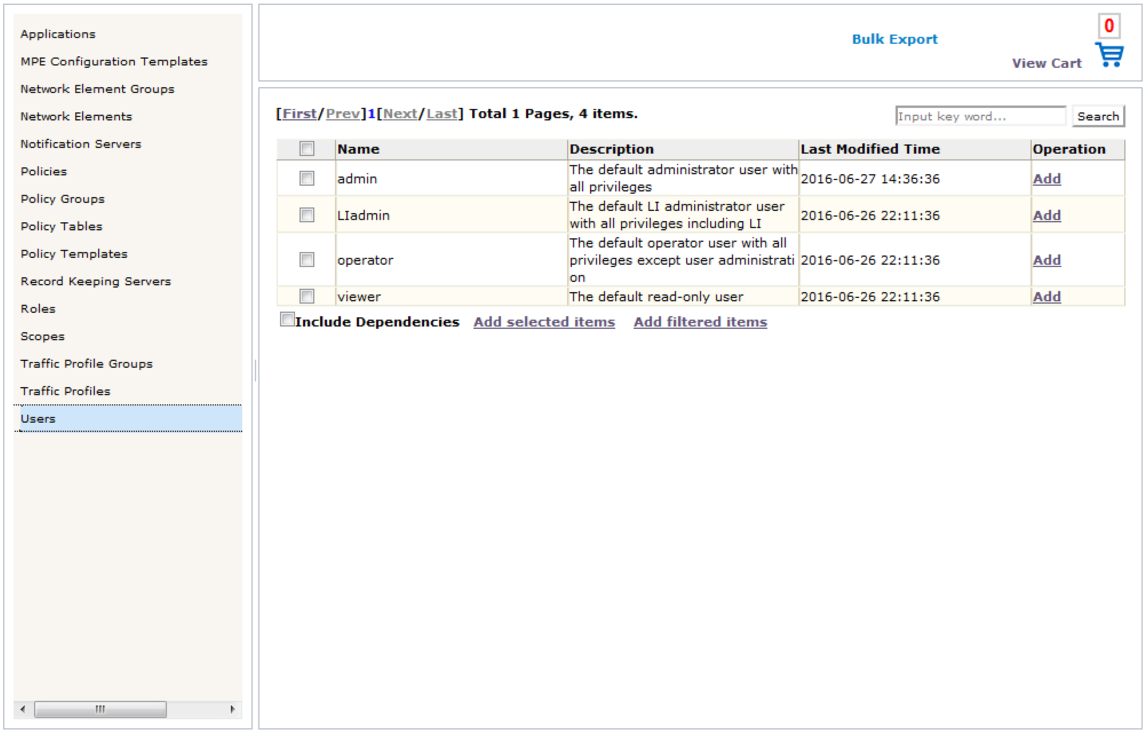

[First/Prev]1[Next/Last] Total 1 Pages, 4 items.

□	Name	Description	Last Modified Time	Operation
□	admin	The default administrator user with all privileges	2016-06-27 14:36:36	Add
□	LIadmin	The default LI administrator user with all privileges including LI	2016-06-26 22:11:36	Add
□	operator	The default operator user with all privileges except user administrati on	2016-06-26 22:11:36	Add
□	viewer	The default read-only user	2016-06-26 22:11:36	Add

□ Include Dependencies   Add selected items   Add filtered items

When operating in Debug mode, the Accounts and Tiers objects are included in the Object List.

Note: The Accounts object depends on Network Elements and Tiers objects.

About Adding Objects for Export

The Export action has different methods for selecting export objects:

- **Add**

  Adds a single object to the export cart. See Exporting a Configuration Object.

- **Add selected items**

  Adds selected objects to the export cart. See Exporting Multiple Selected Objects.

- **Add filtered items**

  Adds the results of a search to the export cart. See Exporting Multiple Filtered Objects.

After configuration objects have been added to the export cart, they are ready for export.

See Exporting All Configuration Objects for details on exporting all configurable objects.

Exporting a Configuration Object

To export a configuration object:

1. From the **System Administration** section of the navigation pane, expand **Import/Export**.

2. Select **Export**.

   The Bulk Export page opens with a listing of exportable CMP object types in the left pane and a work area for selecting objects for export on the right.

3. Select the object type from the left pane (for example, **Users**).

   A list of configured objects for the specified type displays in the work area.

4. To include dependent objects in the export file, select **Include Dependencies**.

   Any object dependencies for the selected objects will also be selected for export.

5. Use the pagination links at the top left of the work area to locate the object for export.

6. Click **Add** for the object to export.

   A message displays indicating the named object has been successfully added to the cart.

7. Click X to close the message.

   > **Note:** If **Include Dependencies** is checked, the total number of items in the export cart may exceed the number of items checked in the list.

8. Click **View Cart** or 🛒 to continue with the export.

   The Check out Page opens.

9. To remove an object from the listing, use any of the following methods:

   - Click **Remove** next to the named object.

- Select the check box next to the named object and click **Remove selected items**.
- Use the **Search** tool to search the list for a specific string to remove and click **Remove filtered items**.

10. After the export list of objects is complete, click **Export Cart**.

    A warning message displays verifying that you want to export the select item in the shopping cart.

11. Click **OK** to proceed with the export.

12. Click **Clear Cart**.

    The contents of the export cart are removed.

A ZIP file is downloaded to your local computer. The filename uses the following naming convention: `CMP_export_`*`yyyyMMddhhmmss`*`.zip`.

Exporting Multiple Selected Objects

To export multiple selected objects:

1. From the **System Administration** section of the navigation pane, expand **Import/ Export**.

2. Select **Export**.

   The Bulk Export page opens with a list of the exportable CMP object types in the left pane and a work area for selecting objects for export on the right.

3. Select the object type from the left pane (for example, **Users**).

   A list of configured objects of the specified type displays in the work area.

4. To select all objects on the page for export:

   a. Select the check box next to the Name heading of the table.

      All objects listed on the page are selected.

      > **Note:** Your selections only apply to the current page. Objects on other pages are not selected.

   b. Click **Add selected items** to add all objects in the filtered list to the export cart.

   A message displays indicating the selected objects have been successfully added to the export cart.

5. To select several objects from multiple pages for export:

   a. Use the pagination links at the top left of the work area to locate the first page of objects for export.

   b. Select the check box next to the name of each of the objects to be exported.

   c. Click **Add selected items** to add the selected objects to the export cart.

   d. Use the pagination links at the top left of the work area to locate the next page of objects for export.

e. Select the check box next to the name of each of the objects to be exported.

f. Click **Add selected items** to add the selected objects to the export cart.

A message displays indicating the selected objects have been successfully added to the export cart.

6. Click the ✕ to the right of the message to close the message.

7. Click **View Cart** or 🛒 to continue with the export.

   The Check out Page opens.

8. Review the listing of objects.

9. To remove an object from the listing, use one of the following methods:

   - Click **Remove** next to the named object.
   - Select the check box next to the named object and click **Remove selected items**.
   - Use the **Search** tool to search the list for a specific string and click **Remove filtered items**.

   Note: Any checked items will become unchecked if you change pages.

10. To view a filtered list of export objects by object type:

    a. Click the **Type** list and select the object types.

    b. Click **Search**.

    All of the selected object types display in the list.

    Note: The remaining object types are still in the export cart. To view the entire contents, select **All** from the **Type** list and click **Search**.

11. When the list of objects is ready, click **Export Cart**.

    A warning message displays verifying that you want to export the select items in the shopping cart.

12. Click **OK** to proceed with the export.

13. Click **Clear Cart**.

    The contents of the export cart are removed and you can proceed with another export.

A ZIP file is downloaded to your local computer. The filename uses the following naming convention: `CMP_export_`*`yyyyMMddhhmmss`*`.zip`.

Exporting Multiple Filtered Objects

To export multiple filtered objects:

1. From the **System Administration** section of the navigation pane, expand **Import/ Export**.

2. Select **Export**.

   The Bulk Export page opens with a listing of exportable CMP object types in the left pane and a work area for selecting objects for export on the right.

3. Select the object type from the left pane (for example, **Users**).

   A list of configured objects of the specified type displays in the work area.

4. To include dependent objects in the export file, select **Include Dependencies**.

   Any object dependencies for selected objects will also be selected for export.

5. To export a subset of objects, use the **Search** tool to search the list.

   Note: The text box accepts the wildcard (%) and underscore (_) characters.

   The list of objects is filtered to a subset of objects matching the search criteria.

6. Click **Add filtered items** to add all objects in the filtered list to the export cart.

   A message displays indicating the objects has been successfully added to the export cart.

7. Repeat the search for other filter criteria.

8. Click the X to the right of the message to close the message.

9. Click **View Cart** or 🛒 to continue with the export.

   The Check out Page opens.

10. Review the listing of objects.

11. To remove an object from the listing, use one of the following methods:

    - Click **Remove** next to the named object.
    - Select the check box next to the named object and click **Remove selected items**.
    - Use the **Search** tool to search the list for a specific string and click **Remove filtered items**.

    Note: Any checked items will become unchecked if you change pages.

12. To view a filtered list of export objects by object type, click the **Type** list, select the object types, and click **Search**.

    All of the selected object types are shown in the list.

    Note: The remaining object types are still in the export cart. To view the entire contents, select **All** from the **Type** list and click **Search**.

13. After the export list of objects is complete, click **Export Cart**.

A warning message displays verifying that you want to export the select items in the shopping cart.

14. Click **OK** to proceed with the export.
15. Click **Clear Cart**.

    The contents of the export cart are removed and you can proceed with another export.

A ZIP file is downloaded to your local computer. The filename uses the following naming convention: `CMP_export_`*`yyyyMMddhhmmss`*`.zip`.

Exporting Configuration Object Groups

To export configuration object groups:

1. From the **System Administration** section of the navigation pane, expand **Import/ Export**.
2. Select **Export**.

   The Bulk Export page opens with a listing of exportable CMP object types in the left pane and a work area for selecting objects for export on the right.

3. Select the group object type from the left pane (for example, **Network Element Groups**).

   A list of object groups are shown in the work area.

4. To include dependent objects in the export file, select **Include Dependencies**.

   Any object dependencies for the selected group will also be selected for export.

5. For each group you want to export, click the check box next to the group Name.
6. Click **Add selected items** to add the selected groups to the export cart.

   A message displays indicating the objects has been successfully added to the export cart.

7. Click the X to the right of the message to close the message.

   > **Note:** If **Include Dependencies** is checked, the total number of items in the export cart may exceed the number of items checked in the list.

8. Click **View Cart** or 🛒 to continue with the export.

   The Check out Page opens.

9. Review the listing of objects.
10. To remove an object from the listing, use one of the following methods:

    - Click **Remove** next to the named object.
    - Select the check box next to the named object and click **Remove selected items**.

- Use the **Search** tool to search the list for a specific string and click **Remove filtered items**.

11. To view a filtered list of export objects by object type, click the **Type** list, select the object types, and click **Search**.

    All of the selected object types are shown in the list.

    > **Note:** The remaining object types are still in the export cart. To view the entire contents, select **All** from the **Type** list and click **Search**.

12. After the list of objects is ready, click **Export Cart**.

    A warning message displays verifying that you want to export the select items in the shopping cart.

13. Click **OK** to proceed with the export.

14. Click **Clear Cart**.

    The contents of the export cart are removed and you can proceed with another export.

A ZIP file is downloaded to your local computer. The filename uses the following naming convention: `CMP_export_`*`yyyyMMddhhmmss`*`.zip`.

Exporting All Configuration Objects

The **Export All** function exports a ZIP file to the local computer that contains all of the exportable types of configuration objects.

> **Note:** **Export All** does not include associations between policies and MPE templates.

To export all configuration objects:

1. From the **System Administration** section of the navigation pane, expand **Import/Export**.

2. Select **Export All**.

   The Bulk Export page appears.

3. Click **Export All**.

A ZIP file is downloaded to your local computer. The filename uses the following naming convention: `CMP_export_`*`yyyyMMddhhmmss`*`.zip`.

About the Manager Reports

The Manager Reports function provides information about the CMP cluster itself. This information is similar to the Cluster Information Report for MPEclusters. The display is refreshed every ten seconds.

Viewing Manager Reports

To view Manager Reports:

1. From the **System Administration** section of the navigation pane, select **Reports**.

   The Manager Reports page opens in the work area.

2. To manage the reports:

   - To pause refreshing the display, click **Pause**.

     The report does not update.

   - To resume refreshing the display, click **Resume**.

     The report refreshes every 10 seconds.

   - To reset the display counters, such as the number of blade failures, click **Reset Counters**.

     The counters in the report are reset to zero.

   The fields that are displayed in the Manager Reports include the following:

   - **Cluster Name and Designation**

     The name of the cluster and whether it is the primary **(P)** or secondary **(S)** site.

   - **Mode**

     The status of the cluster:

     – **Active**

       The cluster is managing the Policy Management network.

     – **Standby**

       The cluster is not currently managing the Policy Management network.

     – **Paused**

       The report is paused. Click **Resume** to return the **Mode** to **Active**.

   - **Cluster Status**

     The status of the servers within the cluster:

     – **On-line**

       If one server, it is active; if two servers, one is active and one is standby.

     – **Degraded**

       One server is active, but the other server is not available.

     – **Out-Of-Service**

       Neither server is active.

     – **No Data**

       The CMP system cannot reach the server.

Also within the Manager Reports is a listing of the **Blades** (that is, the servers) contained within the cluster. A symbol (🔌) indicates which server currently has the external connection (that is, which server is the active server). The report also lists the following server-specific information:

- **Overall** section

  Displays the current topology:

  – **State**

    Active, Standby, or Forced Standby

  – **Blade Failures**

    Counters

  – **Uptime**

    The time providing active or standby service

    For the definitions of these states, see Server Status.

- **Utilization**

  Displays the current usage statistics:

  – **Disk**

    Percentage utilization of disk (of the `/var/camiant` file system)

  – **CPU**

    Average value for the CPU utilization

  – **Memory**

    Average value for Memory use

The **Actions** links let you **Restart** the CMP software on the server or **Reboot** the server.

About the Trace Log

The trace log is part of system administration records notifications for management activity on the CMP system. For information on configuring the cluster-level messages written to the trace log, see Configuring Log Settings. For information on configuring the system-level messages written to the trace log, see Configuring the Trace Log.

Configuring the Trace Log

You can configure the trace log severity message levels for the CMP system.

Note: For information on configuring debug logs, see #unique_313.

To configure the trace log:

1. From the **System Administration** section of the navigation pane, select **Trace Log**.

   The Trace Log page opens in the work area.

2. Click **Modify**.

   The Modify Trace Log Settings page opens.

3. Select the **Trace Log Level** from the list.

   This setting indicates the minimum severity of messages that are recorded in the trace log. These severity levels correspond to the syslog message severities from RFC 3164. Adjusting this setting allows new notifications, at or above the configured severity, to be recorded in the trace log. The levels are:

   - **Emergency**

     Provides the least amount of logging, recording only notification of events causing the system to be unusable.

   - **Alert**

     Action must be taken immediately in order to prevent an unusable system.

   - **Critical**

     Events causing service impact to operations.

   - **Error**

     Designates error events which may or may not be fatal to the application.

   - **Warning** (default)

     Designates potentially harmful situations.

   - **Notice**

     Provides messages that may be of significant interest that occur during normal operation.

   - **Info**

     Designates informational messages highlighting overall progress of the application.

   - **Debug**

     Designates information events of lower importance.

   > **Caution:** Before changing the default logging level, consider the implications. Lowering the log level setting from its default value (for example, from **Warning** to **Info**) causes more notifications to be recorded in the log and can adversely affect performance. Similarly, raising the log level setting (for example, from **Warning** to **Alert**) causes fewer notifications to be recorded in the log and may cause you to miss important notifications.

4. Click **Save**.

The trace log is configured.

Viewing the Trace Log

To view the Trace Log:

1. From the **System Administration** section of the navigation pane, select **Trace Log**.

   The Trace Log page opens in the work area.

2. Click **View Trace Log**.

   The Trace Log Viewer window opens. While data is being retrieved, the progress message indicating that the trace logs are being scanned appears.

   All events contain the following information:

   - **Date/Time —** Event timestamp. This time is relative to the server time.
   - **Code —** The event code. For information about event codes and messages, see the *Troubleshooting Reference*.
   - **Severity —** Severity level of the event. Application-level trace log entries are not logged at a higher level than error.
   - **Message —** The message associated with the event. If additional information is available, the event entry shows as a link. Click on the link to see additional detail in the frame below.

   By default, the window displays 25 events per page. You can change this to 50, 75, or 100 events per page by selecting a value from the **Display results per page** list.

   Events that occur after the Trace Log Viewer starts are not visible until you refresh the display. To refresh the display, click one of the following buttons:

   - **Show Most Recent —** Applies filter settings and refreshes the display. This displays the most recent log entries that fit the filtering criteria.
   - **Next/Prev** — When the number of trace log entries exceeds the page limit, pagination is applied. Use the **Prev** or **Next** buttons to navigate through the trace log entries. When the **Next** button is not visible, you have reached the most recent log entries; when the **Prev** button is not visible, you have reached the oldest log entries.
   - **First/Last** — When the number of trace log entries exceeds the page limit, pagination is applied. Use the **First** and **Last** buttons to navigate to the beginning or end of the trace log. When the **Last** button is not visible, you have reached the end; when the **First** button is not visible, you have reached the beginning.

3. To view the trace log for a different server, select from the **Trace Log Viewer for Server** and click **Search**.

   The trace log for the selected server displays.

4. Click **Close**.

Filtering the Trace Log

The Trace Log can contain a large number of messages. To reduce the number, the log can be filtered using several criteria.

To filter the trace log information:

1. From the **System Administration** section of the navigation pane, select **Trace Log**.

   The Trace Log page opens in the work area.

2. Click **View Trace Log**.

   The Trace Log Viewer window opens. While the data is being retrieved, the a progress message displays.

3. To view the trace log for a different server, select from the **Trace Log Viewer for Server** and click **Search**.

   The trace log for the selected server displays.

4. Specify the filtering parameters using any of the following fields.

   - **Start Date/Time**

     Click 📅 (calendar icon), specify a date and time, and then click **Enter**.

   - **End Date/Time**

     Click 📅 (calendar icon), specify a date and time, and then click **Enter**.

   - **Trace Codes**

     Enter one or a comma-separated list of trace code IDs. Trace code IDs are integers up to 10 digits long.

   - **Use timezone of remote server for Start Date/Time**

     Select to use the time of a remote server (if it is in a different time zone) instead of the time of the CMP server.

   - **Severity**

     Filter by severity level. Events with the selected severity and higher are displayed. For example, if the severity selected is **Warning**, the trace log displays events with the severity level Warning.

   - **Contains**

     Enter a text string to search. For example, if you enter `connection`, all events containing the word `connection` display. This field does not use wildcards and is not case specific.

   Note: The **Start Date/Time** setting overrides the **Contains** setting. For example, if you search for events happening this month, and search for a string in events last month and this month, only results from this month are listed.

5. Click **Search**.

   The filtered log displays.

6. Click **Close**.

   The Trace Log Viewer window closes.

Modifying the Trace Log Configuration

To configure the trace log display:

1. From the **System Administration** section of the navigation pane, select **Trace Log**.

The Trace Log page opens in the work area, displaying the current trace log configuration.

2. Click **Modify**.

   The Modify Trace Log Settings page opens.

3. Define the settings.

   For a description of the settings, see Configuring Log Settings.

4. Click **Save**.

The trace log configuration is modified.

Viewing the Audit Log

The CMP lets you track and view configuration changes within the system. Using the audit log, you can track and monitor each configuration event, providing you better system control. The audit log is stored in the database, so it is backed up and can be restored.

To view the audit log:

1. From the **System Administration** section of the navigation pane, select **Audit Log**.

   The Audit Log page opens in the work area.

2. Click **Show All**.

   The **Audit Log** opens. (Figure 15-3 shows an example.)

   Figure 15-3 Audit Log

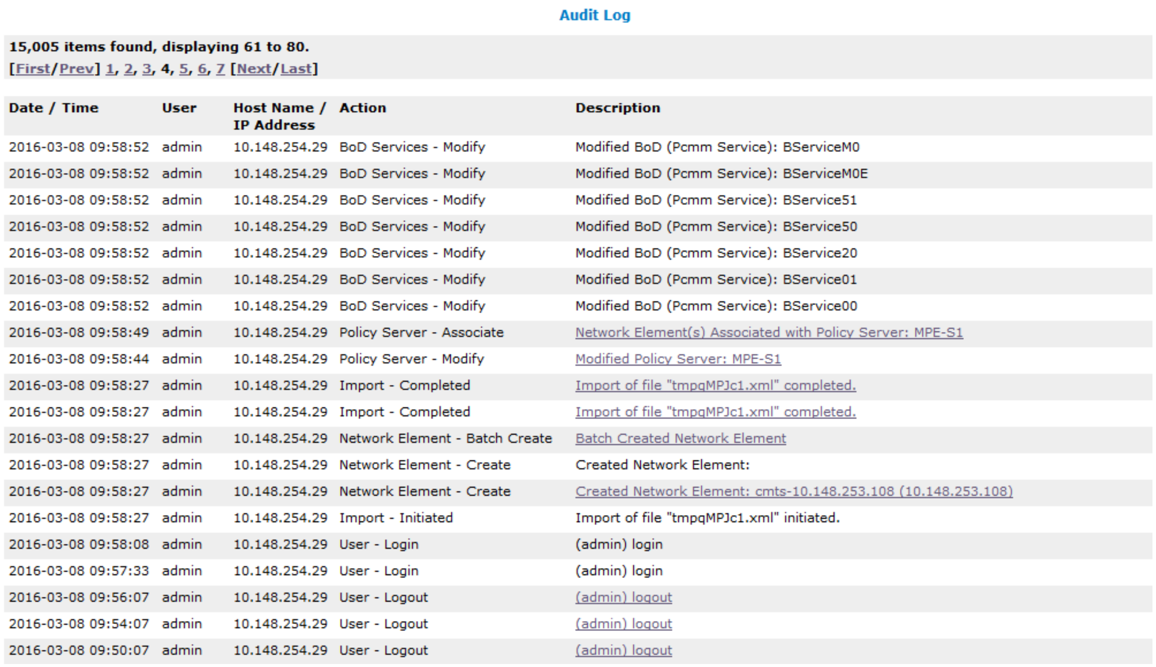

Audit Log

15,005 items found, displaying 61 to 80.
[First/Prev] 1, 2, 3, 4, 5, 6, 7 [Next/Last]

Date / Time	User	Host Name / IP Address	Action	Description
2016-03-08 09:58:52	admin	10.148.254.29	BoD Services - Modify	Modified BoD (Pcmm Service): BServiceM0
2016-03-08 09:58:52	admin	10.148.254.29	BoD Services - Modify	Modified BoD (Pcmm Service): BServiceM0E
2016-03-08 09:58:52	admin	10.148.254.29	BoD Services - Modify	Modified BoD (Pcmm Service): BService51
2016-03-08 09:58:52	admin	10.148.254.29	BoD Services - Modify	Modified BoD (Pcmm Service): BService50
2016-03-08 09:58:52	admin	10.148.254.29	BoD Services - Modify	Modified BoD (Pcmm Service): BService20
2016-03-08 09:58:52	admin	10.148.254.29	BoD Services - Modify	Modified BoD (Pcmm Service): BService01
2016-03-08 09:58:52	admin	10.148.254.29	BoD Services - Modify	Modified BoD (Pcmm Service): BService00
2016-03-08 09:58:49	admin	10.148.254.29	Policy Server - Associate	Network Element(s) Associated with Policy Server: MPE-S1
2016-03-08 09:58:44	admin	10.148.254.29	Policy Server - Modify	Modified Policy Server: MPE-S1
2016-03-08 09:58:27	admin	10.148.254.29	Import - Completed	Import of file "tmpqMPJc1.xml" completed.
2016-03-08 09:58:27	admin	10.148.254.29	Import - Completed	Import of file "tmpqMPJc1.xml" completed.
2016-03-08 09:58:27	admin	10.148.254.29	Network Element - Batch Create	Batch Created Network Element
2016-03-08 09:58:27	admin	10.148.254.29	Network Element - Create	Created Network Element:
2016-03-08 09:58:27	admin	10.148.254.29	Network Element - Create	Created Network Element: cmts-10.148.253.108 (10.148.253.108)
2016-03-08 09:58:27	admin	10.148.254.29	Import - Initiated	Import of file "tmpqMPJc1.xml" initiated.
2016-03-08 09:58:08	admin	10.148.254.29	User - Login	(admin) login
2016-03-08 09:57:33	admin	10.148.254.29	User - Login	(admin) login
2016-03-08 09:56:07	admin	10.148.254.29	User - Logout	(admin) logout
2016-03-08 09:54:07	admin	10.148.254.29	User - Logout	(admin) logout
2016-03-08 09:50:07	admin	10.148.254.29	User - Logout	(admin) logout

3. (Optional) Click **Refine Search** located at the bottom of the page to filter the search results. (See Searching for Audit Log Entries.)

4. (Optional) Click the underlined description for a detailed description of an item.

   The details of the event display. (Figure 15-4 shows an example.)

Figure 15-4 Audit Log Details

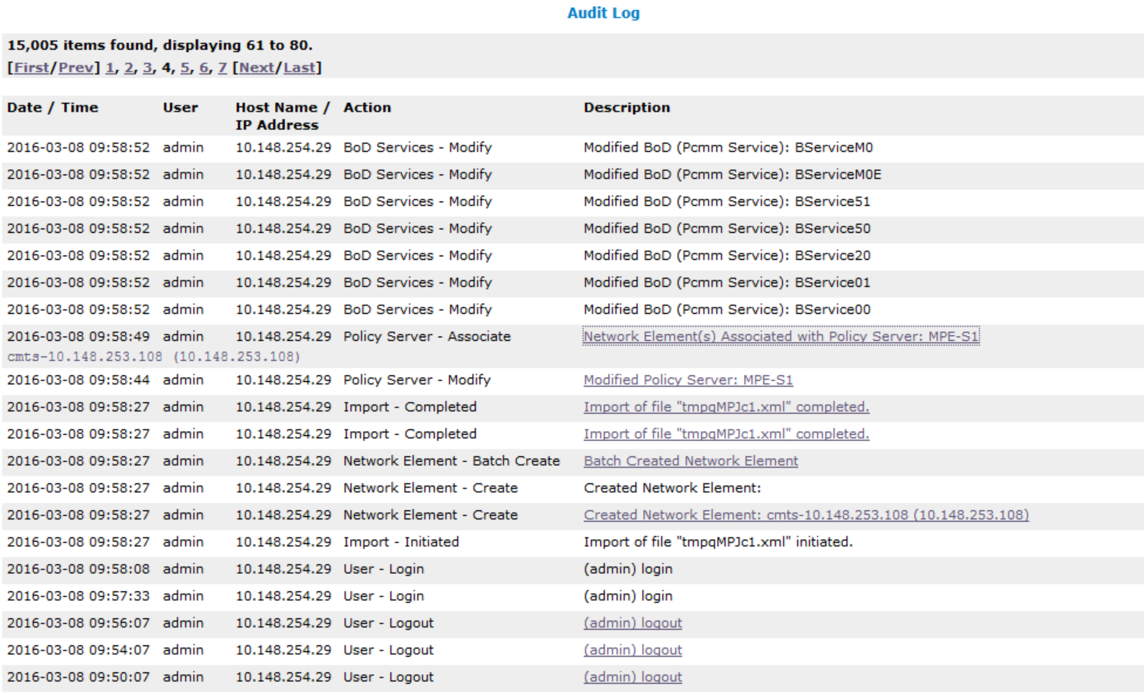

Audit Log

15,005 items found, displaying 61 to 80.
[First/Prev] 1, 2, 3, 4, 5, 6, 7 [Next/Last]

Date / Time	User	Host Name / IP Address	Action	Description
2016-03-08 09:58:52	admin	10.148.254.29	BoD Services - Modify	Modified BoD (Pcmm Service): BServiceM0
2016-03-08 09:58:52	admin	10.148.254.29	BoD Services - Modify	Modified BoD (Pcmm Service): BServiceM0E
2016-03-08 09:58:52	admin	10.148.254.29	BoD Services - Modify	Modified BoD (Pcmm Service): BService51
2016-03-08 09:58:52	admin	10.148.254.29	BoD Services - Modify	Modified BoD (Pcmm Service): BService50
2016-03-08 09:58:52	admin	10.148.254.29	BoD Services - Modify	Modified BoD (Pcmm Service): BService20
2016-03-08 09:58:52	admin	10.148.254.29	BoD Services - Modify	Modified BoD (Pcmm Service): BService01
2016-03-08 09:58:52	admin	10.148.254.29	BoD Services - Modify	Modified BoD (Pcmm Service): BService00
2016-03-08 09:58:49	admin	10.148.254.29	Policy Server - Associate	Network Element(s) Associated with Policy Server: MPE-S1
cmts-10.148.253.108 (10.148.253.108)				
2016-03-08 09:58:44	admin	10.148.254.29	Policy Server - Modify	Modified Policy Server: MPE-S1
2016-03-08 09:58:27	admin	10.148.254.29	Import - Completed	Import of file "tmpqMPJc1.xml" completed.
2016-03-08 09:58:27	admin	10.148.254.29	Import - Completed	Import of file "tmpqMPJc1.xml" completed.
2016-03-08 09:58:27	admin	10.148.254.29	Network Element - Batch Create	Batch Created Network Element
2016-03-08 09:58:27	admin	10.148.254.29	Network Element - Create	Created Network Element:
2016-03-08 09:58:27	admin	10.148.254.29	Network Element - Create	Created Network Element: cmts-10.148.253.108 (10.148.253.108)
2016-03-08 09:58:27	admin	10.148.254.29	Import - Initiated	Import of file "tmpqMPJc1.xml" initiated.
2016-03-08 09:58:08	admin	10.148.254.29	User - Login	(admin) login
2016-03-08 09:57:33	admin	10.148.254.29	User - Login	(admin) login
2016-03-08 09:56:07	admin	10.148.254.29	User - Logout	(admin) logout
2016-03-08 09:54:07	admin	10.148.254.29	User - Logout	(admin) logout
2016-03-08 09:50:07	admin	10.148.254.29	User - Logout	(admin) logout

Searching for Audit Log Entries

To search for the Audit Log entries:

1. From the **System Administration** section of the navigation pane, select **Audit Log**.

   The Audit Log page opens in the work area.

2. Click **Search**.

   The Audit Log Search Restrictions page opens.

3. Define the following items, depending on how restrictive you want the audit log search to be:

   - **From/To** — Click 📅 (calendar icon), specify a date and time then click **Enter**.
   - **Action by User Name(s)** — Enter the **User Name** of the user or users to audit.
   - **Action on Policy Server(s)** — Enter the name of the Policy Management device to audit.
   - **Audit Log Items to Show** — Specifies the category of items to audit:

     a. When you select some categories, a **Name** field displays, which lets you enter a search string.

     b. Leave the **Name** field blank to include all items.

c. When you select a category, an **Actions** link displays, which lets you select individual audit log items within the category.

   By default all items in the category are selected, but you can select individual items instead.

By default you can specify three item categories. Click **More Lines** to add an additional audit log item category.

- **Results Forms** — Specifies the number of items per page to display, including which data to display (most recent or oldest items).

4. Click **Search**.

   The Audit Log displays search results.

Exporting or Purging Audit Log Data

You can export the audit log to a text file; the file name is `AuditLogExport.txt`.

Exporting Audit Log Data

You can export audit log data to a text file. The default file name is `AuditLogExport.txt`.

To export audit log data:

1. From the **System Administration** section of the navigation pane, select **Audit Log**.

   The Audit Log page opens in the work area.

2. Click **Export/Purge**.

   The Export and Purge Audit Log Items page opens.

3. In the **Items to Export** section, select one of the following options:

   a. **Export All Items** — Writes all audit log entries.

   b. **Export Through Date** — Click 📅 (calendar icon), and select a date.

4. Click **Export**.

   A standard File Download window opens; you can open or save the export file.

The audit log is exported.

Purging Audit Log Data

To purge data from the audit log:

1. From the **System Administration** section of the navigation pane, select **Audit Log**.

   The Audit Log page opens in the work area.

2. Click **Export/Purge**.

   The Export and Purge Audit Log Items page opens.

3. In the **Items to Purge** section, click 📅 (calendar icon) and select a date.

4. Click **Purge**.

   You are prompted with a confirmation message.

5. Click **OK**.

The data is purged from the audit log.

About the Manager Log

The Manager log is a series of files that records information about the operation of CMP components and subcomponents. Log data is appended to component and subcomponent logs (for example, tomcat.log and HTTP.log). When the maximum file size is reached a new file is started. When the maximum number of files is reached the oldest file is deleted.

The Manager Log page is available when the CMP system operates in debug mode. Contact My Oracle Support for more information about enabling debug mode. The Manager Log page lets you configure system-wide default values for the available debug logs for Policy Management components and subcomponents.

Configuring the Manager Log

Note: To view the **Manager Log** menu in the navigation pane, you must enable **Debug Mode**. Contact My Oracle Support for more information.

You can use the Manager Log to configure the default debug logs severity message levels for the CMP system subcomponents.

Note: For information on configuring debug logs, see #unique_313.

To configure the Manager Log:

1. From the **System Administration** section of the navigation pane, select **Manager Log**.

   The Manager Log page opens in the work area.

2. Click **Modify**.

   The Manager Log page becomes editable.

3. In the **Tomcat Log Configuration** section of the page:

   a. Enter the **Scan Period (Seconds)**.

      The default value is 20 seconds.

   b. Select the **Root Log Level**:

      Available options are:

      - **OFF**

        Turns off logging.

- **ERROR**

  Designates error events which may or may not be fatal to the application.

- **WARN** (default)

  Designates potentially harmful situations.

- **INFO**

  Designates informational messages highlighting overall progress of the application.

- **DEBUG**

  Designates information events of lower importance.

- **TRACE**

  Designates informational events of very low importance.

- **ALL**

  Records all logging levels.

> **Caution:** Before changing any default logging level, consider the implications. Lowering the log level setting from its default value (for example, from **WARN** to **INFO**) causes more notifications to be recorded in the log and can adversely affect performance. Similarly, raising the log level setting (for example, from **WARN** to **ERROR**) causes fewer notifications to be recorded in the log and may cause you to miss important notifications.

4. In the **File Apppender Configuration** section of the page, for each appender file listed:

   a. Enter the **Maximum File Size**.

      The file size is in megabytes.

   b. Enter the **Maximum File Count**.

5. In the **DC Log Configuration** section of the page:

   a. Enter the **Scan Period (Seconds)**.

      The default value is 20 seconds.

   b. Select the **Root Log Level**.

6. In the **File Apppender Configuration** section of the page, for each appender file listed:

   a. Enter the **Maximum File Size**.

      The file size is in megabytes.

   b. Enter the **Maximum File Count**.

7. Click **Save**.

The Manager Log is configured.

Managing Scheduled Tasks

The CMP system runs batch jobs to complete certain operations. These tasks are scheduled to run at regular intervals, with some tasks scheduled to run in a certain order. You can change the scheduling, reschedule, enable or disable these tasks to better manage network load or to propagate a network element change to the Policy Management devices on demand. You can also abort a running task.

> **Caution:** Oracle recommends that you follow the order in which scheduled tasks are listed. Serious system problems can occur if the order is changed. Consult My Oracle Support before changing the order of task execution.

The tasks include:

Alert Aging
Ensures that alerts age out and are eventually removed from the CMP database. (The valid range is 1 to 365 days.)

Health Checker
Periodically checks the MPE devices to ensure that they are online.

OM Statistics
Periodically retrieves Operational Measurement (OM) statistics from all Policy Management devices.

The Operational Measurements XML interface retrieves operational counters from the system. The OM interface requires that the OM Statistics scheduled task be running on the CMP system. After the specified Stats Collection Period, this task collects the operational counters from the Policy Management devices in the network and records them in the CMP database; the data is then available for query via the OM XML interface. You can configure the task to poll at intervals between 5 minutes and 24 hours, with a default value of 15 minutes; the system keeps the data available for query for 1 to 30 days, with a default value of 7 days. The recommended settings for this task will vary depending on the volume of data you are collecting.

When you request OM statistics, the data for the response is taken from the information that has been collected by this task. You must gather data using the OM Statistics scheduled task if you want data available for subsequent OM queries.

Most values returned as part of the response are presented as the positive change between the start time and end time. To calculate a response, you must have a minimum of two recorded values available; thus you must run the OM Statistics task at least twice in a given time period before you can obtain any statistical data from the OSSI XML interface. *OSSI XML Interface Definitions Reference* describes the OM Interface and the OM Statistics in detail.

PM Statistics Files Uploading
Uploads Performance Management (PM) statistics to the remote FTP server.

PM Statistics
Queries statistics data from the OSSI/XML interface or the TPD platform and writes the data to an XML file.

Legacy OM Statistics
Periodically retrieves OM statistics from MPE devices executing the previous release of Policy Management software. This task should be run only during migration between software releases.

OSSI Distributor Task
(optional) Reads from the database topology and subscriber data that has entered the CMP database using the OSSI Interface, and distributes the data to the MA servers.

Subnet SNMP Collector
Collects all subnet information residing on the CMTS devices by polling, via SNMP, all CMTS devices for all subnets and then stores them in the local database.

Service Class SNMP Collector
Polls, using SNMP, all CMTS devices for the configured service classes and then stores them in the local database.

Subscriber SNMP Collector
Polls, using SNMP, all CMTS devices for the configured subscribers and then stores them in the local database.

CMTS Distributor
Reads CMTS topology data from the CMP database and then distributes it to the appropriate Policy Management devices within the system.

Subscriber Distributor
Reads subscriber data from the CMP database and then distributes it to the appropriate Policy Management devices within the system.

CMTS MA Collector
(optional) Polls all of the MA Servers in the system for subnet and service class data on each CMTS.

PCMM Routing Distribution
Detects changes in the CMTS subnet information, and then forwards this information to any upstream MPE devices configured in a routing hierarchy.

Replication Statistics
Generates replication statistics for MPE and BoD servers.

Configuring a Scheduled Task

To configure an individual task:

1. From the **System Administration** section of the navigation pane, select **Scheduled Tasks**.

   The Scheduled Task Administration page opens showing the settings for the available scheduled tasks.

2. Click **Refresh** to update the page.

3. To configure a scheduled task, click the task name.

   The page displays the current settings and status for the selected scheduled task.

4. Click **Settings**.

Note: The **Health Checker** task has no configurable settings.

Note: The **Subnet Overlap Detector** task has no configurable settings.

5. To configure a **Stats File Synchronizations** task (maximum of 4) for CMP to store copies of stats files to a remote external system:

   a. Enter the **Host Name / IP Address** for the remote server.

      You can use either the FQDN or IP address (either IPv4 or IPv6 format).

   b. Enter the SSH **User Name** required to access the remote server.

   c. Enter the SSH **Password**.

   d. Enter the **Path of Remote Repository** for storing the stats files.

   e. Enter the **Retry Limit** for the number of times to retry synchronizing the stats files (default is 3 with a maximum of 5) if the remote server is unreachable.

6. To configure the **OM Statistics** or **Replication Statistics** task:

   a. Enter the **Number of days to keep statistical data**.

      This is the number of days to retain the specified statistics file. The valid range is from 1 to 30 days with a default of seven days.

7. To configure the **Stats Files Generator**:

   a. Enter the path for the **Local Repository**.

      This is the path on the external server where stats files are replicated. The default root directory for the repository is `/var/camiant/stats_export`. Two levels of subdirectories will be created under this root directory. An MPE or BoD cluster will have a corresponding first level subdirectory. Each stats file type will have its own subdirectory under the cluster-level directory. All the stats files will be created under the stats type subdirectory.

   b. Enter the **Maximum age to keep files (hours)**.

      The default is 72 hours.

   c. Select the **File Format** from the list.

      Available formats are **CSV** or **XML** (default).

   d. Select the **Stats Type** from the list:

      See Generated Statistics for a list of available statistics.

      - Click **Select All** to have the generator collect statistics for all stats types.
      - Click **Inverse All** to deselect any selected stats types and select any unselected stats types.
      - Select each individual stats types.

8. If **Wireless-C** and **CMPP** modes are enabled, to configure the **PM Statistics** task:

a. Enter the **Number of days to keep statistical data**.

   This is the number of days to retain the specified statistics file. The valid range is from 1 to 7 days with a default of seven days.

b. Enter the **Number of Max TPS Capacity**.

   This is the maximum transactions per second. The default value is 5000

9. If **Wireless-C** and **CMPP** modes are enabled, to configure the **PM Statistics Files Uploading** task:

   a. Enter the **Host Name / IP Address** for the remote server.

      You can use either the FQDN or IP address (either IPv4 or IPv6 format).

   b. Enter the FTP **User Name** required to access the remote server.

   c. Enter the FTP **Password**.

   d. Enter the **Path of Remote Repository** for storing the stats files.

   e. Enter the **Retry Limit** for the number of times to retry synchronizing the stats files (default is 3 with a maximum of 5) if the remote server is unreachable.

   f. Select to enable **Security FTP**.

10. If **Wireless-C** and **CMPP** modes are enabled, to configure the **SMS Statistics Files Uploading** task:

    a. Enter the **Host Name / IP Address** for the remote server.

       You can use either the FQDN or IP address (either IPv4 or IPv6 format).

    b. Enter the FTP **User Name** required to access the remote server.

    c. Enter the FTP **Password**.

    d. Enter the **Path of Remote Repository** for storing the stats files.

    e. Enter the **Retry Limit** for the number of times to retry synchronizing the stats files (default is 3 with a maximum of 5) if the remote server is unreachable.

    f. Select to enable **Security FTP**.

11. Click **Save**.

    a. Click **Save**.

The scheduled task is configured.

Rescheduling a Task

To reschedule a scheduled task:

1. From the **System Administration** section of the navigation pane, select **Scheduled Tasks**.

   The Scheduled Task Administration page opens showing the settings for the available scheduled tasks.

2. To configure a scheduled task, click the task name.

   The page displays the current settings and status for the selected scheduled task.

3. Click **Reschedule**.

   The reschedule configuration settings appear.

4. To **Schedule by Interval**:

   a. Select the date and time for the **Next Run Time**.

   b. For the Run Interval, select the **Hours** or **Minutes**.

      Valid intervals are from 0 to 24 hours in 5-minute increments.

5. To schedule the task **Following Another Task**, select the **Task to Follow** from the list.

6. Click **Save**.

The scheduled task is rescheduled.

Enabling or Disabling a Scheduled Task

To enable or disable an individual task:

1. From the **System Administration** section of the navigation pane, select **Scheduled Tasks**.

   The Scheduled Task Administration page opens showing the settings for the available scheduled tasks.

2. To configure a scheduled task, click the task name.

   The page displays the current settings and status for the selected scheduled task.

3. To disable an enabled task, click **Disable**.

   The scheduled task is disabled and the button text changes to **Enable**.

4. To enable a disabled task, click **Enable**.

   The scheduled task is enabled and statistics files will be generated based on the task's configuration settings. The button text changes to **Disable**.

5. Click **OK** to acknowledge the change.

The specified task is disabled or enabled.

About Managing Users

The CMP system lets you configure the following user attributes:

Roles
Determines the actions (and the access level) a user can perform within the CMP system. See About User Roles for details.

Scopes
Determines the network element groups and Policy Management device groups a user can perform actions on and providing a context for a role. See About User Scopes for details.

Users
After you define roles and scopes, you can assign them to user profiles. See About User Profiles for details.

External Authentication
Enables the CMP system to authenticate users using either RADIUS or SANE Authentication. These users must match the RADIUS Server account information before access is permitted. See About External Authentication for details.

About User Roles

The CMP system uses roles to configure what a user can do within the CMP system. Assigning roles to the various users that access the CMP system lets you control who can configure and access features within the CMP system. The default roles are:

Administrator
Permits full read/write access to all functions. You cannot delete the **Administrator** role.

Operator
Permits full read/write access to all Policy Management server management and configuration functions. Access is also permitted to all system administration functions except **User Management**.

Viewer
Permits read-only access to functions associated with Policy Management server management and configuration. Full access is also permitted to some of the system administration functions, such as **Change Password**.

> **Note:** When you create a new role, ensure that it has appropriate access to all functions you intend the role to use. For example, if you create a role for third-party access to OSSI functions, but it does not have the system administration privilege **Import/Export** set to **Show**, a user given that role cannot perform OSSI queries.

The CMP system lets you perform the following role management actions:

- Creating a Role
- Modifying a Role
- Deleting a Role

Creating a Role

To create a role:

1. From the **System Administration** section of the navigation pane, select **User Management**.

The content tree displays the **User Management** group.

2. From the content tree, select the **Roles** group.

   The Role Administration page opens in the work area.

3. Click **Create Role**.

   By default, all access for privileges are set to either **Hide** (that is, the functions do not appear to users of the role, so access must be explicitly granted) or **Read-Only** (that is, information can be displayed but not changed).

   The New Role page opens.

4. Enter the **Name** for the new role.

   The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_).

5. Enter a **Description/Location** (optional).

   Free-form text.

6. **Policy Server Privileges—**Defines access to the following MPE device management functions (with the access **Hide**, **Read-Only**, or **Read-Write**):

   - **Configuration**
   - **Applications**
   - **Quota Profiles & Conventions**
   - **PRA Lists**
   - **Traffic Profiles**
   - **Roaming Profile**
   - **Management Agents**
   - **Custom AVP Definitions**
   - **AVP Definitions**
   - **Custom VSA Definitions**
   - **Notification Server**
   - **SMS Gateway**
   - **Global Configuration Settings**
   - **Bulk Operation**

7. **Network Privileges—**Defines access to the network management functions (with the access **Hide**, **Read-Only**, or **Read-Write**):

   - **Network Elements**
   - **Topology**

8. **Policy Management Privileges—**Defines access to the policy management functions:
   - **Policy Library** (with the access **Hide**, **Read-Only**, **Read and Deploy**, or **Read, Deploy, and Write**)
   - **Template Library** (with the access **Hide**, **Read-Only**, or **Read-Write**)
   - **Policy Table Library** (with the access **Hide**, **Read-Only**, or **Read-Write**)
   - **Policy Checkpoint** (with the access **Hide**, **Read-Only**, or **Read-Write**)
9. **System Wide Reports Privileges—**Defines access to the system-wide reports functions:
10. **Platform Setting Privileges—**Defines access to the platform setting functions:
    - **Topology Settings** (with the access **Hide**, **Read-Only**, or **Read-Write**)
    - **Server Operation** (with the access **Hide** or **Read-Write**)
11. **Upgrade Manager Privileges—**Defines access to software upgrade functions:
    - **ISO Maintenance** (with the access **Hide**, **Read-Only**, or **Read-Write**)
12. **System Administration Privileges—**Defines access to system administration functions:
    - **Import / Export** (with the access **Hide** or **Show**)
    - **Operational Measurements** (with the access **Hide** or **Read-Only**)
    - **User Management** (with the access **Hide**, **Read-Only**, or **Read-Write**)
    - **Scheduled Tasks** (with the access **Hide** or **Read-Write**)
    - **Trace Log of CMP** (with the access **Hide**, **Read-Only**, or **Read-Write**)
    - **Audit Log** (with the access **Hide**, **Read-Only**, or **Read-Write**)
    - **Audit Log User Info** (with the access **Hide** or **Show**)
    - **Alarms** (with the access **Hide**, **Read-Only**, or **Read-Write**)
    - **Password Strength** (with the access **Read-Only** or **Read-Write**)
    - **Push Method for Statistics** (with the access **Read-Only** or **Read-Write**)

      If set to **Read-Only**, the following fields are displayed for the **Stats File Generator** (see Managing Scheduled Tasks) setting:
      - **Name**
      - **Description**
      - **Last Exit Status**
      - **Current State**
      - **Last Start Time**
      - **Last End Time**

  - **Follows Task**

  Task Settings

  - **Local Repository—**Root directory of the local repository.
  - **Maximum age to keep files (hours)—**Stats file retention period. Default is 72 hours.
  - **File Format—**Either CSV or XML (default).
  - **Stats Type—**Any stats type can be selected to generate stats. If you do not select a stats type, the task will not run normally.

- **Debug Options** (with the access **Read-Only** or **Read-Write**)

> **Note:** By default, the **Read-Write** privilege for this is only available for members of the Administrator role. Other roles have **Read-Only** access by default.

New tasks are created to synchronize stats files. These tasks perform a retry if a remote server is unreachable. The following fields are displayed for the Stats Files Synchronization setting:

- **Remove Server Information**
  - **Host Name/IP Address**
  - **User Name**
  - **Password**
  - **Path of Remote Repository**
- **Retry Limit—**You have a limit of three retries in one-minute intervals.

> **Note:** There are a total of four synchronized tasks which are supported but cannot be edited.

13. Click **Save**.

The role is created.

Modifying a Role

To modify a role:

1. From the **System Administration** section of the navigation pane, select **User Management**.

   The content tree displays the **User Management** group.

2. From the content tree, select the **Roles** group.

   The Role Administration page opens in the work area.

3. Select the role to modify.

   The Role Administration page displays the configuration of the role.

4. Click **Modify**.

   The Modify Role page opens.

5. Modify role information as necessary.

   See Creating a Role for a description of the fields on this page.

6. Click **Save**.

The role is modified.

Deleting a Role

> **Note:** You can delete any role except the **Administrator** role.

You cannot delete a role that is in use. You must remove any users assigned to the role before deleting it.

To delete a role:

1. From the **System Administration** section of the navigation pane, select **User Management**.

   The content tree displays the **User Management** group.

2. From the content tree, select the **Roles** group.

   The Role Administration page opens in the work area.

3. Delete the role using one of the following methods:

   - From the work area, click the 🗑 (Delete icon) located next to the role to delete.
   - From the content tree, select the role to delete (role information displays in the work area), then click **Delete**.

   A confirmation message displays.

4. Click **OK**.

The role's information is deleted from the CMP database.

About User Scopes

The CMP system uses scopes to define the network element groups and Policy Management device groups that a user can access, which provides operational context for a role.

> **Note:** You can assign a user more than one scope.

The CMP system allows you to perform the following scope management actions:

- Creating a Scope for CMP Servers
- Modifying a Scope

- Deleting a Scope

Creating a Scope for CMP Servers

Scopes allow you to control what areas or devices in a network a user can manage. The default scope, **Global**, contains all items defined within the CMP database. After you define a scope you can assign it to users.

To create a scope:

1. From the **System Administration** section of the navigation pane, select **User Management**.

   The content tree displays the **User Management** group.

2. In the content tree, select **Scopes**.

   The Scope Administration page opens in the work area.

3. Click **Create Scope**.

   The New Scope page opens.

4. Enter the **Name** for the new scope.

   The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_).

5. Enter the **Description/Location** (optional).

   Free-form text.

6. Select the **Policy Server Group**s included in this scope.

7. Select the **Network Element Groups** included in this scope.

8. Click **Save**.

The scope is created.

Modifying a Scope

To modify a scope:

1. From the **System Administration** section of the navigation pane, select **User Management**.

   The content tree displays the **User Management** group.

2. In the content tree, select **Scopes**.

   The Scope Administration page opens in the work area.

3. Select the scope you want to modify.

   The scope configuration appears.

4. Click **Modify**.

   The Modify Scope page opens.

Note: See Creating a Scope for CMP Servers for descriptions of the fields on this page.

5. Modify the scope as needed.
6. Click **Save**.

The scope is modified.

Deleting a Scope

Note: You cannot delete the **Global** scope.

To delete a scope:

1. From the **System Administration** section of the navigation pane, select **User Management**.

   The content tree displays the **User Management** group.

2. From the content tree, select **Scopes**.

   The Scope Administration page opens in the work area.

3. Delete the scope using one of the following methods:

   - From the work area, click 🗑 (Delete icon) located to the right of the scope you want to delete.
   - From the content tree, select the scope to delete and click **Delete**.

   A confirmation message appears.

4. Click **OK**.

The scope is deleted.

About User Profiles

User Management includes functions to create, modify, or delete user profiles. A user profile defines a user with a role and one or more scopes.

The CMP system is configured initially with the following default user profiles and passwords:

- `admin` (you cannot delete this profile)

Note: Oracle recommends changing the password after your first log in to the CMP system.

- `operator`
- `viewer`

The `admin` user is the only profile that cannot be deleted or have its username modified. The `admin` user is the only user that can create, modify, or delete other users, as well as log off all users.

Note: When logging in, the username is not case sensitive; however, the password is case sensitive.

The CMP system lets you perform the following user management actions:

- Creating a User Profile
- Modifying a User Profile
- Deleting a User Profile

Creating a User Profile

To create a user profile:

1. Log in to the CMP system as `admin`.
2. From the **System Administration** section of the navigation pane, select **User Management**.

   The content tree displays the **User Management** group.
3. In the content tree, select **Users**.

   The User Administration page opens in the work area.

   Note: The **Log Out All Users** button is visible only to the `admin` user.
4. Click **Create User**.

   The New User page opens.
5. Enter the **Username**.

   The name can only contain the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underline (_).

   Note: This value is not case sensitive.
6. Enter a **Description/Location** (optional).

   Free-form text.
7. Enter the **Password**.

   This value is case sensitive and must contain at least six characters; alphabetic, numeric, and special characters are allowed. This value must conform to the password strength rules. See Changing a Password for details on configuring password strength rules.
8. Enter to **Confirm Password** the **Password**.
9. Enter the number of days for the **Password Expiration Period(days; 0=never)**.

   Enter a value from `7` to `365`, or `0` to indicate that the password never expires. The default value is the system setting.

> **Note:** This setting overrides the system setting. For details on configuring password system settings, see #unique_338.

10. Select to **Force to Change Password**.

    If selected, this user must change passwords during the next log in. The default value is enabled.

11. Select a **Role** from the list.

12. Select one or more **Scopes** to assign to the user profile.

13. Click **Save**.

The user profile is created.

Creating a Customer User Management System Profile

To support identity management (IDM), the CMP system can accept HTTP or HTTPS connection requests from an external Customer User Management system to create, update, query, and delete user profiles. Requests and responses consist of XML documents.

For more information on the XML application programming interface, see *OSSI XML Interface Definitions Reference*.

To create a user profile for an external Customer User Management system:

1. Create a user profile as described in Creating a User Profile.
2. Assign the user profile a **Role** that includes the following privileges:
    - **Show** access for **Import/Export** privilege
    - **Read-Write** access for **User Management** privilege
3. Assign the user profile to the default **Global** scope.
4. Click **Save**.
    a. Click **Save**.

The user profile for the Customer User Management System is saved.

Modifying a User Profile

To modify a user profile:

1. Log in to the CMP system as `admin`.
2. From the **System Administration** section of the navigation pane, select **User Management**.

    The content tree displays the **User Management** group.

3. In the content tree, select **Users**.

    The User Administration page opens in the work area.

4. Select the user profile from the content tree.

The profile information page opens.

5. Click **Modify**.

   The Modify User page opens.

6. Modify the user profile.

   For field descriptions, see Creating a User Profile.

7. Click **Save**.

The user profile is modified.

Deleting a User Profile

Note: You cannot delete the `admin` user profile.

To delete a user profile:

1. Log in to the CMP system as `admin`.
2. From the **System Administration** section of the navigation pane, select **User Management**.

   The content tree displays the **User Management** group.

3. In the content tree, select **Users**.

   The User Administration page opens in the work area.

4. Delete the user profile using one of the following methods:

   - From the work area, select 🗑 (Delete icon) located to the right of the profile.
   - From the content tree, select the user profile and click **Delete**.

   A confirmation message displays.

5. Click **OK**.

The user profile is deleted.

About Locking and Unlocking User Profiles

A user is locked out after exceeding the login failure threshold, or if the `admin` user locks the user out.

A locked-out user sees the following message on the login page when attempting to log in: `"Your account is locked. Please contact the Administrator."`

Note: The admin user cannot lock the `admin` user account.

The CMP system includes the following functions:

- Locking a User

- Unlocking a User

Locking a User

To lock a user profile:

1. Log in to the CMP system as `admin`.
2. From the **System Administration** section of the navigation pane, select **User Management**.

   The content tree displays the **User Management** group.
3. In the content tree, select **Users**.

   The User Administration page opens in the work area.
4. Select the user profile from the content tree.

   The User Administration page displays configuration information about the user.
5. Click **Lock**.

   A confirmation message appears.
6. Click **OK**.

   - The user profile is locked.
   - The page displays a message indicating the account was locked successfully.
   - The **Lock** button becomes an **Unlock** button.
   - On the User Administration page, the **Locked Status** for the user shows `Locked`.

Unlocking a User

To unlock a user profile:

1. Log in to the CMP system as `admin`.
2. From the System Administration section of the navigation pane, select **User Management**.

   The content tree displays the **User Management** group.
3. In the content tree, select **Users**.

   The User Administration page opens in the work area.
4. Select the user profile from the content tree.

   The User Administration page displays configuration information about the user.
5. Click **Unlock**.

   A confirmation message appears.
6. Click **OK**.

- The user profile is unlocked.
- The page displays a message indicating that the user account was unlocked successfully.
- The **Unlock** button becomes a **Lock** button.
- On the User Administration page, the **Locked Status** for the user shows `Unlocked by Admin`.

Logging Out All Users

Note: Only the `admin` user can log out all users that are currently logged in to the CMP system. The `admin` user will not be logged out.

To log out all other users:

1. Log in to the CMP system as `admin`.
2. From the **System Administration** section of the navigation pane, select **User Management**.

   The content tree displays the **User Management** group.
3. In the content tree, select **Users**.

   The User Administration page opens in the work area.
4. Click **Log Out All Users**.

   A confirmation message appears.
5. Click **OK**.

Logged-in users are logged out from the CMP system.

About External Authentication

In addition to the built-in authentication functions, you can configure external authentication, RADIUS authentication, and SANE authentication of CMP users.

In the CMP system, you can manage the RADIUS Authentication and Account or the SANE Authentication external authentication method.

About RADIUS Authentication and Accounting

The CMP system supports RADIUS Authentication and Accounting. You can configure the CMP system to operate in a network environment including multiple authentication servers, one authentication server, or no servers.

If both primary and secondary authentication servers are defined, the authentication process is as follows:

1. The CMP system contacts the primary RADIUS server.

   If it responds with Accept or Reject, that action is followed.
2. If the primary server does not respond within a specified number of retries or before a timeout value, the CMP system contacts the secondary RADIUS server (if defined).

If it responds with Accept or Reject, that action is followed.

3. If the secondary server does not respond, the CMP system authenticates against its local database (if enabled).

4. If local authentication is not enabled, authentication fails.

5. The`admin` user is always authenticated locally, regardless of configuration settings.

This process provides a fail-safe mechanism for accessing the CMP system even in the face of misconfiguration or network problems that cause the RADIUS servers to become inaccessible.

RADIUS configuration involves the following steps:

1. See About Configuring the RADIUS Server for details on configuring the RADIUS server to accept authentication (and accounting, if used).

2. See About Defining CMP Users to the RADIUS Server for details on defining CMP users in the RADIUS server.

3. See About Associating Roles and Scopes for details on associating CMP users' roles and scopes with users on the CMP system

4. See About Defining the CMP System as a RADIUS Client for details on configuring the CMP system to work with the RADIUS server.

About Configuring the RADIUS Server

The RADIUS servers must be configured to authenticate clients and users on the CMP system. Some of the configuration values must be consistent with configuration parameters on the CMP system. (The RADIUS administrator is aware of the names and locations of the configuration files.)

See Enabling and Configuring RADIUS on the CMP System for details.

About Defining the CMP System as a RADIUS Client

The client file identifies the systems that use the RADIUS server to authenticate user access. A client should be defined as a single device. For example:

```
client 10.0.10.22 {
        secret = example
        shortname = MPE5
}
client 10.0.10.23 {
        secret = example
        shortname = CMP56
}
```

The best practice is to define IP addresses rather than FQDNs. If a netmask is not given, the default is `/32`. The shared `secret` (in this example, `example`) must be defined on both the RADIUS server and entered into the CMP configuration (see Enabling and Configuring RADIUS on the CMP System). The `shortname` is used as an alias.

About Defining CMP Users to the RADIUS Server

The RADIUS server can use either a database or a simple flat file as its repository of user information. The following example uses a flat file to demonstrate a minimum user configuration. The users file contains authentication and configuration information for each user. It begins with the username and the authentication (that is, the password) that is required from the user. The user/password line is followed by indented lines that are attributes to be passed back to the requesting server.

Figure 15-5 Sample RADIUS User Information Flat File

```
Jeff      Cleartext-Password:="garbage"
          Class="Administrator",
          Oracle-MI-role="Administrator",
          Oracle-MI-scope="Global"

Paul      Cleartext-Password:="apr6279"
          Class="Viewer",
          Oracle-MI-role="Viewer",
          Oracle-MI-scope="Global"
```

When the RADIUS server has authenticated a user, it sends back various attributes with the authentication acceptance message. The CMP system uses these attributes to determine what actions the user can perform.

The best practice is to use a vendor-specific attribute (VSA) dictionary file to define what attributes to send back to the client. Figure 15-6 shows a sample file. The local RADIUS administrator is responsible for incorporating the VSA dictionary file onto the RADIUS server.

Figure 15-6 Sample VSA Dictionary File For RADIUS

```
========== dictionary.oracle ===================
# Oracle Communications VSA's, from RFC 2548
# The filename given here should be an absolute path.
#
# Place additional attributes or $INCLUDEs here.

VENDOR Oracle 21274
BEGIN-VENDOR Oracle
ATTRIBUTE Oracle-MI-role 1 string
ATTRIBUTE Oracle-MI-scope 3 string
END-VENDOR Oracle
======================
```

The attributes `Oracle-MI-role` and `Oracle-MI-scope` are for access to the CMP system. Both a scope and a role are associated with a user. The responses sent back from the RADIUS server should match what is configured in the CMP system. The defaults for the role, in ascending order of capability, are `Viewer`, `Operator`, and `Administrator`, but the system administrator can create other roles or remove any role except that of `Administrator`.

The default scope is `Global`, and the administrator can create other scopes within the CMP system.

About Associating Roles and Scopes

The CMP system assigns two attributes to a user, a role and a scope. Users that authenticate against a RADIUS server are assigned roles and scopes by matching against the attribute values returned by the RADIUS server.

The best practice is to provide role and scope values using the VSA dictionary, by defining the attributes `Oracle-MI-role` and `Oracle-MI-scope`. The flexibility of roles and scopes can be supported by the RADIUS server if the VSA dictionary is integrated.

The following example defines users who have access at different role levels:

Figure 15-7 Sample RADIUS User Information

```
Jeff      Cleartext-Password:="garbage"
          Class="Administrator",
          Oracle-MI-role="Administrator",
          Oracle-MI-scope="Global"

Paul      Cleartext-Password:="apr6279"
          Class="Viewer",
          Oracle-MI-role="Viewer",
          Oracle-MI-scope="Global"
```

In this example, the user `Jeff` has access to the CMP system as an `Administrator`, and the user `Paul` has access to the CMP system as a `Viewer` (read-only access).

However, if Oracle VSAs are not included in the RADIUS dictionary, then they cannot be defined in the user file, and only a `Class` attribute can be returned on a RADIUS authentication. The CMP system can use the `Class` attribute for RADIUS authentication.

To accept the `Class` attribute for CMP login, define a scope and a role that matches what the RADIUS server returns as the `Class` attribute. The CMP system uses the `Class` attribute for both of the required role and scope credentials. For example, consider this user defined in the RADIUS server:

Figure 15-8 Sample RADIUS User Information - No Role or Scope

```
Dawn      Cleartext-Password:="kkmk4813"
          Class="Viewer"
```

`Dawn` can get access to the CMP system if you have defined both a role named `Viewer` and a scope named `Viewer`; the user interface matches the one returned `Class` value to both of the required role and scope credentials.

Enabling and Configuring RADIUS on the CMP System

By default, RADIUS Authentication is disabled in the CMP system. Enabling authentication requires admin privileges. The `admin` user is always authenticated against the local database record; thus, the `admin` user is best suited to setting up RADIUS authentication (see #unique_350).

Two configuration parameters must match with the configuration that was put on the RADIUS server:

- **Source of User Credentials** must match up with the user configuration in the RADIUS server, but this will also depend on what is configured in the next parameter.

- If **Action if missing credentials** is set to **Use following defaults**, then a user will be authenticated as long as the password is correct. This user could log in even though the `Class` is not valid:

Figure 15-9 Sample User for RADIUS Server

```
test      Cleartext-Password := "2931txy"
          Class = "noone"
```

– If **Action if missing credentials** is set to **Reject**, then the configuration of the user will depend on the configuration of **Source of User Credentials**.

To enable RADIUS authentication and accounting:

1. Log in to the CMP system as `admin`.
2. From the **System Administration** section of the navigation pane, select **User Management**.

   The content tree displays the **User Management** group.
3. From the content tree, select **External Authentication**.

   The External Authentication page opens. By default, external authentication is disabled.
4. Click **Modify**.

   The External Authentication page becomes editable.
5. In the **Configuration** section, select **Enable RADIUS Authentication**.

   Configuration and RADIUS Services configuration fields appear.
6. Select to **Enable RADIUS Accounting**.

   This feature is disabled by default. When enabled, the CMP system sends an Accounting-Start message to the accounting server when a user logs in, and an Accounting-Stop message when the user logs out. These messages contain a session ID attribute that uniquely identifies the user session so that it can be matched between Start and Stop.
7. Select the **Destination for Accounting Messages** from the list.

   Available options include:

   - **Both Primary and Secondary** (default)

     Specifies that accounting messages generated for each user session are sent to both the primary and (when configured) secondary RADIUS servers.
   - **Primary (Secondary on error)**

     Accounting messages are sent only to the primary server, as long as it is reachable. If the primary accounting server is unreachable, messages are sent to the secondary accounting server.
8. Enter the **NAS IP Address** (required).

   The IP address, in IPv4 or IPv6 format, of the network access server. By default, this is the local host address.
9. Select when to **Use local authentication** from the list.

   Available options include:

- **When RADIUS servers timeout** (default)
- **When both RADIUS servers timeout or reject**
- **Never**

> **Note:** Fallback to local authentication is never used. However, the `admin` user is always authenticated locally.

10. Select the **Source of User Credentials** from the list.

    Available options include:

    - **RADIUS Class**

      The value of the `Class` attribute returned by the server determines both the role and scope.

    - **Oracle VSAs**

      The value of Oracle VSAs returned by the server determines the role and scope.

11. Select an **Action if Missing Credentials**.

    Available options include:

    - **Reject**

      If you select this option, a user whose login credentials are missing is not logged in.

    - **Use following defaults**

      Select a setting for each of the following attributes:

      – **Default Role**

        The role assigned if the user credentials are missing or mismatched. The default role is **Viewer**.

      – **Default Scope**

        The scope assigned if the user credentials are missing or mismatched. The default scope is **Global**.

12. In the **RADIUS Services** section, edit the following fields:

    a. Configure the **Primary RADIUS Authentication Server**:

    - **Server**

      The FQDN or IP address (in IPv4 or IPv6 format) assigned to the primary authentication server.

    > **Note:** To disable the primary server, delete its IP address.

    - **Port**

      The IP port number of the primary server. The default value is port 1812.

- **Timeout (seconds)**

  The length of time the CMP system waits for a response from the server. The default value is 3 seconds.

- **Retries**

  The number of times the CMP system tries to send a message to the server. The default value is 3 times.

- **Shared Secret**

  A password-like string that must exactly match between the CMP system and the `secret` attribute configured in the entry for this CMP system in the `clients.conf` file in the RADIUS server.

Note: If the two values do not match, the server ignores all messages from the CMP system.

b. Configure the**Secondary RADIUS Authentication Server**:

If configured, the secondary authentication server uses the same fields as the primary authentication server.

c. Configure the **Primary RADIUS Accounting Server**:

- **Server**

  The FQDN or IP address (in IPv4 or IPv6 format) assigned to the primary accounting server.

- **Port**

  The IP port number of the Primary RADIUS Accounting server. The default value is port 1813.

- **Timeout (seconds)**

  The length of time the CMP system waits for a response from the server. The default value is 3 seconds.

- **Retries**

  The number of times the CMP system tries to send a message to the server. The default value is 3 times.

- **Shared Secret**

  A password-like string that must exactly match between the CMP system and the `secret` attribute configured in the entry for this CMP system in the `clients.conf` file in the RADIUS server.

Note: If the two values do not match, the server ignores all messages from the CMP system.

d. **Secondary RADIUS Accounting Server**

If configured, the secondary accounting server uses the same fields as the primary accounting server.

13. Click **Save**.

RADIUS Authentication and Accounting is configured.

About SANE Authentication

The CMP system supports Secure Access to Network Elements (SANE) Authentication and Authorization. You can configure the CMP system to operate in a SANE network environment so that a user elsewhere in the network can gain single sign-on (SSO) access. When the CMP system is configured to authenticate using SANE, users can log in using a SANE client.

Note: Usage of a SANE client is outside the scope of this document.

See Enabling SANE Authentication on the CMP System for details.

The `admin` user profile is treated separately. An `admin` user can log in to the CMP system using any supported browser.

The authentication process is as follows:

1. From a SANE client user interface, the user selects the CMP system in a web browser.
2. An encrypted SANE authentication artifact is sent to the CMP system through the browser.
3. The CMP system forwards the artifact to a SANE server (also called, the SANE responder).

Note: The `admin` user is always authenticated locally, regardless of SANE configuration settings.

- If the SANE server verifies the artifact, it returns an assigned role and scope for the user and the CMP system allows the user to log in to the system.
- If the SANE server does not verify the artifact, the CMP system rejects the login request.

Enabling SANE Authentication on the CMP System

By default, SANE Authentication is disabled in the CMP system. Enabling authentication requires `admin` privileges. The user `admin` is always authenticated against the local database account; thus, the `admin` user is best suited to setting up SANE authentication (see Creating a User Profile).

To enable SANE authentication:

1. Log in to the CMP system as `admin`.
2. From the **System Administration** section of the navigation pane, select **User Management**.

   The content tree displays the **User Management** group.
3. From the content tree, select **External Authentication**.

The External Authentication page opens, displaying the current configuration information. By default, external authentication is disabled.

4. Click **Modify**.

   The External Authentication page becomes editable.

5. In the **Configuration** section, select **Enable SANE Authentication**.

   Configuration and SANE Servers configuration fields appear.

6. Enter the **Artifact Parameter Name**.

   The name of the artifact parameter. Enter an alphanumeric string. The default value is `artifact`.

7. Select the **Verification for Account** setting from the list.

   Available options are:

   - **On login only** (default) — The CMP system authenticates the user once on login. The user is considered authenticated until logout.

   - **On each request** — The CMP system authenticates the user on login, and then for each HTTP or HTTPS request. If any request is not authenticated, the user is immediately logged out.

8. Select the **Action if Missing Credentials**.

   The available options are:

   - **Reject —** If you select this option, a user login is rejected even if the authentication is successful.

   - **Use following defaults —** If you select this option, a user with missing credentials is allowed to log in, but the system assigns a default role and scope:

     – **Default Role —** Default role assigned to the user. The default role is **Viewer**.

     – **Default Scope —** Default scope assigned to the user. The default scope is **Global**.

9. In the **SANE Servers** section, enter the **SAML Service Name**.

   The name of the Security Assertion Markup Language service registered with the UDDI server. Enter an alphanumeric string.

10. Enter the **UDDI Inquiry URL**.

    The Universal Description, Discovery and Integration URL, in HTTP or HTTPS format, for the inquiry.

11. Click **Save**.

SANE authentication is enabled.

Changing a Password

The Change Password option lets users change their password. This system administration function is available to all users.

Note: The `admin` user can change the password for any user.

If the system administrator has configured your account for password expiration, you will receive a warning when you log in that you must change your password.

Note: To reset the administrator password, it is recommended to contact My Oracle Support.

To change a password:

1. From the **System Administration** section of the navigation pane, select **Change Password**.

   The Account Management page opens. If your account is set up with a password expiration period, the expiration date is displayed.

2. Enter your **Current Password**.

3. Enter your **New Password**.

   The password is case sensitive. Depending on your system settings, the password must meet the following requirements:

   - Cannot contain the username.
   - Must contain the minimum number of characters (default is six).
   - Must contain characters as specified from the following categories:
     - Must contain at least the specified number of lower case letters.
     - Must contain at least the specified number of upper case letters.
     - Must contain at least the specified number of numbers.
     - Must contain at least the specified number of symbols.

   Note: Consult with your system administrator to obtain the password criteria for your system.

4. Re-enter your new password to **Confirm Password**.

   Note: If your new password does not conform to the password strength rules configured for your system, a validation error message appears that includes valid password criteria. Enter and confirm another password that conforms to the criteria.

5. Click **Change Password**.

Your password is changed.

Changing the MySQL Password

To change the MySQL password:

1. Verify that HA and MySQL servers are master.
2. Log into the CMP GUI.
3. Clear any critical alarms and MySQL alarms.

   MySQL related alarms are:

   - 70020—QP Master database is outdated
   - 70021—QP slave database is unconnected to the master
   - 70022—QP Slave database failed to synchronize
   - 70023—QP Slave database lagging the master
   - 70024—QP Slave database is prevented from synchronizing with the master
   - 70025—QP Slave database is a different version than the master
4. Set any secondary and tertiary servers (server-b or server-c) to forced standby.
5. Locate the master MySQL node.

   The master MySQL node is the active MA or CMP in primary site for. There are two ways to find the master MySQl node:

   - Login into the CMP GUI, and find the active server by selecting **Platform Settings** > **Topology Settings**.
   - Use the `wbAccess mysqlState` command in the CLI.
6. Using the CLI , enter `manageMySQL ModifyMySQLRootPWD` to modify the MySQL password.

   The password is 1 to 32 characters in length.

   Because of MySQL cluster replication, this change is replicated to all slave MySQL servers, then the password in the database of all the MySQL servers is changed synchronously.
7. Remove the forced standby from all secondary and tertiary servers (server-b or server-c).

Your MySQL password is changed.

A

CMP Modes

The functions available in the CMP system are determined by the operating modes and sub-modes selected when the software is installed. Functions that can change include:

- Items on the navigation pane
- Tabs on the Policy Server Administration page
- Protocols supported
- Configuration options
- Policy options available in the policy wizard
- Reports available

The mode selection process is not normally available. At initial configuration, and if it becomes necessary to replace a server or reinstall Policy Management software, the Mode Settings page becomes visible, and you must reset the operating modes as appropriate for your environment before you can use the product.

This appendix briefly describes the modes and sub-modes available.

Caution: CMP modes should only be set in consultation with My Oracle Support. Setting or changing modes inappropriately could result in the loss of network element connectivity, policy function, statistical data, and cluster redundancy.

About Mode Settings

When you use a web browser to connect to a CMP system after the software is first installed, the Mode Settings page opens (Figure A-1). Select the needed modes, functions, and management options for your installation and then click **OK**. The browser page closes and you are automatically logged out. When you next log in, the CMP system reopens in the selected mode.

Restriction

Important: Options marked as **Restricted** are for use within specific environments and should not be enabled without authorization. For more information about the restricted function or feature, contact My Oracle Support.

Modes and Functions

The following modes and functions are available:

Cable
Enables support of a cable carrier environment.

Note: If Cable Mode is enabled, you must also enable **Manage Direct Link**.

Cable functions include:

PCMM
Supports PacketCable MultiMedia functions.

DQOS
Supports Dynamic Quality of Service functions.

Note: This function is restricted. See Restriction for more information.

Diameter AF
Supports Diameter AF functions.

Wireless
Enables support of a wireless carrier environment. Functions are described in the *Configuration Management Platform Wireless User's Guide*.
Wireless functions include:

Diameter 3GPP
Supports Diameter 3GPP protocol.

Diameter 3GPP2
Supports Diameter 3GPP2 protocol.

Note: This function is restricted. See Restriction for more information.

PCC Extensions
Supports Policy and Charging Control functions.

Note: This function is restricted. See Restriction for more information.

Quotas Gx
Supports a subscriber quota environment using the Diameter Gx protocol. The Gx protocol supports deep packet inspection (DPI) devices.

Quotas Gy
Supports a subscriber quota environment using the Diameter Gy protocol.

Note: This function is restricted. See Restriction for more information.

LI
Supports Lawful Intercept functions.

Note: This function is restricted. See Restriction for more information.

SCE-Gx
Supports the Cisco Service Control Engine Gx protocol. If this mode is selected, **Diameter 3GPP** and **RADIUS** must also be selected, and other Gx sub-modes must not be selected.

Note: This function is restricted. See Restriction for more information.

Gx-Lite
Supports the Gx-Lite protocol, a simplified version of 3GPP Gx for use by non-GGSN PCEF vendors that do not have access to network-level information.

Note: This function is restricted. See Restriction for more information.

Cisco Gx
Supports the Cisco Gx protocol.

Note: This function is restricted. See Restriction for more information.

DSR
Supports Policy Management network segmentation using a Oracle Communications Diameter Signaling Router (DSR).

Note: This function is restricted. See Restriction for more information.

Wireless-C
Supports a wireless system supporting a Mediation server; SMS Notification Statistics; and SCTP counters.

Note: This function is restricted. See Restriction for more information.

SMS
Enables support of SMS servers. Functions are described in the *Configuration Management Platform Wireless User's Guide*.
SMS Mode functions include:

SMPP
Supports SMS using SMPP protocol.

CMPP
Supports SMS using CMPP protocol.

Note: This function is restricted. See Restriction for more information.

XML

Supports SMS using XML.

Note: This function is restricted. See Restriction for more information.

SPR

Enables support of subscriber database management. Select only one sub-mode. Functions are described in the Subscriber Data Management documentation.

SPR Mode functions include:

Subscriber Profiles

Supports subscriber profile functions.

Note: This function is restricted. See Restriction for more information.

Quota

Supports subscriber quotas.

Note: This function is restricted. See Restriction for more information.

Wireline

Enables support of a wireline carrier environment. Functions are described in the *Configuration Management Platform Wireline User's Guide*.

Note: This function is restricted. See Restriction for more information.

SPC

Enables the COPS Application Manager product, which accepts service provisioning requests from a Session Border Controller over the Common Open Policy Service (COPS) protocol. Functions are described in the *Service Provisioning over COPS Application Manager User's Guide*.

Note: This function is restricted. See Restriction for more information.

RADIUS

Enables support of RADIUS Change of Authorization.

Note: This function is restricted. See Restriction for more information.

BoD

Enables the Bandwidth on Demand Application Manager (BoD-AM), which support video on demand (VoD) servers. Functions are described in the *Bandwidth on Demand Application Manager Cable User's Guide*.

PCMM
Supports a network creating PacketCable MultiMedia (PCMM) sessions.

Diameter
Supports a network creating Diameter sessions.

Note: This function is restricted. See Restriction for more information.

RDR
Supports a network containing Service Control Engine (SCE) devices transmitting Raw Data Records (RDRs).

Note: This function is restricted. See Restriction for more information.

Server Management Options

The management options for servers are:

Manage Policy Servers
Manages MPE devices.

Manage MA Servers
Manages Management Agent servers.

Manage Policies
Enables the Policy Wizard.

Manage MRAs
Manages Multi-Protocol Routing Agent servers.

Manage BoDs
Manages Bandwidth on Demand Application Manager servers.

Manage Mediation Servers
Manages Message Distribution Function servers.

Manage SPR Subscriber Data
Manages Subscriber Profile Repository servers.

Manage Geo-Redundant
Manages georedundant MPE, MRA, BoD, MDF, or Mediation clusters.

Manager is HA (clustered)
Enables High Availability features.

Manage Analytic Data
Enables output of policy event records.

Figure A-1 Mode Settings Page

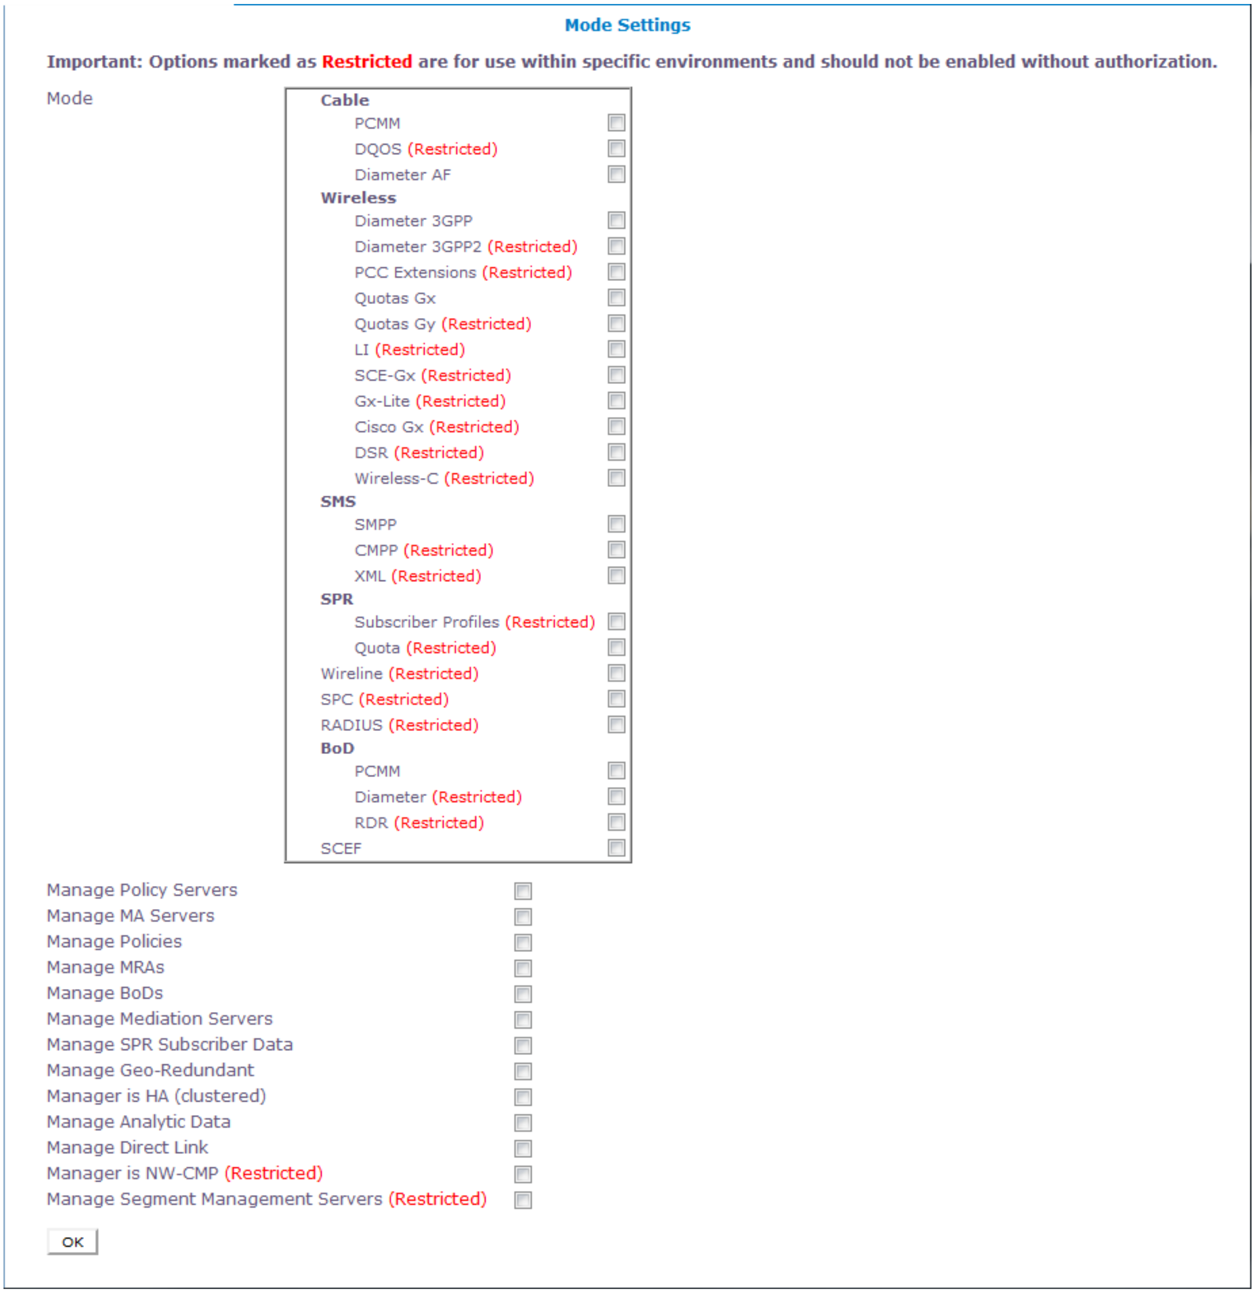

Mode Settings

Important: Options marked as Restricted are for use within specific environments and should not be enabled without authorization.

Mode

Cable
- PCMM □
- DQOS (Restricted) □
- Diameter AF □

Wireless
- Diameter 3GPP □
- Diameter 3GPP2 (Restricted) □
- PCC Extensions (Restricted) □
- Quotas Gx □
- Quotas Gy (Restricted) □
- LI (Restricted) □
- SCE-Gx (Restricted) □
- Gx-Lite (Restricted) □
- Cisco Gx (Restricted) □
- DSR (Restricted) □
- Wireless-C (Restricted) □

SMS
- SMPP □
- CMPP (Restricted) □
- XML (Restricted) □

SPR
- Subscriber Profiles (Restricted) □
- Quota (Restricted) □

Wireline (Restricted) □

SPC (Restricted) □

RADIUS (Restricted) □

BoD
- PCMM □
- Diameter (Restricted) □
- RDR (Restricted) □

SCEF □

Manage Policy Servers □
Manage MA Servers □
Manage Policies □
Manage MRAs □
Manage BoDs □
Manage Mediation Servers □
Manage SPR Subscriber Data □
Manage Geo-Redundant □
Manager is HA (clustered) □
Manage Analytic Data □
Manage Direct Link □
Manager is NW-CMP (Restricted) □
Manage Segment Management Servers (Restricted) □

B

Generated Statistics

This appendix lists the available statistics for generated scheduled tasks.

List of Generated Statistics

The following is a list of generated statistics available for scheduled tasks:

- DiameterAfStats
- DiameterAfPeerStats
- DiameterAfLatencyStats
- DiameterAfPeerLatencyStats
- ProtocolErrorStats
- ConnectionErrorStats
- IntervalStats
- TrafficProfileStats
- StaleSessionStats
- PolicyStats
- TopologyUpdateStats
- PolicyServerStats
- KpiStats
- TpsStats
- RadiusAccountingStats
- RadiusAccountingNetworkElementStats