

**Oracle® Communications
Policy Management**

CMP Cable User's Guide

Release 11.5.2

E65491 Revision 01

August 2015

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

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

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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

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

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

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

Table of Contents

Chapter 1: About this Guide.....	14
Introduction.....	15
How This Guide is Organized.....	15
Scope and Audience.....	15
Documentation Admonishments.....	16
Related Publications.....	16
Locate Product Documentation on the Oracle Technology Network Site.....	17
Customer Training.....	17
My Oracle Support (MOS).....	17
Emergency Response.....	18
Chapter 2: The Oracle Communications Policy Management Solution.....	19
The Multimedia Policy Engine.....	20
Overview.....	20
The Bandwidth on Demand Application Manager.....	22
The Management Agent.....	23
The Oracle Communications Policy Management Configuration Management Platform.....	23
Specifications for Using the GUI.....	23
Logging In.....	24
GUI Overview.....	25
GUI Icons.....	26
Changing a Password.....	27
Overview of Major Tasks.....	27
Chapter 3: Configuring the Policy Management Topology.....	29
About the Policy Management Topology.....	30
High Availability.....	31
Spare Servers.....	33
MPE-S and BoD Georedundancy.....	33
CMP Georedundancy.....	34
Primary and Secondary Sites.....	35

Cluster Preferences.....	36
Server Status.....	36
Setting Up the Topology.....	37
Setting Up a CMP Cluster.....	37
Setting Up a Site.....	39
Setting Up an MPE or BoD Cluster.....	40
Setting Up an MA Cluster.....	43
Setting Up a Georedundant Cluster.....	45
Example: Setting Up Georedundancy.....	48
Modifying the Topology.....	54
Modifying a Site.....	54
Removing a Site from the Topology.....	54
Modifying an MPE, MA, or BoD Cluster.....	55
Modifying a CMP Cluster.....	55
Removing a Cluster from the Topology.....	56
Reversing Cluster Preference.....	56
Demoting a CMP Cluster.....	57
Forcing a Server into Standby Status.....	58
Configuring SNMP Settings.....	59
Configuring the Upsync Log Alarm Threshold.....	61

Chapter 4: Managing Multimedia Policy Engine Devices.....62

Policy Server Profiles.....	63
Creating a Policy Server Profile.....	63
Configuring or Modifying a Policy Server Profile.....	64
Deleting a Policy Server Profile.....	64
Configuring Protocol Options on the Policy Server.....	65
Associations Configuration Options.....	66
General Configuration Options.....	66
PCMM Configuration Options.....	66
Diameter Configuration Options.....	67
Diameter AF Default Profiles Configuration Options.....	68
Configuring MPE Advanced Settings.....	68
Configuring Session Clean Up Options.....	69
Configuring Configuration Keys.....	70
Configuring Load Shedding Rules.....	71
Configuring Data Source Interfaces.....	73
Configuring a DHCP Data Source.....	74
Working with Policy Server Groups.....	75
Creating a Policy Server Group.....	75

Adding a Policy Server to a Policy Server Group.....	76
Creating a Policy Server Sub-group.....	76
Renaming a Policy Server Group.....	76
Removing a Policy Server Profile from a Policy Server Group.....	77
Deleting a Policy Server Group.....	77
About Reapplying a Configuration.....	78
Reapplying the Configuration to a Single Device.....	78
Reapplying the Configuration to a Group of Devices.....	78
Resetting Counters.....	78
Checking the Status of an MPE Server.....	79
Policy Server Reports.....	80
Cluster Information Report.....	81
Policy Statistics.....	81
Session Cleanup Statistics.....	82
Protocol Statistics.....	82
Latency Statistics.....	83
Error Statistics.....	84
Data Source Statistics.....	84
Database Statistics.....	84
KPI Interval Statistics.....	85
Viewing Policy Server Logs.....	92
Viewing the Trace Log.....	93
Syslog Support.....	94
Configuring Log Settings.....	95
Chapter 5: Configuring Protocol Routing.....	97
PCMM Routing Architectures.....	98
Configuring PCMM Routing.....	98
Configuring Rx-to-PCMM Routing.....	99
Chapter 6: Managing Network Elements.....	101
About Network Elements.....	102
Defining a Network Element.....	102
Modifying a Network Element.....	103
Deleting Network Elements.....	103
Deleting Multiple Network Elements.....	104
Finding a Network Element.....	104
Configuring Options for Network Elements.....	105
Configuring CMTS.....	105
Associating a Network Element with an MPE Device.....	106

Working with Network Element Groups.....	107
Creating a Network Element Group.....	107
Adding a Network Element to a Network Element Group.....	107
Creating a Network Element Sub-group.....	108
Deleting a Network Element from a Network Element Group.....	109
Modifying a Network Element Group.....	109
Deleting a Network Element Group or Sub-group.....	110
Chapter 7: Managing Record Keeping Servers.....	111
About Record Keeping Servers.....	112
Creating a Record Keeping Server Profile.....	112
Modifying a Record Keeping Server Profile.....	113
Deleting a Record Keeping Server Profile.....	113
Chapter 8: Managing Event Messaging.....	114
About Event Messaging.....	115
Configuring Global Settings for Event Messaging.....	116
Configuring Local Settings for Event Messaging.....	117
Chapter 9: Managing Management Agent Servers.....	119
About Management Agent Servers.....	120
Creating a Management Agent Profile.....	120
Modifying a Management Agent Profile.....	120
Deleting a Management Agent Profile.....	121
Reapplying a Management Agent Profile Configuration.....	121
Management Agent Tasks.....	121
Managing Management Agent Tasks.....	122
Viewing Task Status on a Management Agent.....	122
Chapter 10: Managing Bandwidth on Demand.....	124
Creating a BoD Server.....	125
Creating a BoD Server Group.....	125
Viewing and Modifying BoD Server Topology Information.....	126
Viewing BoD Topology Information.....	126
Modifying BoD Topology Information.....	127
Configuring a Server with Multiple Destination IP Addresses.....	127
Deleting a BoD Server.....	127
Reapplying the Configuration to a BoD Server.....	128

Viewing and Modifying BoD Server Settings.....	128
Viewing BoD Server Settings.....	128
Modifying BoD Server Information.....	130
Configuring BoD Server Configuration Key Information.....	131
Managing PCMM Services.....	132
Viewing Services.....	132
Creating a Service.....	134
Modifying a Service.....	138
Deactivating a Service.....	147
Deleting a Service.....	147
Importing a Service.....	147
Exporting a Service.....	148
Chapter 11: System-Wide Reports.....	149
KPI Dashboard.....	150
Mapping Display to KPIs.....	151
Configuring the KPI Threshold	153
Viewing the Trending Reports.....	153
Viewing PCMM Transaction Per Second.....	154
Viewing Rx Transaction Per Second.....	154
Viewing Session Count.....	155
Custom Trending Reports.....	156
Alarm Reports.....	159
About Active Alarm Reports.....	159
Viewing Alarms.....	160
Viewing Active Alarms.....	160
Working with the Active Alarms Report.....	161
Viewing the Alarm History Report.....	161
Viewing Other Reports.....	163
Viewing the Connection Status Report.....	163
Viewing the Protocol Errors Report.....	164
Viewing the Policy Statistics Report.....	165
Viewing the MPE/BoD Rep Stats.....	166
Chapter 12: Upgrade Manager.....	168
About ISO Files on Servers.....	169
ISO Maintenance Page Elements.....	169
Viewing the ISO Status of Servers.....	170
Pushing a Script to a Server	170
Adding an ISO File to a Server.....	171

Deleting an ISO File from a Server.....	172
Preparing for an Upgrade.....	172
Performing an Upgrade.....	172
System Maintenance Elements.....	174
Viewing Upgrade Status of Servers.....	176
About Rolling Back an Upgrade.....	176
Chapter 13: Defining Global Configuration Settings.....	178
Configuring Subnet Settings.....	179
About Routing by CMTS IP.....	180
Configuring CMTS IP Routing.....	180
Configuring Stats Settings.....	181
Chapter 14: System Administration.....	183
Configuring System Settings.....	184
Importing and Exporting Configurable Objects.....	186
Using the OSSI XML Interface.....	186
Importing an XML File to Input Objects.....	187
Exporting an XML File.....	188
The Manager Report.....	189
The Trace Log.....	190
Viewing the Trace Log.....	190
Filtering the Trace Log.....	191
Configuring the Trace Log.....	191
Viewing the Audit Log.....	192
Searching for Audit Log Entries.....	194
Exporting or Purging Audit Log Data.....	195
About Scheduled Tasks.....	196
Configuring a Task.....	197
User Management.....	199
Creating a Customer User Management System Profile.....	199
About User Roles.....	199
About User Scope.....	202
About User Profiles.....	203
External Authentication.....	207
Changing a Password.....	214
Appendix A: CMP Modes.....	215
The Mode Settings Page.....	216

Glossary.....	219
----------------------	------------

List of Figures

Figure 1: The CMP and MPE Devices.....	21
Figure 2: CMP Login Page.....	25
Figure 3: Structure of the CMP Cable GUI.....	25
Figure 4: Policy Management Topology.....	31
Figure 5: High Availability.....	32
Figure 6: Cluster with Active, Standby, and Spare Servers.....	33
Figure 7: MPE-S or BoD Georedundant Configuration.....	34
Figure 8: CMP Georedundancy.....	35
Figure 9: Cluster Settings Page for CMP Cluster.....	38
Figure 10: Cluster Settings Page for an MPE Cluster.....	43
Figure 11: Cluster Settings Page for an MA Cluster.....	45
Figure 12: Example of Primary Site Settings.....	51
Figure 13: Group View	80
Figure 14: Sample Protocol Statistics.....	83
Figure 15: Sample Error Statistics.....	84
Figure 16: Policy Server Administration, Logs Tab — Cable.....	93
Figure 17: Modify Routing Configuration Page.....	99
Figure 18: Select Network Elements.....	106
Figure 19: Add Network Element Page.....	108
Figure 20: Modify Event Messaging Page.....	117
Figure 21: BoD Administration Page for the ALL Group.....	125
Figure 22: BoD Administration Page.....	126

Figure 23: Other Advanced Configuration Settings Page.....	131
Figure 24: Services Administration Page.....	132
Figure 25: View Service Page.....	133
Figure 26: Create Service Page.....	134
Figure 27: Upstream Profile.....	135
Figure 28: GateSpec Section.....	135
Figure 29: Classifier Table with Drop-down Menu.....	136
Figure 30: Add Classifier Screen (Standard Classifier selected).....	136
Figure 31: Copy Service Page.....	138
Figure 32: Services Import/Export Page.....	148
Figure 33: Example of KPI Dashboard.....	150
Figure 34: Trending Report Definition Configuration Page.....	157
Figure 35: Sample Active Alarms Report.....	160
Figure 36: Alert Details.....	163
Figure 37: Sample Password Strength Policy.....	186
Figure 38: Audit Log.....	193
Figure 39: Audit Log Details.....	194
Figure 40: Schedule Task Administration - OM Statistics.....	197
Figure 41: Scheduled Task Administration.....	198
Figure 42: Scheduled Task Administration.....	198
Figure 43: Sample VSA Dictionary File For RADIUS.....	208
Figure 44: RADIUS External Authentication Configuration Page.....	212
Figure 45: Mode Settings Page.....	218

List of Tables

Table 1: Admonishments.....	16
Table 2: SNMP Attributes.....	59
Table 3: Associations Configuration Options.....	66
Table 4: General Configuration Options.....	66
Table 5: PCMM Configuration Options.....	66
Table 6: Diameter Configuration Options.....	67
Table 7: Diameter AF Default Profiles Configuration Options.....	68
Table 8: Session Clean Up Options.....	69
Table 9: Default MPE Device Busyness Levels.....	71
Table 10: PCMM (PacketCable MultiMedia) Protocol Statistics.....	86
Table 11: Record Keeping Servers Protocol Statistics.....	87
Table 12: CMTS with Lost Connections Statistics.....	87
Table 13: MGPI Statistics.....	88
Table 14: Diameter AF (Application Function) Statistics.....	88
Table 15: Latency Statistics.....	90
Table 16: Protocol Error Statistics.....	90
Table 17: Connection Error Statistics.....	90
Table 18: KPI Interval Statistics.....	91
Table 19: Policy Statistics.....	91
Table 20: Correct Parameter - Classifier Order	137
Table 21: Multiple Classifier Compatibility Matrix	137
Table 22: Service Field Definitions and Ranges.....	139

Table 23: KPI Definitions for MPE Devices.....	152
Table 24: ISO Maintenance Page Elements.....	169
Table 25: System Maintenance Elements.....	174
Table 26: System Maintenance Operations.....	175
Table 27: Cable Functions.....	216
Table 28: Wireless Functions.....	216
Table 29: SMS Mode functions.....	217
Table 30: SPR Mode functions.....	217

Chapter 1

About this Guide

Topics:

- *Introduction.....15*
- *How This Guide is Organized.....15*
- *Scope and Audience.....15*
- *Documentation Admonishments.....16*
- *Related Publications.....16*
- *Locate Product Documentation on the Oracle Technology Network Site.....17*
- *Customer Training.....17*
- *My Oracle Support (MOS).....17*
- *Emergency Response.....18*

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 how to use the Oracle Communications Policy Management Configuration Management Platform (CMP) system to configure and manage Policy Management devices in a cable network.

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.
- [The Oracle Communications Policy Management Solution](#) 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 Oracle Communications Policy Management Configuration Management Platform (CMP), 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.
- [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 Manager](#) describes the Upgrade Manager pages and the elements found on the pages.
- [Defining Global Configuration Settings](#) 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.

Scope and Audience

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

- Network operators, who configure, operate, monitor, and maintain Policy Management systems in a carrier network
- System administrators, who maintain the accounts of users of CMP systems

Documentation Admonishments

Admonishments are icons and text throughout this manual that alert the reader to assure personal safety, to minimize possible service interruptions, and to warn of the potential for equipment damage.

Table 1: Admonishments

Icon	Description
 DANGER	<p>Danger:</p> <p>(This icon and text indicate the possibility of <i>personal injury</i>.)</p>
 WARNING	<p>Warning:</p> <p>(This icon and text indicate the possibility of <i>equipment damage</i>.)</p>
 CAUTION	<p>Caution:</p> <p>(This icon and text indicate the possibility of <i>service interruption</i>.)</p>
 TOPPLE	<p>Topple:</p> <p>(This icon and text indicate the possibility of <i>personal injury</i> and <i>equipment damage</i>.)</p>

Related Publications

For information about additional publications that are related to this document, refer to the *Related Publications Reference* document, which is published as a separate document on the Oracle Technology Network (OTN) site. See [Locate Product Documentation on the Oracle Technology Network Site](#) for more information.

Locate Product Documentation on the Oracle Technology Network Site

Oracle customer documentation is available on the web at the Oracle Technology Network (OTN) 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 Technology Network site at <http://docs.oracle.com>.
2. Click **Industries**.
3. Under the Oracle Communications subheading, click the **Oracle Communications documentation** link.
The Oracle Communications Documentation page appears with Tekelec shown near the top.
4. Click the **Oracle Communications Documentation for Tekelec Products** link.
5. Navigate to your Product and then the Release Number, and click the **View** link (the Download link will retrieve the entire documentation set).
A list of the entire documentation set for the selected product and release appears.
6. To download a file to your location, right-click the **PDF** link, select **Save target as**, 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 (MOS)

MOS (<https://support.oracle.com>) is your initial point of contact for all product support and training needs. A representative at Customer Access Support (CAS) can assist you with MOS registration.

Call the CAS 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 will be connected to a live agent who can assist you with MOS registration and opening a support ticket.

MOS 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 (CAS) 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.

Chapter

2

The Oracle Communications Policy Management Solution

Topics:

- *The Multimedia Policy Engine.....20*
- *Overview.....20*
- *The Bandwidth on Demand Application Manager.....22*
- *The Management Agent.....23*
- *The Oracle Communications Policy Management Configuration Management Platform.....23*
- *Overview of Major Tasks.....27*

The Oracle Communications Policy Management Solution provides an overview of the major elements of the Policy Management solution. The major elements include:

- The Oracle Communications Policy Management Configuration Management Platform (CMP) system controls MPE devices and associated applications.
- The Oracle Communications Policy Management Multimedia Policy Engine (MPE) device manages multiple network-based client sessions.

The Multimedia Policy Engine

The Multimedia Policy Engine (MPE) device includes a simple, powerful, and flexible policy rules engine. Through the use of policy rules, you can modify the behavior of an MPE device dynamically as it processes protocol messages.

Overview

The core function of the MPE device network is to establish service flows between the subscribers and application servers that provide multimedia services, as shown in *Figure 1: The CMP and MPE Devices*.

The Oracle Communications Policy Management Solution

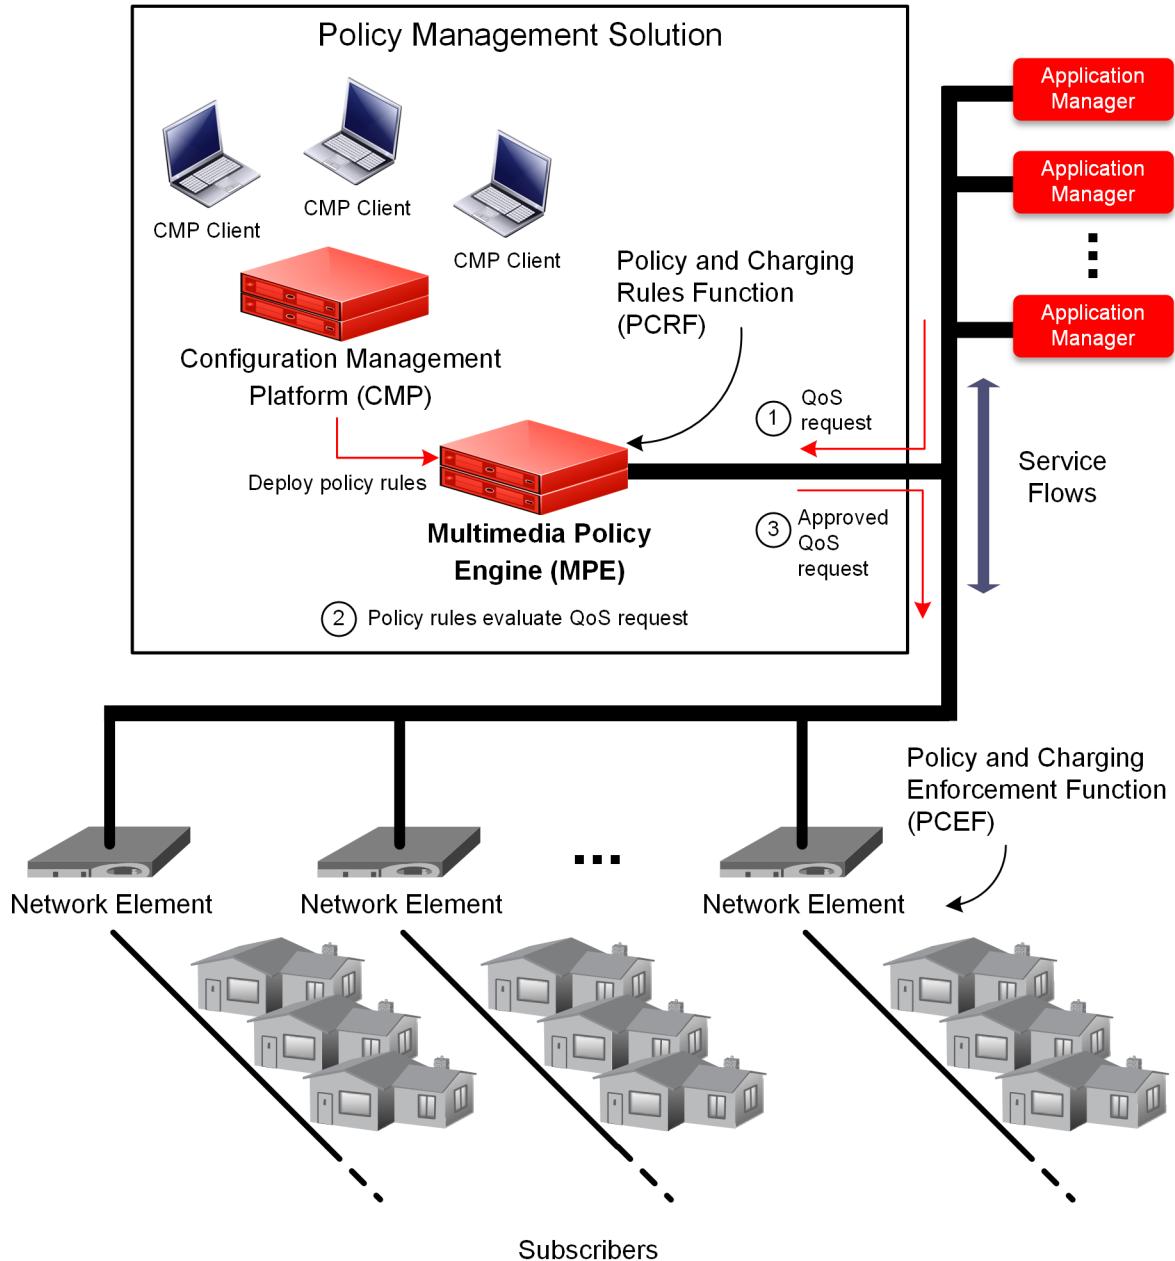


Figure 1: The CMP and MPE Devices

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.

For example, when a subscriber wishes to open an IP-streaming session, the following actions occur:

1. An application receives the subscriber's request and sends a QoS request to the MPE device for the associated network element, requesting that certain network resources be provisioned in order to be used for the application.

2. 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. The MPE device then makes a decision based on the defined policy rules to accept or reject the request.
3. 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.

The MPE device can function in a two-tier hierarchical environment. As a Tier-2 device (called an MPE-S device), it statefully services subscriber flows. As a Tier-1 device (called an MPE-R device), it statelessly routes subscriber flows to MPE-S devices.

The Bandwidth on Demand Application Manager

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

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 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, 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 be defined that simply 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, 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 also allows a calling of an application to specify the duration of QoS resource allocation. The application may choose to completely manage the lifecycle 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 once an application has completed its function. Alternatively, the application may simply tell the BoD AM to keep the resources active for a specified time, or until there is inactivity for a defined period.

BoD devices are configured and managed through the CMP system. For information on using the BoD AM, see the *Bandwidth on Demand Application Manager User Guide*.

The Management Agent

The Management Agent (MA) 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 message routing and policy decisions.

The MA server sits between the Oracle Communications Policy Management Configuration Management Platform (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.

The Oracle Communications Policy Management Configuration Management Platform

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

- Configuration and management of MPE devices
- Definition of network elements
- Creation, modification, deletion, and deployment of policy rules
- Creation, modification, and deletion 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 GUI

You interact with the CMP system through a web browser graphical user interface (GUI). To take best advantage of the GUI, Oracle recommends the following:

Web Browsers

- Mozilla Firefox® release 10.0 or higher
- Microsoft Internet Explorer® 10.0 or higher
- Google Chrome version 20.0 or higher

Monitor Use a resolution of 1024 x 768 or higher

Note: When using the CMP system for the first time, Oracle recommends that you change the default username and password to a self-assigned value. See [Changing a Password](#) for details.

Logging In

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

Note: For information on setting up the login process, see the appropriate CMP user's guide.

Before logging in, you need to know the following:

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

Note: As delivered, the profile **admin** provides full access privileges, and is the assumed profile used in all procedures described in this document. The default username of this profile is **admin** and the default password is **policies**. You cannot delete this user profile, but you should immediately change the password. See [Changing a Password](#).

To log in:

1. Open a web browser and enter the IP address for the CMP system.
The login page opens ([Figure 2: CMP Login Page](#) shows an example).

Note: The title and text on the login page are configurable. For information on changing this page, see [Configuring System Settings](#).

2. Enter the following information in the appropriate fields:
 - a) **Username**
 - b) **Password**
3. Click **Login**.
The main page opens.

You are logged in.

The Oracle Communications Policy Management Solution

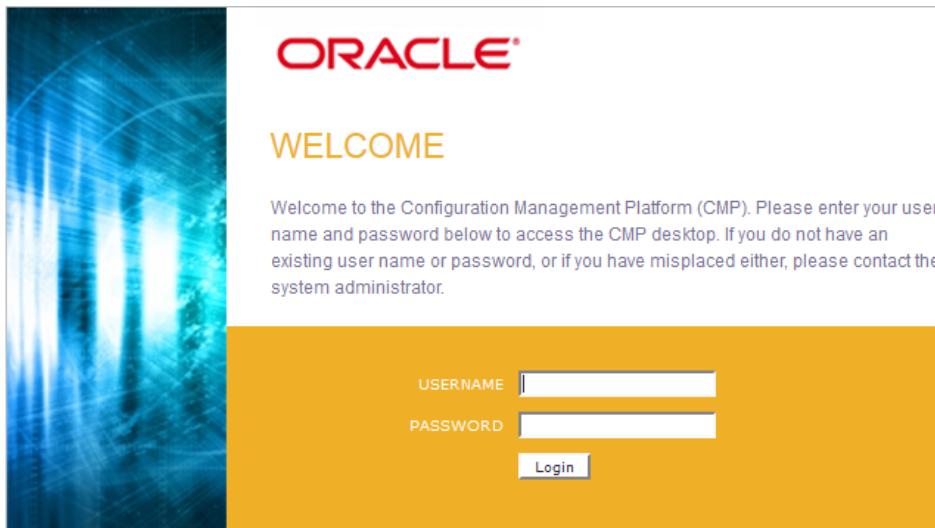


Figure 2: CMP Login Page

GUI Overview

You interact with the CMP system through an intuitive and highly portable graphical user interface (GUI) supporting industry-standard web technologies (SSL, HTTP, HTTPS, IPv4, IPv6, and XML).

Figure 3: Structure of the CMP Cable GUI shows the layout of the CMP GUI.

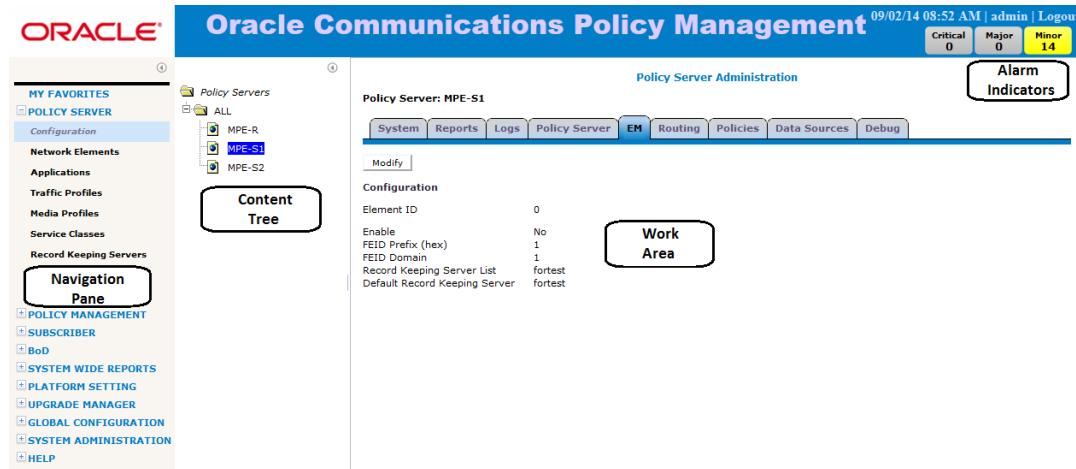


Figure 3: Structure of the CMP Cable GUI

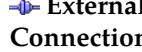
Navigation Pane Provides access to the various available options configured within the CMP system.

You can bookmark options in the navigation pane by right-clicking the option and selecting **Add to Favorite**. Access the bookmarks 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**.

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. The content tree section is not visible with all navigation selections.
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 Mode Indicator	Indicates the current CMP mode. NW-CMP for Network mode or S-CMP for System mode. If there is not a mode indicated, the mode is CMP .

GUI Icons

The CMP GUI 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, time.
 Clone	Use this icon to duplicate a selection in a list.
 Critical error	Displays in reports to indicate a critical error during the blade replication process.
 Delete	When visible in the work area, selecting the Delete icon deletes an item, removing it from the MPE 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.
 >Delete	When visible in the work area, selecting the Delete icon deletes an item, removing it from the MPE 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	The 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	The gear icon displays when a policy references another policy or policy group.
 Hide	When visible in the work area, selecting the hide 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 system.

 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 blade replication process.
 Minor error	Displays in reports to indicate a minor error during the blade replication process.
 Up/Down	The up and down arrow icons are displayed when you can change the sequential order of items in a list.
 Left/Right	The left and right arrow icons are displayed when it is possible to move an item from one list to another.
 OK status	Displays in reports to indicate 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 MPE device PS_1 from policy server group PS_Group2, PS_1 still displays in the ALL group.
 Selection	This icon is in the Policy Wizard. The icon is used to select conditions and actions to add to the policy rule.
 Synch broken	When visible in the Upgrade Manager, indicates that the CMP system 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.
 View Cart	Displays the list of configurable objects selected for the Export action.

Changing a Password

The CMP GUI lets users change their password. Refer to the *CMP User's Guide* for details.

Overview of Major Tasks

The major tasks involved in using MPE devices are configuring, defining manageable elements and profiles, creating and deploying policy rules, managing subscribers and licenses, and administering the authorized CMP users.

The configuration tasks are a series of required steps that must be completed in the following order:

1. Configure the Policy Management topology. This step is described in [Configuring the Policy Management Topology](#).
2. Configure MPE devices by creating Policy Server profiles and then configuring protocol options for each device. This step is described in [Managing Multimedia Policy Engine Devices](#).

3. Configure protocol routing, which enables a Policy Management device to forward requests to other Policy Management devices for further processing. This step is described in [Configuring Protocol Routing](#).
4. Configure BoD devices by creating BoD profiles and then configuring protocol options for each device. This step is described in the [Bandwidth on Demand Application Manager Cable User's Guide](#).

The 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 then specify which MPE device will interact with which network elements. This task is described in [Managing Network Elements](#).
- Create record keeping server profiles. This task is described in [Managing Record Keeping Servers](#).
- Manage events that are sent to record keeping servers. This task is described in [Managing Event Messaging](#).
- Create management agent profiles. This task is described in [Managing Management Agent Servers](#).

The management and administrative tasks, which are optional and performed as needed, are as follows:

1. View reports on the function of the Policy Management systems in your network. This task is described in [System-Wide Reports](#).
2. Manage CMP users, accounts, access, authorization, and operation. This task is described in [System Administration](#).
3. Upgrade software using the Upgrade Manager. These tasks are described in [Upgrade Manager](#).

Chapter 3

Configuring the Policy Management Topology

Topics:

- *About the Policy Management Topology.....30*
- *Setting Up the Topology.....37*
- *Modifying the Topology.....54*
- *Configuring SNMP Settings.....59*
- *Configuring the Upsync Log Alarm Threshold.....61*

Configuring the Policy Management Topology describes how to configure the Policy Management devices into a network and how to configure the CMP system to manage them.

About the Policy Management Topology

You must configure a network topology for the Policy Management products (CMP, MPE, Management Agent (MA), and BoD). 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 loggable incidents and alarms get reported to the CMP system or external network management systems.

Figure 4: Policy Management Topology illustrates a Policy Management topology consisting of a primary (Site 1) and secondary (Site 2) CMP cluster, two georedundant BoD clusters, MA cluster, two Tier-1 (routing) MPE-R clusters, and a series of georedundant Tier-2 MPE-S (serving) clusters.

Note: These terms are defined in subsequent topics.

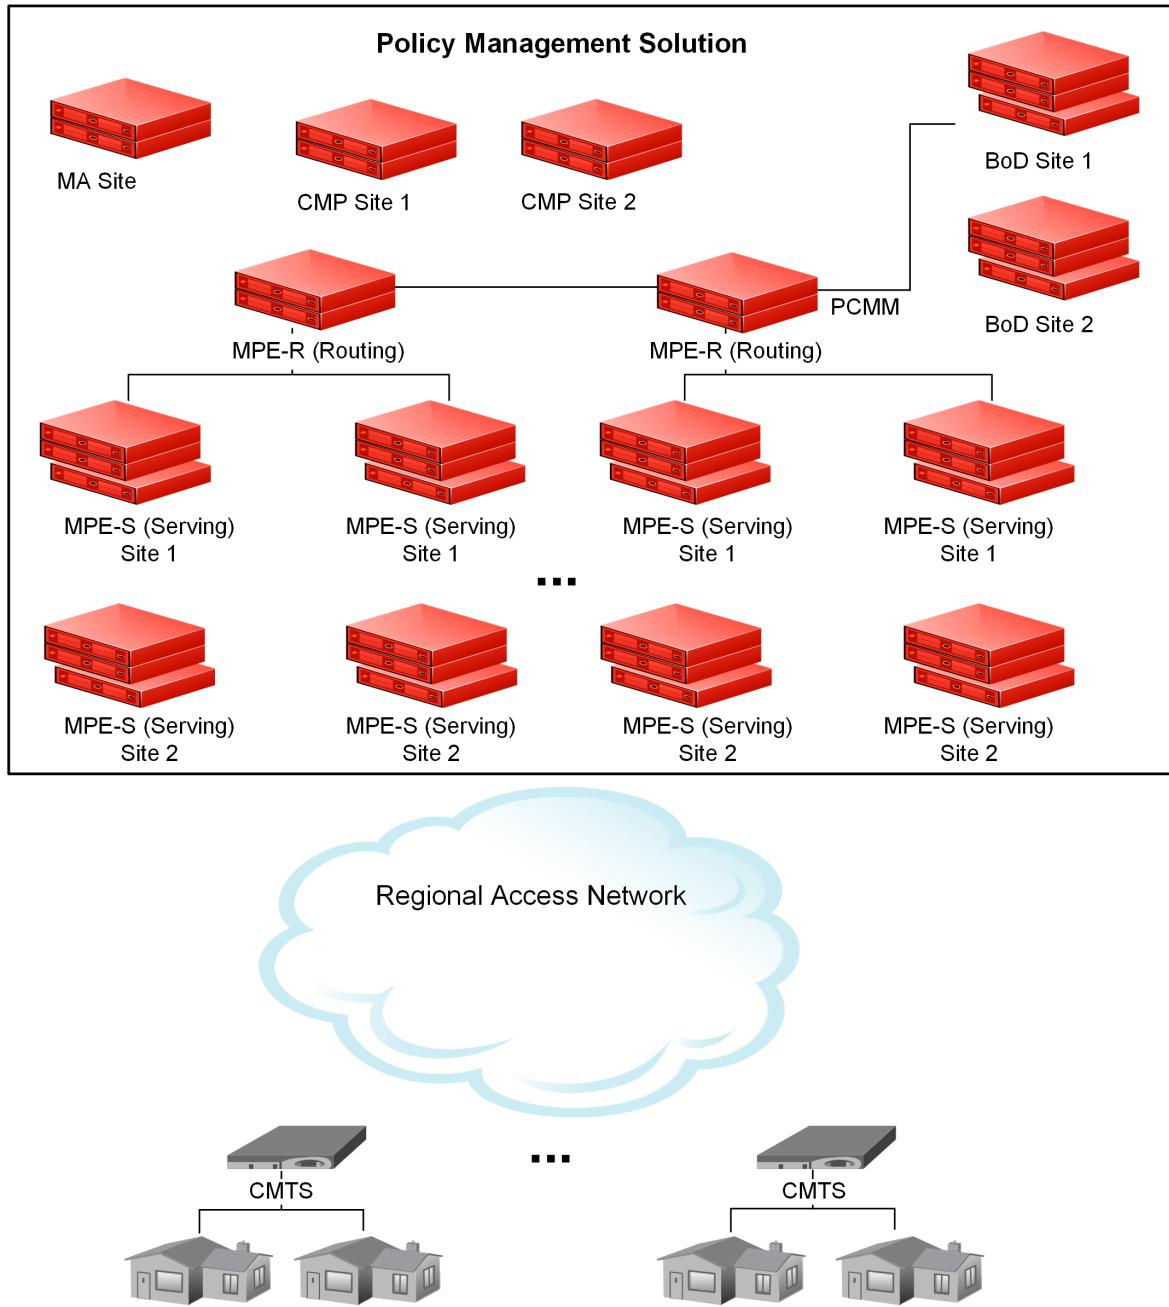


Figure 4: Policy Management Topology

High Availability

High Availability (HA) is provided for all Policy Management cluster configurations. HA is accomplished by using two servers per cluster, an active server and a standby server. Servers are continually monitored by the Communications Core Object Library (COMCOL) in-memory database. As shown in [Figure 5: High Availability](#), 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.

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.
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. COMCOL database runtime processes on each server constantly monitor server status using heartbeat signals.
5. If the active server fails, indicated by skipping a succession of heartbeats, COMCOL instructs the standby server to become the active server and take over the VIP address and connections. Because it has been receiving session state data updates through replication, it 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, or states, can change based on decisions made by the underlying COMCOL database, automatically and at any time. If necessary, the standby server can assume control, at which point it 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, or state, 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; if an active server fails and then recovers, it becomes the standby server, rather than resuming its role as the active server.

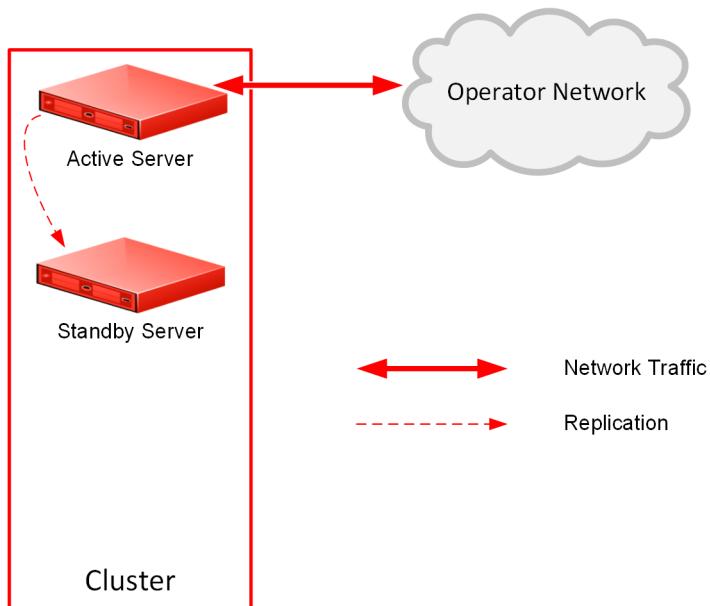


Figure 5: High Availability

Spare Servers

As shown in [Figure 6: Cluster with Active, Standby, and Spare Servers](#), an MPE-S or BoD cluster can contain an additional server, called a spare server. The active server will replicate its database to the spare server as well as the standby 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.

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 COMCOL state database runtime process constantly monitors the status of all servers in the cluster.
5. If the active server fails, it instructs the standby server to take over and become the active server.

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

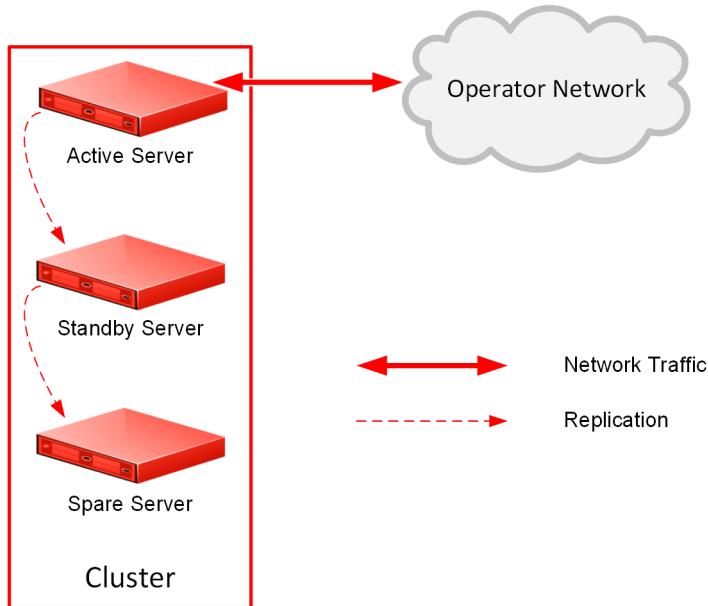


Figure 6: Cluster with Active, Standby, and Spare Servers

MPE-S and BoD Georedundancy

The spare server need not be physically close to the active and standby servers. Georedundancy is an optional configuration provided for MPE-S and BoD clusters in which the spare server can be located in a separate geographical location, as shown in [Figure 7: MPE-S or BoD Georedundant Configuration](#).

If the two servers at one site become unavailable, the third server, located at another site, automatically continues to provide service.

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

1. 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 spare servers communicate using several TCP connections over the OAM network to perform replication, monitor heartbeats, and merge trace-log and alarm data.
2. The servers must share a virtual IP (VIP) cluster address to support failover (georedundancy).
3. The COMCOL database runtime process constantly monitors the status of all servers in the cluster.
4. If the active server fails, it instructs the standby server to take over and become the active server.
5. If both the active and standby servers fail, it instructs the spare server to take over and become the active server.

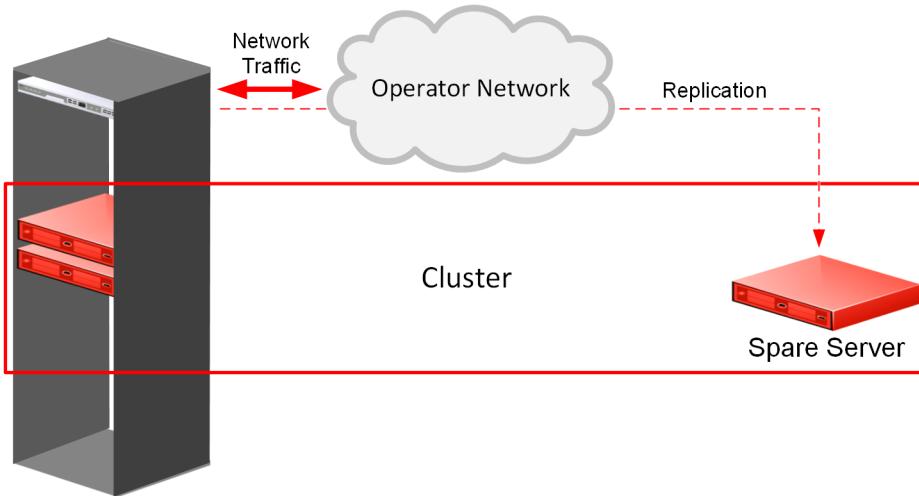


Figure 7: MPE-S or BoD Georedundant Configuration

CMP Georedundancy

As shown in [Figure 8: CMP Georedundancy](#), 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 High Availability, to the active server of the Site 2 cluster.

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, or states, 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. 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").

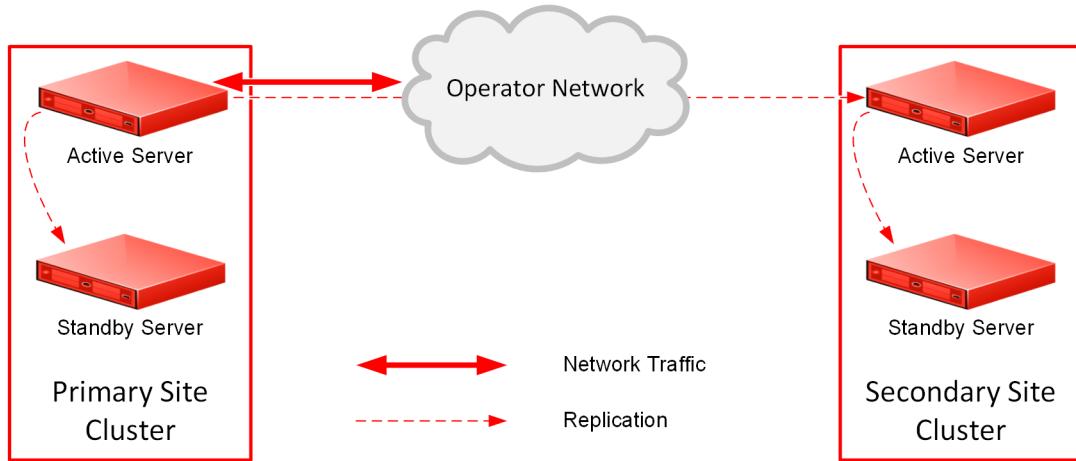


Figure 8: CMP Georedundancy

In a georedundant topology, HP Proliant BL460G6 servers (with a 1x4 mezzanine card) and Netra servers can communicate over a dedicated backup (BKUP) network.

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, server, and connection. 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 is once again able to provide service, then the "active" status transitions back to the server at the primary site. (In contrast, CMP failover is manual. MPE-R and MA clusters do not support failover.)

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

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, 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 Cluster A is also the secondary site of Cluster B, and vice versa.

Cluster 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 (CMP clusters cannot be reversed)

Server Status

You can display the status of a server in the Cluster Information Report (see [Cluster Information Report](#)). The display refreshes every 10 seconds.

The status of a server can be thought of as its current role. The status describes what function the server is currently performing in the 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. (Two servers in the same cluster with the same status is an error condition.)

The status values are as follows:

- **Active:** The active server in a cluster is the server that is externally connected. The active server is the only server that is handling connections and servicing messages and requests. Only the 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 is available to provide service at the primary site.
- **Standby:** The standby server in a cluster is the server 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. Once the previously active server has recovered, it reverts to its former status of standby server.
- **Spare:** The spare server in an MPE-S or BoD cluster is the server that is prepared to take over if no server at the primary site is able to provide service. The spare server has the same replicated data as the servers at the primary site. If there is no server available at the primary site, the spare server becomes active and provides service. As soon as a server in the primary site is available to provide service, that server becomes the active server and the spare server demotes itself and reverts to its former status of spare or standby (depending on the availability of the other servers in the cluster).
- **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 previous Policy Management releases. It can be observed during the upgrade process.

Setting Up the Topology

Topology configuration consists of defining Policy Management sites and clusters, including their addresses and hierarchy. You can add MPE, Management Agent (MA), 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 is as follows:

1. Configure the primary CMP cluster — You start to build a topology by logging in to the active CMP server at the primary site. Configure the CMP cluster settings. The settings are replicated (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. The primary site cannot be deleted from the topology.
2. Configure the secondary CMP cluster (optional) — Use the primary CMP cluster to configure a secondary, or Site 2, CMP cluster. A secondary CMP cluster can provide georedundancy.
3. Configure MPE, MA, and BoD clusters — Enter MPE, MA, and BoD cluster settings on the active CMP server on the primary site. You can define the topology before defining the servers themselves. Once defined, the configuration information is replicated as follows:
 - a. 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.
 - b. 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.)
 - c. The COMCOL 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.

Once 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.

Setting Up a CMP Cluster

You must define at least one CMP cluster before continuing with the topology. The first site you define will be the primary (Site 1) cluster. You can optionally define a secondary CMP cluster.

Before defining the primary (Site 1) cluster, ensure the following:

- The CMP 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
- The CMP server IP connection is active
- The CMP application is running on at least one server

To set up the primary CMP cluster:

1. Log in to the CMP server.
2. From the **Platform Setting** section of the navigation pane, select **Topology Settings**. The **Cluster Configuration** page opens. If a primary cluster is not yet defined, you are prompted, **Initial Configuration Detected. Please add CMP Site 1 Cluster**.
3. From the content tree, select the **All Clusters** group.
4. Click **Add CMP Site1 Cluster**. The **Cluster Settings** page opens. The cluster name and application type are fixed.
5. Enter the following information (*Figure 9: Cluster Settings Page for CMP Cluster* shows an example):
 - a) **HW Type** — Select the type of hardware:
 - **HP ProLiant DL360G8/G9** (default)
 - **NETRA** (for a Netra server)
 - **All other RMS** (for a rack-mounted server)
 - b) **Network VLAN IDs** (appears if you selected **NETRA**) — Enter the Operation, Administration, and Management (OAM), SIG-A, and (optionally) SIG-B virtual LAN (VLAN) IDs, in the range 1–4095. The defaults are 3 for the OAM Virtual IP (VIP) and server IP, 5 for the SIG-A VIP, and 6 for the SIG-B VIP.
 - c) **OAM VIP** (required) — The OAM VIP is the IP address the CMP uses to communicate with a Policy Management cluster. Enter the IPv4 address and mask of the OAM VIP. Enter the address in the standard dot format and the subnet mask in CIDR notation from 0–32 (IPv4).
 - d) **Signaling VIP 1 through Signaling VIP 4** (optional) — Enter up to four IPv4 or IPv6 addresses and masks of the signaling virtual IP (VIP) addresses; for each, select **None**, **SIG-A**, or **SIG-B** to indicate whether the cluster will use an external signaling network. You must enter a Signaling VIP value if you specify either SIG-A or SIG-B. If you enter an IPv4 address, use the standard dot format, and enter the subnet mask in CIDR notation from 0–32. If you enter an IPv6 address, use the standard 8-part colon-separated hexadecimal string format, and enter the subnet mask in CIDR notation from 0–128.

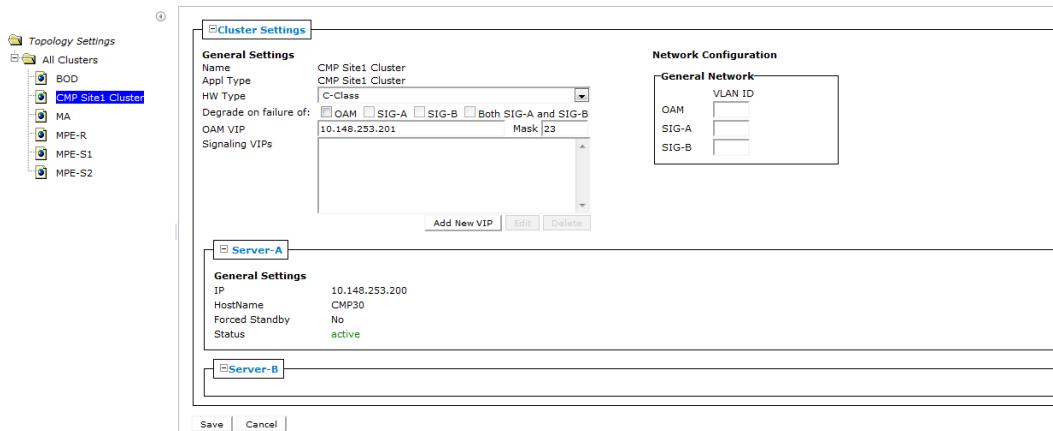


Figure 9: Cluster Settings Page for CMP Cluster

6. Select **Server-A** and enter the following information for the first server of the cluster (which will be the initial active server):
 - a) **IP** (required) — The IP address of the server. Use the standard dot-formatted IP address string.

- b) **HostName** (required) — The name of the server. This must exactly match the host name provisioned for this server (that is, the output of the Linux command `uname -n`).
- c) **Forced Standby** — Select to force this server into standby mode. The flag 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.

7. When you finish, click **Save**.
You are prompted, The VLAN IDs on the page must match the VLAN IDs configured on the server. A mismatch will cause HA to fail. Please confirm that the VLAN IDs are correct before saving.

8. Click **OK**.
You are prompted, Active server will restart and you will be logged out.

9. Click **OK**.
The active server restarts.

10. Log back in to the CMP server.

11. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.
The **Cluster Configuration** page opens.

12. From the content tree, select the **CMP Site1 Cluster**.
The **Topology Configuration** page opens.

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

14. When you finish, click **Save**.

The CMP cluster topology is defined.

Once 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.

Once you define the primary (Site 1) CMP cluster, you can repeat this procedure to define a secondary (Site 2) CMP cluster.

Backup traffic between CMP sites can be sent between sites on the BKUP network.

Setting Up a Site

Sites are available in georedundant mode only. Georedundant sites can contain one or more MPE clusters.

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

To set up a site:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.
The **Cluster Configuration** page opens.
2. From the content tree, select the **All Sites** group.
The **Site Configuration** page opens.
3. Click **Create Site**.
The **New Site** page opens.
4. Enter values for the configuration attributes:
 - a) **Name** (required) — The site name. Enter up to 35 alphanumeric characters, underscores (_), or hyphens (-).

b) **Max Primary Site Failure Threshold** — If the number of cluster pair failures reaches this threshold, a trace log entry and a major alarm are generated.
A pair failure is recorded when both servers at a primary site are either out of service or in forced standby. You can optionally enter a number up to the total number of servers provisioned at this site. The default is no threshold.

5. When you finish, click **Save**.

The site configuration is saved in the CMP database.

To define multiple sites, repeat the procedure starting at [Step 3](#).

Setting Up an MPE or BoD Cluster

Before defining an MPE or BoD cluster, ensure the following:

- The MPE or BoD 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 Operations, Administration, and Maintenance (OAM) IP addresses
- The MPE or BoD server IP connection is active
- The MPE or BoD application is running on at least one server

For information about setting up an MA cluster, see [Setting Up an MA Cluster](#).

To define an MPE or BoD cluster:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.
The **Cluster Configuration** page opens.
2. From the content tree, select the **All Clusters** group.
The **Cluster Configuration** page opens.
3. Click **Add MPE/BoD/MA Cluster**.
The **Topology Configuration** page opens.
4. In the **Cluster Settings** pane, enter the following information ([Figure 10: Cluster Settings Page for an MPE Cluster](#) shows an example):
 - a) **Name** (required) — Name of the cluster. Enter up to 250 characters, excluding quotation marks (") and commas (,).
This field only appears on the page if the CMP system supports georedundancy.
 - b) **Appl Type** — Select **MPE** (the default) or **BoD**.
 - c) **Site Preference** — Select **Normal** (the default) or **Reverse**.
To separate replication traffic onto SigA/B network, The SigA or SigB static IP must be individually specified in Server A, B, and C's **Path Configuration** field. When the **Sig-A** or **Sig-B** option is selected, the primary replication stream connection is set up on a SigA/B bonding interface on a rack mounted server (RMS) or a NETRA server. The replication data is sent to Server C through the SigA/B static path IP.
 - d) **DSCP Marking** — Select the Differentiated Services Control Point (DSCP). The default is **PHB(None)**.
 - e) **Replication Stream Count** — Select from 1 to 8 streams for replications. The default is **1**.
 - f) **Replication & Heartbeat** — Select the network type for replication and heartbeat stream. **None** is the default.

g) **Backup Heartbeat** — Select the network type for the backup heartbeat stream. **None** is the default.

The Backup Heartbeat network option is the secondary path for PCRF-GEO replication heartbeat. For example, if the Server-C primary replication connection fails, another replication heartbeat connection is set up over the Backup Heartbeat network with associated static IP.

5. In the **Primary Site Settings** pane, enter the following information ([Figure 10: Cluster Settings Page for an MPE Cluster](#) shows an example):

- Site Name** — Select **Unspecified** (the default) or the name of a previously defined site. If you select **Unspecified** you create a non-georedundant site, and you cannot subsequently add a secondary site. You can assign multiple clusters to the same site.
- HW Type** — Select **HP ProLiant DL360G8/G9** (the default), **NETRA** (for a configuration in which Signaling and OAM networks are separated onto physically separate equipment), or **All other RMS** (for a rack-mounted server).
- General Network VLAN ID** (appears if you selected **NETRA**) — Enter the Operation, Administration, and Management (OAM), SIG-A, and SIG-B virtual LAN IDs, in the range 1–4095. The defaults are 3 for the OAM VIP and server IP, 5 for the SIG-A VIP, and 6 for the SIG-B VIP.
- OAM VIP** (optional) — Enter the IPv4 address and mask of the OAM virtual IP (VIP) address. The OAM VIP is the IP address the CMP cluster uses to communicate with the MPE or BoD cluster. Enter the address in the standard dot format, and the subnet mask in CIDR notation from 0–32.

Note: This address corresponds to the cluster address in Policy Management systems before V7.5.

- Signaling VIPs** — Click **Add New VIP** to enter up to four IPv4 or IPv6 addresses and masks of the signaling virtual IP (VIP) addresses; for each, select **SIG-A** or **SIG-B** to indicate whether the cluster will use an external signaling network. The Signaling VIP is the IP address a PCEF device uses to communicate with an MPE or BoD cluster. (To support redundant communication channels, an MPE or BoD cluster uses both **SIG-A** and **SIG-B**.) You must enter a Signaling VIP value for SIG-A or SIG-B. If you enter an IPv4 address, use the standard dot format, and enter the subnet mask in CIDR notation from 0–32. If you enter an IPv6 address, use the standard 8-part colon-separated hexadecimal string format, and enter the subnet mask in CIDR notation from 0–128. For a CMP cluster, the Signaling VIP is optional, but for an MPE or BoD cluster, at least one signaling VIP is required (whether it's SIG-A or SIG-B).

6. In the **Server-A** pane enter the following information for the first server of the cluster (which will be the initial active server):

- IP** (required) — The IPv4 address of the server. Enter the standard dot-formatted IPv4 address string.
- HostName** (required) — The name of the server. This must exactly match the host name provisioned for this server (that is, the output of the Linux command `uname -n`).
- Static IP** — Click **Add New** to enter up to four IPv4 or IPv6 addresses and masks of the signaling virtual IP (VIP) addresses; for each, select **SIG-A** or **SIG-B** to indicate whether the cluster will use an external signaling network. The Signaling VIP is the IP address a PCEF device uses to communicate with an MPE or BoD cluster. (To support redundant communication channels, an MPE or BoD cluster uses both **SIG-A** and **SIG-B**.) You must enter a Signaling VIP value for SIG-A or SIG-B. If you enter an IPv4 address, use the standard dot format, and enter the subnet mask in CIDR notation from 0–32. If you enter an IPv6 address, use the standard 8-part colon-separated hexadecimal string format, and enter the subnet mask in CIDR notation from

0–128. For a CMP cluster, the Signaling VIP is optional, but for an MPE or BoD cluster, at least one signaling VIP is required (whether it's SIG-A or SIG-B).

7. (Optional) Click **Add Server-B** and enter the appropriate information for the second server of the cluster.
8. (Optional) **Secondary Site** — For a georedundant cluster, select the name of a previously defined site. The secondary site name must be different from the primary site name.

This section only appears on the page if the CMP system supports georedundancy.

9. (Optional) For a georedundant cluster, click **Add Server-C** and enter the appropriate information for the spare server of the cluster.

This section only appears on the page if the CMP system supports georedundancy. If you define a secondary site, you must define a spare server.

10. When you finish, click **Save**.

You are prompted, *Active server will restart.*

11. Click **OK**.

12. If you are setting up multiple clusters, repeat the above steps.

The MPE or BoD cluster is defined.

Once 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.

For information on setting up a hierarchy of MPE-R and MPE-S clusters, see [Configuring Protocol Routing](#).

Configuring the Policy Management Topology

Topology Configuration

Modify Cluster Settings | Modify Primary Site | Modify Secondary Site | Delete Secondary Site | Back

Cluster Settings

Name	MPE-S01-63-64-8838	DSCP Marking	AF11
App Type	MPE	Replication Stream Count	1
Site Preference	Normal	Replication & Heartbeat	None
Degrade on failure of:	<input type="checkbox"/> OAM <input type="checkbox"/> SIG-A <input type="checkbox"/> SIG-B <input type="checkbox"/> Both SIG-A and SIG-B	Backup Heartbeat	None

Primary Site Settings

General Settings	
Site Name	MPE-S01-Site1
HW Type	RMS
OAM VIP	
Signalling VIPs	<Signalling VIP1> <10.196.145.163/23> <SIG-A>

Server-A

General Settings		Path Configuration
IP	10.148.244.63	Static IP
HostName	MPE83	
Forced Standby	No	
Status	active	

Server-B

General Settings		Path Configuration
IP	10.148.244.64	Static IP
HostName	MPE84	
Forced Standby	No	
Status	standby	

Secondary Site Settings

General Settings	
Site Name	MPE-S01-Site2
HW Type	RMS
OAM VIP	
Signalling VIPs	<Signalling VIP1> <10.196.145.168/23> <SIG-A>

Server-C

General Settings		Path Configuration
IP	10.148.244.68	Static IP
HostName	MPE88	
Forced Standby	No	
Status	Spare	

Figure 10: Cluster Settings Page for an MPE Cluster

Setting Up an MA Cluster

Before defining a Management Agent (MA) cluster, ensure the following:

- The MA 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
- The MA server IP connection is active
- The MA application is running on at least one server

To define an MA cluster:

- From the **Platform Setting** section of the navigation pane, select **Topology Setting**.
The **Cluster Configuration** page opens.
- From the content tree, select the **All Clusters** group.
The **Cluster Configuration** page opens.
- Click **Add MPE/BoD/MA Cluster**.
The **Topology Configuration** page opens.
- In the **Cluster Settings** pane, enter the following information (*Figure 11: Cluster Settings Page for an MA Cluster* shows an example):
 - Name** (required) — Name of the cluster. Enter up to 250 characters, excluding quotation marks (") and commas (,).
 - Appl Type** — Select **MA**.
 - Degrade on failure of** — Select the server degrade criteria based on the connection problems of the network interfaces.
- In the **General Settings** pane, enter the following information (*Figure 11: Cluster Settings Page for an MA Cluster* shows an example):
 - HW Type** — Select **HP ProLiant DL360G8/G9** (the default), **NETRA**, or **All other RMS** (for a rack-mounted server).
 - Network VLAN IDs** (appears if you selected **NETRA**) — Enter the Operation, Administration, and Management (OAM), SIG-A, and SIG-B virtual LAN IDs, in the range 1–4095. The defaults are 3 for the OAM VIP and server IP, 5 for the SIG-A VIP, and 6 for the SIG-B VIP.
 - OAM VIP** (optional) — Enter the IPv4 address and mask of the OAM virtual IP (VIP) address. The OAM VIP is the IP address the CMP cluster uses to communicate with the MA cluster. Enter the address in the standard dot format, and the subnet mask in CIDR notation from 0–32.

Note: This address corresponds to the cluster address in Policy Management systems before V7.5.
 - Signaling VIP 1** through **Signaling VIP 4** — Enter up to four IPv4 or IPv6 addresses and masks of the signaling virtual IP (VIP) addresses; for each, select **None**, **SIG-A**, or **SIG-B** to indicate whether the cluster will use an external signaling network. The Signaling VIP is the IP address a PCEF device uses to communicate with an MA cluster. (To support redundant communication channels, an MA cluster can use both **SIG-A** and **SIG-B**.) You must enter a Signaling VIP value if you specify either **SIG-A** or **SIG-B**. If you enter an IPv4 address, use the standard dot format, and enter the subnet mask in CIDR notation from 0–32. If you enter an IPv6 address, use the standard 8-part colon-separated hexadecimal string format, and enter the subnet mask in CIDR notation from 0–128.

Note: For an MA cluster, at least one signaling VIP is required (either **SIG-A** or **SIG-B**).
- Select **Server-A** and enter the following information for the first server of the cluster (which will be the initial active server):
 - IP** (required) — The IPv4 address of the server. Enter the standard dot-formatted IPv4 address string.
 - HostName** (required) — The name of the server. This must exactly match the host name provisioned for this server (that is, the output of the Linux command **uname -n**).
- (Optional) Click **Add Server-B** and enter the appropriate information for the second server of the cluster.
- When you finish, click **Save**.
You are prompted, **Active server will restart**.

9. Click **OK**.
10. If you are setting up multiple clusters, repeat the above steps.

The MA cluster is defined.

Once 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.



Figure 11: Cluster Settings Page for an MA Cluster

Setting Up a Georedundant Cluster

This procedure contains the steps for setting up a non-CMP cluster:

- MPE
- BoD

Before defining a 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
- The server IP connection is active
- The server application is running on at least one server

If your system is not set up for georedundancy, see [Setting Up an MPE or BoD Cluster](#).

To define a cluster in a georedundant system:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**. The **Topology Configuration** page opens.
2. From the content tree, select the **All Clusters** folder. The defined clusters are listed.
3. Click **Add MPE/BoD Cluster**. The **Topology Configuration** page opens. Each section of the **Topology Configuration** page can be collapsed or expanded.
4. Define the general settings for the cluster in the **Cluster Settings** section of the page:

- a) **Name** (required) — Name of the cluster. Enter up to 250 characters, excluding quotation marks (") and commas (,).
- b) **Appl Type** — Select the application type:
 - **MPE** (default)
 - **BoD**
- c) **Site Preference** — Select **Normal** (default) or **Reverse**.
- d) **DSCP Marking** — Select the type of Differentiated Services Code Point (DSCP) marking for replication traffic. The valid code points are **AF11, AF12, AF13, AF21, AF22, AF23, AF31, AF32, AF33, AF41, AF42, AF43** (assured forwarding), **CS1, CS2, CS3, CS4, CS5, CS6, CS7** (class selector), **EF** (expedited forwarding), or **PHB(None)** (default for no marking). For information on DSCP marking, see [Setting Up an MPE or BoD Cluster](#).
- e) **Replication Stream Count** — Select 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) **Replication & Heartbeat** — Select a network to carry inter-site replication and heartbeat traffic. This field only appears if the system supports georedundancy:
 - **None** (default)
 - **OAM**
 - **SIG-A**
 - **SIG-B**

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

- g) **Backup Heartbeat** — Select a network to carry inter-site backup heartbeat traffic. This field only appears if the system supports georedundancy:
 - **None** (default)
 - **OAM**
 - **SIG-A**
 - **SIG-B**

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

5. Define the primary site settings in the **Primary Site Settings** section of the page:
 - a) **Site Name** — Select **Unspecified** (default) or the name of a previously defined site. If you select **Unspecified**, you create a non-georedundant site, and cannot add a secondary site. You can assign multiple clusters to the same site.
 - b) **HW Type** — Select the hardware type:
 - **HP Proliant DL360G8/G9**
 - **NETRA** (Netra server)
 - **All other RMS**
 - c) **OAM VIP** (optional) — The OAM VIP is the address the CMP cluster uses to communicate with the cluster. Enter the IPv4 address and mask of the OAM virtual IP (VIP) address. Enter the address in the standard dot format, and enter the subnet mask in CIDR notation from 0–32.

d) **Signaling VIPs** — 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.

At least one signaling VIP is required.

Click **Add New VIP** to add a VIP to the system. You can enter up to four IPv4 or IPv6 addresses and masks of the signaling VIP addresses.

For each new VIP, enter the address and mask in the New Signaling VIP dialog. Select **SIG-A** or **SIG-B** to indicate whether the cluster will use an external signaling network.

For an IPv4 address, use the standard dot format, and enter the subnet mask in CIDR notation from 0–32. For an IPv6 address, use the standard 8-part colon-separated hexadecimal string format, and enter the subnet mask in CIDR notation from 0–128.

e) **General Network VLAN ID** — This field appears if you selected **NETRA**. Enter the **OAM**, **SIG-A**, and **SIG-B** VLAN IDs, in the range 1–4095. The defaults are 3 for the OAM network and server IP, 5 for the SIG-A network, and 6 for the SIG-B network.

6. Define Server-A in the **Server-A** section of the page:

- IP** (required) — The IPv4 address of the server. Enter the standard IP dot-formatted IPv4 address string.
- HostName** — The name of the server. This must exactly match the host name provisioned for this server (the output of the Linux command `uname -n`).
- Forced Standby** — Select to put Server A into forced standby. (By default, Server A will be the initial active server of the cluster.)
- Static IP** — If an alternate replication path and secondary HA heartbeat path is used, then a server address must be entered in this field. Click **Add New**. In the New Path dialog, enter an IP address and mask, and select the network:
 - **SIG-A**
 - **SIG-B**
 - **BKUP** (if the hardware type is **C-Class(Segregated Traffic)** or **NETRA**)

7. (Optional) Define Server-B in the **Server-B** section of the page. Click **Add Server-B** and enter the standby server information for the cluster:

- IP** (required) — The IPv4 address of the server. Enter the standard IP dot-formatted IPv4 address string.
- HostName** — The name of the server. This must exactly match the host name provisioned for this server (the output of the Linux command `uname -n`).
- Forced Standby** — Select to put Server A into forced standby. (By default, Server A will be the initial active server of the cluster.)
- Static IP** — If an alternate replication path and secondary HA heartbeat path is used, then a server address must be entered in this field. Click **Add New**. In the New Path dialog, enter an IP address and mask, and select the network:
 - **SIG-A**
 - **SIG-B**

8. Define the secondary site information in the **Secondary Site Settings** section of the page:

- Site Name** — Select **Unspecified** (default) or the name of a previously defined site. This site name must be different from the primary site name. If you select **Unspecified**, you create a

non-georedundant site, and cannot add a secondary site. You can assign multiple clusters to the same site.

b) **HW Type** — Select the hardware type:

- **HP Proliant DL360G8/G9**
- **NETRA** (Netra server)
- **All other RMS**

c) **OAM VIP** (optional) — Enter the IPv4 address and mask of the OAM virtual IP (VIP) address. The OAM VIP is the address the CMP cluster uses to communicate with the cluster.

Enter the address in the standard dot format, and enter the subnet mask in CIDR notation from 0–32.

d) **Signaling VIPs** — The signaling VIP is the IP address a PCEF device uses to communicate with a cluster. Clusters support redundant communication channels, named SIG-A and SIG-B, for carriers that use redundant signaling channels.

At least one signaling VIP is required.

Click **Add New VIP** to add a VIP to the system. You can enter up to four IPv4 or IPv6 addresses and masks of the signaling VIP addresses.

For each new VIP, enter the address and mask in the New Signaling VIP dialog. Select **SIG-A** or **SIG-B** to indicate whether the cluster will use an external signaling network.

For an IPv4 address, use the standard dot format, and enter the subnet mask in CIDR notation from 0–32. For an IPv6 address, use the standard 8-part colon-separated hexadecimal string format, and enter the subnet mask in CIDR notation from 0–128.

e) **General Network VLAN ID** — This field appears if you selected **NETRA**. Enter the **OAM**, **SIG-A**, and **SIG-B** VLAN IDs, in the range 1–4095. The defaults are 3 for the OAM network and server IP, 5 for the SIG-A network, and 6 for the SIG-B network.

9. When you finish, click **Save**.

You are prompted, The VLAN IDs on the page must match the VLAN IDs configured on the server. A mismatch will cause HA to fail. Please confirm that the VLAN IDs are correct before saving..

10. Click **OK**.

11. If you are setting up multiple clusters, repeat the procedure.

The cluster is defined.

Example: Setting Up Georedundancy

This topic describes how to add a secondary site, Site-2, to a Policy Management topology, and a third server, located at Site-2, to an existing active/standby MPE cluster located at the primary site, Site-1, to create a two-site (Site-1 and Site-2), three-system (active, standby, and spare, or Server-A, Server-B, and Server-C) mated georedundant MPE cluster. If the primary site were to fail, the spare server would assume the active role. The procedure 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 systems in the topology are running 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

Before beginning the procedure, you will need to collect or provide the following information (to collect information, see [Setting Up a Georedundant Cluster](#)).

Tip: This information can be collected at any time before beginning the procedure without interrupting service.

- The names of existing clusters
- Names for the sites (this procedure uses **site-1** and **site-2**)
- The maximum primary site failure threshold, to record site failures (0 is recommended)
- The OAM VIP address of the existing Site-1 CMP system and, if applicable, the georedundant CMP system
- (Optional) a designated network path, either OAM, SIG-A or SIG-B, for backup (secondary) HA heartbeats between Site-1 and Site-2
- (Optional) a designated network path, either OAM, SIG-A or SIG-B, for WAN replication traffic between Site-1 and Site-2
- Initial provisioning information for Server-C:
 - A hostname (this procedure uses **Server-C**)
 - For CMP access, an OAM IPv4 or IPv6 address and subnet mask
 - An OAM IPv4/IPv6 default route
 - A list of network time protocol (NTP) server IP addresses
 - A list of domain name system (DNS) server IP addresses
 - VLAN IDs for OAM, SIG-A, and SIG-B network paths
 - 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:
 - If the SIG-A/SIG-B network is used for either WAN replication traffic or backup (secondary) HA heartbeats, a VLAN ID for the SIG-A/SIG-B network path
 - If the SIG-A/SIG-B network is used for either WAN replication traffic or backup (secondary) HA heartbeats, an IPv4 static address and subnet mask on the SIG-A network for Server-A
 - If the SIG-A/SIG-B network is used for either WAN replication traffic or backup (secondary) HA heartbeats, an IPv4 static address and subnet mask on the SIG-B network for Server-B
 - Verify that firewall rules are correctly provisioned (for more information, see the *Platform Configuration User's Guide*)
- 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

To create a secondary site and a georedundant MPE cluster, follow these steps.



Caution: This procedure interrupts service.

1. Using the Platform Management & Configuration utility, install the MPE application on Server-C. For more information, refer to the PM&C documentation, or contact MOS for support.

- Using the Platform Configuration utility, provision Server-C with the following configuration information.

For more information, see the *Platform Configuration User's Guide*.

- HostName
- OAM Real IP Address
- OAM Default Route
- NTP Server
- DNS Server A
- DNS Server B (optional)
- DNS Search
- Device
- OAM VLAN Id
- SIG A VLAN Id
- SIG B VLAN Id (optional)

- Using the Platform Configuration utility, export routing configuration information from Server-A or Server-B and import it into Server-C.

For more information, see the *Platform Configuration User's Guide*.

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

Note: Unless otherwise noted, the remaining steps are performed within the CMP system.

- If this is the first georedundant cluster in your topology, set the CMP system to manage georedundant MPE/MRA/BoD systems.

See [The Mode Settings Page](#).

On the content tree of the **Topology Configuration** page, the **All Sites** group becomes available.

- Define the two sites.

See [Setting Up a Site](#).

The sites become visible on the **Site Configuration** page:

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

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

The **Cluster Configuration** page opens, displaying the defined clusters.

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

The **Topology Configuration** page opens for the MPE cluster.

- Click **Modify Primary Site**.

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

10. 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 values in the **HW Type** field, **Network Configuration** section, and **Signaling VIPs** field.
11. In the **Server-A** section of the page:
 - a) Confirm the values in the **General Settings** section.
 - b) In the **Path Configuration** section, click **Add New**, enter the Static IP address and subnet mask for the **SIG-A** network in the pop-up window, and click **Save**.
12. Repeat [Step 11](#) for Server-B.

The primary site settings are defined; for example:

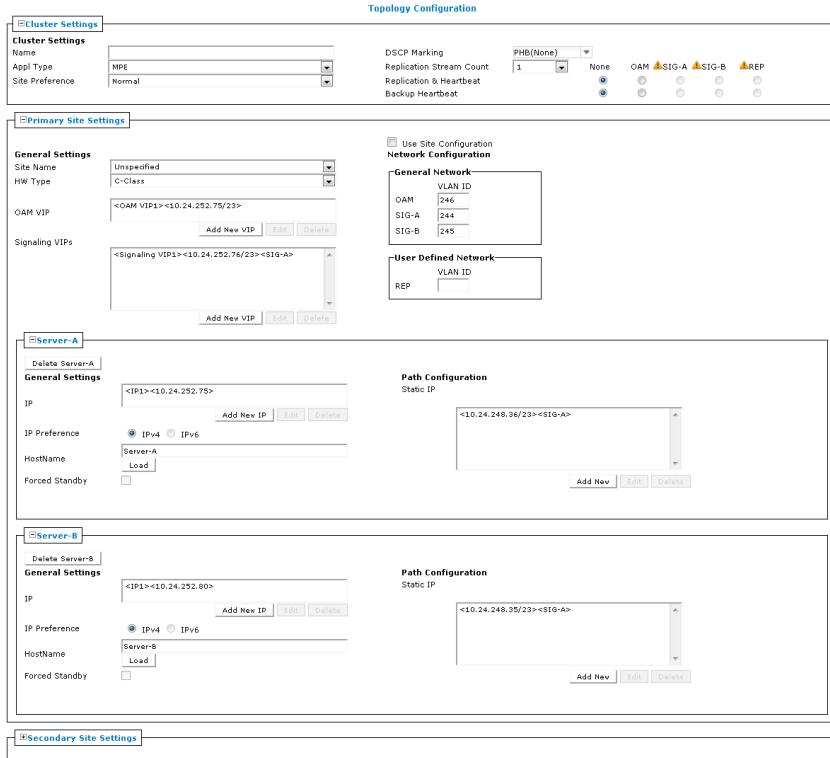


Figure 12: Example of Primary Site Settings

13. Click **Save** (at the bottom of the page).
You are prompted, **Active server will restart**.
14. Click **OK**.
Server-A restarts. You must now define the Site-2 and Server-C configuration.
15. From the **Platform Setting** section of the navigation pane, select **Topology Settings**. The **Cluster Configuration** page opens.
16. From the content tree, select the **All Clusters** group.
The **Cluster Configuration** page opens, displaying the defined clusters.
17. On the **Cluster Configuration** page, for the MPE cluster you are expanding, click the operation **View**.

The **Topology Configuration** page opens for the MPE cluster.

18. Click Modify Secondary Site.

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

19. In the Secondary Site Settings section of the page:

- a) In the **Site Name** field, select the secondary site name (**Site-2** in this example).
- b) Confirm the values in the **HW Type** field, **Network Configuration** section, and **Signaling VIPs** field.

20. In the Server-C section of the page:

- a) Click **Add Server-C**.
- b) In the **IP** field, enter the OAM IP address.
- c) In the **IP Preference** field, enter the preferred IP version, either **IPv4** or **IPv6**. If IPv6 is selected, the server will prefer to use the IPv6 address for communication. If neither an OAM IPv6 IP nor a static IP address defined, the IPv6 radio button cannot be selected here. Similarly, If neither an IPv4 OAM IP nor a static IP address is defined, the IPv4 radio button isn't accessible.
- d) In the **HostName** field, enter the host name.
- e) In the **Path Configuration** section, click **Add New**, enter the Static IP address and subnet mask for the SIG-A network in the window, and click **Save**.

Site-2 and Server-C are defined, and Server-C is placed in Force Standby status; for example:

21. Click Save (at the bottom of the page).

You are prompted, **Active server will restart**.

22. Click OK.

Server-A restarts.

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

23. Click the status of Server-C.

The status changes to **Spare**.

24. Click Save.

The configuration is saved.

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

The **Cluster Configuration** page opens.

26. From the content tree, select the All Clusters group.

The **Cluster Configuration** page opens, displaying the defined clusters.

27. On the Cluster Configuration page, for the MPE cluster you are expanding, click the operation View.

The **Topology Configuration** page opens for the MPE cluster.

28. Click Modify Cluster Settings.

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

29. In the Cluster Settings section of the page:

- a) If DSCP marking is used, in the **DSCP Marking** field, select the type of marking.
- b) If replication streams are used, in the **Replication Stream Count** field, select the number of streams.
- c) In the **Replication & Heartbeat** field, select the network used (or **None** to return to the system default).
- d) If the backup (secondary) heartbeat feature is used, in the **Backup Heartbeat** field, select the network used (or **None** to disable the feature).

30. Click **Save.**

The configuration is saved.

31. Verify the status of Server-C by viewing the cluster.

Server-C appears as part of the cluster in the Force Standby state with replication on.

32. 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) appears in the report, but with severity Clear.

33. On Server-C, using the Platform Configuration utility, exchange SSH keys with the other servers of the cluster.

This step is not completed using the CMP software. See the *Platform Configuration User's Guide*.

34. On the CMP system, using the Platform Configuration utility, exchange SSH keys with all other CMP systems in the topology.

This step is not completed using the CMP software. See the *Platform Configuration User's Guide*.

35. Modify the cluster configuration to cancel the Force Standby state of Server-C.

The state of Server-C changes to Spare.

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

For more information, see [KPI Dashboard](#).

Server-C appears as part of the cluster, in the Spare state.

37. (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, see the *Policy Wizard Reference*.

38. Use the **Data Sources function to configure routes on Server-C to existing data sources.**

For more information, see [Configuring Data Source Interfaces](#).

39. Use the **Topology Settings function to force Server-A and Server-B to standby status to verify that Server-C is functioning normally:**

a) Select the MPE cluster and click **Modify Primary Site**.

b) In the **Server-A** section of the page, select **Forced Standby**.

c) In the **Server-B** section of the page, select **Forced Standby**.

d) Click **Save** (at the bottom of the page). You are prompted, *Active server will restart*.

e) Click **OK**.

f) Use the **System Maintenance** function to verify that Server-C has become the active server.

g) For more information, see [Data Source Statistics](#).

40. Use the **Topology Settings function to cancel the Force Standby state of Server-A and Server-B.**
On the **System Maintenance** page, the state of Server-C changes to Spare.

Note: Either Server-A or Server-B may assume the Active role. Oracle recommends not attempting to force Server-A back into the Active role, as doing so would interrupt service.

The two sites, and the georedundant MPE cluster, are defined, and the normal function of all servers is verified.

If your topology includes MRA systems, add additional routes on the system to reach Server-C in the case of a cluster restart, and add the georedundant MPE cluster to an MPE pool. For more information, refer to the *Policy Front End Wireless User's Guide*.

Modifying the Topology

After the topology is configured, 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

You can modify a cluster even if the standby or spare server is off line. However, you cannot modify or delete the active server of a cluster.

Modifying a Site

To modify a site:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.
The **Cluster Configuration** page opens, displaying information about the clusters in the Policy Management network topology.
2. From the content tree, select the site you want to modify.
The **Site Configuration** page displays information about the site.
3. Click **Modify**.
The **Modify Site** page opens.
4. Modify site information.
For a description of the fields contained on this page, see [Setting Up a Site](#).
5. When you finish, click **Save**.

The site is modified.

Removing a Site from the Topology

You can remove a site from a georedundant topology. You can only remove a site if the site is not referenced by a C-level cluster. When a site is in use by a cluster, you will receive the following message if you try to delete the site: *Site cannot be deleted because it is referred in following clusters: cluster1[, cluster2[, . . .]]*.

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.
2. Select the **All Sites** group.

The **Site Configuration** page opens, displaying the configured sites.

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

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

A confirmation message displays.

4. Click **Delete**.

The site is removed from the topology.

Modifying an MPE, MA, or BoD Cluster

To modify an MPE, MA, or BoD cluster:

1. From the **Platform Setting** section of the navigation pane, select **Topology Setting**.
The **Topology Configuration** page opens.

2. From the content tree, select the cluster.

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

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

- To modify cluster settings, click **Modify Cluster Settings**.
- To modify the primary site configuration, click **Modify Primary Site**.
- To modify the secondary site configuration, click **Modify Secondary Site**.
- To delete the secondary site configuration, click **Delete Secondary Site**.

The fields on the **Topology Configuration** page become editable.

4. Make the changes.

You must make changes to each section individually. You can remove some servers from a cluster, but not all of them. You can select **Forced Standby** on all servers of an MPE, MA, or BoD cluster.

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

5. When you finish, click **Save**.

You are prompted, **Warning:** You may need to restart the application or reboot the server for the new topology configuration to take effect.

6. Click **OK**.

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

Modifying a CMP Cluster

To modify a CMP cluster:

1. From the **Platform Setting** section of the navigation pane, select **Topology Setting**.
The **Cluster Configuration** page opens.

2. From the content tree, select the cluster.

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

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

- To modify cluster settings, click **Modify Cluster Settings**.

- To modify the configuration of the first server defined in the cluster, click **Modify Server-A**.
- To modify the configuration of the second server defined in the cluster, click **Modify Server-B**.

The fields on the **Topology Configuration** page become editable. For information on configurable fields, see [Setting Up a CMP Cluster](#).

4. Make changes.

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.

5. When you finish, click **Save**.

The following message displays: **Warning:** You may need to restart the application or reboot the server for the new topology configuration to take effect.

6. Click **OK**.

The cluster is modified. 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 an MPE, BoD, MA, or Site 2 CMP cluster from the topology. (You cannot remove the Site 1 (primary) CMP cluster from the topology.) Before removing an MPE, BoD, MA , or Site 2 CMP cluster, remove the profiles of its servers; see [Deleting a Policy Server Profile](#).

To remove a cluster from the topology:

1. From the **Platform Setting** section of the navigation pane, select **Topology Setting**.
The **Topology Configuration** page opens.
2. From the content tree, select the **All Clusters** folder.
The **Cluster Configuration** page opens, displaying a table that lists the cluster settings information for the clusters defined in the topology.
3. In the topology configuration table, in the row listing the cluster, click **Delete**.
A confirmation message displays.
4. Click **Delete**.

The cluster is removed from the topology and the following message displays: The cluster *cluster_name* was successfully deleted. Go to each server and su - platcfg -> Policy Configuration -> Cluster Configuration Removal -> Cluster information cleanup

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

Reversing Cluster Preference

You can change the preference, or predilection, of the servers in a cluster to be active or spare. This setting is only available when the system has been configured for georedundancy.

To reverse cluster preference:

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

The Cluster Configuration page opens.

2. Select the cluster from the content tree.

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

3. Click **Modify Cluster Settings**.

The fields become editable.

4. In the **Cluster Settings** section of the page, toggle the **Site Preference** between **Normal** and **Reverse**.
5. Click **Save**.

The cluster preferences are reversed.

Demoting a 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.

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



Caution: Perform cluster demotion before cluster promotion. 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 CMP cluster:

1. Log in to the currently active georedundant CMP cluster.
2. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.

The Cluster Configuration page opens, displaying a cluster settings table listing information about the clusters defined in the topology. 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 options to **View** and **Demote**.

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

The page displays the message: **This server you signed in is the 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**.

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

The Cluster Configuration page opens, displaying a cluster settings table listing information about the clusters defined in the topology. You should see options to **View** and **Promote**.



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

5. Return to the browser window logged in to the primary CMP cluster.

You should still be on the Cluster Configuration page.

6. In the Cluster Settings table, in the row listing the primary CMP cluster, click **Demote**.
You are prompted, *Are you sure you want to demote this Cluster?*
7. Click **OK**.
The page displays the message *Demote cluster successfully*.
8. Log out of the CMP system for the cluster you have just demoted.
9. Return to the browser window logged in to the secondary CMP cluster.
You should still be on the **Cluster Configuration** page.
10. Wait two minutes.
11. In the Cluster Settings table, in the row listing the secondary CMP cluster, click **Promote**.
You are prompted, *Are you sure you want to promote this Cluster?*
12. Click **OK**.
The page displays the message *Promote cluster successfully*.
13. Log out of the CMP system for the cluster you have just promoted.
14. Log in to the CMP system for the cluster you have just promoted.
15. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.
The **Cluster Configuration** page opens, displaying a cluster settings table listing information about the clusters defined in the topology. The primary cluster is marked with (P), and the name of the secondary cluster is marked with (S). The old primary cluster may briefly display as off-line.
Note: You should see options to **View** and **Demote**. All functions available from the primary CMP cluster should now appear and be accessible.
16. Wait ten minutes and then use the **Cluster Configuration** page to verify that both the primary and secondary CMP clusters are available and have the correct status.

The primary CMP cluster is demoted, and the secondary cluster is promoted to primary status.

Forcing a Server into Standby Status

You can change the status of an active or spare server in a cluster to Standby. You would do this, for example, to the active server prior to performing maintenance on it.

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

- If the server is active, it demotes itself.
- The server does not assume the active role, regardless of the status or roles of the other servers in the cluster.
- The server continues as part of its cluster, and reports its status as Forced-Standby.

To force a server into standby status:

1. From the **Platform Setting** section of the navigation pane, select **Topology Setting**.
The **Topology Configuration** page opens, displaying a cluster settings table listing information about the clusters defined in the topology.
2. In the cluster settings table, in the row listing the cluster containing the server you want to force into standby status, click **View**.
The **Topology Configuration** page displays information about the cluster.
3. Select the server:
 - For a CMP cluster, click **Modify Server-A** or **Modify Server-B**.

- For an MPE, MA, or BoD cluster, click the site containing the server, either **Modify Primary Site** or **Modify Secondary Site**.

4. Select **Forced Standby**.
5. When you finish, click **Save**.

The server is placed in standby status.

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, or such that each server generates and delivers its own traps.

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

To configure SNMP settings:

1. Log in to the CMP system from its server address as a user with administrator privileges.
The navigation pane opens.
2. From the **Platform Setting** section of the navigation pane, select **SNMP Setting**.
The **SNMP Settings** page opens.
3. Click **Modify**.
The **Edit SNMP Settings** page opens.
4. Edit the settings.
5. When you finish, click **Save**.

Table 2: SNMP Attributes describes the SNMP attributes that can be edited.

Table 2: SNMP Attributes

Field Name	Description
Manager 1-5	<p>This field is required for an NMS to receive traps and send SNMP requests. Each Manager field can be filled with either a valid host name or an IPv4 address. These fields have the following restrictions:</p> <ul style="list-style-type: none"> • A hostname should include only alphanumeric characters. • Maximum length is 20 characters. • Case insensitive (uppercase and lowercase are treated as the same). • This field can contain an IP address. An IP address should be in a standard dot-formatted IP address string. <p>By default, these fields are blank.</p> <p>Note: The IPv6 address is not supported.</p>

Field Name	Description
Enabled Versions	<p>Supported SNMP versions:</p> <ul style="list-style-type: none"> • SNMPv2c • SNMPv3 • SNMPv2c and SNMPv3 (default)
Traps Enabled	<p>Enable the sending SNMPv2 traps. The default is enabled.</p> <p>Note: This option must be selected to use the SNMP Trap Forwarding feature</p> <p>Clear the checkbox to disable sending SNMPv2 traps.</p>
Traps from individual Servers	<p>Enable sending traps from an individual server. The default is disabled.</p> <p>Note: To use the SNMP Trap Forwarding feature, ensure that this option is not selected.</p> <p>Clear the checkbox to send traps from the active CMP system only.</p>
SNMPv2c Community Name	<p>The SNMP read-write community string. This field has the following restrictions:</p> <ul style="list-style-type: none"> • The field is required if SNMPv2c is enabled. • The name can contain alphanumeric characters and cannot exceed 31 characters in length. • The name cannot be either private or public. <p>The default value is snmppublic.</p>
SNMPv3 Engine ID	<p>Configured Engine ID for SNMPv3. This field has the following restrictions:</p> <ul style="list-style-type: none"> • The field is required if SNMPv3 is enabled. • The Engine ID includes only hexadecimal digits (0-9 and a-f). • The length can be from 10 to 64 digits. <p>The default is no value (empty).</p>
SNMPv3 Security Level	<p>SNMPv3 Authentication and Privacy options are:</p> <ul style="list-style-type: none"> • No Auth No Priv — Authenticate using the Username. No Privacy. • Auth No Priv — Authentication using MD5 or SHA1 protocol. • Auth Priv — Authenticate using MD5 or SHA1 protocol. Encrypt using the AES and DES protocol. <p>The default value is Auth Priv.</p>
SNMPv3 Authentication Type	<p>Authentication protocol for SNMPv3. Options are:</p> <ul style="list-style-type: none"> • SHA-1 — Use Secure Hash Algorithm authentication. • MD5 — Use Message Digest authentication. <p>The default value is SHA-1.</p>

Field Name	Description
SNMPv3 Privacy Type	Privacy Protocol for SNMPv3. Options are: <ul style="list-style-type: none"> • AES — Use Advanced Encryption Standard privacy. • DES — Use Data Encryption Standard privacy. The default value is AES .
SNMPv3 Username	The SNMPv3 User Name. This field has the following restrictions: <ul style="list-style-type: none"> • The field is required if SNMPv3 is enabled. • The name must contain alphanumeric characters and cannot exceed 32 characters in length. The default value is TekSNMPUser .
SNMPv3 Password	Authentication password for SNMPv3. This value is also used for msgPrivacyParameters. This field has the following restrictions: <ul style="list-style-type: none"> • The field is required if SNMPv3 is enabled. • The length of the password must be between 8 and 64 characters and can include any character. The default value is snmpv3password .

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 upsSync 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 Rep Stats](#) for more information.

To configure the upsSync log alarm threshold:

1. From the **Platform Setting** section of the navigation pane, select **Platform Configuration Setting**. The **Platform Configuration** page is displayed.
2. Click **Modify**.
3. Enter the threshold.
4. When you finish, click **Save**.

Chapter 4

Managing Multimedia Policy Engine Devices

Topics:

- *Policy Server Profiles*.....63
- *Configuring Protocol Options on the Policy Server*.....65
- *Configuring MPE Advanced Settings*.....68
- *Configuring Data Source Interfaces*.....73
- *Working with Policy Server Groups*.....75
- *About Reapplying a Configuration*.....78
- *Resetting Counters*.....78
- *Checking the Status of an MPE Server*.....79
- *Policy Server Reports*.....80
- *Viewing Policy Server Logs*.....92

Managing Multimedia Policy Engine Devices describes how to use the CMP system to configure and manage the Multimedia Policy Engine (MPE) devices in a network.

Note: The MPE device is the Policy Management policy server. The terms policy server and MPE device are synonymous.

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 define profiles, you deploy them to MPE devices across the network.

The following subsections describe how to manage policy server profiles. For information on deploying defined policies to an MPE device, see the *Policy Wizard Reference*.

Creating a Policy Server Profile

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. Enter values for the configuration attributes:
 - a) **Associated Cluster** (required) — Select the cluster with which to associate this MPE device.
 - b) **Name** — Name of this MPE device. The default is the associated cluster name. A name is subject to the following rules:
 - Case insensitive (uppercase and lowercase are treated as the same)
 - Must be no longer than 255 characters
 - Must not contain quotation marks ("") or commas (,)
 - c) **Description / Location** (optional) — Information that defines the function or location of this MPE device.
 - d) **Secure Connection** — Designates whether or not to use the HTTPS protocol for communication between Policy Management devices. If selected, devices communicate over port 8443.
 - e) **Type** — Defines the policy server type:
 - **Oracle** (default) — The policy server is an MPE device and can be fully managed by the CMP.
 - **Unmanaged** — The policy server is not an MPE device and therefore cannot be actively managed by the CMP. This selection is useful when an MPE device is routing traffic to a non-Oracle policy server.
5. When you finish, click **Save**.

The profile appears in the list of policy servers. You have defined the policy server profile.

For most protocols to function correctly, once a policy server profile is created, you must configure attribute information on the **Policy Server** tab (see [Configuring Protocol Options on the Policy Server](#)).

Once you have defined policy server profiles for the MPE devices in your Policy Management network, you can associate network elements with them (see [Managing Network Elements](#)).

Configuring or Modifying a Policy Server Profile

To configure or modify a policy server profile:

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 in the work area.

The 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).
- **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.
- **Logs** — Displays the Trace Log and Syslog configurations.
- **Policy Server** — Lets you associate applications and network elements with the MPE device and configure protocol information.
- **EM** — Lets you view and configure event messages.
- **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.
- **Policies** — Lets you manage policies that are deployed on the policy server.
- **Data Sources** — Lets you configure interfaces to DHCP (Dynamic Host Configuration Protocol) systems.

3. Select the tab that contains the information you want to modify and click **Modify**.
4. Edit the configuration.
5. When you finish your modifications, click **Save**.

Deleting a Policy Server Profile

Deleting a policy server (MPE device) profile from the ALL group also deletes it from any associated group.

To delete an MPE device profile:

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, displaying all defined MPE devices; for example:

Policy Server	Status
mpe26-098	On-line
mpe26-100	On-line
mpe26-101	On-line
mpe26-103	On-line
mpe26-104	On-line
mpe26-105	Degraded
mpe26-107	On-line

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, select the MPE device. The **Policy Server Administration** page opens. Click the **System** tab, and then click **Delete**.

A confirmation message displays.

4. Click **OK** to delete the MPE device profile.
The profile is removed from the list.

The policy server profile is deleted.

Configuring Protocol Options on the Policy Server

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 policy server groups; the initial group is **ALL**.
2. From the content tree, select the desired 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 is used.

The following tables define the available options. (The options you see vary depending on the mode configuration of your system.)

- *Associations Configuration Options*
- *General Configuration Options*
- *PCMM Configuration Options*
- *Diameter Configuration Options*
- *Diameter AF Default Profiles Configuration Options*

- When you finish, click **Save**.

You have defined the protocol options for this MPE device.

Associations Configuration Options

Table 3: 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 the <i>Policy Wizard Reference</i> .
Network Elements	The network elements associated with this MPE device. To modify this list, click Manage . For more information on network elements, see Managing Network Elements .
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 Managing Network Elements .

General Configuration Options

Table 4: General Configuration Options

Attribute	Description
Management Agent	Visible if your network contains management agents. For more information, see Managing Management Agent Servers .

PCMM Configuration Options

Table 5: 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.

Attribute	Description
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 Configuration Options

Table 6: Diameter Configuration Options

Attribute	Description
Diameter Realm	The domain of responsibility (for example, <code>galactel.com</code>) for the MPE device.
Diameter Identity	The fully qualified domain name (FQDN) of the MPE device (for example, <code>mpe3.galactel.com</code>).
Diameter PCMM AMID	This is the AMID used when requests are received from an Application Function (AF) that are translated to PCMM. This AMID must be unique among all the AMIDs that are used by any PCMM Application Managers (AMs) in your network. The default is 3472.
Diameter PCMM Classifier Priority	The default classifier priority for the PCMM gate. The default is 64.
Validate user	If enabled, sessions for unknown users are rejected.
Protocol Timer Profile	The timer profile to use.
Prevent Overlapping Rule Names	If selected, rule names that are dynamically generated on the primary and spare MPE devices in the same Gx session are unique.

Attribute	Description
Allow Multiple Rx Connections with the same Origin-host Id	When enabled, the MPE device accepts multiple Rx connections with the same Origin-Host Attribute Value Pair (AVP) and source IP address.
Timers	Rx-to-PCMM gate timers. Enter values in seconds for T1 (authorized, default 1 second), T2 (reserved, default 300 seconds), and T3 (committed, default 300 seconds).

Diameter AF Default Profiles Configuration Options

Table 7: 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.

Configuring MPE Advanced Settings

The Advanced configuration page provides access to factory-default attribute settings that are not normally changed.



Caution: Do not attempt to change a service override without first consulting with My Oracle Support.

CAUTION

The MPE advanced settings are used for the following:

- [Configuring Session Clean Up Options](#)
- [Configuring Configuration Keys](#)
- [Configuring Load Shedding Rules](#)

Configuring Session Clean Up Options

Session cleanup options are used to configure the methods used for cleaning up stale sessions and how often cleanup occurs.

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.
3. Select the **Policy Server** tab.
The configuration settings for the policy server are displayed.
4. Click **Advanced**.
Advanced configuration settings, including the session clean up options are displayed and can be edited.

Table 8: Session Clean Up Options

Attribute	Description
Max Session Cleanup Rate (sessions/sec)	The rate at which the cleanup task attempts to clean stale sessions. The default is 50 sessions/sec. Valid range is 1–5000 sessions/sec. Do not modify this setting without consulting My Oracle Support.
Max Session Iteration Rate (sessions/sec)	The maximum rate at which the cleanup task iterates through the sessions database. Default value is 1000. Valid range is 1–100000. Do not modify this setting without consulting My Oracle Support.
Max Duration For Session Iteration (hours)	The maximum duration to iterate through the sessions. Default value is 2 hours. Valid range is 1–2 hours. Do not modify this setting without consulting My Oracle Support.
Session Cleanup Start Time	Defines the time of day when the cleanup task occurs. Specify either Start Time or Interval for defining when session cleanup occurs by clicking the associated radio button and entering/selecting a value. Time can be specified in 24-hour format from the list. No default value is defined.
Session Cleanup Interval (hours)	Defines the interval, in hours, at which the cleanup task occurs. Specify either Start Time or Interval for defining when session cleanup occurs by clicking the associated radio button and entering/selecting a value. The default is 5 hours. Valid range

	is 0–6 hours. A value of 0 disables cleanup. Do not modify this setting without consulting My Oracle Support.
Override Cleanup Audit	Select to turn override clean up audit on. When selected, the cleanup task bypasses the audit process and deletes all sessions that are stale for the session validity time. The default is deselected.
Cleanup Stale Rx Sessions	Determines whether the RxSessionCleanUp task should clean up stale Rx sessions. The default is true.
Rx Session Validity Time (hours)	The amount of time, hours, after which an Rx session is declared as stale. The default is 24 hours.
Cleanup Stale PCMM Sessions	Determines whether the CleanupStalePcmmSessions task should clean up stale PCMM sessions. The default is true.
PCMM Session Validity Time (hours)	The amount of time, hours, after which a PCMM session is declared as stale. The default is 24 hours.

- When you finish, click **Save**.

The settings are applied to the selected MPE device.

Configuring Configuration Keys



Caution: Do not attempt to change a service override without first consulting with My Oracle Support.

CAUTION

Configuration key changes are made using the Other Advanced Configuration Settings section of the Advanced configuration page.

Make configuration key changes as follows:

- 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**.
- From the content tree, select the MPE device.
The **Policy Server Administration** page opens.
- Select the **Policy Server** tab.
The configuration settings for the policy server are displayed.
- Click **Advanced**.
The advanced configuration displays.
- Configuration key changes are made using the Other Advanced Configuration Settings section.
 - To add a key to the table** — 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 device is undetermined.

Enter the following values:

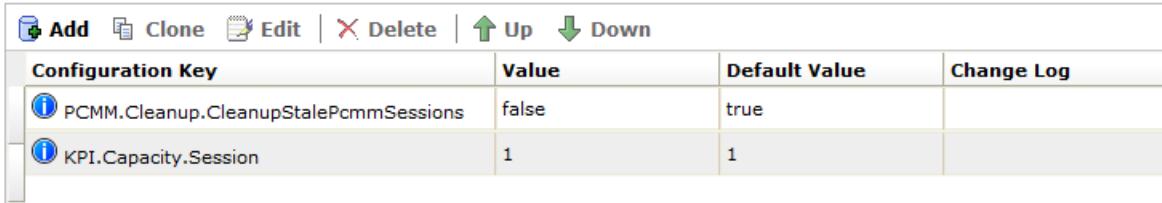
- **Configuration Key** — The attribute to set
- **Value** — The attribute value (up to 255 characters)

When you finish, click **OK**. The key is displayed in the table with its defined and default values.

- **To clone a key in the table** — Select an existing key in the table and click **Clone**; the Clone Configuration Key Value window opens with that key's information filled in. Make changes as required. When you finish, click **Save**.
- **To edit a key in the table** — Select an existing key in the table and click **Edit**; the Edit Configuration Key Value window opens with that key's information. Make changes as required. When you finish, click **Save**.
- **To delete a key from the table** — Select an existing key in the table and click **Delete**; you are prompted with a confirmation message. Click **Delete** to remove the key.

The following figure shows an example configuration key.

Other Advanced Configuration Settings



Configuration Key	Value	Default Value	Change Log
PCMM.Cleanup.CleanupStalePcmmSessions	false	true	
KPI.Capacity.Session	1	1	

6. When you finish, click **Save**.

The settings are applied to the selected MPE device.

Configuring Load Shedding Rules

You can configure load shedding rules to determine how an MPE device reacts to a processing backlog. This state is called “busyness.” By default there are three levels of busyness, from Level 1, the least busy, to Level 3, the most busy. With each successive level, the MPE 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 silently discarded without further processing, since the originator would have already given up on that request. [Table 9: Default MPE Device Busyness Levels](#) shows the default load-shedding rules for an MPE device.

Note: MPE-R has enough capacity to handle route message. It is recommended that Load Shedding is disabled for the for MPE-R.

Table 9: Default MPE Device Busyness Levels

Busyness Level	Rule Name	Actions
Level 1	DefaultRule14	Reject Gate messages with PCMM_TOO_BUSY This message applies to GateSet and GateDelete only and has an error code of 127 with a subcode of 127.
Level 2	DefaultRule7	Reject Rx AAR messages with DIAMETER_TOO_BUSY
	DefaultRule15	Reject Gate messages with PCMM_TOO_BUSY

Busyness Level	Rule Name	Actions
Level 3	DefaultRule11	Reject Rx AAR messages with DIAMETER_TOO_BUSY
	DefaultRule16	Reject Gate messages with PCMM_TOO_BUSY

Use the **Load Shedding Configuration** section of the **Advanced Configuration** page to edit, reorder, or add new rules at each of the three levels of busyness for an MPE-S device based on the amount of backlog. To reach a configured level of busyness:

- The backlog of outstanding messages in a node crosses a pre-defined threshold for the level.
- The backlog has been above the busyness level threshold for a minimum amount of time.

At each level, the MPE-S device can be configured to take one of the following actions (referred to as rules) until the busyness level clears:

- Reject new messages with a specific result code (the default is DIAMETER_TOO_BUSY).
- Drop the message.

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. 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-S device.
The **Policy Server Administration** page opens.
3. Select the **Policy Server** tab.
The Policy Server configuration settings are displayed.
4. Click **Advanced**.
Advanced configuration settings are displayed and can be edited.
5. In the **Load Shedding Configuration** section of the page, **Enabled** is selected by default.
6. Configure the rules for the busyness levels:
 - a) Click  (arrow next to the level) to expand the level.
 - b) Click **Add** and select either **Diameter** or **PCMM**.
The **Add Load Shedding Rule** dialog appears.
 - c) Enter the values for the load shedding rule:
 - **Name** — Name of the rule.
 - **Application** — Select the application the rule applies to. You can select **Rx**.
 - **Message** — Type of message the rule applies to (which depends on the application chosen).
 - **Request Types** (available depending on the message selected) — Select the Request-Type attribute-value pairs (AVPs) that the message must contain. You can select **Initial**, **Update**, and/or **Terminate**. The default request type varies depending on the configured mode.
 - **APNs** (available depending on the message selected) — Enter a CSV list of one or more access point names that the message must contain.
 - **Action** — Select the action to be taken if the criteria are met for the busyness level.

Drop (drop the message)

Answer With (select a code from the drop-down list)

Answer With Code (enter a code) (available depending on the message selected)

Vendor ID (enter a vendor ID) (available depending on the message selected)

- d) Click **Save**.

The rule is displayed in the table.

7. Once a rule is defined, you can clone ()**Clone**), edit ()**Edit**), or delete (x) the rule by selecting the rule and clicking the appropriate button.
8. When you finish, click **Save**.

The settings are applied to the selected MPE-S device.

Configuring Data Source Interfaces

Before the MPE device can communicate with any external data sources, you must configure the interface. To configure a data source interface:

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, listing 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) When you finish, click **Save**.
6. (Optional) Add, modify, delete, or reorder data sources.
 - Cloning an entry in the table
 1. Select an entry in the table.
 2. Click  **Clone**. The **Clone** window opens with the information for the entry.
 3. Make changes as required.
 4. When you finish, click **Save**. The entry is added to the table

- Editing an entry in the table
 1. Select the entry in the table.
 2. Click  **Edit**. The **Edit Response** window opens, displaying the information for the entry.
 3. Make changes as required.
 4. When you finish, click **Save**. The entry is updated in the table.
- Deleting a value from the table
 1. Select the entry in the table.
 2. Click  **Delete**. A confirmation message displays.
 3. 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:

 1. Select an entry.
 2. Click  **Up** or  **Down**. The search order is changed.

7. When you finish, click **Save**.

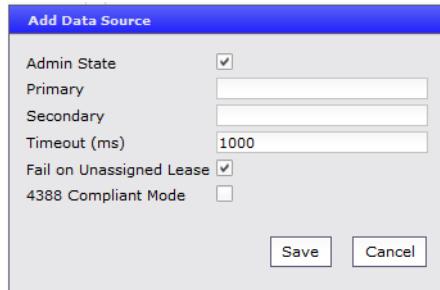
8. The following general settings are available:

- **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. When you finish, click **Save**.

Configuring a DHCP Data Source

For DHCP, you can configure connections to one or two DHCP servers. In the **Add Data Source** window, enter the following:



1. **Admin State** — Select to enable this data source.
Selected by default.
2. **Primary** — FQDN or IP address in IPv4 or IPv6 format of primary DHCP server.
3. **Secondary** — FQDN or IP address in IPv4 or IPv6 format of secondary DHCP server.
4. **Timeout (ms)** — Length of time to wait before a DCHP request times out.

The default timeout is 1000 ms (one second).

5. **Fail on Unassigned Lease** — Action to take if the DHCP server returns an unassigned lease.
By default, the action fails.
6. **4388 Compliant Mode** — Compliant with RFC4388 (“Dynamic Host Configuration Protocol (DHCP) Leasequery”).
7. When you finish, click **Save**.

The DHCP data source is defined.

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 (,).

The screenshot shows a web-based administrative interface for policy server groups. The main title is 'Policy Server Administration'. A sub-section titled 'Create Group' is active. Below it, a table labeled 'Information' has one row. The 'Name' column contains the value 'Denver' in a text input field. At the bottom of the dialog are two buttons: 'Save' and 'Cancel'.

5. When you finish, 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. When you finish, 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 cannot contain quotation marks ("") or commas (,).
5. When you finish, click **Save**.

The sub-group is added to the selected group.

Renaming a Policy Server Group

To modify the name assigned to 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.
3. Click **Modify**.
The **Modify Group** page opens.

4. Enter the new name in the **Name** field.
The name cannot contain quotation marks ("") or commas (,).
5. When you finish, click **Save**.

The group is renamed.

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 [Deleting a Policy Server Profile](#).

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 policy server is removed immediately; there is no 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 **Remove**.

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 apply the Rediscover Cluster operation to the CMP system to re-initialize the Cluster Information Report for the device, thereby clearing out status of the failed (see [Policy Server Reports](#) 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 its **System** tab with the notation **Config 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

The device is synchronized with the CMP system.

Reapplying the Configuration to a Group of Devices

All of the servers in a group are synchronized with the CMP system.

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 device's current operating status (**On-line** or **Off-line**) and profile configuration.

Figure 13: Group View shows an example of a Group View in which one of the servers is degraded.

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

Figure 13: Group View

Trash can icon Click Delete (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 All Counters** — Resets all counters under Policy Statistics and Protocol Statistics back to initial values except for “Session count” and “Downstream Bandwidth” under Network Elements.
 - **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 filesystem), 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 finishing.
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 0–20, 20–40, 40–60, 60–80, 80–100, 100–150, 150–200, 200–250, and >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 the *Policy Wizard Reference*.

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 14: Sample Protocol Statistics](#) shows a sample.

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	

Figure 14: Sample Protocol Statistics

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.

Note:

1. Statistical information is returned from the MPE device 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 device is brought on line, 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:

Reset Counters	Resets all latency counters.
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 15: Sample Error Statistics shows a sample.

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

Figure 15: Sample Error Statistics

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 MaxTransactionsPerSecond	The highest value of the counter MaxTransactionsPerSecond during the previous interval.
Interval MaxMRABindingCount	The highest value of the counter MaxMRABindingCount during the previous interval. (This value is 0 on MPE clusters.)
Interval MaxSessionCount	The highest value of the counter MaxSessionCount during the previous interval.
Interval MaxPDNConnectionCount	The highest value of the counter MaxPDNConnectionCount 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 10: PCMM \(PacketCable MultiMedia\) Protocol Statistics](#) shows information for these protocols:
 - PCMM CMTS (Cable Modem Termination System)
 - PCMM AM (Application Manager)
 - PCMM DPS

- *Table 11: Record Keeping Servers Protocol Statistics* shows information for Record Keeping Servers (RKSs).
- *Table 12: CMTS with Lost Connections Statistics* shows information for individual CMTS systems with lost connections.
- *Table 13: MGPI Statistics* shows information for the MGPI protocol.
- *Table 14: Diameter AF (Application Function) Statistics* shows information for the Diameter AF protocol.
- *Table 15: Latency Statistics* shows information for these statistics:
 - Diameter AF
 - PCMM AM
 - PCMM CMTS
 - PCMM DPS
- *Table 16: Protocol Error Statistics* shows information for these statistics:
 - Diameter
 - PCMM
- *Table 17: Connection Error Statistics* shows information for these statistics:
 - Diameter
 - PCMM
- *Table 18: KPI Interval Statistics* shows information for the KPI collection interval.
- *Table 19: Policy Statistics* shows information for policy execution.

Table 10: 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 11: 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 12: CMTS with Lost Connections Statistics

Reports Display Name	OSSI XML Name
CMTS Name	
CMTS IP Address	
Last Connection Time	
Last Disconnection Time	

Table 13: MGPI Statistics

Reports Display Name	OSSI XML Name
Total flows	
Actual gates	
Multi-flow gates	
Effective gates	

Table 14: 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
Rx-Pcmm Messages Timeout	
Last stats reset time	
Currently active sessions	Active Session Count

Reports Display Name	OSSI XML Name
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
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 15: 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 [<i>timeframe</i>] (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 16: 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	
Diameter Protocol Errors on each error codes	(see specific errors listed in GUI)

Table 17: 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	

Reports Display Name	OSSI XML Name
Protocol Errors on each error codes	(see specific errors listed in GUI)

Table 18: 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 19: Policy Statistics

Reports Display Name	OSSI XML Name
Peg Count	
Evaluated	
Executed	
Ignored	
Policy Details Stats:	
Name	
Evaluated	Eval Count
Executed	Trigger Count
Ignored	
Policy write state on session create	
Name	
Evaluated	
Executed	
Ignored	
Total Execution Time (ms)	
Max Execution Time (ms)	

Reports Display Name	OSSI XML Name
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.

Log information, including the log levels, is displayed. Refer to example for [Figure 16: Policy Server Administration, Logs Tab — Cable](#). 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.

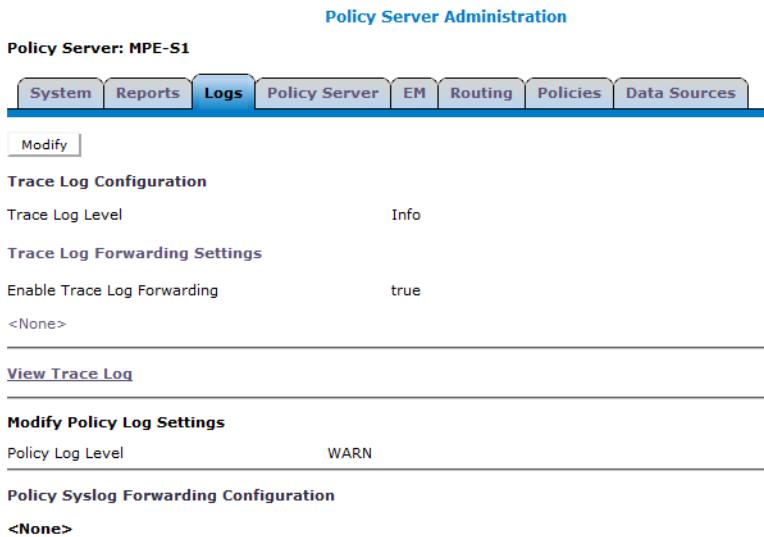


Figure 16: Policy Server Administration, Logs Tab — Cable

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, which is a database table, at eight levels: Emergency (ID 4560), Alert (ID 4561), Critical (4562), Error (ID 4563), Warning (ID 4564), Notice (ID 4565) Info (ID 4566), and Debug (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” 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 *Policy Management Troubleshooting Guide*.
- **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 that appeared in events last month and this month, only results from this month appear.

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** — Once 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** — Once 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.

When you are finished viewing the trace log, click **Close**.

Syslog Support

Notifications generated by policy actions are sent to the standard UNIX syslog. No other notifications are forwarded to syslog. For information on policy actions, see the *Policy Wizard Reference*.

Note: This feature is separate from TPD syslog support.

You can define multiple destinations for notifications, and filter notifications by severity level. For more information, see [Configuring Log Settings](#).

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 > ALL**
The **Policy Server Administration** page opens in the work area and displays a list of MPE devices.

2. 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.

3. Select the **Logs** tab.

The **Policy Server Administration** page opens and details the logs configuration settings for the specified device.

4. To edit the logs configuration settings, click **Modify**.

The editable fields open in the work area.

5. 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** — 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

Caution: Before changing the default logging level, consider the implications. Lowering the trace log level setting from its default value (for example, from **Warning** to **Info**) causes more notifications to be recorded in the trace 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 trace log, and may cause you to miss important notifications.

6. You can enable and configure Trace Log Forwarding for individual clusters.

Note: The CMP system provides log forwarding configuration for all products that have trace logs: MPE, MRA, MA, BoD, and the CMP itself.

For each cluster, enter the following:

- a) Select the **Enable Trace Log Forwarding** checkbox in the **Modify Trace Log Forwarding Settings** section of the page.
Trace Log Forwarding fields 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 hostname 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 5](#) for a description of each level.

7. 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.

8. When you finish, click **Save**.

The log settings are configured.

Chapter 5

Configuring Protocol Routing

Topics:

- *PCMM Routing Architectures.....98*
- *Configuring PCMM Routing.....98*
- *Configuring Rx-to-PCMM Routing.....99*

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 architecture, 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 17: Modify Routing Configuration Page* 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. When you finish, click **Save**.

PCMM routing is configured.

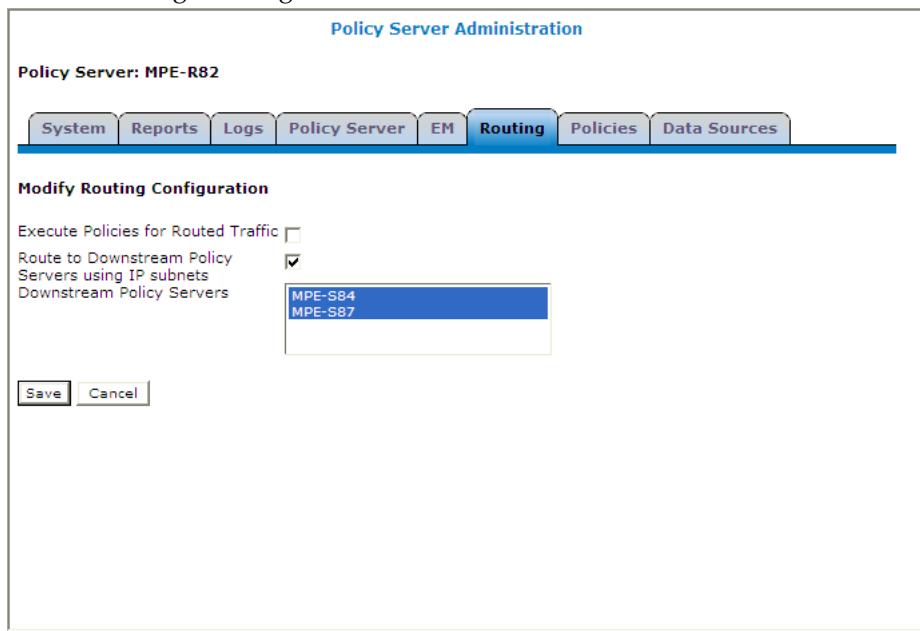


Figure 17: Modify Routing Configuration Page

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. When you finish, click **Save**.

Rx-to-PCMM routing is configured.

Chapter

6

Managing Network Elements

Topics:

- *About Network Elements.....102*
- *Defining a Network Element.....102*
- *Configuring Options for Network Elements....105*
- *Associating a Network Element with an MPE Device.....106*
- *Working with Network Element Groups.....107*

Managing Network Elements 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

Once 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 once 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, once you define network elements, you can combine them into network element groups.

Defining a Network Element

You must define a network element for each device associated with any of the MPE devices within the network. To define 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 **Network Elements**.
The **Network Element Administration** page opens.
3. Click **Create Network Element**.
The **New Network Element** page opens.
4. Enter information for the network element:
 - a) **Name** (required) — The name you assign to the network element.
Enter up to 250 alphanumeric characters. The name can include underscores (_), hyphens (-), colons (:), and periods (.)
 - b) **Host Name/IP Address** (required) — 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) **Type** (required) — Select the type of network element.
The supported types are:
 - **CMTS** (default) — Cable Modem Termination System
- 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.
- h) **Network Element Groups which contain this Network Element** — Specifies the links to other network elements.

5. In **Policy Servers associated with this Network Element**, select one or more policy servers (MPE devices) to associate with this network element.
6. In **Network Element Groups which contain this Network Element** select the group (see [Adding a Network Element to a Network Element Group](#)).
7. When you finish, click **Save**.

You have created the definition for a network element and the network element is listed on the **Network Element Administration** 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 [Defining a Network Element](#).
5. When you finish, 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**.

- From the content tree, select the **ALL** group.

The **Network Element Administration** page opens in the work area, displaying all defined network elements.

- From the work area, click  (trash can icon) located to the right of the network element.
A confirmation message displays.
- Click **OK** to delete the network element.

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:

- 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**.

- From the content tree, select **ALL**.

The **Network Element Administration** page opens in the work area.

- Click **Bulk Delete**.

The **Bulk Delete Network Elements** page opens.

- Select the network elements or network element groups to delete.

- (Optional) Filter the search by entering a search pattern (for example, **cmts***) and click **Filter**.

By default, the **Search Pattern** entry box contains an asterisk (*) to match all network elements.

- Click **Bulk Delete**.

A confirmation message displays.

- Click **OK** to delete the network element.

The selected network element or group definitions are deleted from the CMP database and all associated MPE devices.

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:

- 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**.

- From the content tree, select **ALL**.

The **Network Element Administration** page opens in the work area.

- Click **Search**.

The **Network Element Search Criteria** window opens.

- 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.

Configuring Options for Network Elements

The following subsections 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 [About Scheduled Tasks](#) for more information.

Configuring CMTS

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 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.

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. When you finish, click **Save**.

The CMTS device is configured.

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:

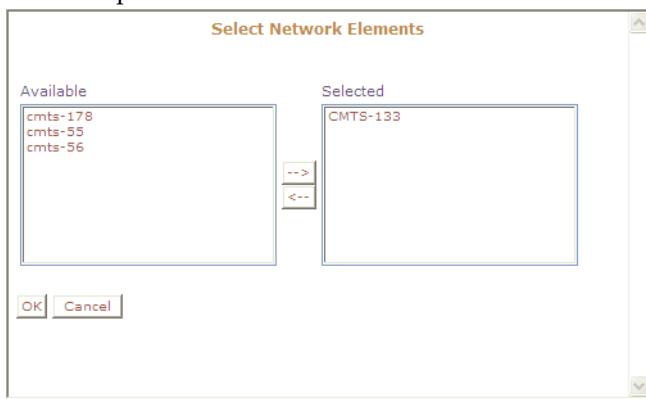


Figure 18: Select Network Elements

6. Select the network elements in the **Available** list and click **-->**.

To disassociate a network element from the MPE device, select the network element from the **Selected** list and click <--. To select entries, press the Ctrl or Shift key and select the entries.

7. When you finish, 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. When you finish, 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

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 be up to 250 characters long and must not contain quotation marks ("") or commas (,).
5. Enter a text **Description/Location** of the network group.
6. When you finish, click **Save**.

You have created a network element group.

Adding a Network Element to a Network Element Group

Once a network element group is created, you can add individual network elements to it. To add a network element to a network element group:

1. From the 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. ([Figure 19: Add Network Element Page](#) shows an example.)
- If there are more than 25 network elements defined, the page does not display any of them. 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. 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. When you finish, click **Save**.

The network element is added to the selected group, and a message indicates the change.

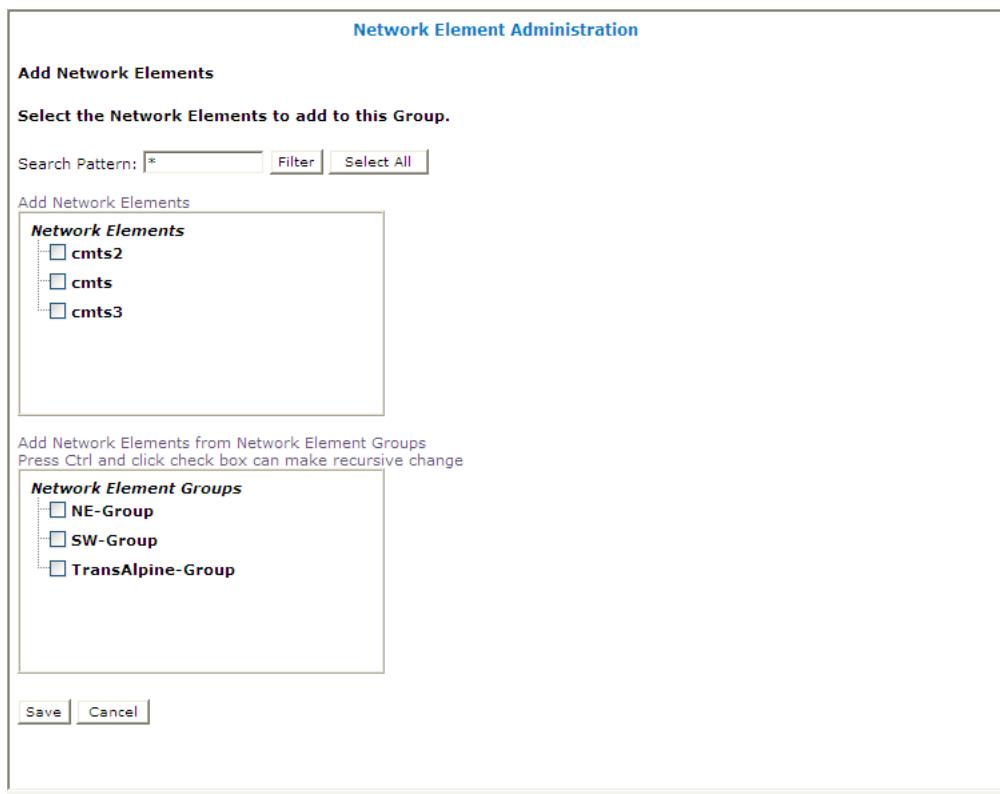


Figure 19: Add Network Element Page

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 cannot contain quotation marks ("") or commas (,).
5. Enter a text description of the sub-group.
6. When you finish, 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 of 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

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 or description.
5. When you finish, click **Save**.

The 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.

Chapter

7

Managing Record Keeping Servers

Topics:

- *About Record Keeping Servers.....112*
- *Creating a Record Keeping Server Profile.....112*
- *Modifying a Record Keeping Server Profile.....113*
- *Deleting a Record Keeping Server Profile.....113*

Managing Record Keeping Servers 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 is a repository for PacketCable event messages. It gathers billing event messages and passes them on to back-office support systems. To use event messaging, you must configure profiles for one or more Record Keeping Servers, 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 Record Keeping Server, note that a single Record Keeping Server can correspond to a single external server, but it can also correspond to a pair of external servers. This depends on how the Record Keeping Server 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 Record Keeping Server. 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 Record Keeping Server (that acts as a backup pair) or you could configure them as two separate Record Keeping Servers, each with a primary address/port and no secondary address/port. However, if a Record Keeping Server does not have a secondary address/port, then that Record Keeping Server will not be able to participate in the Record Keeping Server 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. When you finish, 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. When you finish, 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.

Chapter 8

Managing Event Messaging

Topics:

- *About Event Messaging.....115*
- *Configuring Global Settings for Event Messaging.....116*
- *Configuring Local Settings for Event Messaging.....117*

Managing Event Messaging 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 20: Modify Event Messaging Page* 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. When you finish, click **Save**.

The global event message settings are defined.

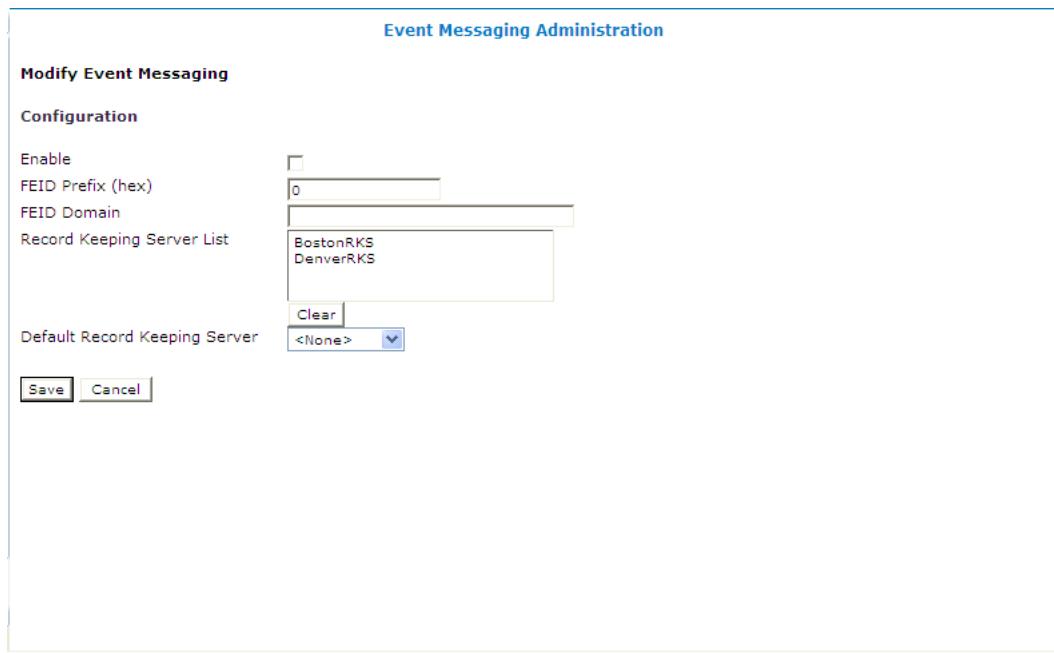


Figure 20: Modify Event Messaging Page

Configuring Local Settings for Event Messaging

The CMP system lets you configure how event messages are handled for a specific MPE device. Local event messaging settings override global settings.

To configure local 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 are displayed.
4. Click **Modify**.
The **Modify Event Messaging** page opens.
5. Configure the attributes as follows. Select **Overrides** to configure a value only for this MPE device.
 - a) **Element ID** — This attribute is set for each MPE device. The Element ID identifies event messages sent from this MPE device.
Type 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.
 - b) **Enable** — Indicates whether event messaging is enabled.
If this value is set to **Yes**, event messages can be sent from the MPE device. If this value is set to **No**, event messages are not sent.

- c) **FEID Prefix (hex)** — 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.
- d) **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 MSO, which uniquely identifies the operator for billing and settlement purposes. This domain name is limited to 239 characters.
- e) **Record Keeping Server 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.
- f) **Default Record Keeping Server** — Defines the default Record Keeping Server for event messaging.

6. When you finish, click **Save**.

Local settings are defined for this MPE device.

Chapter 9

Managing Management Agent Servers

Topics:

- *About Management Agent Servers.....120*
- *Management Agent Tasks.....121*

Managing Management Agent Servers 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. When you finish, 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. When you finish, 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 [About Scheduled Tasks](#).

Viewing Task Status on a Management Agent

To view the status and the current execution for scheduled 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.
Detailed information is displayed.

The following options are available on this page:

- **Reschedule** — Reschedules when the task process starts:
 - Click  (calendar icon), select the date and time, and then click **Enter**.
 - Define the run interval. Valid values are from 0 to 24 hours and 0 to 55 minutes (in 5-minute increments).
 - Define the task, if any, that this task follows.
 - When you finish, click **Save** to save the information to the Management Agent.
- **Run Now** — Runs the task process immediately.
- **Disable or Enable** — Disables or enables this feature.

- **Refresh** — Refreshes the current page.
- **Cancel** — Ignores any information added and closes this page.

Chapter 10

Managing Bandwidth on Demand

Topics:

- *Creating a BoD Server.....125*
- *Creating a BoD Server Group.....125*
- *Viewing and Modifying BoD Server Topology Information.....126*
- *Viewing and Modifying BoD Server Settings..128*
- *Managing PCMM Services.....132*

Managing Bandwidth on Demand 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.

Creating a BoD Server

To create a BoD AM server:

1. From the **BoD** section of the navigation pane, select **Configuration**.
The content tree displays a list of the Bandwidth on Demand servers.
2. Select the **ALL** folder.
The **BoD Administration** page opens in the work area.

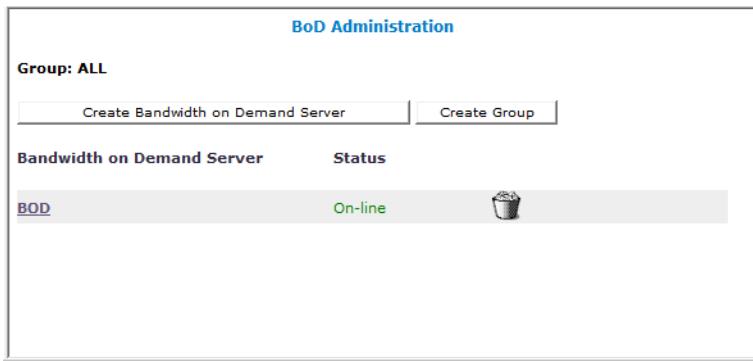


Figure 21: BoD Administration Page for the ALL Group

3. Click **Create Bandwidth on Demand Server**.

The **New BoD** page opens in the work area.

Edit the fields to create the server:

- **Associated Cluster** — Select the name of the cluster where the BoD server associates.
- **Name** — The name of the BoD server.
- **Description/Location** — Descriptive text helping to identify the BoD server.
- **Secure Connection** — Select to require a secure connection for the BoD server.

4. When you finish, click **Save**.

The BoD server is created and added to the list in the content tree.

Creating a BoD Server Group

You can create groups for BoD servers. Updates made to a group also apply to all servers in a group.

To create a BoD server group:

1. From the **BoD** section of the navigation pane, select **Configuration**.
The content tree displays a list of the BoD servers.
2. Select the **ALL** folder.
The **BoD Administration** page opens in the work area.
3. Click **Create Group**.

The **Create Group** page opens in the work area.

4. Enter the name of the group in the **Name** field.
5. When you finish, click **Save**.

A BoD server group is created.

Viewing and Modifying BoD Server Topology Information

The **System** tab on the **BoD Administration** page allows you to view and modify BoD server information. You can also use this tab to delete BoD servers from the CMP system and re-apply a configuration to a server.

Viewing BoD Topology Information

To view the server topology information:

1. From the **BoD** section of the navigation pane, select **Configuration**.
The content tree displays a list of the servers.
2. From the content tree, select a server.
The **BoD Administration** page opens in the work area and displays topology information for the selected server.

Note: The **BoD Administration** page automatically opens to the **System** tab.

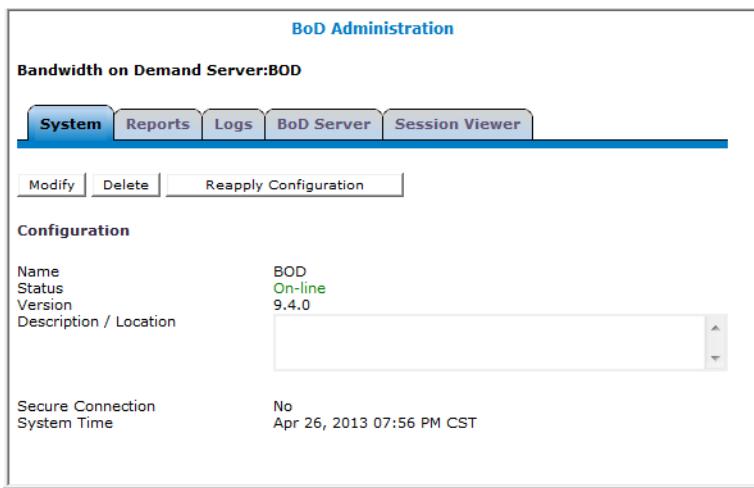


Figure 22: BoD Administration Page

The following settings are displayed:

- **Name** — The name of the server.
- **Status** — The status of the server.
- **Version** — The version of the software running on the server.
- **Description/Location** — Descriptive text for the server.

- **Secure Connection** — Whether the server has a secure connection.
- **System Time** — The current date and time.

Modifying BoD Topology Information

To modify BoD server topology information:

1. From the **BoD** section of the navigation pane, select **Configuration**.
The content tree displays a list of the BoD servers.
2. Select a BoD server.
The **BoD Administration** page opens in the work area and displays topology information about the server.
3. Click **Modify**.
The **Modify System Settings** page opens in the work area.
4. Update the information. For a detailed description of each setting, refer to [Creating a BoD Server](#).
5. When you finish, click **Save**.

Configuring a Server with Multiple Destination IP Addresses

To configure the BoD server topology information to send notifications to multiple pre-configured servers:

1. From the **BoD** section of the navigation pane, select **Configuration**.
The content tree displays a list of the BoD servers.
2. From the content tree, select a BoD server.
The **BoD Administration** page opens in the work area and displays topology information about the server.
3. Click **Modify**.
The **Modify System Settings** page opens in the work area.
4. Select the server settings. See [Viewing BoD Server Settings](#) for details.
5. Select the HTTP method. See [Viewing BoD Server Settings](#) for details.
6. Configure the server information. Refer to [Viewing BoD Server Settings](#) for details.
7. Click **Add**.
The information appears in the list field.
8. Repeat steps *Step 4* through *Step 7* to add more addresses.
9. When you finish, click **Save**.

Deleting a BoD Server

To delete a BoD server:

1. From the **BoD** section of the navigation pane, select **Configuration**.
The content tree displays a list of the BoD servers.
2. Remove the BoD server 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. A confirmation message appears

3. Click **Ok** to delete the server.

The BoD server is deleted.

Reapplying the Configuration to a BoD Server

Note: Reapplying the configuration pushes the settings on the CMP system to the selected BoD server and overwrites the current settings stored on that server.



Caution: Reapplying the configuration pushes the settings on the CMP system to the selected BoD server and overwrites the current settings stored on that server.

To reapply the configuration to a BoD server:

1. From the **BoD** section of the navigation pane, select **Configuration**.
The content tree displays a list of the BoD servers.
2. Select a BoD server.
The **BoD Administration** page opens in the work area and displays information about the selected server.
3. Click **Reapply Configuration**.
The configuration is reapplied.

Viewing and Modifying BoD Server Settings

The **BoD Server** tab on the BoD Administration page is used to configure customer-specific settings for the BoD server, such as PacketCable Multimedia (PCMM) settings and session status notification.

This information is used globally in all future BoD session activity.

Viewing BoD Server Settings

To view BoD server settings:

1. From the **BoD** section of the navigation pane, select **Configuration**.
The content tree displays a list of the BoD servers.
2. Select a BoD server.
The **BoD Administration** page opens in the work area.
3. Select the **BoD Server** tab.
The **BoD Administration** page displays information for the selected BoD server.
The configuration information includes the following:

- **MAC Translation URL** — The address of the server that provides the IP Address mapping for a request containing a MAC address. To support a design where the MAC address is delivered to the BoD application and a separate system provides the mapping of the MAC address to the corresponding IP address, configurable settings are available to look up these translations on a per-session basis. If this method is used there is a latency introduced for the lookup.
- **Translate MAC to IP in recreate** — If this parameter is enabled, then an additional MAC to IP translation is performed when the BoD server receives a gate report or learns that a gate has closed before the BoD server attempts to re-establish the gates. If the parameter is disabled, then the BoD server attempts to re-establish the gates utilizing the IP address received from the original MAC translation operation.
- **Session Status Notification** — This section of the **BoD Administration** page is used to configure whether the BoD server will notify a third-party calling application of events that change the session status.
 - **Notification server IP or FQDN** — The address of the BoD server that is used to notify a third-party calling application of events that change the session status.
 - **Server listening port** — The number of the port that is used to accept a specific request.

Note: This option is not visible when Notification Strategy is disabled.
 - **Server pathname** — HTTP URL for the global or third-party calling application server.

Note: This option is not visible when Notification Strategy is disabled.
 - **Notification strategy** — Determines whether the BoD server will notify a third-party calling application through HTTP requests of events that change the session status. You can select one of the following:
 - **Disabled** — Notification of session modification does not occur.
 - **Global server** — A global notification server is used to accept all HTTP requests. If this option is selected, then values must be entered in the **Notification server IP or FQDN**, **Server pathname** and **Server listening port** fields.

Note: Use this option if multiple destination IP addresses are used for notifications.
 - **Calling Application Server** — Each notification is sent to the application server that created the session. If this option is selected, then a value between 0 - 65535 must be entered in the **Server listening port** field; and the value entered in the **Server pathname** field must be able to be resolved to a valid HTTP URL when combined with a valid server, port and schema.

Note: Use this option if multiple destination IP addresses are used for notifications.

 - **HTTP GET** — The BoD sends the parameters through the URL using this format: `http://NOTIF_IP/PATH?ssid=SESS_ID&reqIP=REQ_IP&state=S&reason=R&subcode=C`
 - **HTTP POST** — the BoD sends the parameters in the body of HTTP message using this format: `ssid=SESS_ID&reqIP=REQ_IP&state=S&reason=R&subcode=C`
- **PCMM** — This section of the **BoD Administration** page is used to configure PacketCable Multimedia (PCMM) settings.
 - **PCMM Enabled** — Defines whether PacketCable Multimedia is enabled.
 - **PCMM Policy Server Configuration** — This section is used to add one or more policy servers (internal or external), allowing the BoD server to send out PCMM requests to an MPE device.

- **PCMM Application Manager Id** — The identifier for a BoD server. This value is an identifier within the protocol, and serves as a label for business rule enforcement at the policy server. It is also possible, within the specific service definitions, to override this value on a per service basis.
- **PCMM Gate Delete Retry Interval (seconds)** — The period of time to wait before attempting another gate delete.
- **Maximum gate delete retry (times)** — The maximum number of times a gate delete will be attempted.

Modifying BoD Server Information

To modify BoD server information:

1. From the **BoD** section of the navigation pane, select **Configuration**.
The content tree displays a list of the BoD servers.
2. From the content tree, select a BoD server.
The **BoD Administration** page opens in the work area.
3. Select the **BoD Server** tab.
The **BoD Administration** page displays the server information.
4. Click **Modify**.
The **BoD Administration** page displays options to allow you to modify the BoD server information.
5. Enter the information. For a detailed description of each field, refer to [Viewing BoD Server Settings](#).
6. Use the PCMM area to add a new MPE device or clone/edit/delete the settings of an existing MPE device.
7. (Optional) Add, modify, or delete settings.
 - Cloning an entry in the table
 1. Select an entry in the table.
 2. Click  **Clone**. The **Clone** window opens with the information for the entry.
 3. Make changes as required.
 4. When you finish, click **Save**. The entry is added to the table
 - Editing an entry in the table
 1. Select the entry in the table.
 2. Click  **Edit**. The **Edit Response** window opens, displaying the information for the entry.
 3. Make changes as required.
 4. When you finish, click **Save**. The entry is updated in the table.
 - Deleting a value from the table
 1. Select the entry in the table.
 2. Click  **Delete**. A confirmation message displays.
 3. Click **Delete** to remove the entry. The entry is removed from the table.
8. When you finish, from the **BoD Administration** page, click **Save**.

Configuring BoD Server Configuration Key Information

To view or modify BoD server configuration key information:

1. From the **BoD** section of the navigation pane, select **Configuration**.
The content tree displays a list of the BoD servers.
2. From the content tree, select a BoD server.
The **BoD Administration** page opens in the work area.
3. Select the **BoD Server** tab.
The **BoD Administration** page displays the server information.
4. Click **Advanced**.

The **Other Advanced Configuration Settings** page opens.

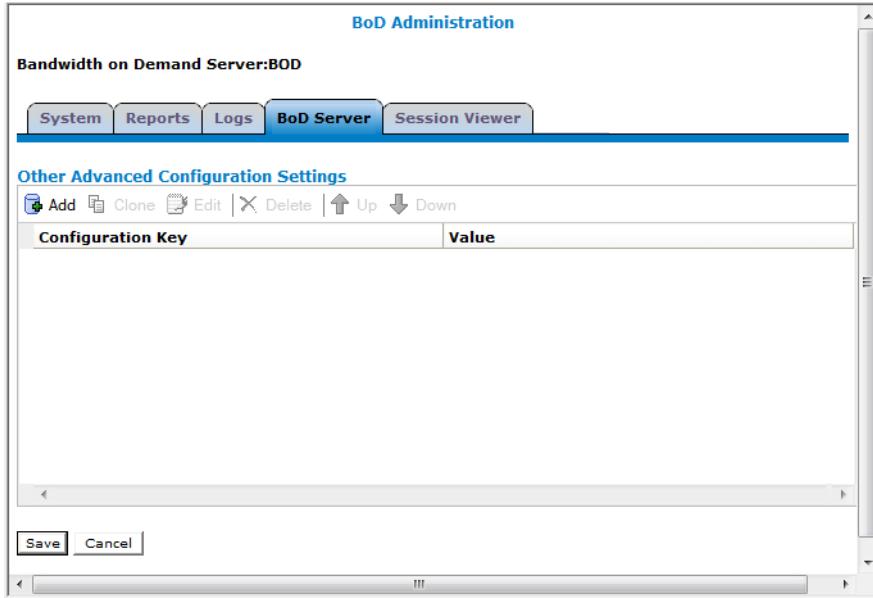


Figure 23: Other Advanced Configuration Settings Page

5. Click **Add** to add a configuration key.
The **Add Configuration Key Value** pop-up appears.
6. When you finish, click **Save**.
7. (Optional) Add, modify, or delete keys.
 - Cloning an entry in the table
 1. Select an entry in the table.
 2. Click **Clone**. The **Clone** window opens with the information for the entry.
 3. Make changes as required.
 4. When you finish, click **Save**. The entry is added to the table
 - Editing an entry in the table
 1. Select the entry in the table.

2. Click  **Edit**. The **Edit Response** window opens, displaying the information for the entry.
3. Make changes as required.
4. When you finish, click **Save**. The entry is updated in the table.

- Deleting a value from the table
 1. Select the entry in the table.
 2. Click  **Delete**. A confirmation message displays.
 3. Click **Delete** to remove the entry. The entry is removed from the table.

8. When you finish, from the **BoD Administration** page, click **Save**.

Managing PCMM Services

The PCMM Services management features allow you to view, create, modify, and delete services.

Note: The services displayed in the examples in this chapter are preconfigured. The services shown in your application may vary.

Viewing Services

The **Services Administration** page displays the names and information for the PCMM services.

From this page, you can modify the parameters contained in a service, copy the parameters to a new service, delete the service from the BoD application, or deactivate the service.

To view currently configured services:

1. From the **BoD** section of the navigation pane, select **Services**.

The **Services Administration** page opens in the work area.

Services Administration				
Policy Server 1/		PCMM	Create Service	Delete All Services
PCMM Services				
Service Name	Upstream Profile Type	Downstream Profile Type	Status	
 audioFlowspec	PCMM Flow Spec		Active	
 beplus	PCMM Best Effort	PCMM Down Stream	Active	
 docsisClassNameService	PCMM DOCSIS		Active	
 mediaDownloadService		PCMM Down Stream	Active	
 mediaUploadService	PCMM Best Effort	PCMM Down Stream	Active	
 netcamService	PCMM Best Effort	PCMM Down Stream	Active	
 realservice	PCMM Best Effort	PCMM Down Stream	Active	
 sandvineService	PCMM Best Effort	PCMM Down Stream	Active	
 snapservice	PCMM Best Effort	PCMM Down Stream	Active	
 snapservice2	PCMM Best Effort	PCMM Down Stream	Active	
 testbeService	PCMM Best Effort		Active	
 turboService	PCMM Best Effort	PCMM Down Stream	Active	
 voiceServiceUnsolGrant	PCMM Unsolicited Grant	PCMM Down Stream	Active	

Figure 24: Services Administration Page

The Services Administration page displays the following fields:

- **Service Name** — User-supplied unique identifiable name for the service.
- **Upstream Profile Type** — Traffic profile associated with a gate through a FlowSpec, DOCSIS Service Class Name, or a DOCSIS specific parameterization scheme.
- **Downstream Profile Type** — Traffic profile associated with a gate through a FlowSpec, DOCSIS Service Class Name, or a DOCSIS specific parameterization scheme.
- **Status** — Status of the service.

2. Click on a service name to display additional properties associated with the service.

The **View Service** page opens in the work area.

Note: The fields shown on the **View Service** page depend on the values selected for the **Classifier Type** and **Scheduling Type** fields. The fields in your service may differ from this example. See [Table 22: Service Field Definitions and Ranges](#) for a description of each field.

Figure 25: View Service Page

View Service

Modify **Copy** **Delete** **Cancel** **Deactivate**

State: Active

Service Name	36431
Duration (minutes)	0
Volume Limit (kilobytes)	0
AM Identifier Override	123
T3 Timer (minutes)	60
Gate Info Interval	1 hour(s) 0 min(s)
Classifier Type	Extended Classifier

Downstream Profile

Scheduling Type	PCMM Down Stream
Traffic Priority	1
Max Sustained Traffic Rate (bps)	80000
Max Traffic Burst	80000
Min Reserved Traffic Rate (bps)	60000
Assumed Min Packet Size (bytes)	255
Max Downstream Latency	60
Downstream Peak Traffic Rate	80000
Required Attribute Mask	0
Forbidden Attribute Mask	0
Attribute Aggregation Rule Mask	2048000
Downstream Resequencing	1
Minimum Buffer	0
Target Buffer	0
Maximum Buffer	0

GateSpec

DSCP/TOS Overwrite Enabled	No
DSCP/TOS Overwrite	
DSCP/TOS Mask	

Classifier

Protocol ID	17
DSCP Tos Field	64
DSCP Tos Mask	64
Source IP / Mask	10.15.0.0 / 16
Source Port Range	80 - 8080
Destination IP / Mask	10.0.1.0 / 24
Destination Port Range	80 - 8080
Priority	1

Creating a Service

You can create a service using the **Create Service** button or by copying an existing service.

In addition, you can create (or modify) a PCMMService that contains multiple Classifiers at the same time. A PCMMService can contain multiple Classifiers which means that you can create sessions with multiple standard /extended/ipv6 Classifiers.

Note: Only extended/ipv6 types can be mixed.

You can also configure Services with multiple Classifiers in each direction (Upstream or Downstream) and add or remove Classifiers from Services.

Creating a Service

To create a service:

1. From the **BoD** section of the navigation pane, select **Services**.

The **Services Administration** page opens in the work area.

2. Click **Create Service**.

The **Create Service** page opens in the work area.

Figure 26: Create Service Page

3. Enter a **Service Name**.
4. Configure the service. See [Table 22: Service Field Definitions and Ranges](#)
5. Configure an **Upstream Profile**. See [Adding Upstream and Downstream Profiles](#).

6. Configure the **GateSpecs**. See [Adding Gate Specifications](#).
7. Add a **Classifier**. See [Adding Classifiers](#).
8. Create the **Upstream Profile**. See [Table 22: Service Field Definitions and Ranges](#).
9. When you finish, click **Save**.

Adding Upstream and Downstream Profiles

To create upstream and downstream profiles for a new service:

Note: The parameters stated here are for both upstream and downstream profiles.

1. From the Create Service screen scroll to the **Upstream Profile** section.

Figure 27: Upstream Profile

Note: For information on entering values, see [Table 22: Service Field Definitions and Ranges](#).

2. Select the **Scheduling Type** from the pull-down menu.
3. Type in a **Traffic Priority**.
4. Set the **Request Transmission Policy**.
5. Set the **Max Sustained Traffic Rate**.
6. Set the **Max Traffic Burst**.
7. Set the **Min Reserved Traffic Rate**.
8. Set the **Assumed Min Packet Size**.
9. Associate the stream with a **Classifier**.

Adding Gate Specifications

To set the gate specifications for a service:

1. From the **Create Service** screen, scroll down to the **GateSpec** section.

Note: The parameters for the gate specifications are the same for both upstream and downstream profiles.

For specific information on each field for a gate, see [Table 22: Service Field Definitions and Ranges](#).

Figure 28: GateSpec Section

2. (Optional) Select if the **DSCP/TOS Overwrite Enabled** is active or not.
3. Enter the **DSCP/TOS Overwrite** period.
4. Enter the **DSCP/TOS Mask** period.

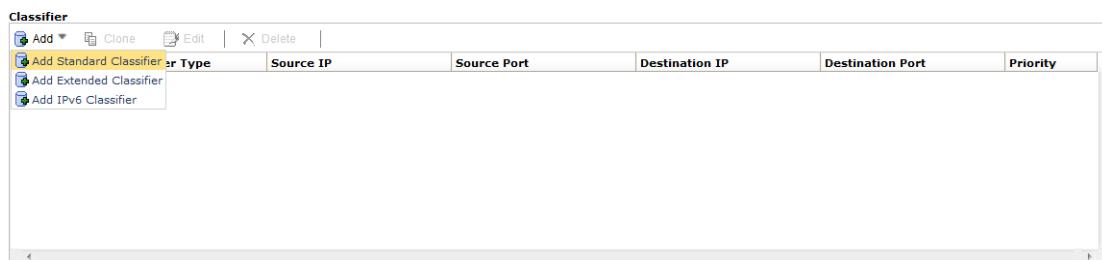
- Associate the **GateSpec** with a classifier(s).

Adding Classifiers

To add one or more classifiers:

- From the **Create Service** page click the **Add** button on the Classifier tool bar. The **Classifier** drop-down menu opens.

Note: This process is for adding both upstream and downstream classifiers.

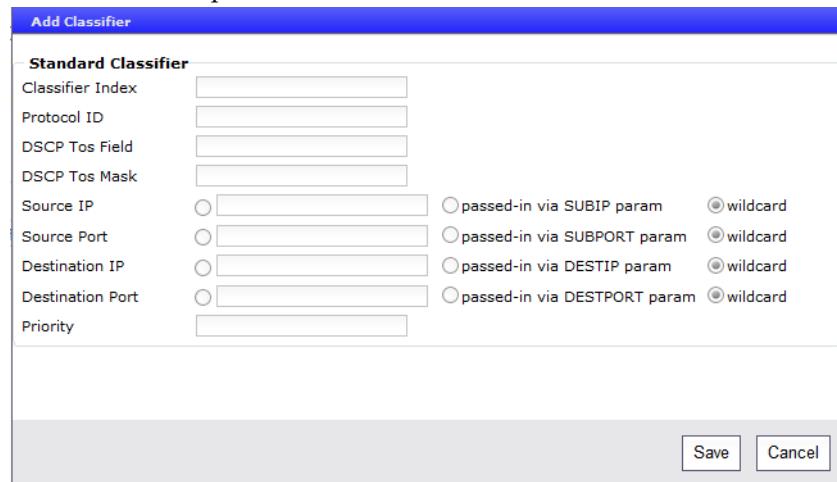


Classifier	Add	Clone	Edit	Delete	Source IP	Source Port	Destination IP	Destination Port	Priority
	Add Standard Classifier								
	Add Extended Classifier								
	Add IPv6 Classifier								

Figure 29: Classifier Table with Drop-down Menu

- Select the **Classifier**.

The **Add Classifier** window opens.



Add Classifier

Standard Classifier

Classifier Index:

Protocol ID:

DSCP Tos Field:

DSCP Tos Mask:

Source IP: passed-in via SUBIP param wildcard

Source Port: passed-in via SUBPORT param wildcard

Destination IP: passed-in via DESTIP param wildcard

Destination Port: passed-in via DESTPORT param wildcard

Priority:

Save **Cancel**

Figure 30: Add Classifier Screen (Standard Classifier selected)

- Type in the appropriate values in each of the fields, See [Table 22: Service Field Definitions and Ranges](#).
- Click the **Save** button to save the settings.

You are returned to the **Create Service** screen.

- Repeat steps [Step 1](#) through [Step 4](#) to add additional classifiers.

Note: For classifier compatibility. See [Compatibility Matrix for Creating Multiple Classifiers](#).

Correct Parameter and Classifier Order

The parameter/classifier order should match the guidelines for the PCMMService. For example, if two standard Classifiers are pre-defined in the PCMMService profile illustrated in this example:

```
http://bodurl/bod/createsession.do?Servicename=123&Subip=x1.x1.x1.x1|x2.x2.x2.x2&Destip=d1.d1.d1.d1|d2.d2.d2.d2&subport=subport1|subport2&destport=destport1|destport2
```

Then the correct parameter order should follow the Classifier order.

Table 20: Correct Parameter - Classifier Order

	Classifier index1	Classifier index2
Subip	x1.x1.x1.x1	x 2 . x 2 . x 2 . x 2
Destip	d1 . d1 . d1 . d1	d2 . d2 . d2 . d2
Subport	subport1	subport2
Destport	destport1	destport2

Stated in a different way:

The parameters from Subip_list[0], Destip[0], subport[0] and destport[0] belong to the Classifier1.

The parameters from Subip_list[1], Destip[1], subport[1] and destport[1] belong to the Classifier2.

Compatibility Matrix for Creating Multiple Classifiers

This table provides the compatibility rules for creating session interfaces with multiple classifiers to the PCMMService:

Note: Parameter values used for multiple Classifiers are formatted and interpreted similarly to the single classifier createSession.do. However, parameters related to the Classifier(s) for the service will include multiple values, separated by a pipe (" | ") character, and ordered corresponding to the Index values in the Service Definition. If a given parameter is not required for one Classifier in a multiple-Classifier Service, then the corresponding parameter should be omitted but the pipe character is still required.

Table 21: Multiple Classifier Compatibility Matrix

Classifier Type	Classifier Type	Classifier Compatibility
Standard	Standard	Yes
Standard	Extended	No
Standard	Ipv6	No
Extended	Extended	Yes
Ipv6	Ipv6	Yes
Extended	Ipv6	Yes

Creating a Service from an Existing Service

To create a service from an existing service:

1. From the **BoD** section of the navigation pane, select **Services**.
The **Service Administration** page opens in the work area.
2. Select a service.
The **View Service** page opens in the work area.
3. Click **Copy**.
The **Copy Service** page opens in the work area.

Figure 31: Copy Service Page

4. Enter a new service name for the service, and edit the fields. For a description of each field and its associated values, see the [Table 22: Service Field Definitions and Ranges](#).
5. When you finish, click **Save**.

Modifying a Service

To modify a service:

1. From the **BoD** section of the navigation page, select **Services**.
The **Services Administration** page opens in the work area.
2. Select a service.
The **View Service** page opens in the work area.

3. Click **Modify**.

The **Edit Service** page opens in the work area.

4. Edit the fields.

See [Table 22: Service Field Definitions and Ranges](#) for descriptions of the fields used to edit a PCMM service. The fields available depend on the values selected for the **Classifier Type** and **Scheduling Type** fields.

Table 22: Service Field Definitions and Ranges

Field	Description	Valid Value
Service Name	A unique name for the PCMM service.	A string of up to 32 characters.
Duration (minutes)	The duration this service is deployed. Enter a value or select passed-in via DUR param or indefinite .	Integer (0 - 35791394)
Volume Limit (kilobytes)	The volume limit. Enter a value or select passed-in via VOLLIMIT param or indefinite .	Long (0 - 9223372036854775807)
AM Identifier Override	An override for the BoD AM ID. This value overrides the value set in the Configure Settings AM Identifier field.	Long (0 - 4294967295)
T3 Timer (minutes)	The T3 inactivity timer value.	Integer (1 - 1092.25)
Gate Info Interval	The gate information interval, in hours and minutes.	4-byte unsigned integer
Classifier Type	Classifier type. Depending on what you select here, other fields may become available on the page.	<ul style="list-style-type: none"> • Standard • Extended Classifier • IPv6 Classifier • Mixed Classifier
Upstream Profile		
Scheduling Type	The traffic profile type. Depending on what you select here, other fields may become available on the page.	<ul style="list-style-type: none"> • NA • PCMM Best Effort • PCMM Real-Time Polling • PCMM Non-Real-Time Polling • PCMM Unsolicited Grant • PCMM Unsolicited Grant with Activity Detection • PCMM Upstream Drop • PCMM Flow Spec • PCMM DOCSIS

Field	Description	Valid Value
Traffic Priority	The relative priority assigned to the service flow in comparison with other flows.	1 byte (0 - 255)
Request Transmission Policy	Specifies which IUC opportunities the CM uses for upstream transmission requests and packet transmissions for this service flow.	4-byte (0 - 4294967295)
Max Sustained Traffic Rate (bps)	The rate parameter for a token-bucket-based rate limit for this service flow. Enter a value or select passed-in via UPBWMAX param.	4-byte (0 - 4294967295)
Max Traffic Burst	The token bucket size, in bytes, for a token-bucket-based rate limit for this service flow.	4-byte (0 - 4294967295)
Min Reserved Traffic Rate (bps)	The minimum rate reserved for this service flow. Enter a value or select passed-in via UPBW param.	4-byte (0 - 4294967295)
Assumed Min Packet Size (bytes)	The assumed minimum packet size for which the Minimum Reserved Traffic rate is provided for this service flow.	2 bytes (0 - 65535). Enter 0 if a specific Assumed Minimum Reserved Traffic Rate Packet size is not required. Upon receipt of a value of 0, the CMTS must utilize its implementation-specific default size for this parameter, not 0 bytes.
Maximum Concatenated Burst	The maximum concatenated burst, in bytes, that a service flow is allowed.	2-byte (0 - 65535)
Service Class Name	The DOCSIS Service Class to be used to describe QoS attributes.	32 characters
Service Number	The service number. A controlled load service must contain only the TSpec token bucket parameters, and not the RSpec. A guaranteed service must contain both the TSpec and the RSpec.	Short (0 - 255) 5 - controlled load 2 - guaranteed

Field	Description	Valid Value
Token Bucket Rate (bytes/sec)	Defines how traffic is injected into the network by the sending application.	Float (0.0 - 3.4028234663852886E38)
Token Bucket Size (bytes)	Controls the maximum amount of data that the flow can send at the peak rate.	Float (0.0 - 3.4028234663852886E38)
Peak Data Rate (bytes/sec)	The peak data rate.	Float (0.0 - 3.4028234663852886E38)
Minimum Policed Unit (bytes)	The minimum size of a packet that can be subject to policing.	Long (0 - 2147483647)
Maximum Policed Size (bytes)	The maximum size of a burst of data that can exceed the given bandwidth limit.	Long (0 - 2147483647)
Rate (bytes/sec)	The rate.	Float (0.0 - 3.4028234663852886E38)
Slack Term (microsec)	The slack term, corresponding to latency or jitter depending on the service.	Long (0 - 2147483647)
Envelope	The envelope types (i.e. Authorized, Reserved, and Committed) that are present in the object.	1 byte (0 - 255) A value of 1 indicates that the envelope type is present in the Traffic Profile.
Unsolicited Grant Size (bytes)	The grant size.	2-byte (0 - 65535) There is no default value.
Grants Per Interval	The number of grants per Nominal Grant Interval.	1 byte (0 - 255) There is no default value. A value of 1 is recommended.
Nominal Grant Interval	The nominal time, in microseconds, between successive data grant opportunities for a service flow.	4-byte (0 - 4294967295)
Tolerated Grant Jitter (microsec)	The maximum amount of time that transmission opportunities can be delayed from the nominal periodic schedule.	4-byte (0 - 4294967295)
Nominal Polling Interval (microsec)	The nominal interval between successive unicast request opportunities for this service flow on the upstream channel.	4-byte (0 - 4294967295)

Field	Description	Valid Value
Tolerated Poll Jitter (microsec)	The maximum amount of time that a polling request can be delayed.	Long (0 - 4294967295)
Required Attribute Mask	Limits the set of channels. Enter a value or select passed-in via UPRAMASK param.	4-byte (0 - 4294967295)
Forbidden Attribute Mask	Limits the set of channels and bonding groups to which the CMTS assigns the service flow by forbidding certain attributes. Enter a value or select passed-in via UPFAMASK param.	4-byte (0 - 4294967295)
Attribute Aggregation Rule Mask	Guides the CMTS on how it can use the attribute masks of individual channels to construct a dynamic bonding group for this service flow. Enter a value or select passed-in via UPAAMASK param.	4-byte (0 - 4294967295)
Upstream Peak Data Rate	The peak traffic rate, in bits per second, that is allowed for a service flow. Enter a value or select passed-in via UP PTR param.	4-byte (0 - 4294967295)
Minimum Buffer	The lower limit for the size of the buffer to be provided for a service flow.	4-byte (0 - 4294967295)
Target Buffer	The desired value for the size of the buffer to be provided for a service flow.	4-byte (0 - 4294967295)
Maximum Buffer	The upper limit for the size of the buffer to be provided for a service flow.	4-byte (0 - 4294967295)
GateSpec		
DSCP/TOS Overwrite Enabled	Enables the DSCP/TOS Overwrite functionality. If this field is set, then the CMTS must mark the packets traversing the CMTS DSCP/TOS value. If the field is cleared, then the CMTS must not perform any marking.	Enabled/Disabled
DSCP/TOS Overwrite	Used to overwrite the DSCP/TOS field of packets	1 byte (0 - 255)

Field	Description	Valid Value
	associated with the DOCSIS Service Flow that corresponds to the Gate	
DSCP/TOS Mask	The bit mask used to identify particular bits within the DSCP/TOS field	1 byte (0 - 255)
Classifier		
Protocol ID	The protocol identifier.	0 - 257
DSCP/Tos Field	The Differentiated Services Code Point (DSCP) or IP Precedence.	1 byte (0 - 255)
DSCP/Tos Mask	The bit mask used to select relevant bits from the accompanying DSCP/Tos field value.	1 byte (0 - 255)
Next Header Type	The desired next header type value for any header or extension header associated with the packet.	A value of 256 matches traffic with any IPv6 next header type value. A value of 257 matches both TCP and UDP traffic.
Traffic Class Mask	A mask defining which of the 8 bits should be used for matching a traffic class value.	1 byte (0 - 255).
Traffic Class Range	Enter a lower and upper value to match a range of traffic class values.	1 byte (0 - 255).
Flow Label	Contains valid data for comparison with the IPv6 Flow Label.	4 bit (0 - 15) This flag must be set to 1 if data is needed. When comparison of the IPv6 Flow Label for this entry is irrelevant then the flag cannot be set (value = 0). When the Flow Label flag is set to 0, the CMTS cannot include the IPV6 Flow Label field in the classifier. All other flags must be set to zero
Source IP (Standard classifier) Source IP / Mask (Extended classifier)	The source IP address used to classify traffic. Enter a value,	4-octet IPv4 address (Standard classifier)

Field	Description	Valid Value
Source IP / Prefix Length (IPv6 classifier)	select that the values are passed in by parameters, or select wildcard .	4-octet IPv4 mask (Extended classifier) 16-octet IPv6 address (IPv6 classifier) Prefix Length: Short (0 - 128)
Source Port (Standard classifier) Source Port Range (Extended and IPv6 classifiers)	The source port range used to classify traffic. Enter a value, select that the values are passed in by parameters, or select wildcard .	2 bytes 0 - 32768 (Standard classifier) 0 - 65535 (Extended classifier)
Destination IP (Standard classifier) Destination IP / Mask (Extended classifier) Destination IP / Prefix Length (IPv6 classifier)	The destination IP address used to classify traffic. Enter a value, select that the values are passed in by parameters, or select wildcard .	4-octet IPv4 address (Standard classifier) 4-octet IPv4 mask (Extended classifier) 16-octet IPv6 address (IPv6 classifier) Prefix Length: Short (0 - 128)
Destination Port (Standard classifier) Destination Port Range (Extended and IPv6 classifiers)	The destination port used to classify traffic. Enter a value, select that the values are passed in by parameters, or select wildcard .	2 bytes 0 - 32768 (Standard classifier) (0 - 65535) Extended classifier
Priority	The priority level for this classifier.	Current DOCSIS supported values are 64 - 191.
Downstream Profile		
Scheduling Type	The downstream profile type. Depending on what you select here, other fields may become available on the page.	<ul style="list-style-type: none"> • NA • PCMM Down Stream • PCMM DOCSIS • PCMM Flow Spec
Traffic Priority	The relative priority assigned to the service flow in comparison with other flows.	1 byte (0 - 255)
Max Sustained Traffic Rate (bps)	The rate parameter for a token-bucket-based rate limit for this service flow. Enter a value or select passed-in via DOWNBWMAX param .	4-byte (0 - 4294967295)
Max Traffic Burst	The token bucket size, in bytes, for a token-bucket-based rate limit for this service flow.	4 bytes (0 - 4294967295)

Field	Description	Valid Value
Downstream Peak Traffic Rate	The rate parameter of a token-bucket-based peak rate limiter for packets of a downstream service flow. Enter a value or select passed-in via DOWNPTR param.	4-byte (0 - 4294967295)
Min Reserved Traffic Rate (bps)	The minimum rate reserved for this service flow. Enter a value or select passed-in via DOWNBW param.	4-byte (0 - 4294967295)
Assumed Min Packet Size (bytes)	The assumed minimum packet size for which the Minimum Reserved Traffic rate is provided for this service flow.	2-byte (0 - 65535). Enter 0 if a specific Assumed Minimum Reserved Traffic Rate Packet size is not required. Upon receipt of a value of 0, the CMTS must utilize its implementation-specific default size for this parameter, not 0 bytes.
Max Downstream Latency	The maximum latency between the receptions of a packet on the CMTS's NSI and the forwarding of the packet on its RF interface.	4 byte (0 - 4294967295)
Service Class Name	The pre-configured service class name associated with a gate.	32 characters
Service Number	The service number. A controlled load service must contain only the TSpec token bucket parameters, and not the RSpec. A guaranteed service must contain both the TSpec and the RSpec.	Short (0 - 255) 5 - controlled load 2 - guaranteed
Token Bucket Rate (bytes/sec)	Defines how traffic is injected into the network by the sending application.	Float (0.0 - 3.4028234663852886E38)
Token Bucket Size (bytes)	Controls the maximum amount of data that the flow can send at the peak rate.	Float (0.0 - 3.4028234663852886E38)
Peak Data Rate (bytes/sec)	The peak data rate.	Float (0.0 - 3.4028234663852886E38)
Minimum Policed Unit (bytes)	The minimum size of a packet that can be subject to policing.	Long (0 - 2147483647)

Field	Description	Valid Value
Maximum Policed Size (bytes)	The maximum size of a burst of data that can exceed the given bandwidth limit.	Long (0 - 2147483647)
Rate (bytes/sec)	The rate.	Float (0.0 - 3.4028234663852886E38)
Slack Term (microsec)	The slack term, corresponding to latency or jitter, depending on the service.	Long (0 - 2147483647)
Required Attribute Mask	Limits the set of channels. Enter a value or select passed-in via DOWNRAMASK param.	4-byte (0 - 4294967295)
Forbidden Attribute Mask	Limits the set of channels. Enter a value or select passed-in via DOWNRAMASK param.	4-byte (0 - 4294967295)
Attribute Aggregation Rule Mask	Guides the CMTS on how it can use the attribute masks of individual channels to construct a dynamic bonding group for this service. Enter a value or select passed-in via DOWNRAMASK param.	4-byte (0 - 4294967295)
Downstream Resequencing	Specifies the use of sequence numbers in downstream DOCSIS 3.0 service flows.	1 byte (0 - 255)
Minimum Buffer	The lower limit for the size of the buffer to be provided for a service flow.	4-byte (0 - 4294967295)
Target Buffer	The desired value for the size of the buffer to be provided for a service flow.	4-byte (0 - 4294967295)
Maximum Buffer	The upper limit for the size of the buffer to be provided for a service flow.	4-byte (0 - 4294967295)

Note: The **passed-in** radio button indicates the field obtains its value from a passed-in HTTP parameter from the create session or add traffic classifier request. The use of the **Wildcard** radio button indicates to accept or match all and on a duration field indicates an indefinite duration. Refer to the PCMM specification for details for each wildcard value.

- When you finish, click **Save**.

Deactivating a Service

Deactivating a service prevents any new requests for this service from establishing QoS. Any existing sessions will not be affected.

To deactivate a service:

1. From the **BoD** section of the navigation pane, select **Services**.

The **Services Administration** page opens in the work area.

2. Click on the desired service.

The **View Service** page opens in the work area.

3. Click **Deactivate**.

The service is deactivated and the **Deactivate** button is replaced by an **Activate** button, which is used to reactivate the service.

Deleting a Service

Deleting a service prevents any new requests for this service from establishing QoS. Any existing sessions will not be affected.

To delete a service:

1. From the **BoD** section of the navigation pane, select **Services**.

The **Services Administration** page opens in the work area.

2. Select a service.

The **View Service** page opens in the work area.

3. Click **Delete**.

A confirmation message appears.

4. Click **OK** to delete the service.

The service is deleted.

Importing a Service

The Import Services page allows an XML file containing services definitions to be imported into the BoD application. By clicking **Browse**, a file chooser dialog box is opened. The operator can then choose a file to import into the BoD application. If the imported file passes validation, all the Service definitions currently defined in the BoD application are replaced with the Service definitions defined in the import file. This feature is typically used in conjunction with the Export Services feature.

Note: Attributes `upRequiredAttrMask`, `upForbiddenAttrMask`, `upAttrAggrRuleMask`, `downRequiredAttrMask`, `downForbiddenAttrMask`, `downAttrAggrRuleMask`, and `downPeakTrafficRate` are expected in a SOAP or HTTP request for gates. If these attributes are not included, the request will fail. Therefore, after importing an XML file from previous versions of the BoD application, ensure that the service has included the new attributes. To configure these attributes in the BoD application, see [Modifying a Service](#).

To import a service:

- From the **BoD** section of the navigation pane, select **Services Import/Export**.

The **Import/Export** page opens in the work area.

Figure 32: Services Import/Export Page

- Enter the filename to import in the **Services Import File Name** field or click **Browse** to locate the desired file.
- To have the file validated before import, click **Validate Service File**.
- After you have entered the desired information, click **Save** to import the file.

Exporting a Service

The Export Services page invokes a secondary browser window containing an XML document that reflects the currently defined service profile definitions configured within the BoD AM application. This content can be saved by choosing the **File > Save As...** menu item in this secondary browser window. This feature can also be used to save snapshots of service definitions, which can be later imported back into the BoD AM application using the Import Services feature.

To export a Service:

- From the **BoD** section of the navigation screen, select **Services Import/Export**.
The Services Import/Export page opens in the work area.
- Click **Export**.
- You are prompted to open or save the file.
- Click **Open**. An XML file appears.
- Select **Save As ...** from the **File** pull-down menu.
- Enter the location to export the file and click **Save**.
The file exports to the entered location.

Chapter 11

System-Wide Reports

Topics:

- *KPI Dashboard.....150*
- *Viewing the Trending Reports.....153*
- *Alarm Reports.....159*
- *Viewing Other Reports.....163*

System-Wide Reports 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 server. 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 periodically refreshed every 10 seconds; this time period is not configurable. The color changing thresholds are user configurable.

Figure 33: Example of KPI Dashboard illustrates the dashboard's contents.

KPI Dashboard (Stats Reset: Manual / Last Refresh:09/09/2014 09:04:07)														Change Thresholds	
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 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 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 of 0	0	0	2	0	0

Figure 33: Example of KPI Dashboard

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.

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 (since 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, see the *Troubleshooting Reference* for this release.

Mapping Display to KPIs

Table 23: KPI Definitions for MPE Devices 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 23: 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
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 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. When you finish, click **Save** to retain your changes.

If you want to revert back to the system defaults, click **Reset**.

The KPI thresholds are configured.

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 34: Trending Report Definition Configuration Page* shows a sample.

Figure 34: 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 contain up to 255 characters, cannot contain double quotes or commas, and cannot begin or end with a space.

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

1. Select an entry in the table.
2. Click  **Clone**. The **Clone** window opens with the information for the entry.
3. Make changes as required.
4. When you finish, click **Save**. The entry is added to the table

- Editing an entry in the table
 1. Select the entry in the table.
 2. Click  **Edit**. The **Edit Response** window opens, displaying the information for the entry.
 3. Make changes as required.
 4. When you finish, click **Save**. The entry is updated in the table.
- Deleting a value from the table
 1. Select the entry in the table.
 2. Click  **Delete**. A confirmation message displays.
 3. Click **Delete** to remove the entry. The entry is removed from the table.

6. When you finish, 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.)

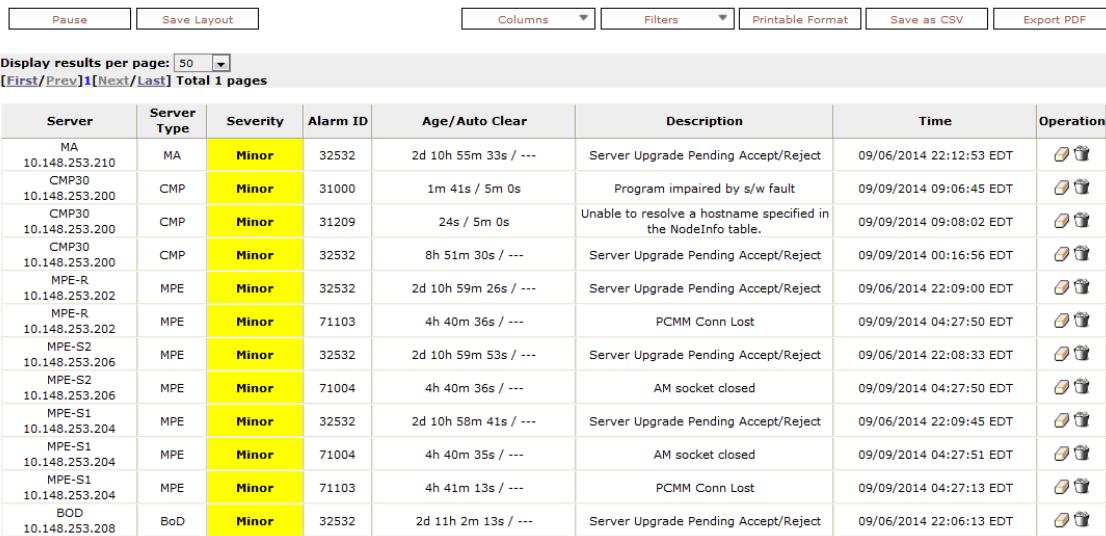
Notifications, which have a severity of Info, are not displayed in the Active Alarms report, but are written to the trace log. For more information, see [The Trace Log](#).

Figure 35: Sample Active Alarms Report 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.



<input type="button" value="Pause"/> <input type="button" value="Save Layout"/> <input type="button" value="Columns"/> <input type="button" value="Filters"/> <input type="button" value="Printable Format"/> <input type="button" value="Save as CSV"/> <input type="button" value="Export PDF"/>							
Display results per page: <input type="button" value="50"/> <input type="button" value="▼"/>							
[First/Prev] <input type="button" value="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	

Figure 35: Sample Active Alarms Report

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 whose alarms you want to view.
- **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. When you finish, click **Close** to close the window.

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 or 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:

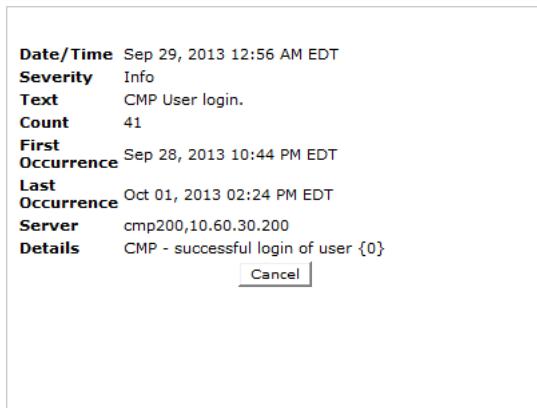


Figure 36: Alert Details

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 Rep Stats

The MPE/BoD replication statistics report provides replication counters for MPE/BoD replication, with one row for each set of servers, letting you gauge the replication of MPE/BoD servers. The display is refreshed every ten seconds. The Description field will display a percentage if the network is slow and the COMCOL UpSyncLog buffer usage exceeds a fixed threshold.

To view the MPE/BoD Replication 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**/**MPE/BoD Rep Stats**.
The MPE/BoD Replication 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** pulldown menu. The following columns are available:

- **Servers** — The replication server group.
- **Cluster Name** — The name of the cluster containing the servers.

- **App Type** — The type of server **MPE** or **BoD**.
- **Total Sent KB** — The total amount of replication data sent during the last five minutes.
- **Peak Sent KB/s** — The peak (highest) amount of data sent during the last five minutes.
- **Avg Sent KB/s** — The average amount of data sent during the last five minutes.
- **Total Recv KB** — The total amount of replication data received during the last five minutes.
- **Peak Recv KB/s** — The peak (highest) amount of data received during the last five minutes.
- **Avg Recv KB/s** — The average amount of data received during the last five minutes.
- **Description** — Describes how the servers are connected.

You can filter results by controlling which table rows appear, using the **Filters** pulldown menu. You can define filtering criteria using the following fields:

- **App Type** — Filter in all types (the default) or one specific type.
- **Server Name** — Filter in all servers (the default) or one specific server.
- **Cluster Name** — Filter in all clusters (the default) or one specific cluster.

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 rep stats 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.

Chapter 12

Upgrade Manager

Topics:

- *About ISO Files on Servers.....169*
- *Preparing for an Upgrade.....172*
- *Performing an Upgrade.....172*
- *About Rolling Back an Upgrade.....176*

The Upgrade Manager allows you to manage upgrade ISOs and perform software upgrades on servers in the topology. During the upgrade process, the **System Maintenance** page displays the upgrade status. Access to these GUI options can be restricted by user role; see *User Management* for more information.

For specific steps on performing an upgrade, contact My Oracle Support. See *My Oracle Support (MOS)* for more information.

About ISO Files on Servers

Policy Management software upgrades are distributed and stored for use as ISO files, which are archive files of optical (DVD) discs.

Use the **ISO Maintenance** option to show the current Policy Management software version executing on servers, and determine what ISO files are available to use for upgrades. Operations performed from here include distributing ISO files to servers, deleting ISO files from servers, and pushing the upgrade script to servers. An audit log is generated for each operation that occurs on this page.

ISO Maintenance Page Elements

On the section of the navigation pane, **ISO Maintenance** is an option. All clusters and their constituent servers in the Policy Management network appear in the table on this page. You can collapse or expand the display of servers by clicking the [-] or [+] icons in the first column of the table. The display is updated every ten seconds.

The following types of elements appear on the **ISO Maintenance** page:

- Check boxes to select clusters or servers on which to perform operations
- The table of filtered clusters and servers
- Lists (**Columns**, **Filters**, and **Operations**) for changing what displays in the table and for selecting operations

Table 24: ISO Maintenance Page Elements describes the elements that appear on the **ISO Maintenance** page.

Table 24: ISO Maintenance Page Elements

Element	Description
<input type="checkbox"/>	(checkbox) Use this column to select the clusters or servers on which an operation is to be performed. If you select a cluster, all servers in that cluster are selected. Note: At least one cluster or server must be selected before you can select an operation from the Operations menu.
Name	Displays the names of all filtered clusters and servers. When a server is receiving an ISO file, a download icon appears next to the name.
Appl Type	Displays the type of application running on each server. The Filters list lets you select the application type: CMP Site1 Cluster , CMP Site2 Cluster , MPE , BoD , MA , or All applications.
Site	Displays the site name, if any, that is associated with each server. The Filters list also lets you display Unspecified or All sites. Note: This column only appears for a georedundant Policy Management network.
IP	Displays the OAM server IP address of each server. The Filters list lets you filter on only a server with a specific IP address or display All servers.

Element	Description
Running Release	Displays the current Policy Management software release of each server. The Filters list lets you filter on only a specific major release only or display All releases.
ISO	Displays the ISO files available on each server. Use the checkbox to select the ISO file to delete during the Delete ISO operation.
Columns	Use the Columns list to change the columns that appear in this table. The Name column is mandatory. By default, all columns appear. To change which columns appear, uncheck the columns to be removed from the page.
Filters	Use the Filters list to select a subset of clusters and servers to appear on this page. On this menu are the following pulldown filter submenus: App Type , IP , and Running Release . By default, the filters are set to All , and all servers appear. Selecting another option from one or more of these filters reduces the number of servers displayed.
Operations	<p>Use the Operations list to select an ISO operation to perform.</p> <p>Note: You must select (in the first column of the table) the cluster(s) or server(s) on which the operation is being performed before you can select an operation. The operations that appear in the list depend on the state of the servers that are selected; that is, if you select more than one server, only the operations that are valid for all selected servers appear.</p> <p>Possible operations are Push Script, Upload ISO, and Delete ISO. As a protective feature, before Push Script or Delete ISO are executed, you are prompted whether you want to execute the operation. If you click OK, the operation is performed. A progress bar displaying the status of the command execution appears in a window.</p> <p>Note: After an operation is confirmed, it cannot be cancelled.</p>

Viewing the ISO Status of Servers

Use this procedure to view the status of in-service servers before, during, and after a software upgrade.

1. From the section of the navigation pane, select **ISO Maintenance**.
The **ISO Maintenance** page appears.
2. (Optional) Click **Filters** and specify the criteria to customize the list of servers that display in the table.
3. (Optional) Click **Columns** and select columns to customize the table.

All in-service servers that meet the filter criteria are listed. Server information is updated every ten seconds.

Pushing a Script to a Server

Before starting this procedure, you must have mounted the ISO file manually and copied the following files to `/opt/camiant/bin` on the CMP system on which you are performing this procedure:

- policyUpgrade.pl
- policyUpgradeHelper.pl
- qpSSHKeyProv.pl
- policySSHKey.pl
- lvm_reclam.pl

Upgrades are controlled by a set of script files. Use this procedure to push upgrade scripts to the remote servers receiving a software upgrade. This procedure is required before a software upgrade can occur on a server.

1. From the section of the navigation pane, select **ISO Maintenance**.
The **ISO Maintenance** page opens.
2. Select the servers receiving the upgrade script.
3. Click the **Operations** list and select **Push Script**.
You are prompted: *Are you sure you want to execute Push Script?*
4. Click **OK**.
A progress bar displays the progress of the operation.

The upgrade scripts are downloaded to the selected servers.

Adding an ISO File to a Server

Before adding an ISO file to a server, you should ensure that the directory `/var/TKLC/upgrade` is empty on that server.

Use this procedure to load an upgrade ISO file onto a remote server for a software upgrade.

1. From the section of the navigation pane, select **ISO Maintenance**.
The **ISO Maintenance** page opens.
2. Select the clusters or servers to receive the ISO file.
3. Click the **Operations** list and select **Upload ISO**.
The **Upload ISO** window opens.
4. Enter the following information for the ISO file (all fields are required):
 - a) **Mode** — Mode used to transfer the ISO file to remote servers. Currently, SCP is available.
 - b) **ISO Server Hostname/IP** — Enter the name or address of the server receiving the ISO file.
 - c) **User** — Enter the root account user name.
 - d) **Password** — Enter the root account password.
 - e) **Source ISO file full path** — Enter the location where the ISO file is to be stored on the remote server.
5. Click **Add** (or **Back** to abandon your request).
The **Upload ISO** window closes, and the transfer process begins to the selected servers. A download icon appears in the **Name** column for the servers receiving the ISO file during the file transfer process. A progress bar displays during the operation. Once the process completes, the icon disappears.

The ISO file is distributed to the servers.

Deleting an ISO File from a Server

Use this procedure to delete an ISO file from a server.

1. From the section of the navigation pane, select **ISO Maintenance**.
The **ISO Maintenance** page opens.
2. Select the clusters or servers.
3. Select the ISO file to be removed.
4. Click the **Operations** list and select **Delete ISO**.
You are prompted: Are you sure you want to execute Delete ISO?
5. Click **OK**.
A progress bar displays the progress of this operation.

The selected ISO files are deleted from the selected clusters or servers.

Preparing for an Upgrade

Upgrading a server requires a large amount of preparation. For detailed information about preparing for an upgrade, please see the My Oracle Support website (<https://support.oracle.com>).



Caution: Contact My Oracle Support and inform them of your upgrade plans prior to beginning this or any upgrade procedure. Before upgrading any system, go to the My Oracle Support website and review any relevant Technical Service Bulletins (TSBs). Use only the upgrade procedure provided by My Oracle Support.



Caution: Once you begin an upgrade, any changes you make to the configuration during the process (such as creating or editing network elements or policies) may be lost.

Note: In Policy Management version 9.3, secure connections used port 443. Before upgrading from version 9.3 to version 11.5, disable **Secure Connection** until all devices are upgraded. For more information, see [Creating a Policy Server Profile](#).

Performing an Upgrade

The information in this section is a general overview of the Upgrade Manager steps you take to upgrade a cluster or servers. Specific details, including the order in which systems are upgraded, are provided by My Oracle Support. See <https://support.oracle.com> for more information.



Caution: Once you begin an upgrade, any changes you make to the configuration during the process (such as creating or editing network elements or policies) may be lost.

A server must display **Forced Standby** in the Server State column on the **System Maintenance** page before a software upgrade can be performed on that server.

Before upgrading any server in any cluster of the Policy Management network:

1. Use **Upload ISO** to obtain upgrade files.
2. Use **Push Script** to distribute upgrade files to each server.

You must upgrade the primary-site CMP cluster first. To upgrade a primary-site CMP cluster:

1. On the active server of the primary-site cluster, execute the command **policyUpgrade.pl --prepareUpgrade**. (For details of this script and how to execute it, contact My Oracle Support. See <https://support.oracle.com> for more information.)
2. Select the standby server of both the primary-site and secondary-site cluster and apply **Force Standby**.
3. Select the forced standby server of the primary-site cluster and apply **Start Upgrade** to begin the upgrade process on that server.
4. Select the primary site and apply **Switch ForceStandby** to make the standby server active and the active server standby. You are logged out of the CMP system.
5. Log in to the CMP system, select the forced standby server, and apply **Start Upgrade** to begin the upgrade process on that server.
6. Select the forced standby server and apply **Cancel Force Standby** to make it standby.
7. Select each server and apply **Upgrade Completion**.

Once you upgrade the primary-site CMP cluster, you can upgrade a secondary-site CMP cluster. To upgrade a secondary-site CMP cluster:

1. Select the forced standby server of the secondary-site cluster and apply **Start Upgrade** to begin the upgrade process on that server.
2. Select the secondary site and apply **Switch Force Standby** to make the standby server active and the active server standby.
3. Select the forced standby server and apply **Start Upgrade** to begin the upgrade process on that server.
4. Select the forced standby server and apply **Cancel Force Standby** to make it standby.
5. Select each server and apply **Upgrade Completion**.

To upgrade a non-georedundant MPE or BoD cluster:

1. Select the active server of the cluster and apply **Turn Off Replication** to stop replication traffic.
2. Select the standby server of the cluster and apply **Force Standby**.
3. Select the forced standby server of the cluster and apply **Start Upgrade** to begin the upgrade process on that server.
4. Select the cluster and apply **Switch Force Standby** to make the standby server active and the active server standby.
5. Select the cluster and apply **Reapply Configuration** (see [About Reapplying a Configuration](#)) to distribute configuration information to it.
6. Select the forced standby server and apply **Start Upgrade** to begin the upgrade process on that server.
7. Select the active server of the cluster and apply **Turn On Replication** to restart replication traffic.
8. Select the standby server and apply **Cancel Force Standby** to make it standby.
9. Select each server and apply **Upgrade Completion**.

To upgrade a georedundant MPE or BoD cluster:

1. Select the active and spare servers of the cluster and apply **Turn Off Replication** to stop replication traffic.

2. Select the standby server of the cluster and apply **Force Standby**.
3. Select the forced standby server of the cluster and apply **Start Upgrade** to begin the upgrade process on that server.
4. Select the spare server of the cluster and apply **Force Standby**.
5. Select the cluster and apply **Switch Force Standby** to make the standby server active and the active server standby.
6. Select the cluster and apply **Reapply Configuration** (see *About Reapplying a Configuration*) to distribute configuration information to it.
7. Select the forced standby server and apply **Start Upgrade** to begin the upgrade process on that server.
8. Select the spare server of the cluster (at the georedundant site) and apply **Force Standby**.
9. Select the forced standby server and apply **Start Upgrade** to begin the upgrade process on that server.
10. Select the active server of the cluster and apply **Turn On Replication** to restart replication traffic.
11. Select the standby server and apply **Cancel Force Standby** to make it standby.
12. Select each server and apply **Upgrade Completion**.

Once the upgrade is accepted, the last step is to select each server and apply **Accept Upgrade**.

Note: An upgrade must be accepted (or rejected) before any subsequent upgrade can occur.

System Maintenance Elements

On the **Upgrade Manager** menu, **System Maintenance** is an option. All servers in the topology appear in the server table on this page. Servers display in groups by cluster; clusters can be collapsed or expanded by clicking the [-] or [+] icons in the first column of the table. Server information is updated every ten seconds.

There are three types of elements that appear on the **Upgrade Manager** GUI page:

- Checkboxes to select servers/ISOs on which to perform operations.
- Table of filtered servers.
- Lists (**Columns**, **Filters**, and **Operations**) for changing what displays in the table and for performing operations.

Table 25: System Maintenance Elements describes all of the elements. *Table 26: System Maintenance Operations* describes all the operations.

Table 25: System Maintenance Elements

Element	Description
<input type="checkbox"/> (checkbox)	<p>Use the checkbox column to select the servers on which an operation is to be performed. If you select a main cluster server, all servers in that cluster are selected.</p> <p>Note: At least one server must be selected before you can select an operation from the Operations list.</p>
Name	Displays the server name of each server. When a server is in the process of being upgraded, a special upgrade icon appears next to the name.

	Likewise, if a server upgrade has failed, a special failed icon appears next to the name.
IP	Displays the IP address of each server. The Filters list allows you to display only the server with a specific IP address or All servers.
Server State	Displays the current state of the server.
ISO	Displays the current ISO files that are loaded on the server.
Running Release	Displays the current Policy Management software release of each server. The Filters list allows you to display a specific release only or All releases.
Upgrade Status	Displays details of last upgrade performed on each server.
Columns	Use the Columns list to change the columns that appear on this page. By default, all columns appear. To change which columns appear, uncheck the columns to be removed from the page. The Name column is mandatory.
Filters	Use the Filters list to select a subset of servers to appear on this page. On this menu are the following pulldown filter submenus: Appl Type , Site , IP , State , Prev Release , and Running Release . These filters are set to All by default, so all servers appear initially. Selecting another option from one or more of these filters reduces the number of servers displayed.
Operations	<p>Use the Operations list to select an upgrade operation to perform.</p> <p>Note: At least one server must be selected before you can select an operation from the Operations list. The operations that appear in the list depend on the state of the servers that are selected. In other words, when more than one server is selected, only the Operations that are available on all of these servers appear.</p> <p>See Table 26: System Maintenance Operations for the possible operations. As a protective feature, when a command is executed, a warning message displays, asking if you are sure you want to execute this operation. If you click OK, a progress bar displays the status of the command completion in a window.</p> <p>Note: Once the operation is confirmed, it cannot be cancelled.</p>

Table 26: System Maintenance Operations

Operation	Description
Push Script	Pushes script to remote server. Upgrade Manager uses the script to communicate with the remote server and to perform the upgrade or backout.
Upload ISO	Adds ISO to the specified Policy Management products.

Force Standby	Forces the selected server(s) into standby status.  CAUTION Note: You cannot force both servers of a CMP cluster into standby status.
Turn Off Replication	Turns off replication for the selected servers.
Prepare Upgrade	Turns off COMCOL replication of database tables.
Upgrade Completion	Turns off legacy replication.
Undo Upgrade Completion	Prepare for a backout of a software upgrade. This process turns on legacy replication for all the clusters.
Switch ForceStandby	Switches the upgraded server to active and the previously active server to forced standby to upgrade it.
Cancel ForceStandby	Cancels the forced standby status of the selected server(s).
Start Upgrade	Begins the upgrade on the selected server(s) with the selected ISO on each server.
Accept Upgrade	Removes backout information. Once the upgrade is accepted, it cannot be rolled back.
Backout	Initiates a backout on the selected server(s).
Export SSD	Exports the SSD file.

Viewing Upgrade Status of Servers

Use this procedure to view the status of in-service servers before, during, and after a software upgrade.

1. From the **Upgrade Manager** section of the navigation pane, select **System Maintenance**. The **System Maintenance** page appears.
2. (Optional) Click **Filters** and specify the information to customize the list of servers that display in the table.
3. (Optional) Click **Columns** to select columns display in the table.

All in-service servers that meet the filter criteria are listed. Server information is updated every ten seconds.

About Rolling Back an Upgrade

It is possible to roll back, or back out, the Policy Management software to the previous version in a production environment until the upgrade is accepted. Procedures and scripts are available to preserve the current state of subscriber data, such as MPE sessions. Before beginning a rollback, contact My Oracle Support and inform them of your plans.

Note: Once an upgrade is accepted with the **Accept Upgrade** operation, it cannot be rolled back.

Rollback is subject to the following limitations:

- You can roll back one version only, regardless of whether the last upgrade was a major, minor, or maintenance release.
- Subscriber sessions affected by new features may be affected.

If you decide to roll back an upgrade, you should do so in reverse order. That is, first roll back the last cluster you upgraded, then the previous one, and so on, and then roll back CMP clusters. Once the systems are rolled back to a previous release, they can be upgraded to another supported version.

Chapter 13

Defining Global Configuration Settings

Topics:

- *Configuring Subnet Settings.....179*
- *About Routing by CMTS IP.....180*
- *Configuring Stats Settings.....181*

This section describes how to configure global CMP settings.

Configuring 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 enable and configure 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.

1. 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.

2. Click **Add** to add the IP address and prefix length. The IPv6 address and prefix length are added to the list of addresses.
3. 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. When you finish, 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.

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 Subnet Settings](#).

Configuring 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 Subnet Settings](#) to configure IPv4 subnet filtering.

To enable and configure CMTS IP routing, follow these steps.

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.

- When you finish, 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.

- 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.

- Click **Create Network Element**.

The **New Network Element** page opens.

- 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.

- When you finish, 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.

To change stats settings, do the following:

- From the **Global Configuration** section of the navigation pane, select **Global Configuration Settings**.

The content tree displays a list of global configuration settings.

- From the content tree, select the **Stats Settings** folder.

The **Stats Settings** page opens in the work group area.

- Click **Modify**.

The fields on the **Stats Settings** page become editable.

- Enter values for the configuration attributes:

- a) **Stats Reset Configuration** — From the list, select **Manual** or **Interval**. The default value is **Manual**.

- When in Manual mode, numeric values can only reset when the system restarts (for example, on failover or initial startup) or when you issue a reset command. Manual mode disables the resetting of numeric fields at regular intervals but does not alter historical data collection.
- When configured for Interval mode, numeric values are reset at regular intervals, controlled by **Stats Collection Period**. In Interval mode, a reset occurs on the hour and then every 5, 10, 15, 20, 30 or 60 minutes afterwards, depending on the value selected in **Stats Collection Period**, providing performance information about the Policy Management system at specific times of day.

b) **Stats Collection Period** — When **Stats Reset Configuration** is set to **Interval**, specify the time interval in minutes to use from the list. Options are: minutes.

- 5 (default)
- 10
- 15
- 20
- 30
- 60

5. When you finish, click **Save**



Caution: Saving changes to the statistics settings causes the historical stats data to be lost.

The Stats Settings attributes are configured.

Chapter 14

System Administration

Topics:

- *Configuring System Settings.....184*
- *Importing and Exporting Configurable Objects.....186*
- *The Manager Report.....189*
- *The Trace Log.....190*
- *Viewing the Audit Log.....192*
- *About Scheduled Tasks.....196*
- *User Management.....199*
- *Changing a Password.....214*

System Administration 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 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 message: Your account already has the maximum number of concurrent sessions.
 - 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, once 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 MPE device using IPv4 format, then define the alert destination using IPv4 format; if you defined the address of this MPE device 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 Text** — The text that displays on the login page. You can enter up to 10,000 characters.
- k) **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.
- l) **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) **Actions upon Crossing Threshold** — The system action to take if a user reaches the invalid login threshold:
 - **Lock user** — 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 (in IPv4 or IPv6 format) from which the login attempts were 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 letters (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 letters (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 numerals (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-alphabetic characters (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 0 (disabled). Select it to require users to conform to a new password policy effective the next time they log in.

6. When you finish, click **Save**.

The system settings are configured.

Figure 37: Sample Password Strength Policy 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.

Password Strength Settings

Lower-case letter

Upper-case letter

Numeral

Non-alphanumeric character

Require at least categories of the above

Require at least lower-case letter(s) (1-64)

Require at least upper-case letter(s) (1-64)

Require at least numeral(s) (1-64)

Require at least non-alphanumeric character(s) (1-64)

Force users with weak password to change password at their next login

0

1

1

0

Save **Cancel**

Figure 37: Sample Password Strength Policy

Importing and Exporting Configurable Objects

In addition to defining manageable objects manually, you can add them to the CMP database using the OSSi XML Interface or by importing them from an XML file. You can also export a list of objects of various types to an XML output file. This section describes the OSSi XML interface and the XML bulk import and export processes.

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.

Using the OSSi XML Interface

The OSSi XML interface provides access to raw data in the system directly via HTTP. The system data is entered and returned as XML documents in accordance with a defined schema. The schema for the input XML is provided to specify exactly which attributes of a manageable object are permitted on import, as well as the formatting for those attributes.

You can also define object groups as part of the XML file and import them within the same file. Groups let you define a logical organization of objects within the CMP database at the time of import. Group structures include not only group attributes, but also relationships between groups, subgroups, and objects.

The OSSi XML interface includes the following:

- **Topology Interface** — Allows you to query and manage network elements within the system
- **Operational Measurements (OM) Interface** — Allows you to retrieve statistical data from the system

- **Policy Tables** — Allows you to export policy tables, and import them to add, edit, replace or delete a table

For detailed information, see the document *OSSI XML Interface Definition*.

Importing an XML File to Input Objects

During the import process, object definitions are read one at a time from the user-specified XML file. Each object is then validated and checked against the existing database for collisions (duplications). Collisions are detected based on the object name, which is a unique database key. If the object already exists within the system, the existing object's attributes are updated (overwritten) by the attributes specified in the XML file being imported. If the object does not exist within the system, the object is created and imported as a new object. A blank element value is replaced with a default or null value, as appropriate.

An XML import is limited to 20,000,000 bytes. If you try to import a file larger than that the import will fail with a result code of 102 (input stream error).

Note: Export the existing database of objects before starting an import operation to ensure that you can recreate the previous state if necessary (see [Exporting an XML File](#)).

To use an XML file to input defined objects:

1. From the **System Administration** section of the navigation pane, select **Import/Export**. The **Import/Export** page opens in the work area.

Note: Do not select **Policy Import/Export**, in the **Policy Management** section; that is a different function.
2. On the **Import/Export** page, enter the file name of the XML import file, or click **Browse** and, from the standard file open window that appears, locate it.
3. Select what to import:
 - * (specifies import all types) (default value)
 - **Network Elements**
 - **Tiers**
 - **Traffic Profiles**
 - **Match Lists**
 - **Policy Table**
 - **Applications**
 - **Roles**
 - **Scopes**
 - **Users**

If you select **Network Elements**, additional filtering fields appear to help you manage the volume of data being imported. You can filter by network element name. Each additional field accepts a string that can include the wildcard characters * (to represent any string) and ? (to represent any character). By default, all elements matching the filter are included. For each field you can select the operators **AND**, **OR**, **AND NOT**, or **OR NOT**; if you select an operator, an additional statement field appears. You can specify up to six logical combinations of filtering statements.

Note: The concatenation of all filters is left associative. For example, C1 AND C2 OR C3 equals (C1 AND C2) OR C3. The NOT operator affects the succeeding statement(s); for example, C1 AND NOT C2 AND C3 equals C1 AND (NOT C2) AND C3.

4. Click Import.

Data from the XML file is imported. If the operation takes more than five seconds, a progress bar appears.

Following the import, status messages provide the total counts of all successful imports, updates, and failures. Click **Details** (the button is below the status messages) to open a window containing detailed warnings and errors for each object. The error messages contain identifying information for the XML structure that caused the error, allowing you to pinpoint and fix problems in the XML file.

For each User element, ensure that Role and Scope data is also defined. The recommended sequence of elements in the XML import file is Network Element, Role, Scope, and then User.

If an imported user password does not satisfy the current password rules, the user will have to change passwords on first login. Password expiration timestamps are imported, so the passwords will expire on the schedule of the CMP system from which they were exported.

Exporting an XML File

The Export feature creates an XML file containing definitions for objects within the CMP database, in the same schema used on import. You can back up data by exporting it to an XML file, and restore it by importing the same file. The export file can also be transferred to a third-party system. To export an XML file:

1. From the **System Administration** section of the navigation pane, select **Import/Export**.
The Import/Export page opens in the work area.

Note: Do not select **Policy Import/Export**, in the **Policy Management** section; that is a different function.

2. Select the type of export:

- **Network Elements** (default)
- **Accounts**
- **Tiers**
- **Traffic Profiles**
- **Match Lists**
- **Applications**
- **Policy Table**
- **Roles**
- **Scopes**
- **Users**

3. If you select **Network Elements**, the **Network Element Name** fields are available to help you manage the volume of potential data being imported.

You can filter by network element name and Diameter identifier. Each additional field accepts a string that can include the wildcard characters * (to represent any string) and ? (to represent any character). By default, all elements matching the filter are included. For each field you can select the operators **AND**, **OR**, **AND NOT**, or **OR NOT**; if you select an operator, an additional statement field appears. You can specify up to six logical combinations of filtering statements.

Note: The concatenation of all filters is left associative. For example, C1 AND C2 OR C3 equals (C1 AND C2) OR C3. The NOT operator affects the succeeding statement(s); for example, C1 AND NOT C2 AND C3 equals C1 AND (NOT C2) AND C3.

4. Click **Export**.

A standard file download window opens, and you are prompted, Do you want to open or save this file?

5. Click **Save** to save the file.

Data is exported to an XML file. If the operation takes more than five seconds, a progress bar appears. User passwords are exported in encrypted text. Password expiration time stamps are retained, so the passwords expire on the schedule of the CMP system from where they were exported.

Note: The user accounts datacollector and _policy_server cannot be exported.

The Manager Report

The Manager Report provides information about the CMP cluster itself. This information is similar to the Cluster Information Report for MPE clusters. The display is refreshed every ten seconds.

To view the Manager Report, select **Reports** from the **System Administration** section of the navigation pane.

The fields that display in the Manager Reports section include the following:

- **Cluster Name and Designation** — The name of the cluster, and also 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.

To reset the display counters, click **Reset All Counters**.

To pause refreshing the display, click **Pause**. To resume refreshing, click **Resume**.

- **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 Report is a listing of 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 GUI service). For the definitions of these states, see [Server Status](#).
- **Utilization** — Displays the percentage utilization of disk (of the /var/camiant filesystem), CPU, and memory.

The **Actions** let you restart the CMP software on the server or restart the server.

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](#).

Viewing the Trace Log

To view log information using the Trace Log Viewer:

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 **Scanning Trace Logs** 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** — Once 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** — Once 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. When you finish, 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 different fields.

To filter the log information using the Trace Log Viewer:

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 progress message **Scanning Trace Logs** appears.
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 that appeared in events last month and this month, only results from this month appear.

5. Click **Search**.
The filtered log displays.
6. When you finish, click **Close**.
The **Trace Log Viewer** window closes.

Configuring the Trace Log

You can configure the trace log severity message levels and forwarding settings for the CMP system.

Note: For information on configuring cluster-level messages written to the trace log, see [Configuring Log Settings](#).

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

Caution: Before changing the default logging level, consider the implications. Lowering the trace log level setting from its default value (for example, from **Warning** to **Info**) causes more notifications to be recorded in the trace 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 trace log, and may cause you to miss important notifications.

4. You can enable and configure Trace Log Forwarding at the system level.

Note: The CMP system provides log forwarding configuration for all products that have trace logs: MPE, MRA, MA, BOD, and the CMP itself.

For each cluster, enter the following:

- a) Select the **Enable Trace Log Forwarding** checkbox in the **Modify Trace Log Forwarding Settings** section of the page.
Trace Log Forwarding is enabled.
- b) Enter a valid **Hostname/IP Address** for each CMP cluster receiving the trace logs.

Note: The system validates the IP address is unique based on the literal value. It does not resolve the hostname 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 3](#) for a description of each level.

5. When you finish, click **Save**.

The system trace log settings are configured.

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, giving you better system control. The audit log is stored in the database, so it is backed up and can be restored.

To display 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 38: Audit Log* shows an example.)

Audit Log			
124 items found, displaying 1 to 20. [First/Prev] 1, 2, 3, 4, 5, 6, 7 [Next/Last]			
Date / Time	User	Host Name / Action IP Address	Description
2012-04-20 14:00:47	admin	10.15.5.15	User - Login
2012-04-20 13:55:40	admin	10.25.170.220	User - Logout
2012-04-20 13:48:45	admin	10.15.5.108	(admin) login
2012-04-20 12:48:20	admin	10.25.170.220	User - Login
2012-04-20 12:29:36	admin	10.33.251.15	User - Logout
2012-04-20 12:07:03	admin	10.15.5.108	(admin) logout
2012-04-20 11:49:13	admin	10.15.5.108	(admin) login
2012-04-20 11:36:12	admin	10.33.251.15	User - Login
2012-04-20 11:32:35	admin	10.15.5.108	User - Logout
2012-04-20 11:01:20	admin	10.15.5.108	(admin) login
2012-04-20 10:07:31	admin	172.31.251.25	User - Logout
2012-04-20 09:58:17	admin	10.26.3.2	User - Login
2012-04-20 09:58:17	admin	10.26.3.2	User - Logout
2012-04-20 09:28:46	admin	10.26.3.2	MRA - Reapply Config
2012-04-20 09:28:30	admin	10.26.3.2	Policy Server - Reapply Config
2012-04-20 09:27:59	admin	10.26.3.2	Policy Group - Associate
2012-04-20 09:27:47	admin	10.26.3.2	Policy Group - Associate
2012-04-20 09:27:14	admin	10.26.3.2	Policy Group - Associate
2012-04-20 09:27:03	admin	10.26.3.2	Import - Completed
2012-04-20 09:27:02	admin	10.26.3.2	Import - Initiated

Figure 38: Audit Log

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 39: Audit Log Details* shows an example.)

Audit Log				
124 items found, displaying 21 to 40. [First/Prev] 1, 2, 3, 4, 5, 6, 7 [Next/Last]				
Date / Time	User	Host Name / Action	Description	
2012-04-20 09:26:39	admin	10.26.3.2 Import - Completed	Import of file "PolicyTableDataExport.xml" completed.	
2012-04-20 09:26:37	admin	10.26.3.2 Policy Table Library - Batch Create	Batch Created Policy Table Library	
2012-04-20 09:26:37	admin	10.26.3.2 Policy Table Library - Create	Created Policy Table Library: martin - O2 Device specific flow or session	
2012-04-20 09:26:33	admin	10.26.3.2 Policy Table Library - Create	Created Policy Table Library: martin - O2 ApnChargingRuleList	
2012-04-20 09:26:29	admin	10.26.3.2 Policy Table Library - Create	Created Policy Table Library: matTable1	
2012-04-20 09:26:24	admin	10.26.3.2 Import - Initiated	Import of file "PolicyTableDataExport.xml" initiated.	
2012-04-20 09:26:17	admin	10.26.3.2 Import - Completed	Import of file "TrafficProfileExport.xml" completed.	
2012-04-20 09:26:16	admin	10.26.3.2 Traffic Profile - Create	Created Traffic Profile: netcom.sp_5	
Name: netcom.sp_5 QosProfileType: Predefined PCC Rule Rule Name: netcom.sp_5 Description:				
2012-04-20 09:26:16	admin	10.26.3.2 Traffic Profile - Create	Created Traffic Profile: netcom.sp_2	
2012-04-20 09:26:16	admin	10.26.3.2 Traffic Profile - Create	Created Traffic Profile: surf.sp_5	
2012-04-20 09:26:16	admin	10.26.3.2 Traffic Profile - Create	Created Traffic Profile: surf.sp_0	
2012-04-20 09:26:16	admin	10.26.3.2 Traffic Profile - Create	Created Traffic Profile: mmapn.sp_5	
2012-04-20 09:26:16	admin	10.26.3.2 Traffic Profile - Create	Created Traffic Profile: mmapn.sp_3	
2012-04-20 09:26:16	admin	10.26.3.2 Traffic Profile - Create	Created Traffic Profile: enigma-test_5	
2012-04-20 09:26:16	admin	10.26.3.2 Traffic Profile - Create	Created Traffic Profile: enigma-test_3	
2012-04-20 09:26:16	admin	10.26.3.2 Traffic Profile - Create	Created Traffic Profile: penigma-test_33	
2012-04-20 09:26:16	admin	10.26.3.2 Traffic Profile - Create	Created Traffic Profile: penigma-test_33	
2012-04-20 09:26:16	admin	10.26.3.2 Traffic Profile - Create	Created Traffic Profile: internet1_5	
2012-04-20 09:26:16	admin	10.26.3.2 Traffic Profile - Create	Created Traffic Profile: internet1_3	
2012-04-20 09:26:16	admin	10.26.3.2 Traffic Profile - Create	Created Traffic Profile: blackberry.net_5	
Refine Search				

Figure 39: Audit Log Details

Searching for Audit Log Entries

To search for entries in the Audit Log:

- From the **System Administration** section of the navigation pane, select **Audit Log**.
The **Audit Log** page opens in the work area.
- Click **Search**.
The **Audit Log Search Restrictions** page opens.
- Define one or more of the following items:
 - From** — Click (calendar icon), specify a date and time then click **Enter**.
 - To** — Click (calendar icon), specify a date and time then click **Enter**.
 - Action by User Names** — Enter the name of the user or users to audit.
 - Action on Policy Servers** — Enter the name of the Policy Management device to audit.
 - Audit Log Items to Show** — Specifies an item to audit for display. By default you can specify three items; click **More Lines** to add an additional item.
 - Policy Server**
 - Network Element**
 - Network Element Group**
 - Application**
 - Policy**
 - Policy Group**
 - Account**
 - Alert**

- User
- Audit
- Alarm
- OM Statistics
- MPE Manager
- Upgrade Manager
- Topology Setting
- Global Configuration Settings
- Trending Report
- User Layout
- Upsync Log Alarm Threshold
- BoD
- BoD Services
- BoD Group

• **Results Forms** — Specifies the number of items to display on each page, along with the chronological order of the items (most recent or oldest).

4. When you finish, click **Search**.
The Audit Log displays search results.

Exporting or Purging Audit Log Data

You can export the audit log to a text file or you can purge items from the system.

Exporting Data

To export data from the audit logs:

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** — Enter a date in the format *mm/dd/yyyy*, or click  (calendar icon), located to the right of the field, to select a date.
4. When you finish, click **Export**.
A standard file download window opens; you can open or save the export file. The default filename is **AuditLogExport.txt**.

The audit log is exported.

Purging 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 **Purge Through Date** field of the **Items to Purge** section, enter a date in the format *mm/dd/yyyy*, or click  (calendar icon), located to the right of the field, to select a date from the pop-up window.

4. When you finish, click **Purge**.

You are prompted: "Click 'OK' to purge all audit log items through: *mm/dd/yyyy*."

5. Click **OK** (or **Cancel** to cancel the request).

The data is purged from the audit log.

About Scheduled Tasks

The CMP 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 of these tasks to better manage network load or to propagate a network element change to the Policy Management devices on demand.



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:

- **Subnet Overlap Detector** — Detects overlapping (duplicate) subnets across all CMTS devices.
- **Health Checker** — Periodically checks the MPE devices to ensure that they are online.
- **OM Statistics** — Periodically retrieves Operational Measurement (OM) statistics from all MPE 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. 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 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 in order to provide any data through the OM XML interface. The *OSSI XML Interface Definition* document describes the OM Interface and the OM Statistics in detail.

- **OSSI Distributor Task** (optional) — Reads from the database topology and subscriber data that has entered the CMP 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, via SNMP, all CMTS devices for the configured service classes and then stores them in the local database.
- **Subscriber SNMP Collector** — Polls, via 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 local database and then distributes it to the appropriate Policy Management devices within the system.
- **Subscriber Distributor** — Reads subscriber data from the CMP local database and then distributes it to the appropriate Policy Management devices within the system.
- **CMTS MA Collector** (optional) — Polls all of the MAs 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 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 in the work area.
2. To display details about a task, click the task name. The current settings and status are displayed.

For example:

Scheduled Task Administration

Name	OM Statistics
Description	The task to retrieve OM statistics.
Last Exit Status	Success
Current State	Idle
Last Start Time	Jun 7, 2013 2:30:00 PM
Last End Time	Jun 7, 2013 2:30:02 PM
Next Run Time	Jun 7, 2013 2:45:00 PM
Run Interval	15 mins 0 sec

Settings

Number of days to keep statistical data (1 - 30)	7
--	---

Server time: Jun 07, 2013 02:32 PM EDT

Figure 40: Schedule Task Administration - OM Statistics

3. The options for this task are as follows:
 - **details** (Subnet Overlap Detector only) — Located in the Exit Status Message section, click to display all the duplicate and overlapping subnets.

For example:



Figure 41: Scheduled Task Administration

- **Reschedule** — Click to reschedule the time that this task is performed on the Policy Management device.

For example:

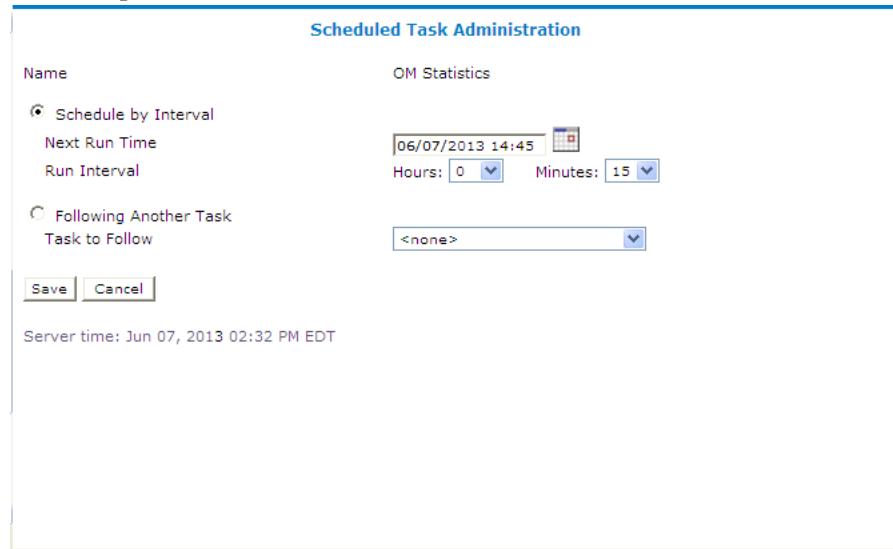


Figure 42: Scheduled Task Administration

- **Schedule by Interval (Next Run Time or Run Interval)** — Defines the run interval for the task to follow.
Valid run intervals are from 0 to 24 hours in 5-minute increments.
- **Following Another Task** — Schedules the task run time as following the completion of another scheduled task selected from the list.
- **Settings** — Number of days to keep data; the default is seven days. Available for the OM Statistics and Replication Statistics tasks only.
- **Run Now** — Runs the process immediately.
You are prompted with a confirmation message. Click **OK** to run the task.
- **Disable or Enable** — Disables or enables the next scheduled execution of this process.
If you click **Disable**, a confirmation message displays. Click **OK**. The task is disabled and will not run at the next scheduled time, and the button changes to **Enable**.

- **Refresh** — Refreshes the page.
- **Cancel** — Returns to the previous page.

User Management

The CMP system lets you configure the following user attributes:

- **Roles** — What a user can do within the CMP system.
- **Scopes** — Network element groups and Policy Management device groups that provide a context for a role.
- **Users** — Once you define roles and scopes, you can apply them to user profiles.

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 accounts. Requests and responses consist of XML documents. You must define a user profile for the external system. The profile is a regular CMP user profile with specific roles and scope.

Assign the profile a role that includes the following privileges:

- Show privilege for XML Import/Export
- Read-Write privilege for User Management

Assign the profile the default scope, Global.

For information on creating a user profile, see [Creating a User Profile](#). For more information on the XML application programming interface, see the [OSSI XML Interface Definitions Guide](#).

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 device management and configuration functions. Access is also permitted to all system administration functions except user administration.
Viewer	Permits read-only access to functions associated with Policy Management device management and configuration. Full access is also permitted to some of the system administration functions, such as Change Password.

The CMP system allows you to 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, displaying existing roles.
3. On the Role Administration page, click **Create Role**.
The **New Role** page opens. By default, all privileges are set to **Hide** (functions do not appear to users of the role, so access must be explicitly granted) or **Read-Only** (function is displayed but not changed).
4. Enter the following general role information:
 - a) **Name** — The name for the new role (up to 64 characters long)
 - b) **Description/Location** (optional) — Free-form text
5. **Policy Server Privileges** — Defines access to the following MPE device management functions (assigning each the privilege **Hide**, **Read-Only**, or **Read-Write**):
 - **Configuration**
 - **Network Element**
 - **Application**
 - **Traffic Profiles**
 - **Media Profile**
 - **Service Class**
 - **Record Keeping Server and Event Messaging**
 - **Management Agent**
 - **Custom AVP Definition**
 - **Global Configuration Settings**
 - **Bulk Operation**
6. **Network Privileges** — Defines access to the network management Paths function (assigning the privilege **Hide**, **Read-Only**, or **Read-Write**):

Topology

7. **BoD Privileges** — Defines access to the Bandwidth on Demand (BoD) Application Manager (with the privileges **Hide**, **Read-Only**, or **Read-Write**).
 - **Configuration**
 - **Services**
 - **Service Import/Export**
8. **Policy Management Privileges** — Defines access to the Policy Management functions:
 - **Policy Library** (with the privileges **Hide**, **Read-Only**, **Read and Deploy**, or **Read, Deploy, and Write**)
 - **Template Library** (with the privileges **Hide**, **Read-Only**, or **Read-Write**)
 - **Policy Table Library** (with the privileges **Hide**, **Read-Only**, or **Read-Write**)
 - **Policy Import/Export** (with the privileges **Hide**, **Read-Only**, or **Read-Write**)

9. **System Wide Reports Privileges** — Defines access to the system-wide reports functions (with the privileges **Hide**, **Read-Only**, or **Read-Write**).
10. **Platform Setting Privileges** — Defines access to the platform setting functions:
 - **Topology Configuration** (with the privileges **Hide**, **Read-Only**, or **Read-Write**)
 - **Server Operation** (with the privileges **Hide** or **Read-Write**)
11. **Upgrade Manager Privileges** — Defines access to software upgrade functions:
 - **ISO Maintenance** (with the privileges **Hide**, **Read-Only**, or **Read-Write**)
 - **System Maintenance** (with the privileges **Hide**, **Read-Only**, or **Read-Write**)
12. **System Administration Privileges** — Defines access to system administration functions:
 - **XML Import/Export** (with the privileges **Hide** or **Show**)
 - **Reports** (with the privileges **Hide** or **Show**)
 - **Operational Measurements** (with the privileges **Hide** or **Read-Only**)
 - **User Management** (with the privileges **Hide**, **Read-Only**, or **Read-Write**)
 - **Scheduled Tasks** (with the privileges **Hide** or **Read-Write**)
 - **Trace Log** (with the privileges **Hide**, **Read-Only**, or **Read-Write**)
 - **Trace Log of CMP** (with the privileges **Hide**, **Read-Only**, or **Read-Write**)
 - **Audit Log** (with the privileges **Hide**, **Read-Only**, or **Read-Write**)
 - **Audit Log User Info** (with the privileges **Hide** or **Show**)
 - **Alarms** (with the privileges **Hide**, **Read-Only**, or **Read-Write**)
 - **Password Strength** (with the privileges **Read-Only** or **Read-Write**)
 - **Push Method for Statistics** (with the privileges **Read-Only** or **Read-Write**)
13. When you finish, click **Save**.

Privileges are assigned to the role.

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, displaying existing roles.
3. Select the role to modify.
The **Role** page opens.
4. Click **Modify**.
The **Modify Role** page opens.
5. Modify the role information.
See [Creating a Role](#) for a description of the fields.
6. When you finish, click **Save**.

The role is modified.

Deleting a Role

You can delete any role except the Administrator role. You cannot delete a role that is in use.

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, displaying existing roles.
3. Delete the role using one of the following methods:
 - From the work area, click  (trash can icon) located next to the role to delete.
 - From the content tree, select the role (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 Scope

Scope defines which network element groups and Policy Management device groups that a user has access to, which provides operational context for a role.

The CMP system allows you to perform the following scope management actions:

- *Creating a Scope*
- *Modifying a Scope*
- *Deleting a Scope*

Creating a Scope

The CMP lets you configure scopes. Scopes contain selections of network element groups and Policy Management device groups that provide a context for a role. This lets you control what areas or devices in a network a user can manage. The default scope, **Global**, contains all items defined within the CMP. Once you define a scope you can apply it to a user.

To configure a new 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, click **Scopes**.
The **Scope Administration** page opens in the work area, displaying existing scopes. The default scope is **Global**.
3. Click **Create Scope**.
The **New Scope** page opens.
4. Enter the following information:
 - a) **Name** — The name for the new scope.
 - b) **Description/Location** (optional) — Free-form text.

5. Select the policy server groups included in this scope.
6. Select the network element groups included in this scope.
7. Select the BoD groups included in this scope.
8. When you finish, click **Save** to create the scope.

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, click **Scopes**.
The **Scope Administration** page opens in the work area, displaying existing scopes. The default scope is **Global**.
3. Select the scope.
The scope description opens.
4. Click **Modify**.
The **Modify Scope** page opens. *Creating a Scope* describes the fields on this page.
5. Modify the scope information.
6. When you finish, click **Save**.

The scope is modified.

Deleting a Scope

You can delete any scope except **Global**. 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, click **Scopes**.
The **Scope Administration** page opens in the work area, displaying existing scopes.
3. Delete the role using one of the following methods:
 - From the work area, click  (trash can icon) located to the right of the role.
 - From the content tree, select the role. The information for the role displays in the work area. Click **Delete**.

A confirmation message displays.
4. Click **OK**.

The scope is deleted.

About User Profiles

A user profile defines a user with a role and one or more scopes.

The CMP system allows you to perform the following scope management actions:

- *Creating a User Profile*

- *Modifying a User Profile*
- *Deleting a User Profile*

Creating a User Profile

The User Management functions include the tools necessary to create, modify, or delete system user profiles.

The CMP system is configured initially with the following default user profiles and passwords:

- admin/policies (you cannot delete this profile)
- operator/policies
- viewer/policies

Each default user profile has an associated role assigned to it. The **admin** user is the only profile that cannot be deleted or have its username modified. Also, the **admin** user is the only user who can create, modify, or delete other users. The password assigned to the **admin** user can be changed. For security reasons, it is recommended that you change this value from the default value as soon as the system is installed.

Note: When logging in, the username is not case sensitive; however, the password is case sensitive.

To create a new user profile:

1. Log into 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, click **Users**.
The **User Administration** page opens in the work area, displaying existing users.
- Note:** The **Log Out All Users** button is visible only to the **admin** user.
4. Click **Create User**.
The **New User** page opens.
5. Define the following information:
 - a) **Username** — Assign a name to the user profile of up to 64 characters (this value is not case sensitive).
 - b) **Description/Location** (optional) — Free-form text.
 - c) **Password** — Assign a password to the user profile.
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.
 - d) **Confirm Password** — Re-enter the password to confirm the value entered above.
 - e) **Password Expiration Period(days; 0=never)** — The number of days a password can be used before it expires. (This overrides the system setting.)
Enter a value from 7 to 365, or 0 to indicate that the password never expires. The default is the system setting.
 - f) **Force to Change Password** — If selected, this user must change passwords when he or she next logs in.
 - g) **Role** — Select a role from the drop list to assign to the user profile.
 - h) **Scopes** — Select one or more scopes to assign to the user profile.
6. When you finish, click **Save**.

The user profile is created and stored in the **Users** group.

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, click **Users**.
The **User Administration** page opens in the work area, displaying existing users.
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. When you finish, click **Save**.

The user profile is modified.

Deleting a User Profile

You can delete any user profile except **admin**. 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, click **Users**.
The **User Administration** page opens in the work area, displaying existing users.
4. Delete the user profile using one of the following methods:
 - From the work area, select  (trash can icon) located to the right of the profile.
 - From the content tree, select the user profile (profile information displays in the work area), then click **Delete**.

A confirmation message displays.

5. Click **OK** to delete the user profile.

The user profile is deleted.

Locking and Unlocking User Accounts

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** account cannot lock the **admin** account.

Locking an Account

To lock a user account:

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, click **Users**.
The **User Administration** page opens in the work area, displaying existing users.
4. Select the user profile from the content tree.
The **User Administration** page displays the information for the user profile.
5. Click **Lock**.
A confirmation message displays.
6. Click **OK**.

The account is locked. The page displays: **User account locked successfully**. The **Lock** button becomes an **Unlock** button. On the **User Administration** page, the Locked Status for the user changes to **Locked**.

Unlocking an Account

To unlock a user account:

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, click **Users**.
The **User Administration** page opens in the work area, displaying existing users.
4. Select the user profile from the content tree.
The **User Administration** page opens.
5. Click **Unlock**.
A confirmation message displays.
6. Click **OK**.

The account is unlocked. The page displays: **User account unlocked successfully**. The **Unlock** button becomes a **Lock** button. On the **User Administration** page, the Locked status for the user changes to **Unlocked by Admin**.

Logging Out All Users

You can log out all users (except the **admin** account) from the CMP system. To log out all 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, click **Users**.
The **User Administration** page opens in the work area, displaying existing users.
4. Click **Log Out All Users**
A confirmation message appears.
5. Click **OK** to log out all users.

Users are logged out.

External Authentication

In addition to the built-in authentication functions, you can configure external authentication, RADIUS authentication, and SANE authentication of CMP users.

The CMP system allows you to manage the following external authentication methods:

- [RADIUS Authentication and Accounting](#)
- [SANE Authentication](#)

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 user **admin** 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 three steps:

1. [Configuring the RADIUS Server](#) to accept authentication (and accounting, if used)
2. [Defining CMP Users to the RADIUS Server](#) and [Associating Roles and Scopes](#) on the CMP system
3. [Defining the CMP System as a RADIUS Client](#) to work with RADIUS

Configuring the RADIUS Server

The RADIUS server 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 RADIUS on the CMP System](#) for details.

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 = oracle
    shortname = MPE5
```

```

}
client 10.0.10.23 {
    secret = oracle
    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, `oracle`) must be defined on both the RADIUS server and entered into the CMP configuration (see [Enabling RADIUS on the CMP System](#)). The shortname is used as an alias.

Defining CMP Users to the RADIUS Server

RADIUS 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 (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.

When RADIUS has authenticated a user, it sends back various attributes with the authentication acceptance message. The CMP system uses these attributes to determine what the user can do. The best practice is to use a vendor-specific attribute (VSA) dictionary file to define what attributes to send back to the client. [Figure 43: Sample VSA Dictionary File For RADIUS](#) shows a sample file. The local RADIUS administrator is responsible for incorporating the VSA dictionary file onto the RADIUS server.

```

=====
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
=====

```

Figure 43: Sample VSA Dictionary File For RADIUS

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.

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.

It is easiest 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 RADIUS if the VSA dictionary is integrated.

The following example defines users who have access at different role levels:

```
Jeff      "garbage"
Class="Administrator",
Oracle-MI-role="Administrator",
Oracle-MI-scope="Global"

Paul     "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 required credentials. For example, consider this user defined in RADIUS:

```
Dawn      "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 GUI matches the one returned value to both of the required credentials.

Enabling RADIUS on the CMP System

By default, RADIUS 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 RADIUS authentication (see [Creating a User Profile](#)).

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:

```
test      Cleartext-Password := "camiant"
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, displaying the current configuration information. By default, external authentication is disabled.
4. Click **Modify**.
The page opens with the available authentication options.
5. In the **Configuration** section, select **Enable RADIUS Authentication**.
Additional fields appear (*Figure 44: RADIUS External Authentication Configuration Page*).
6. Edit the following fields:
 - a) **Enable RADIUS Accounting** — Enables RADIUS accounting on the CMP system.
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.
 - b) **Destination for Accounting Messages** — Choose the following from the list:
 - **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.
 - c) **NAS IP Address** (required) — IP address, in IPv4 or IPv6 format, of the network access server.
By default, this is the local host address.
 - d) **Use local authentication** — Choose when to use local authentication:
 - **When RADIUS servers timeout** (default)
 - **When RADIUS servers timeout or reject**
 - **Never** — Fallback to local authentication is never used (however, the user **admin** is always authenticated locally)
 - e) **Source of User Credentials** — Choose the following from the list:
 - **RADIUS Class** (default) — If selected, the value of the Class attribute returned by the server determines both the role and scope.
 - **Camiant VSAs** — If selected, the value of Camiant VSAs returned by the server determines the role and scope.
 - f) **Action if Missing Credentials**:
 - **Reject** — If you select this option, a user whose login credentials are missing or mismatched is not logged in.
 - **Use following defaults** — If you select this option, a user whose login credentials are missing or mismatched is assigned a default role and scope:
 1. **Default Role** — Role assigned if the user credentials are missing or mismatched. The default is **Viewer**.
 2. **Default Scope** — Scope assigned if the user credentials are missing or mismatched. The default is **Global**.

7. In the **RADIUS Servers** section, edit the following fields:

a) **Primary RADIUS Authentication Server**

- **Server** — 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** — IP port number of the primary server. The default is port 1812.
- **Timeout (seconds)** — How long the CMP system waits for a response from the server. The default is 3 seconds.
- **Retries** — How many times the CMP system tries to send a message to the server. The default is 3.
- **Shared Secret** — A password-like string that must match between the CMP system and the server. If it does not match, the server ignores all messages from the CMP system.

b) **Secondary RADIUS Authentication Server**

If configured, the secondary authentication server uses the same fields as the primary server.

c) **Primary RADIUS Accounting Server**

- **Server** — FQDN or IP address (in IPv4 or IPv6 format) assigned to the primary accounting server.
- **Port** — IP port number of the primary server. The default is port 1813.
- **Timeout (seconds)** — How long the CMP system waits for a response from the server. The default is 3 seconds.
- **Retries** — How many times the CMP system tries to send a message to the server. The default is 3.
- **Shared Secret** — A password-like string that must match between the CMP system and the server. If it does 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 server.

8. When you finish, click **Save**.

RADIUS Authentication and Accounting is configured.

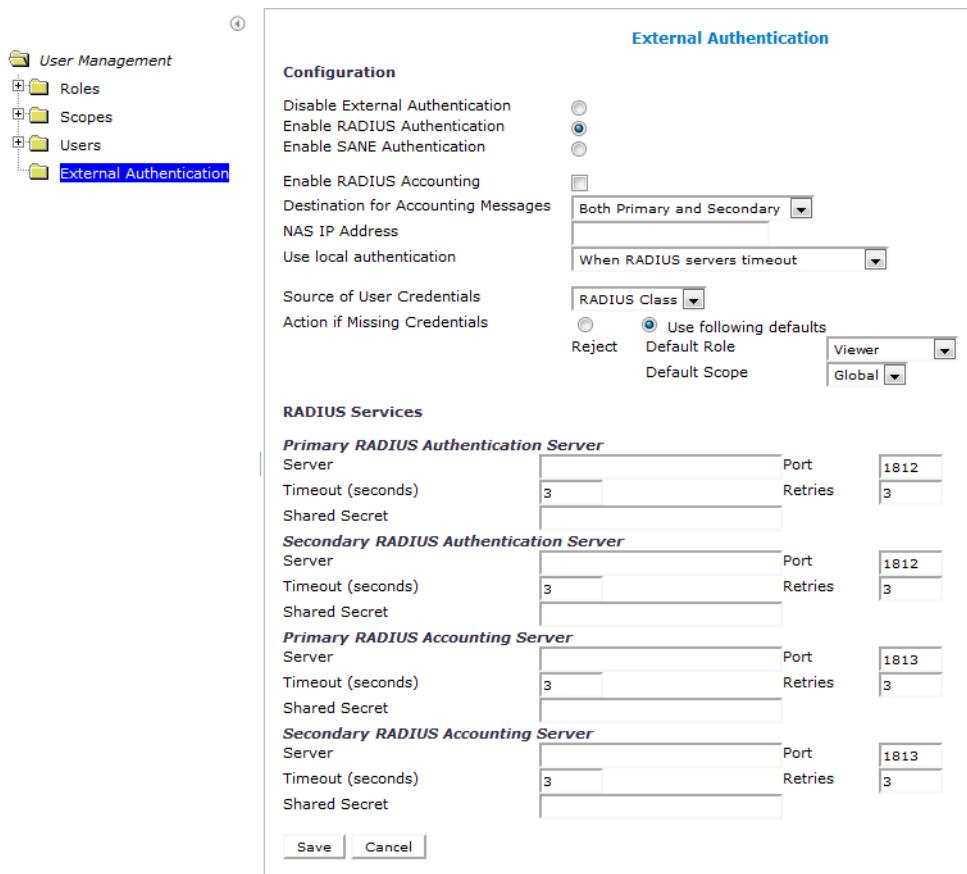


Figure 44: RADIUS External Authentication Configuration Page

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 such 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. (Usage of a SANE client is outside the scope of this document.) See [Enabling and Configuring SANE Authentication on the CMP System](#) for details.

The **admin** account is treated separately. An admin user enters the CMP URL in any supported browser to log in.

The authentication process is as follows:

1. From a SANE client GUI, the user selects the CMP system. A web browser session is launched. An encrypted SANE authentication artifact is sent to the CMP system through the browser.
2. The CMP system forwards the artifact to a SANE server (the SANE responder).
3. 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 accordingly. Otherwise, the CMP system rejects the login request.
4. The user **admin** is always authenticated locally, regardless of configuration settings. (That user clicks on the **Login** link.)

Enabling and Configuring 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 modify page opens.
5. In the **Configuration** section, select **Enable SANE Authentication**.
Additional fields appear.
6. Edit the following fields:
 - a) **Artifact Parameter Name** — Name of the artifact parameter. Enter an alphanumeric string. The default is **artifact**.
 - b) **Verification for Account** — Choose the following from the list:
 - **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 again for each HTTP or HTTPS request. If any request is not authenticated, the user is immediately logged out.
 - c) **Action if Missing Credentials**:
 - **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:
 1. **Default Role** — Default role assigned to the user. The default role is **Viewer**.
 2. **Default Scope** — Default scope assigned to the user. The default scope is **Global**.
7. In the **SANE Servers** section, edit the following fields:
 - a) **SAML Service Name** — Name of the Security Assertion Markup Language service registered with the UDDI server. Enter an alphanumeric string.
 - b) **UDDI Inquiry URL** — Universal Description, Discovery and Integration URL, in HTTP or HTTPS format, for the inquiry.
8. When you finish, click **Save**.

SANE authentication is configured on the CMP system.

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 a system administrator has configured your account for password expiration, you receive a warning when you log in that you must change your password.

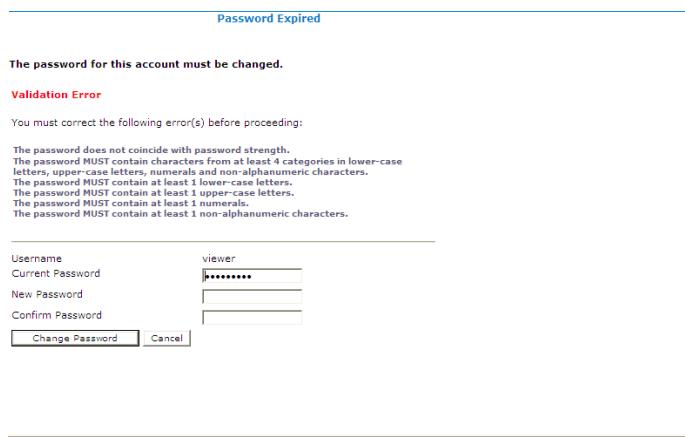
Note: To reset the administrator password, contact My Oracle Support.

To change your password:

1. From the **System Administration** section of the navigation pane, select **Change Password**. The **Change Password** page opens. If your account is set up with a password expiration period, the expiration date is displayed.
2. Enter the following information:
 - a) **Current Password** — The present value of the password.
 - b) **New Password** — The value of the new password.

This value is case sensitive and must conform to the password strength rules. The password cannot contain the user name.
- c) **Confirm Password** — Enter the new password value again.

If your new password does not conform to the password strength rules, a validation error message appears; for example:



The screenshot shows the 'Change Password' dialog box with the following content:

Username: viewer
Current Password: [REDACTED]
New Password: [REDACTED]
Confirm Password: [REDACTED]

Validation Error:

The password for this account must be changed.

You must correct the following error(s) before proceeding:

The password does not coincide with password strength.
 The password MUST contain characters from at least 4 categories in lower-case letters, upper-case letters, numerals and non-alphanumeric characters.
 The password MUST contain at least 1 lower-case letters.
 The password MUST contain at least 1 upper-case letters.
 The password MUST contain at least 1 numerals.
 The password MUST contain at least 1 non-alphanumeric characters.

3. When you finish, click **Change Password**.

Your password is changed.

Appendix A

CMP Modes

Topics:

- *The Mode Settings Page.....216*

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

Normally, servers are pre-configured before delivery. However, if it becomes necessary to replace a server or reinstall the software in the field, the mode selection screen becomes visible, and you must reset the operational 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 modes inappropriately could result in the loss of network element connectivity, policy function, statistical data, and cluster redundancy.

The Mode Settings Page

When you use a web browser to connect to a CMP system after the software is first installed, the **Mode Settings** page opens ([Figure 45: Mode Settings Page](#)). Select modes, functions, and management options, 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.

The modes and functions available are:

Cable Mode Enables support of a cable carrier environment. Functions are described in the *Configuration Management Platform Cable User's Guide*.

Table 27: Cable Functions

Function	Description
PCMM	Supports PacketCable MultiMedia functions.
DQOS	Supports Dynamic Quality of Service functions.
Diameter AF	Supports Diameter AF functions.

Wireless Mode Enables support of a wireless carrier environment. Functions are described in the *Configuration Management Platform Wireless User's Guide*.

Table 28: Wireless Functions

Function	Description
Diameter 3GPP	Supports Diameter 3GPP protocol.
Diameter 3GPP2	Supports Diameter 3GPP2 protocol.
PCC Extensions	Supports Policy and Charging Control functions.
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
LI	Supports Lawful Intercept functions. Described in the <i>Configuring Lawful Intercept Application Note</i> .
SCE-Gx	Supports the Cisco Service Control Engine Gx protocol. If this mode is selected, Diameter 3GPP and RADIUS

Function	Description
	must also be selected, and other Gx sub-modes must not be selected.
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.
Cisco Gx	Supports the Cisco Gx protocol.
DSR	Supports Policy Management network segmentation using a Diameter Signaling Router.
KT	Supports the KT custom features.

SMS Mode Enables support of SMS servers. Functions are described in the *Configuration Management Platform Wireless User's Guide*.

Table 29: SMS Mode functions

Function	Description
SMPP	Supports SMS using SMPP protocol.
XML	Supports SMS using XML.

SPR Mode Enables support of subscriber database management. Select only one sub-mode. Functions are described in the Subscriber Data Management documentation.

Table 30: SPR Mode functions

Function	Description
Subscriber Profiles	Supports subscriber profile functions.
Quota	Supports subscriber quotas.

Wireline Mode Enables support of a wireline carrier environment. Functions are described in the *Configuration Management Platform Wireline User's Guide*.

SPC Mode 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*.

RADIUS Mode Enables support of RADIUS AAA.

The management options are as follows:

- **Manage Policy Servers** — Manage MPE devices.

- **Manage SIP-AM Servers** — Manage Session Initiation Protocol Application Manager (SIP-AM) servers.
- **Manage CD-AM Servers** — Manage Content Distribution Network servers.
- **Manage MA Servers** — Manage Management Agent servers.
- **Manage Policies** — Enable the policy wizard.
- **Manage MRAs** — Manage Multi-Protocol Routing Agent servers.
- **Manage BoDs** — Manage Bandwidth on Demand Application Manager servers.
- **Manage Geo-Redundant MPE/MRA/BoD** — Manage georedundant MPE, MRA, or BoD clusters.
- **Manager is HA (clustered)** — Enable High Availability features.
- **Manage Analytic Data** — Enable output of policy event records.

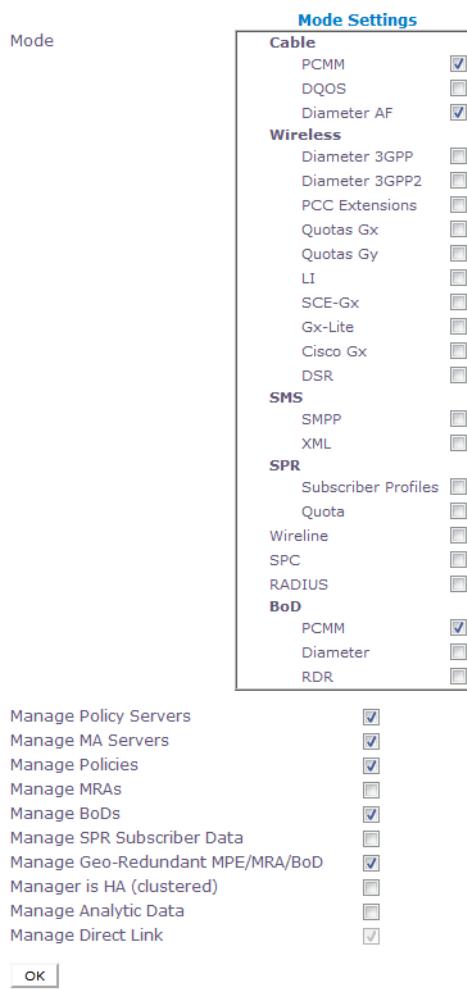


Figure 45: Mode Settings Page

Glossary

#

3GPP 3rd Generation Partnership Project
The standards body for wireless communications.

3GPP2 3rd Generation Partnership Project 2

A

AM Application Manager
A server within a network that is responsible for establishing and managing subscriber sessions associated with a specific application.

AMID Application Manager ID

application The telecommunications software that is hosted on the platform. A service provided to subscribers to a network; for example, voice over IP (VoIP), video on demand (VoD), video conferencing, or gaming.

architecture Used to conceptually describe the function, interaction, and connectivity of hardware, software, and/or system components within a network.

B

Bandwidth on Demand See BoD.

B

BoD Bandwidth on Demand

An application that provides dynamic allocation of bandwidth; for example, a broadband speed promotion.

C

CMP Configuration Management Platform

A centralized management interface to create policies, maintain policy libraries, configure, provision, and manage multiple distributed MPE policy server devices, and deploy policy rules to MPE devices. The CMP has a web-based interface.

CMTS

COMCOL Communications Core Object Library

A suite of re-usable C++ libraries, as well as processes and procedures available for use in Tekelec products. Many of its features are focused toward the communications area of software developments, although its purpose is not intended to restrict its functionality to any particular area

COPS Common Open Policy Service

A protocol that is part of the internet protocol suite as defined by the IETF's RFC 2748. COPS specifies a simple client/server model for supporting policy control over Quality of Service (QoS) signaling protocols (e.g., RSVP).

D

D

Diameter	Diameter can also be used as a signaling protocol for mobility management which is typically associated with an IMS or wireless type of environment. Diameter is the successor to the RADIUS protocol. The MPE device supports a range of Diameter interfaces, including Rx, Gx, Gy, and Ty.
	Protocol that provides an Authentication, Authorization, and Accounting (AAA) framework for applications such as network access or IP mobility. Diameter works in both local and roaming AAA situations. Diameter can also be used as a signaling protocol for mobility management which is typically associated with an IMS or wireless type of environment.
DOCSIS	Data Over Cable Service Interface Specification - An international telecommunications standard for adding high-speed data transfer to an existing cable TV system. Employed by many cable television operators to provide Internet access over their existing infrastructure.
DSCP	Differentiated Service Code Point Differentiated Services Code Point Provides a framework and building blocks to enable deployment of scalable service discrimination in the internet. The differentiated services are realized by mapping the code point contained in a field in the IP packet header to a particular forwarding treatment or per-hop behavior (PHB). Differentiated services or DiffServ is a computer networking architecture that specifies a simple, scalable and coarse-grained

D

mechanism for classifying and managing network traffic and providing quality of service (QoS) on modern IP networks.

E

event

In Policy Management, an expected incident that is logged. Events can be used for debugging purposes.

F

FEID

Financial Entity ID

Identifies the paying entity. The first 8 bytes of the ID designate the multiple-service operator (MSO). By default, the first 8 bytes are zero filled. The 9th byte of the field contains the domain name to uniquely identify the MSO for billing and settlement purposes. The domain name for the MSO is limited to 239 bytes.

FQDN

Fully Qualified Domain Name

The complete domain name for a specific computer on the Internet (i.e., www.oracle.com).

A domain name that specifies its exact location in the tree hierarchy of the DNS.

G

georedundancy

Redundancy between two geographically separate CMP systems.

GUI

Graphical User Interface

The term given to that set of items and facilities which provide the user with a graphic means for

G

manipulating screen data rather than being limited to character based commands.

H**HA**

High Availability

High Availability refers to a system or component that operates on a continuous basis by utilizing redundant connectivity, thereby circumventing unplanned outages.

HTTP

Hypertext Transfer Protocol

I**IPv4**

Identifies an Internet Protocol version 4 address composed of 4 bytes in a dotted decimal format (for example, nnn.nn.nnn.nn).

IPv6

Identifies an Internet Protocol version 6 address composed of 8 groups of colon-separated 4 hexadecimal digits.

IP

Internet Protocol - IP specifies the format of packets, also called datagrams, and the addressing scheme. The network layer for the TCP/IP protocol suite widely used on Ethernet networks, defined in STD 5, RFC 791. IP is a connectionless, best-effort packet switching protocol. It provides packet routing, fragmentation and re-assembly through the data link layer.

IPv4

Internet Protocol version 4

I

IPv6	Internet Protocol version 6
------	-----------------------------

M

MA	Management Agent
----	------------------

MAC	Media Access Control Address The unique serial number burned into the Ethernet adapter that identifies that network card from all others.
-----	--

MPE	Multimedia Policy Engine A high-performance, high-availability platform for operators to deliver and manage differentiated services over high-speed data networks. The MPE includes a protocol-independent policy rules engine that provides authorization for services based on policy conditions such as subscriber information, application information, time of day, and edge resource utilization.
-----	--

MSO	Multiple-service operator
-----	---------------------------

N

network topology	A map of physical equipment or logical entities in a network.
------------------	---

NTP	Network Time Protocol
-----	-----------------------

O

OSS	Operations Support System Computer systems used by telecommunications service providers, supporting processes
-----	--

O

such as maintaining network inventory, provisioning services, configuring network components, and managing faults.

OSSI

Operation Support System Interface

An interface to a “back-end” (office) system. The Configuration Management Platform includes an OSSi XML interface.

P**PCEF**

Policy and Charging Enforcement Function

Maintains rules regarding a subscriber’s use of network resources. Responds to CCR and AAR messages. Periodically sends RAR messages. All policy sessions for a given subscriber, originating anywhere in the network, must be processed by the same PCRF.

Policy and charging enforcement function

A system responsible for enforcing policies on network subscriber authentication, authorization, accounting, and mobility. A PCEF device, such as a CMTS or GGSN, communicates with a PCRF device, such as a policy server.

PCMM

PacketCable MultiMedia

PCEF

Policy and charging enforcement function

A system responsible for enforcing policies on network subscriber authentication, authorization, accounting, and mobility. A PCEF

P

device, such as a CMTS or GGSN, communicates with a PCRF device, such as a policy server.

Q

QoS

Quality of Service

Control mechanisms that guarantee a certain level of performance to a data flow.

R

RKS

Record Keeping Server

S

SANE

Secure Access to Network Elements

Verizon Wireless's central authentication and authorization system for network elements. It provides single-sign-on capability to network elements, for user of the SANE GUI client, and it allows network element vendors to use open-source, open-protocol methodologies to integrate clients into the Verizon Wireless security infrastructure.

Secure Access to Network Elements

See SANE.

server

In Policy Management, a computer running Policy Management software, or a computer providing data to a Policy Management system.

SMPP

Short Message Peer-to-Peer Protocol

An open, industry standard protocol that provides a flexible

S

data communications interface for transfer of short message data.

SNMP

Simple Network Management Protocol.

An industry-wide standard protocol used for network management. The SNMP agent maintains data variables that represent aspects of the network. These variables are called managed objects and are stored in a management information base (MIB). The SNMP protocol arranges managed objects into groups.

SOAP

Simple Object Access Protocol

SSO

Single Sign-On

V**VIP**

Virtual IP Address

Virtual IP is a layer-3 concept employed to provide HA at a host level. A VIP enables two or more IP hosts to operate in an active/standby HA manner. From the perspective of the IP network, these IP hosts appear as a single host.

X**XML**

eXtensible Markup Language

A version of the Standard Generalized Markup Language (SGML) that allows Web developers to create customized tags for additional functionality.