Oracle® Communications Policy Management

CMP Cable User Guide

Release 9.4

910-6736-001 Revision C

April 2014

ORACLE®

Oracle® Communications CMP Cable User Guide, Release 9.4

Table of Contents

List of Figures

List of Tables

Chapter 1

About This Guide

Topics:

This chapter describes the organization of the document and provides other information that could be useful to the reader.

Introduction

This guide describes how to use the Configuration Management Platform (CMP) product to configure and manage Policy Management devices in a cable network.

Conventions

The following conventions are used throughout this guide:

- **Bold text** in procedures indicates icons, buttons, links, or menu items that you click on.
- *Italic text* indicates variables.
- `Monospace text` indicates text displayed on screen.
- **`Monospace bold text`** indicates text that you enter exactly as shown.

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 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 Configuration Management Platform (CMP), which controls MPE devices and associated applications.
- *Configuring the Policy Management Topology* describes how to set the topology configuration.
- *Managing MPE 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 Application Profiles* describes how to manage application profiles.
- *Managing Traffic Profiles* describes how to manage traffic profiles.
- *Managing Media Profiles* describes how to manage media profiles.
- *Managing Service Classes* describes how to manage service classes.
- *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.
- *Understanding and Creating Policy Rules* describes policy rules, which dynamically control how an MPE device processes protocol messages as they pass through it.
- *Managing Policy Rules* describes how to manage your library of policy rules and policy groups.
- *Managing Policy Tables* describes how to manage policy tables.
- *System-Wide Reports* describes the reports available on the function of Policy Management systems in your network.

- *Upgrade Manager* describes the purpose of the Upgrade Manager GUI page and the elements found on that page.
- *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 MPE devices:

- System operators
- System administrators

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**: (This icon and text indicate the possibility of *personal injury*.)
	**Warning**: (This icon and text indicate the possibility of *equipment damage*.)
	**Caution**: (This icon and text indicate the possibility of *service interruption*.)
	**Topple**: (This icon and text indicate the possibility of *personal injury* and *equipment damage*.)

Customer Care Center

The Tekelec Customer Care Center is your initial point of contact for all product support needs. A representative takes your call or email, creates a Customer Service Request (CSR) and directs your requests to the Tekelec Technical Assistance Center (TAC). Each CSR includes an individual tracking number. Together with TAC Engineers, the representative will help you resolve your request.

The Customer Care Center is available 24 hours a day, 7 days a week, 365 days a year, and is linked to TAC Engineers around the globe.

Tekelec TAC Engineers are available to provide solutions to your technical questions and issues 7 days a week, 24 hours a day. After a CSR is issued, the TAC Engineer determines the classification of the trouble. If a critical problem exists, emergency procedures are initiated. If the problem is not critical, normal support procedures apply. A primary Technical Engineer is assigned to work on the CSR and provide a solution to the problem. The CSR is closed when the problem is resolved.

Tekelec Technical Assistance Centers are located around the globe in the following locations:

Tekelec - Global

Email (All Regions): support@tekelec.com

- **USA and Canada**

  Phone:

  1-888-FOR-TKLC or 1-888-367-8552 (toll-free, within continental USA and Canada)

  1-919-460-2150 (outside continental USA and Canada)

  TAC Regional Support Office Hours:

  8:00 a.m. through 5:00 p.m. (GMT minus 5 hours), Monday through Friday, excluding holidays

- **Caribbean and Latin America (CALA)**

  Phone:

  +1-919-460-2150

  TAC Regional Support Office Hours (except Brazil):

  10:00 a.m. through 7:00 p.m. (GMT minus 6 hours), Monday through Friday, excluding holidays

  - **Argentina**

    Phone:

    0-800-555-5246 (toll-free)

  - **Brazil**

    Phone:

    0-800-891-4341 (toll-free)

    TAC Regional Support Office Hours:

    8:00 a.m. through 5:48 p.m. (GMT minus 3 hours), Monday through Friday, excluding holidays

- **Chile**

  Phone:

  1230-020-555-5468

- **Colombia**

  Phone:

  01-800-912-0537

- **Dominican Republic**

  Phone:

  1-888-367-8552

- **Mexico**

  Phone:

  001-888-367-8552

- **Peru**

  Phone:

  0800-53-087

- **Puerto Rico**

  Phone:

  1-888-367-8552 (1-888-FOR-TKLC)

- **Venezuela**

  Phone:

  0800-176-6497

- **Europe, Middle East, and Africa**

  Regional Office Hours:

  8:30 a.m. through 5:00 p.m. (GMT), Monday through Friday, excluding holidays

  - **Signaling**

    Phone:

    +44 1784 467 804 (within UK)

  - **Software Solutions**

    Phone:

    +33 3 89 33 54 00

- **Asia**

  - **India**

    Phone:

+91-124-465-5098 or +1-919-460-2150

TAC Regional Support Office Hours:

10:00 a.m. through 7:00 p.m. (GMT plus 5 1/2 hours), Monday through Saturday, excluding holidays

- **Singapore**

  Phone:

  +65 6796 2288

  TAC Regional Support Office Hours:

  9:00 a.m. through 6:00 p.m. (GMT plus 8 hours), Monday through Friday, excluding holidays

Emergency Response

In the event of a critical service situation, emergency response is offered by the Tekelec Customer Care Center 24 hours a day, 7 days a week. 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 the Tekelec Customer Care Center.

Related Publications

The Policy Management product set includes the following publications, which provide information for the configuration and use of Policy Management products in the following environments:

Cable

- *Feature Notice*
- *Cable Release Notice*
- *Roadmap to Hardware Documentation*
- *Platform Configuration User Guide*
- *CMP Cable User Guide*

- *Bandwidth on Demand Application Manager User Guide*
- *Troubleshooting Reference Guide*
- *SNMP User Guide*
- *OSSI XML Interface Definitions Reference Guide*

Wireless

- *Feature Notice*
- *Wireless Release Notice*
- *Roadmap to Hardware Documentation*
- *Platform Configuration User Guide*
- *CMP Wireless User Guide*
- *Multi-Protocol Routing Agent User Guide*
- *Troubleshooting Reference Guide*
- *SNMP User Guide*
- *OSSI XML Interface Definitions Reference Guide*
- *Analytics Data Stream Reference*

Wireline

- *Feature Notice*
- *Wireline Release Notice*
- *Roadmap to Hardware Documentation*
- *Platform Configuration User Guide*
- *CMP Wireline User Guide*
- *Troubleshooting Reference Guide*
- *SNMP User Guide*
- *OSSI XML Interface Definitions Reference Guide*

The following documents are useful for reference:

- PCMM specifications PKT-SP-MM-I06
- PKT-SP-DQOS-I12-050812 - PacketCable™ Dynamic Quality-of-Service Specification

Locate Product Documentation on the Customer Support Site

Access to Tekelec's Customer Support site is restricted to current Tekelec customers only. This section describes how to log into the Tekelec Customer Support site and locate a document. Viewing the document requires Adobe Acrobat Reader, which can be downloaded at www.adobe.com.

1. Log into the *Tekelec Customer Support* site.

   Note: If you have not registered for this new site, click the **Register Here** link. Have your customer number available. The response time for registration requests is 24 to 48 hours.

2. Click the **Product Support** tab.
3. Use the Search field to locate a document by its part number, release number, document name, or document type. The Search field accepts both full and partial entries.
4. Click a subject folder to browse through a list of related files.
5. To download a file to your location, right-click the file name and select **Save Target As**.

Chapter 2

The Policy Management Solution

Topics:

The Policy Management Solution provides an overview of the Multimedia Policy Engine (MPE) device, which manages multiple network-based client sessions; the network in which MPE devices operate; policies; and the Configuration Management Platform (CMP) system, which controls MPE devices and associated applications.

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*.

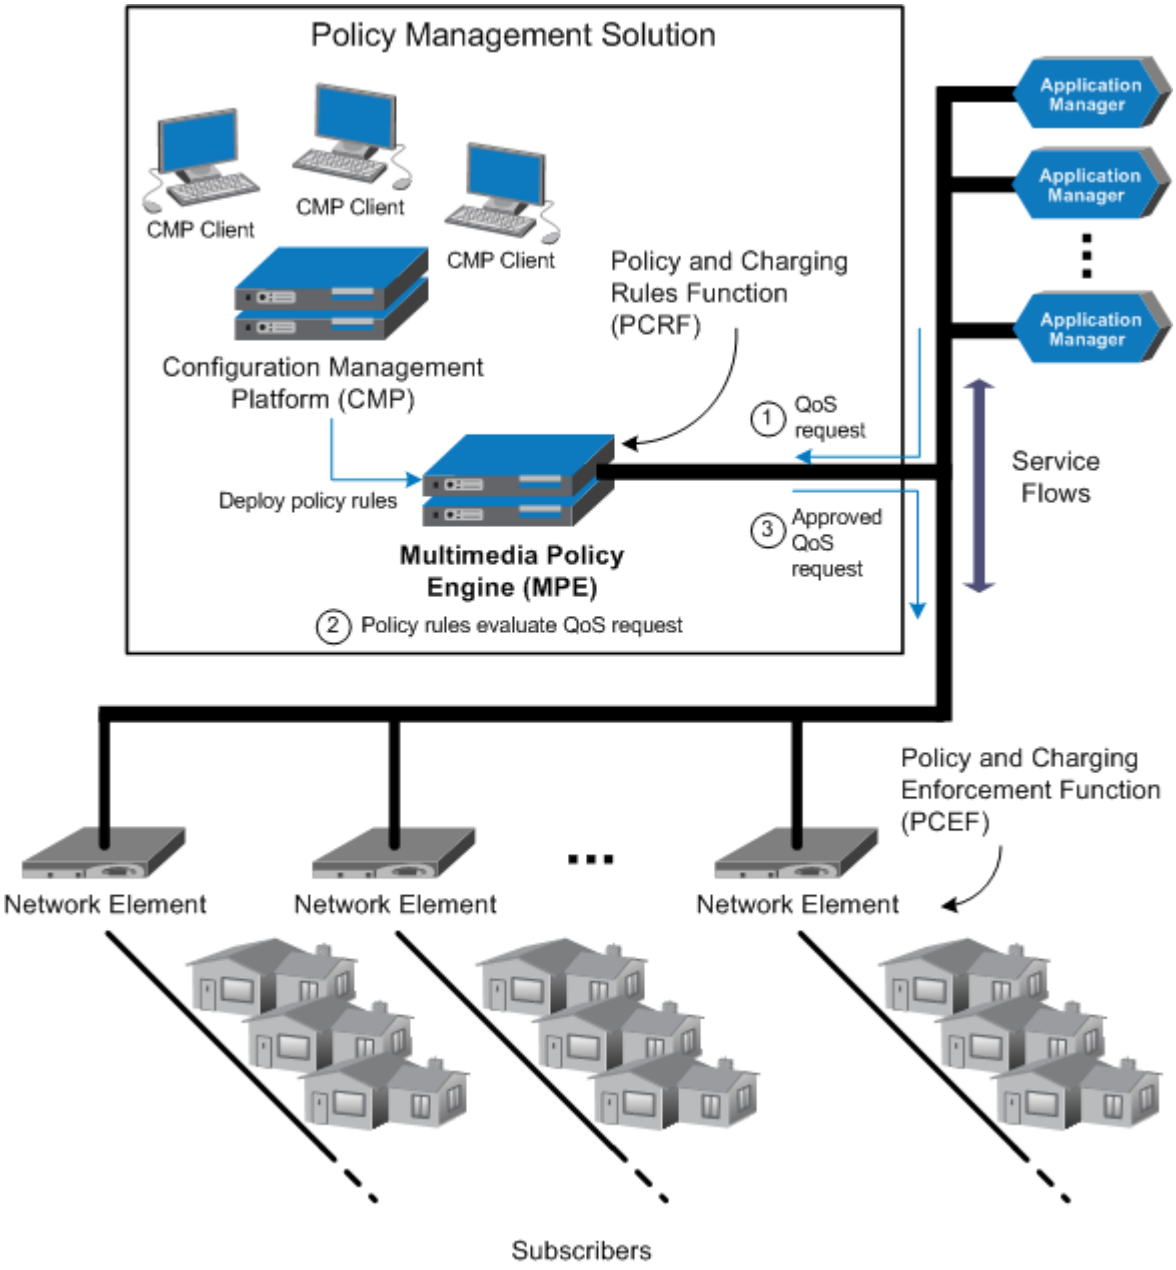

Figure 1: The CMP and MPE Devices

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.

Understanding Policy Rules

A policy rule is an if-then statement that has a set of conditions and actions. If the conditions are met, the actions are performed. You create policy rules within the CMP database, using a policy wizard that organizes a large number of conditions and actions to assist you in the construction of policy rules. Once you create policy rules, you deploy them to MPE devices.

You can combine policy rules to provide additional power and flexibility. When there are multiple policy rules, the order in which the policy rules are evaluated can also influence MPE device behavior, so the order of evaluation is also configurable through the CMP system. You can also organize policy rules into groups to simplify the management of policy rules. You can cause groups of rules to be executed.

The following are sample scenarios for which you might use policy rules:

- You can modify the contents of protocol messages using policy rules. For example, you could use a policy rule to override the requested bandwidth parameters in a request.
- You can create policy rules that track the use of resources for devices in the network and implement limits on how those resources are used. For example, some cable modems have limits on the number of dynamic flows that they can support. Using policy rules, you can ensure that a cable modem does not exceed this limit.
- Some protocols allow for the provisioning of default QoS parameters for subscribers. With these protocols, policy rules can implement subscriber tiers where different subscribers have different bandwidth available.
- You can configure policy rules to monitor the reservation of bandwidth on network elements and notify operators when an element exceeds certain threshold levels.
- In many protocols, the policy server acts as an intermediary between the Application Managers (AMs) and the QoS enforcement devices. Many of these QoS enforcement devices implement proprietary features that are activated through the use of standard (or non-standard) fields in protocol messages. Using policy rules, you can activate these proprietary features on behalf of the AMs, thus allowing them to use these features without modification.

The Bandwidth on Demand Application Manager

The Bandwidth on Demand (BoD) Application Manager 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 product, 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 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 Configuration Management Platform

The 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 web-based and supports the following features and functions:

- Configuration and management of MPE and MRABandwidth on Demand (BoD) 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 MPE and CMP software

Organizing Policy Rules

The CMP system includes features to simplify the management of multiple policy rules.

The order in which rules are evaluated is important. The CMP system lets you configure the evaluation order of policies. See *Structure and Evaluation of Policy Rules*.

The CMP system provides a policy template feature to simplify the creation of multiple policy rules that have similar conditions and actions. Once you create a policy template, you can use it to create additional rules. See *Creating a Policy Template*.

The CMP system also provides a policy rule grouping feature. Policy rules can be organized into groups and the groups can be used to simplify the process of deploying policies to MPE devices. See *Creating a Policy Group*. Policy rule groups can be executed with a single action. See *Structure and Evaluation of Policy Rules*.

Policies with similar conditions or actions can be consolidated into tabular form. See *Managing Policy Tables*.

Specifications for Using the GUI

Tekelec recommends the following:

- **Web Browsers** —
  - Mozilla Firefox release 4.0.1 or higher
  - Microsoft Internet Explorer 8.4 or higher, on Windows XP
- **Monitor** — 1024 x 768 or higher

Note: When using the CMP system for the first time, it is recommended that you change the default username and password to a self-assigned value. See *Changing a Password* for information on this procedure.

Logging In

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

Note: The CMP system also supports carrier-specific network authentication and authorization environments. For information on setting up an alternate login process, see *System Administration*.

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 *Creating a User Profile* for information about user profiles.

To log in:

1. Open a web browser and enter the IP address of 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.

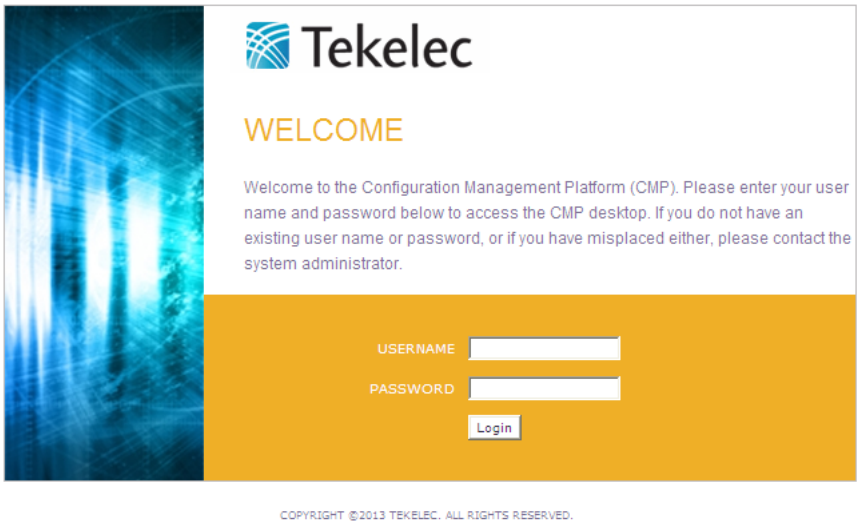

Figure 2: CMP Login Page

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 GUI* shows the structure of the CMP GUI.

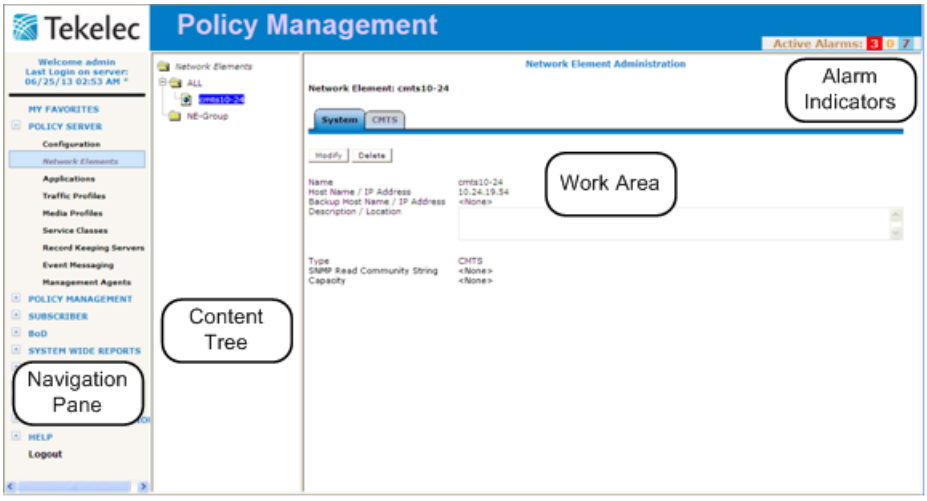

Figure 3: Structure of the CMP GUI

- **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**. Bookmarked options can be accessed from 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 on 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 in which you perform all work.
- **Alarm Indicators** — Provides visual indicators that show the number of active alarms.

GUI Icons

The CMP GUI provides icons for removing, deleting, or changing the sequential order of items displayed in a list:

 **Remove icon** — When visible in the work area, selecting the Remove icon removes an item from the group it is associated with. The item is still listed in the ALL group and any other group that it is currently associated with. For example, if you remove MPE device PS_1 from policy server group PS_Group2, PS_1 still displays in the ALL group.

 **Delete icon** — 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.

 **Move icon** — The up/down arrow icons are displayed when it is possible to change the sequential order of items in a list.

Shortcut Selection Keys

The CMP GUI supports the following standard browser techniques for selecting multiple items from a list:

- **Shift/click** — selects two or more consecutive items. To do this, select the first item, then Shift/click on a second item to select both items and all items in between.
- **Control/click** — selects two or more non-consecutive items. To do this, hold down the Ctrl key as you click on each item.

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 any user's password.

If a system administrator has configured your account for password expiration, you will receive a warning when you log in that you will need to change your password.

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** — Retype the new password.

      If your new password does not conform to the password strength rules, a validation error message appears; for example:

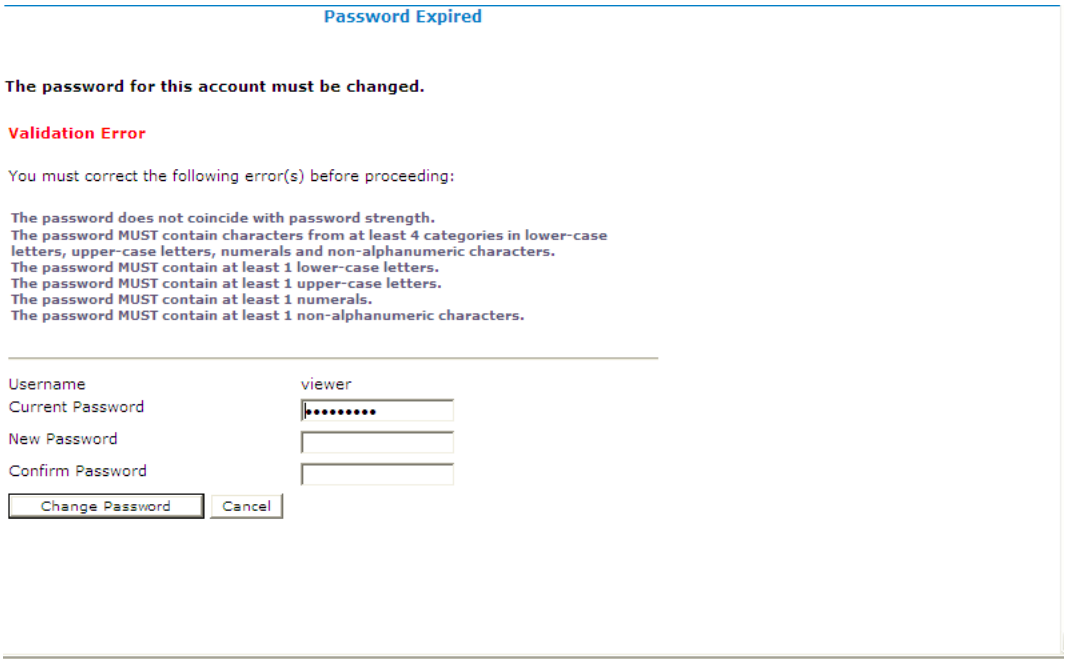

3. When you finish, click **Change Password**.

Your password is changed.

Overview of Major Tasks

The major tasks involved in using MPE devices are configuration, 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 on each one. This step is described in *Managing MPE 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 on each one. This step is described in the *Bandwidth on Demand Application Manager User Guide*.

The element and profile definition tasks you need to perform depend on what exists in your network. They can be done in any order at any time as needed. The complete set of tasks are as follows:

- Create network element profiles, including protocol options, for each network element with which the MPE devices interact. This task is described in *Managing Network Elements*.
- Specify which MPE device will interact with which network element(s). This task is described in *Managing Network Elements*.
- Create application profiles, which specify protocol information to associate each request with an application. This task is described in *Managing Application Profiles*.
- Create traffic profiles, which define default settings for protocol messages. This task is described in *Managing Traffic Profiles*.
- Create media profiles, which describe audio and video CODECs supported for Rx-to-PCMM translation. This task is described in *Managing Media Profiles*.
- Create service classes, which correspond to Data-Over-Cable Service Interface Specification (DOCSIS) traffic descriptions defined in cable modem termination systems (CMTSs). This task is described in *Managing Service Classes*.

The steps to create and deploy policy rules are required and must be done in the following order:

1. Create policy rules on the CMP device. This step is described in *Understanding and Creating Policy Rules*.
2. Deploy the policy rules from the CMP to MPE devices. This step is described in *Managing Policy Rules*.

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

1. Manage CMP users, accounts, access, authorization, and operation. This task is described in *System Administration*.

Chapter 3

Configuring the Policy Management Topology

Topics:

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 need to 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, two MA clusters, 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.

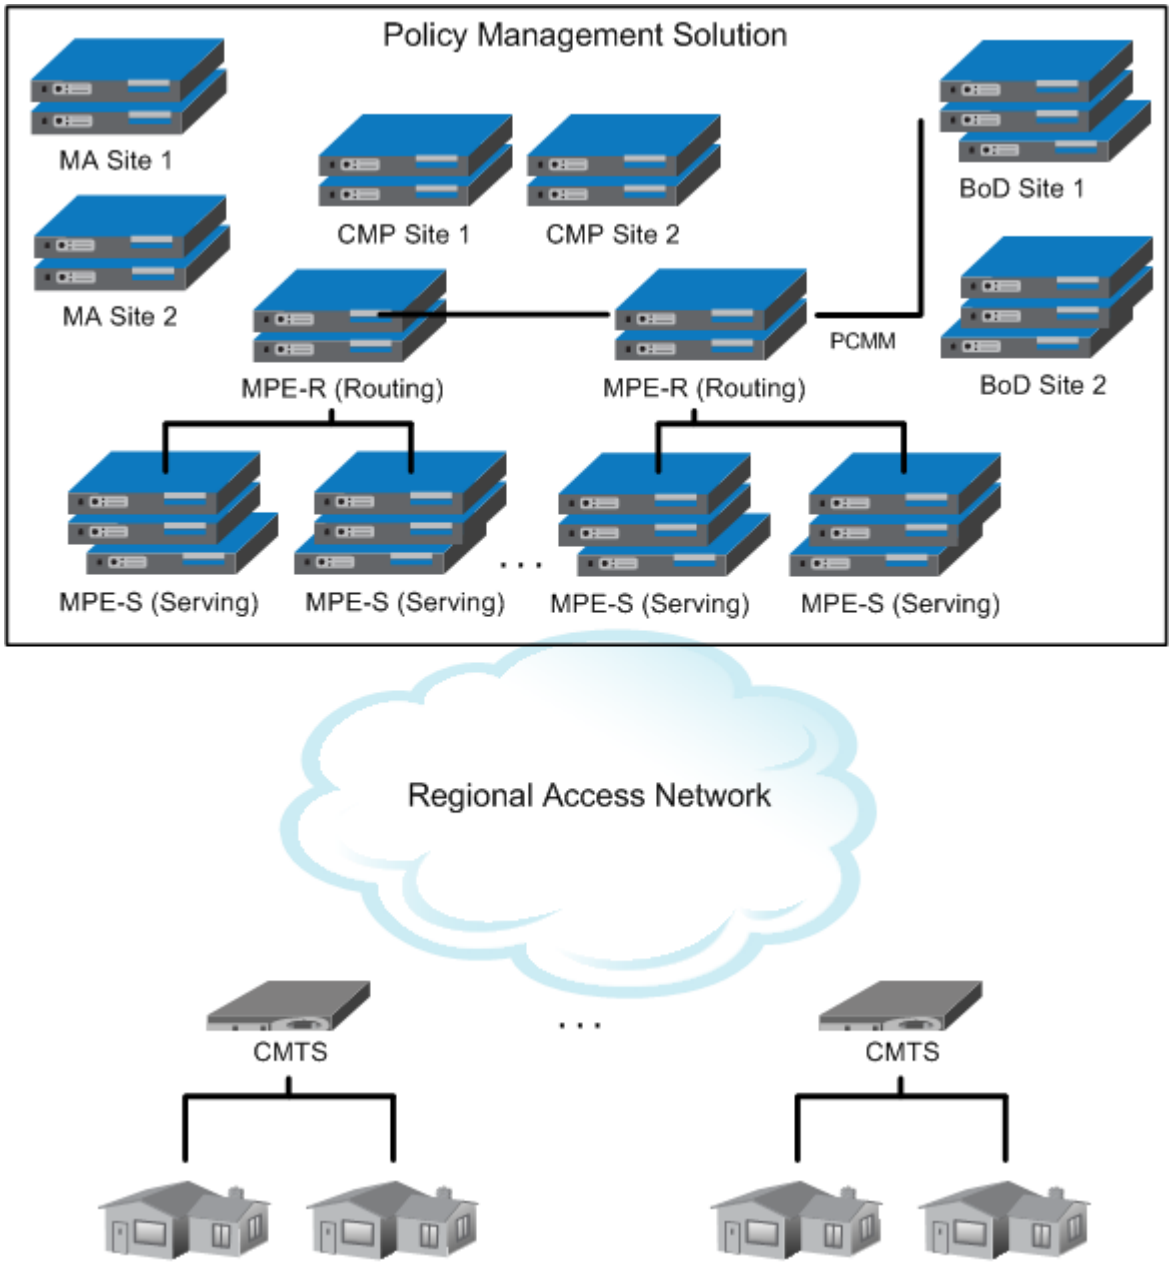

Figure 4: Policy Management Topology

High Availability

High Availability is provided for all Policy Management cluster configurations. High Availability is afforded by using two servers per cluster, an active server and a standby server. 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 perform replication, monitor server heartbeats, and merge trace logs and alarms.
2. The servers share a virtual IP (VIP) cluster address to support automatic failover.
3. The COMCOL database runtime process constantly monitors the status of both servers in the cluster.
4. If the active server fails, it instructs the standby server to take over and become the active server.

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.

Note: Some Policy Management product clusters can also include a spare server that can be in a physically separate location. The role of the spare server in high availability is described in *MPE-S and BoD Georedundancy*.

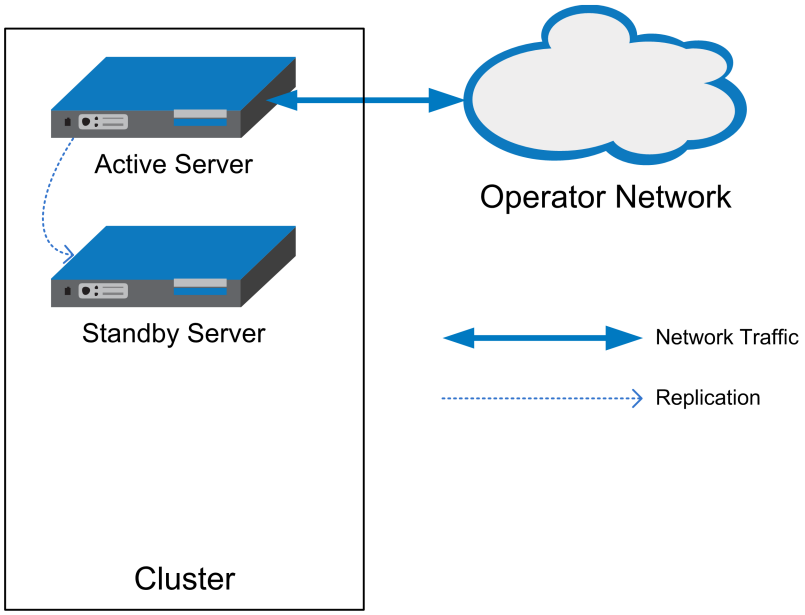

Figure 5: High Availability

MPE-S and BoD Georedundancy

As shown in *Figure 6: MPE-S or BoD 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.

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.

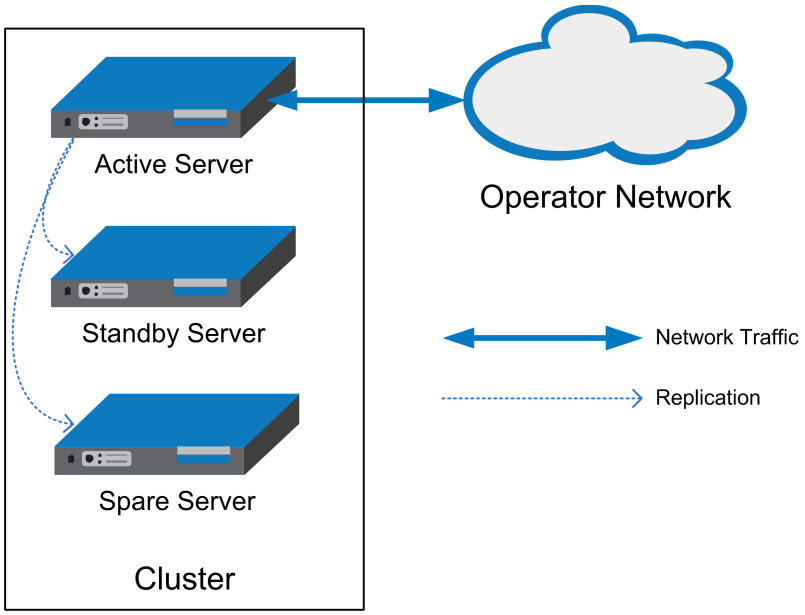

Figure 6: MPE-S or BoD Cluster with Active, Standby, and Spare Servers

The additional (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 share a virtual IP (VIP) cluster address to support automatic failover.
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.

Note: The CMP supports MPE-S and BoD georedundancy as an optional configuration mode. This mode must be configured before your CMP system will display MPE-S and BoD georedundancy options. Contact Tekelec Customer Support to change an existing CMP system to support georedundant MPE-S or BoD clusters.

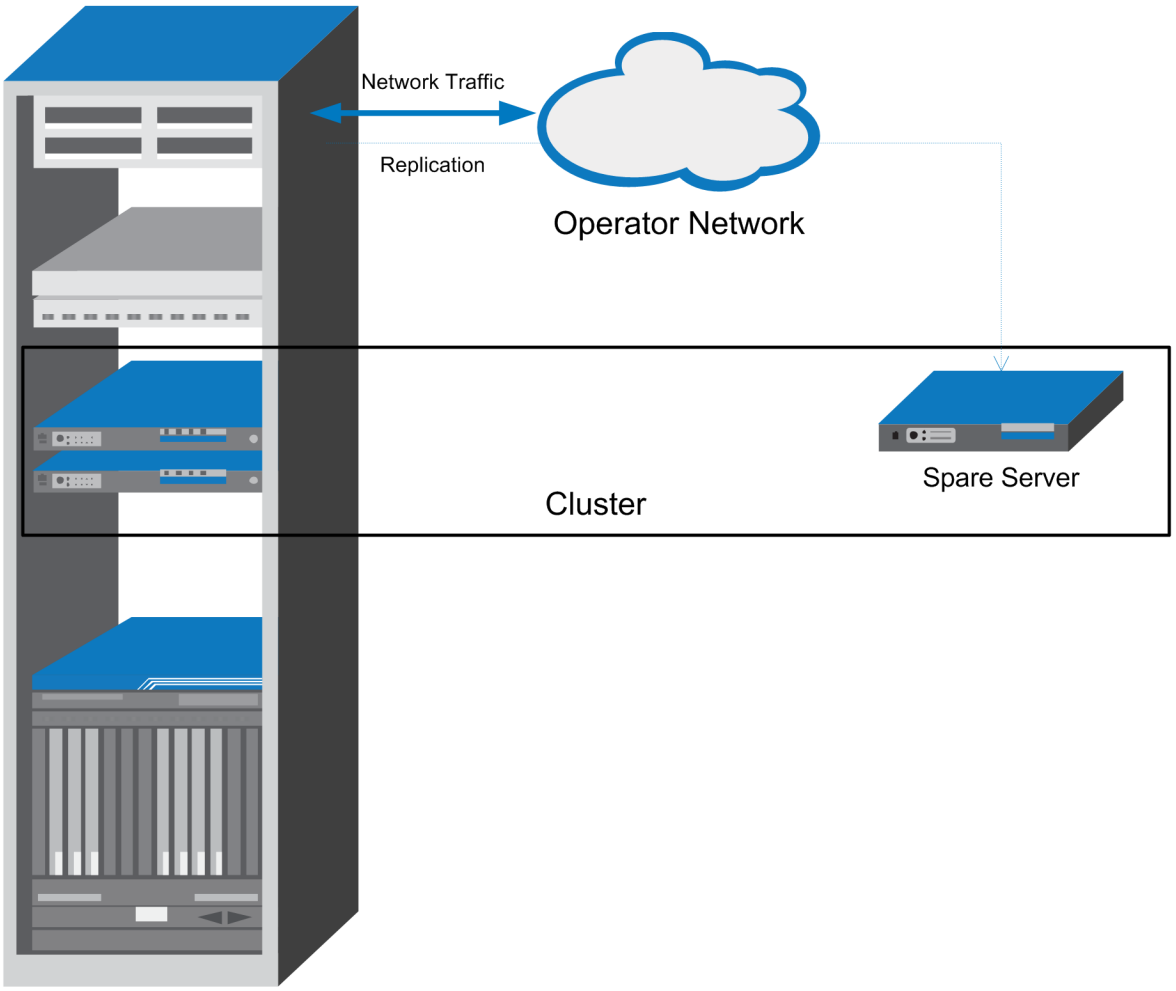

Figure 7: MPE-S or BoD Georedundant Configuration

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 topology and application data to active server of the Site 2 cluster.

The secondary cluster need not 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.

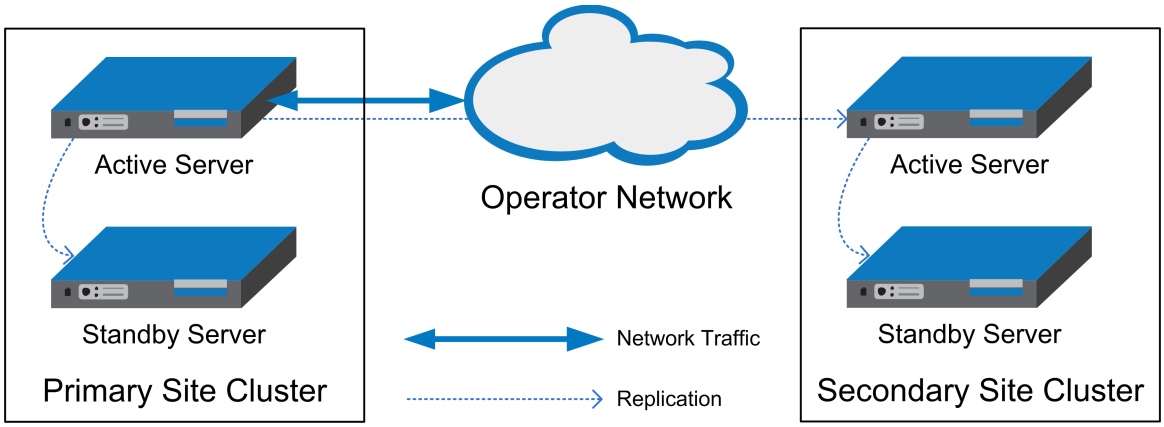

Figure 8: CMP Georedundancy

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, if necessary, 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. Once 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 on MPE-S or BoD cluster preferences under the heading "Site Preference." A cluster preference is either "Normal" or "Reverse" (or "N/A" for CMP clusters, which 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 the 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 become 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. The topology configuration, including the cluster settings, is replicated to all servers. These servers form an MPE, MA, or BoD cluster based on the topology configuration.
   b. The servers share a virtual IP (VIP) cluster address to support automatic failover.
   c. The COMCOL database runtime process constantly monitors the status of the servers in each cluster. If an active server fails, it instructs the standby server to take over and become the active server. In a georedundant cluster, if the active and standby servers in a cluster both fail, the COMCOL database 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 *Reapplying the Configuration to a Policy Server* 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 Setting**.
   The Topology 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.
   The Cluster Configuration page opens.
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 **C-Class** (the default), **C-Class(Segregated Traffic)** (for a configuration in which Signaling and OAM networks are separated onto physically separate equipment), or **RMS** (for a rack-mounted server).
   b) **Network VLAN IDs** (appears if you selected **C-Class** or **C-Class(Segregated Traffic)**) — 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) — Enter the IPv4 address and mask of the OAM VIP. The OAM VIP is the IP address the CMP uses to communicate with a Policy Management 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.

   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.
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. Enter 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. Once you define a Server A, you can select **Server-B** to enter the appropriate information for the second server of the cluster.
8. When you finish, click **Save** (or **Cancel** to discard your changes).
   You are prompted, "Active server will restart and you will be logged out." The active server restarts.

The CMP cluster topology is defined.

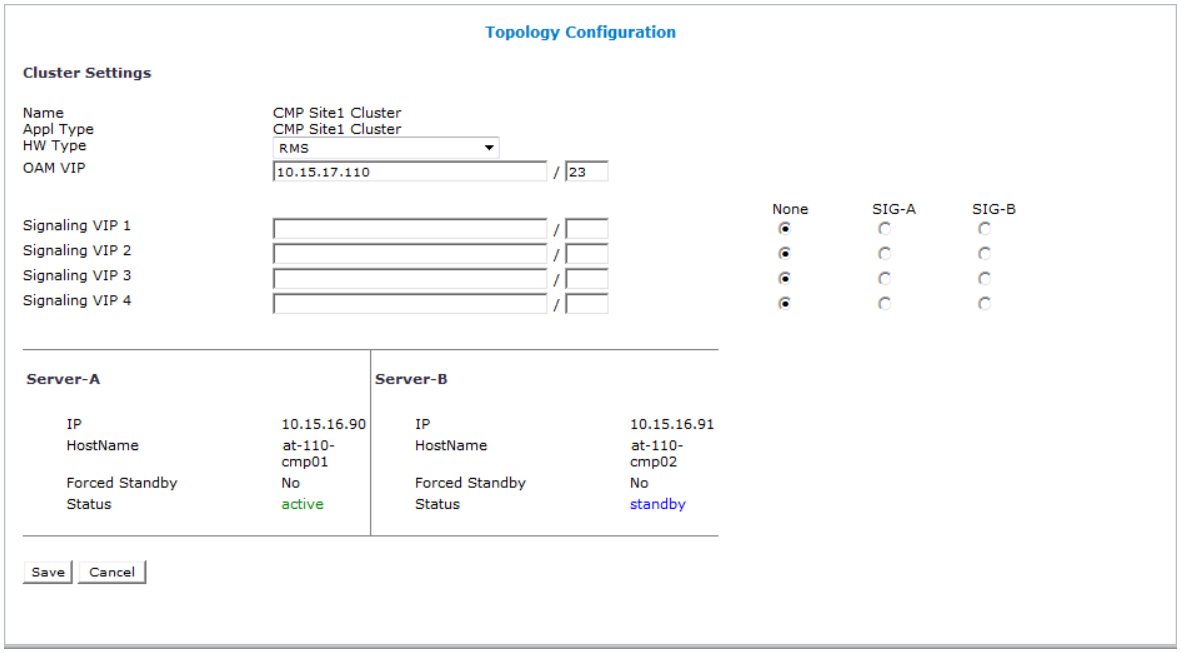

Figure 9: Cluster Settings Page for CMP Cluster

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

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

Setting Up a Site

Georedundant sites can contain one or more MPE-S or BoD clusters. 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 Topology Configuration page opens.
2. From the content tree, select the **All Sites** group.
   The Site Configuration page opens.
3. On the Site Configuration page, 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** (or **Cancel** to abandon your request).
   The site configuration is saved in the CMP database.

The site is defined.

If you need to define multiple sites, repeat steps 3 through 5 as necessary.

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 OAM IP addresses
- The MPE or BoD server IP connection is active
- The MPE or BoD application is running on at least one server

To define an MPE or BoD 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 **All Clusters** group.
   The Cluster Configuration page opens.
3. Click **Add MPE/BoD/MA Cluster**.
   The Topology Configuration page opens.
4. 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 (,).
   b) **Appl Type** — Select **MPE** (the default) or **BoD**.
   c) **Site Preference** — Select **Normal** (the default) or **Reverse**.
      This field only appears on the page if the CMP system supports georedundancy.
   d) **Primary Site** — 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.
   e) **HW Type** — Select **C-Class** (the default), **C-Class(Segregated Traffic)** (for a configuration in which Signaling and OAM networks are separated onto physically separate equipment), **HP ProLiant DL360G6/G7**, or **RMS** (for a rack-mounted server).
   f) **Network VLAN IDs** (appears if you selected **C-Class** or **C-Class(Segregated Traffic)**) — 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.
   g) **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.

h) **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 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 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. 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).

5. 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 IPv4 address of the server. Enter the standard dot-formatted IPv4 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`**).
6. Once you define Server A, you can optionally click **Add Server-B** and enter the appropriate information for the second server of the cluster.
7. (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.
8. (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.
9. When you finish, click **Save** (or **Cancel** to discard your changes).
   You are prompted, "Active server will restart." Click OK or Cancel.
10. If you are setting up multiple clusters, repeat the above steps as often as necessary.

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 *Reapplying the Configuration to a Policy Server* for more information.

For information on setting up a hierarchy of MPE-R and MPE-S clusters, see *Configuring Protocol Routing*.

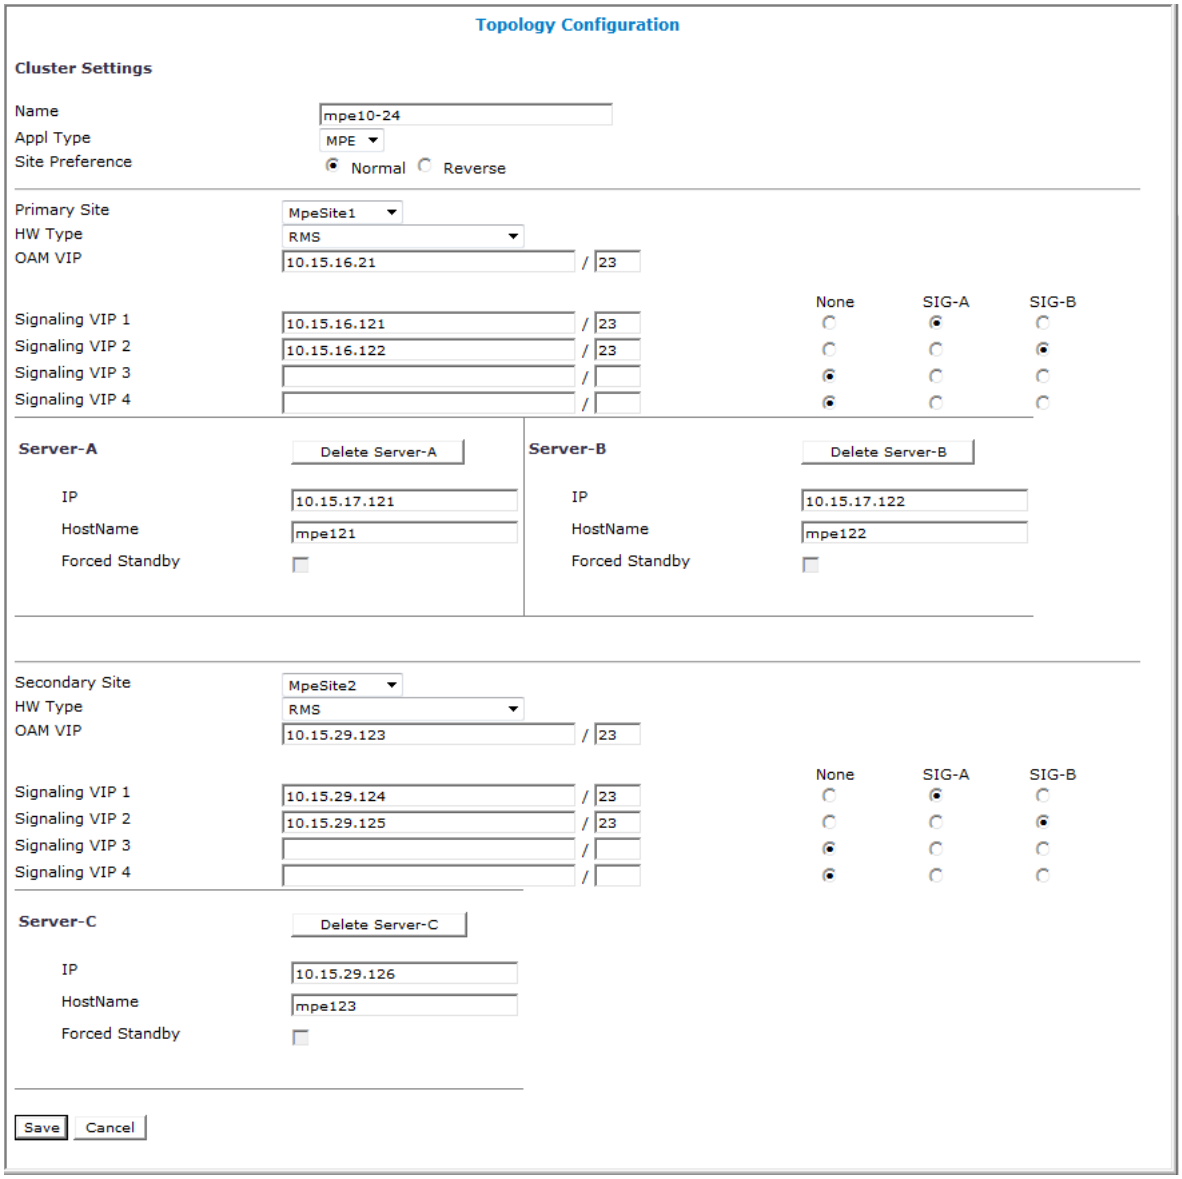

Figure 10: Cluster Settings Page for an MPE Cluster

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:

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 **All Clusters** group.
   The Cluster Configuration page opens.
3. Click **Add MPE/BoD/MA Cluster**.
   The Topology Configuration page opens.
4. Enter the following information (*Figure 11: Cluster Settings Page for an MA Cluster* shows an example):
   a) **Name** (required) — Name of the cluster. Enter up to 250 characters, excluding quotation marks (") and commas (,).
   b) **Appl Type** — Select **MA**.
   c) **HW Type** — Select **C-Class** (the default), **C-Class(Segregated Traffic)** (for a configuration in which Signaling and OAM networks are separated onto physically separate equipment), **HP ProLiant DL360G6/G7**, or **RMS** (for a rack-mounted server).
   d) **Network VLAN IDs** (appears if you selected **C-Class** or **C-Class(Segregated Traffic)**) — 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.
   e) **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.

   f) **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. For a CMP cluster, the Signaling VIP is optional, but for an MA cluster, at least one signaling VIP is required (either SIG-A or SIG-B).
5. 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 IPv4 address of the server. Enter the standard dot-formatted IPv4 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`**).
6. Once you define Server A, you can optionally click **Add Server-B** and enter the appropriate information for the second server of the cluster.
7. When you finish, click **Save** (or **Cancel** to discard your changes).
   You are prompted, "Active server will restart." Click OK or Cancel.
8. If you are setting up multiple clusters, repeat the above steps as often as necessary.

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 *Reapplying the Configuration to a Policy Server* for more information.

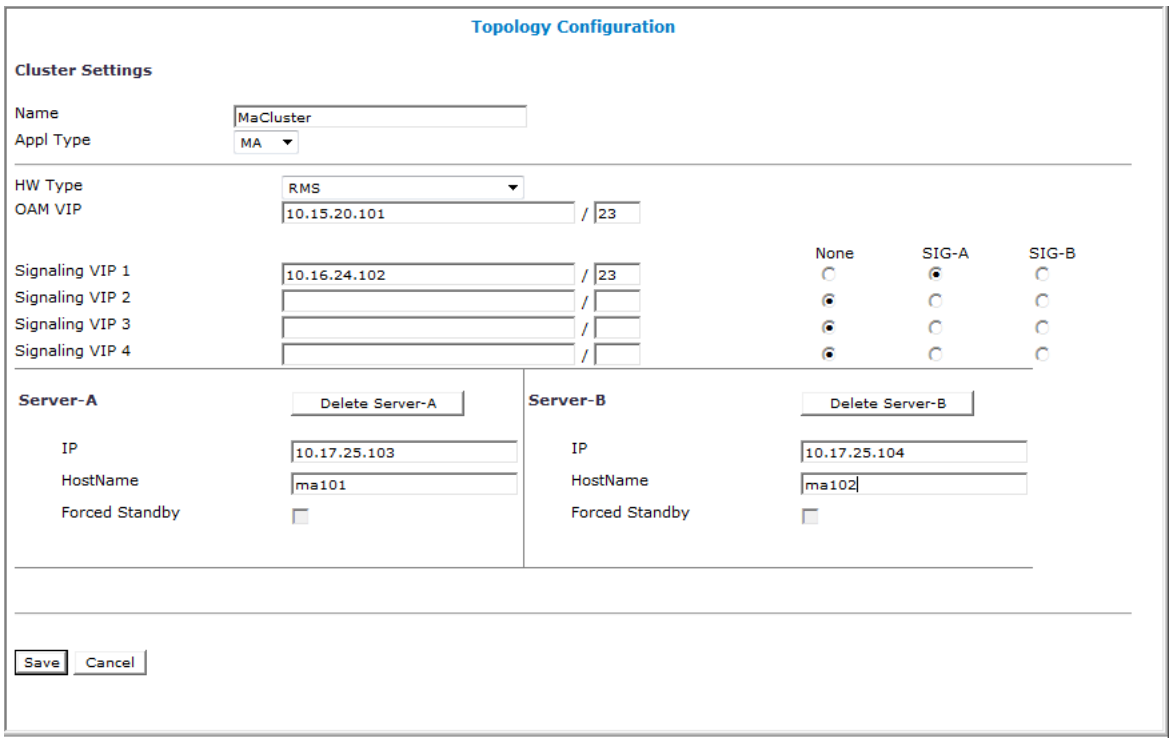

Figure 11: Cluster Settings Page for an MA Cluster

Modifying the Topology

Once the topology is configured, you can change it as necessary, to correct errors, add a server to a cluster, define new clusters, or 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 the topology is described in the following topics:

- *Modifying a Site*
- *Removing a Site from the Topology*
- *Modifying an MPE, MA, or BoD Cluster*
- *Modifying a CMP Cluster*
- *Removing a Cluster from the Topology*
- *Reversing Cluster Preference*
- *Demoting a CMP Cluster*
- *Forcing a Server into Standby Status*

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. On the Site Configuration page, click **Modify**.
   The Modify Site page opens.
4. Modify site information as required.
   For a description of the fields contained on this page, see *Setting Up a Site*.
5. When you finish, click **Save** (or **Cancel** to discard your changes).

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 it is not referenced by an MPE or BoD cluster. Once it is in use by a cluster, if you try to delete it, you are prompted, "*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 Topology 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 the Delete icon, located to the right of the site you wish to delete.
   - From the content tree, select the site and click **Delete**.

   You are prompted, "Are you sure you want to delete this Site?"
4. Click **Delete** (or **Cancel** to abandon your request).
   The page closes.

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 you want to modify.
   The Topology Configuration page opens, displaying information about the cluster.
3. On the Topology Configuration page, click the appropriate 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 appropriate fields on the Topology Configuration page become editable.

4. Make changes as required.

   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 will restart.

5. When you finish, click **Save** (or **Cancel** to discard your changes).
   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** (or **Cancel** to discard your changes).

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 Topology Configuration page opens.
2. From the content tree, select the cluster you want to modify.
   The Topology Configuration page opens, displaying information about the cluster.
3. On the Topology Configuration page, click the appropriate 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 appropriate fields on the Topology Configuration page become editable. For information on configurable fields, see .

4. Make changes as required.

   You must make changes to each section individually. You can remove either server from the cluster, but not both. You can select **Forced Standby** on either server of the cluster, but not both, and not at all if the cluster has only one server.

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

5. When you finish, click **Save** (or **Cancel** to discard your changes).
   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** (or **Cancel** to discard your changes).

The cluster is modified. You can determine if there is a topology mismatch by using the System tab of 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 cluster settings table listing information about the clusters defined in the topology.
3. In the topology configuration table, in the row listing the cluster you want to remove, click **Delete**.
   You are prompted, "Are you sure you want to delete this Cluster?"
4. Click **Delete** (or **Cancel** to abandon your request).
   The page closes.

The cluster is removed from the topology.

Reversing Cluster Preference

You can change the preference, or predilection, of the servers in a cluster to be active or spare.

To reverse cluster preference:

1. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.
   The Topology 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** (or **Cancel** to abandon your change).

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 Topology 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."
4. From the **Platform Setting** section of the navigation pane, select **Topology Settings**.
   The Topology 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 both steps 2 and 4, stop this procedure and do not try to change the current active georedundant cluster. Contact Tekelec Support before proceeding.

5. Return to the browser window logged in to the primary CMP cluster.
   You should still be on the Topology 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** (or **Cancel** to abandon your request).
   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 Topology 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** (or **Cancel** to abandon your request).
    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 Topology Configuration page opens, displaying a cluster settings table listing information about the clusters defined in the topology. The 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 Topology 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 will 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**, as appropriate.
   - 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. Click **Save** (or **Cancel** to abandon your request).
   The page closes.

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.

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 is displayed.

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

The SNMP Settings attributes are displayed.

3. Click **Modify**.

   The **SNMP Settings** page opens.

4. Edit the settings that need to be entered or changed.
5. When you finish, click **Save** (or **Cancel** to discard your changes).

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

Table 2: SNMP Attributes

Field Name	Description
Manager 1-5	SNMP Manager to receive traps and send SNMP requests. Each Manager field can be filled as either a valid host name or an IPv4 address. A hostname should include only alphanumeric characters. Maximum length is 20 characters, and it is not case-sensitive. This field can also be an IP address. An IP address should be in a standard dot-formatted IP address string. By default, these fields are empty. **Note:** The IPv6 address is not supported.
Enabled Versions	Supported SNMP versions: • SNMPv2c • SNMPv3 • SNMPv2c and SNMPv3 (default)
Traps Enabled	Enable sending SNMPv2 traps (default is box check marked) Disable sending SNMPv2 traps (box not check marked)
Traps from individual Servers	Enable sending traps from an individual server (box check marked). Send traps only from the active CMP system (default is box not check marked)
SNMPv2c Community Name	The SNMP read-write community string. 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”. The default value is “snmppublic”.

Field Name	Description
SNMPv3 Engine ID	Configured Engine ID for SNMPv3. 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. The default is no value (empty).
SNMPv3 Security Level	SNMPv3 Authentication and Privacy options. 1. "No Auth No Priv" - Authenticate using the Username. No Privacy. 2. "Auth No Priv" - Authentication using MD5 or SHA1 protocol. 3. "Auth Priv" - Authenticate using MD5 or SHA1 protocol. Encrypt using the AES and DES protocol. The default value is "Auth Priv".
SNMPv3 Authentication Type	Authentication protocol for SNMPv3. Options are: 1. "SHA-1" - Use Secure Hash Algorithm authentication. 2. "MD5" - Use Message Digest authentication. The default value is "SHA-1".
SNMPv3 Privacy Type	Privacy Protocol for SNMPv3. Options are: 1. "AES": Use Advanced Encryption Standard privacy. 2. "DES": Use Data Encryption Standard privacy. The default value is "AES".
SNMPv3 Username	The SNMPv3 User Name. The field is required if SNMPv3 is enabled. The name must contain alphanumeric characters and cannot not exceed 32 characters in length. The default value is "TekSNMPUser."
SNMPv3 Password	Authentication password for SNMPv3. This value is also used for msgPrivacyParameters. The field is required If SNMPv3 is enabled. The length of the password must be between 8 and 64 characters; it can include any character. The default value is "snmpv3password".

Defining Global Configuration Settings

This section describes how to configure global CMP settings.

Setting IPv6 Settings

You can define whether aggregation and/or filtering for IPv6 prefixes is enabled. Aggregation allows multiple IPv6 prefixes to be aggregated into a single entry. Filtering allows IPv6 prefixes that match the configured criteria to be discarded before data is routed to the CMP system and MPE devices. Both functions allow reduction in the data set that is handled.

IPv6 prefix filters and aggregation configuration changes are shown in the audit log (see *Viewing the Audit Log*). Prefix filtering and aggregation functionality for each CMTS and the net result of the functionality for all CMTSs are shown in the trace log (see *The Trace Log*).

To change IPv6 subnet settings, do the following:

1. From the **Global Configuration** section of the navigation pane, select **Global Configuration Settings**.
   The content tree displays a list of global configuration settings.
2. From the content tree, select the **IPv6 Subnet Settings** folder.
   The IPv6 Subnet Settings page opens in the work group area.
3. From the IPv6 Subnet Settings page, click **Modify**.
4. Enter values for the IPv6 settings:
   a) **IPv6 Subnet Aggregation Enable** — Click to enable the IPv6 aggregation functionality.
   b) **IPv6 Subnet Filtering Enable** — Click to enable the IPv6 filtering functionality.
      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. The IP string must be a valid IPv6 address and is case insensitive.

         To filter out all IPv6 prefixes, enter * for the IP address and leave the prefix length field blank.

         To filter out all prefixes with a specific prefix length, enter * for the IP address and the appropriate value for the prefix length.
      2. Click **Add** to add the IPv6 address and prefix length.

         The IPv6 address and prefix length are added to the list of addresses.
      3. To remove an IPv6 address/prefix length from the list, select the IPv6 address/prefix length, and click **Delete**. Click **Delete All** to remove all addresses and prefix lengths from the list.

5. When you finish, click **Save** (or **Cancel** to discard your changes).

   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 IPv6 subnet settings are configured.◀

Setting Stats Settings

You can define when and how measurement statistic values are reset.

To change stats settings, do the following:

1. From the **Global Configuration** section of the navigation pane, select **Global Configuration Settings**.
   The content tree displays a list of global configuration settings.
2. From the content tree, select the **Stats Settings** folder.
   The Stats Settings page opens in the work group area.
3. On the Stats Settings page, click **Modify**.
   The Modify Stats Settings page opens.
4. Enter values for the configuration attributes:
   a) **Stats Reset Configuration** — From the pulldown menu, select **Manual** or **Interval**. 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 the Stats Collection Period variable. 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 a better idea of the performance of the Policy Management system at specific times of day. The default value is Manual.
   b) **Stats Collection Period** — When the Stats Reset Configuration variable is set to Interval, specify the time interval to use from the pulldown menu. Options are 5, 10, 15, 20, 30, and 60 minutes.
5. When you finish, click **Save** (or **Cancel** to discard your changes).
   The Stats Settings page closes.

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

The Stats Settings attributes are configured.

Chapter 4

Managing MPE Devices

Topics:

Managing MPE 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. Once 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 *Deploying a Policy or Policy Group to MPE Devices*.

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 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. On the Policy Server Administration page, 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:
      - Is 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.
   e) **Type** — Defines the policy server type:
      - **Tekelec** (the 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 Tekelec policy server.
5. When you finish, click **Save** (or **Cancel** to discard your changes).
   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** — 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. When you finish your modifications, click **Save** (or **Cancel** to discard your changes).

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:

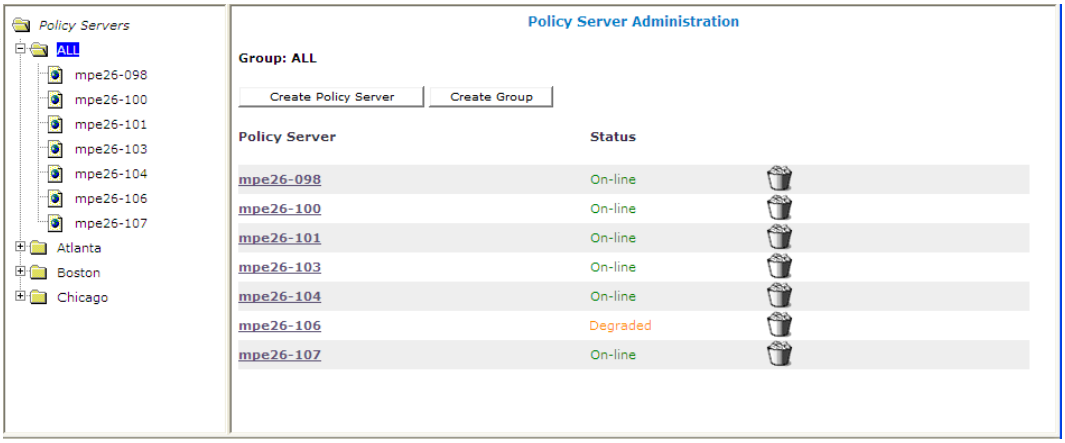

3. Use one of the following methods to select the MPE device profile to delete:
   - From the work area, click the **Delete** icon located next to the MPE device profile you want to delete.
   - From the policy server group tree, select the MPE device; the Policy Server Administration page opens. Click the System tab; the System tab opens. Click **Delete**.

   You are prompted, "Are you sure you want to delete this Policy Server?"
4. Click **OK** to delete the MPE device profile (or **Cancel** to cancel the request).
   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. On the Policy Server Administration page, select the **Policy Server** tab.
   The current configuration options are displayed.
4. Click **Modify** and define options as necessary.
   Table 3: Policy Server Protocol Configuration Options defines available options. (The options you see may vary depending on the mode in which your system is configured.)
5. When you finish, click **Save** (or **Cancel** to discard your changes).

Table 3: Policy Server Protocol Configuration Options

Attribute	Description
Associations	

Attribute	Description
Applications	The applications associated with this MPE device. To modify this list, click **Manage**.
Network Elements	The network elements associated with this MPE device. To modify this list, click **Manage**.
Network Element Groups	The network element groups associated with this MPE device. To modify this list, select or deselect groups.
Configuration	
Management Agent	Visible if your network contains management agents. For more information, see *Managing Traffic Profiles*.
PCMM	
Validate the application	When enabled, all PCMM requests are checked to ensure that there is an application defined that can be associated with the request (typically by matching the application manager ID, or AMID, in the request). If there is no such application, the MPE device rejects the request.
Validate the service class	When enabled, any PCMM requests that refer to a Service Class Name in a traffic profile are checked to ensure that the service class is known to be valid for the destination CMTS.
Validate the gate ID	When enabled, all PCMM requests that refer to an existing gate are checked against the MPE device's database of existing gates. If the request refers to a gate ID that does not exist, then it is rejected without forwarding to the CMTS.
Validate traffic profile envelopes	When enabled, all PCMM requests that include traffic profiles are checked to ensure that the parameters for the Authorized, Reserved, and Committed envelopes are valid, as defined in the PCMM Specification.
Enable MGPI	Enable Multiple Grants Per Interval (MGPI) for all Rx applications. By default, not selected (that is, MGPI is disabled). For more information, see *Configuring Protocol Routing*. **Note:** If MGPI is enabled, flow aggregation begins with the next call that creates or modifies an application flow.
Upstream Flow Limit for Triggering MGPI	The number of upstream service flows above which MGPI is triggered. A value from 1 through 99; the default is 8 flows.
Maximum Number of Grants per Interval	The maximum number of grants per interval allowed on one gate (that is, the maximum number of sub-flows aggregated on one service flow). A value from 2 through 99; the default is 8 grants.

Attribute	Description
Default Local Time Mode	Select the time used within a user's session from the pulldown menu: **System Local Time** to use the local time of the MPE device (the 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	
Diameter Realm	The domain of responsibility (for example, **`galactel.com`**) for the MPE device.
Diameter Identity	The fully qualified domain name (FQDN) of the MPE device (for example, **`mpe3.galactel.com`**).
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.
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	
	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: **Default**, **Audio**, **Video**, **Data**, **Application**, **Control**, **Text**, **Message**, and **Other**. (The **Default** profile is used when a profile for a media type is not defined.) To specify values, create Diameter profiles in the general profile configuration.
Load Shedding Configuration	
Enabled	Select to enable Call Admission Control on managed MPE devices, which implements and enforces load shedding. You can enable or disable load shedding on individual MPE devices.

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 configuration key without first consulting with Tekelec Technical Support.

To configure an advanced setting 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. On the **Policy Server Administration** page, select the **Policy Server** tab.
   The Policy Server configuration settings are displayed.
4. Click **Advanced**.
   Advanced configuration settings are displayed and can be edited.

   - **Session Clean Up Settings**

     Table 4: Session Clean Up Options

Attribute	Description
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 pulldown menu. 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 Tekelec Technical Support.
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.

- **Other Advanced Configuration Settings**— Configuration Key changes are made using this table.
  - **To add a key to the table** — Click **Add**; the Add Configuration Key Value window opens. Enter the following values:
    - **Configuration Key** — The attribute to set
    - **Value** — The attribute value

    For example:

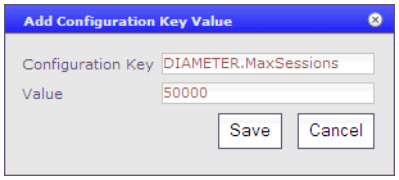

    When you finish, click **Save** (or **Cancel** to discard your changes).

    Caution: There is no input validation on keys or values. Also, if you overwrite a setting that is already configurable using the CMP GUI, the value adopted by the MPE device is undetermined.

  - **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** (or **Cancel** to discard your changes).
  - **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** (or **Cancel** to discard your changes).
  - **To delete a key from the table** — Select an existing key in the table and click **Delete**; you are prompted, "Are you sure you want to delete the selected Configuration Key Value(s)?" Click **Delete** to remove the key (or **Cancel** to cancel your request).

5. When finished making changes, click **Save** (or **Cancel** to discard changes).
   The settings are applied to the selected MPE 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 desired policy server.
   The Policy Server Administration page opens.
3. On the Policy Server Administration page, select the Data Sources tab.
   The current data sources are displayed, listing the administrative state, name, role, type, primary host, and secondary host.
4. To modify the list of data sources, click **Modify**.

The Modify Data Sources page opens. The functions available from this table are as follows:

- **To add a data source to the table** — Select the data source type from the Add pulldown list; the appropriate Add Data Source window opens. Configure values as appropriate.
- **To clone a data source in the table** — Select an existing data source in the table and click **Clone**; the Clone Data Source window opens with that data source's information filled in. Make changes as required.
- **To edit a data source in the table** — Select the data source in the table and click **Edit**; the Edit Data Source window opens, displaying the data source's information. Change the configuration values as required.
- **To delete a data source from the table** — Select the data source in the table and click **Delete**; you are prompted, "Are you sure you want to delete the selected data source(s)?" Click **Delete** to remove the data source entry (or **Cancel** to cancel your request).
- **To change the order of the list** — If you define multiple data sources, they are searched in the order displayed in this list. To change the order, select a data source and click the Up or Down arrows.

When you finish, click **Save** (or **Cancel** to discard your changes).

5. 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.
6. When you finish, click **Save** (or **Cancel** to discard your changes).

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:

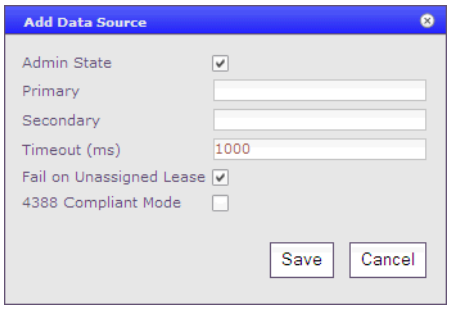

1. **Admin State** — Select to enable this data source.
   Selected by default.
2. **Primary** — FQDN or IP address in IPv4 or IPv6 format of primary DHCP server.
3. **Secondary** — FQDN or IP address in IPv4 or IPv6 format of secondary DHCP server.
4. **Timeout (ms)** — Length of time to wait before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

When you finish, click **Save** (or **Cancel** to abandon your changes). The DHCP data source is defined.

Policy Server Groups

For organizational purposes, you can aggregate the 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. On the Policy Server Administration page, click **Create Group**.
   The Create Group page opens.
4. Enter the name of the new policy server group.
   The name cannot contain quotation marks (") or commas (,).

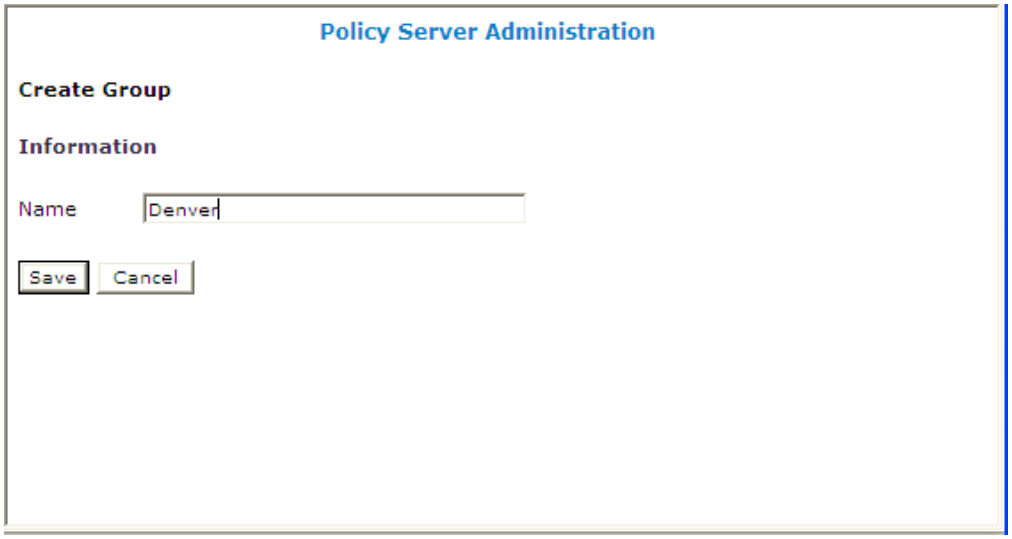

5. When you finish, click **Save** (or **Cancel** to discard your changes).
   The new group appears in the content tree.

You have created a policy server group.

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 desired policy server group.
   The Policy Server Administration page opens in the work area, displaying the contents of the selected policy server group.
3. On the Policy Server Administration page, click **Add Policy Server**.
   The Add Policy Server page opens, displaying the policy servers not already part of the group.
4. Click on the policy server you want to add; use Ctrl or Shift-Ctrl to select multiple policy servers.
5. When you finish, click **Save** (or **Cancel** to cancel the request).

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 desired policy server group.
   The Policy Server Administration page opens in the work area, displaying the contents of the selected policy server group.
3. On the Policy Server Administration page, 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** (or **Cancel** to discard your changes).
   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 desired policy server group or sub-group.
   The Policy Server Administration page opens in the work area.
3. On the Policy Server Administration page, 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** (or **Cancel** to cancel the request).
   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 desired 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 desired policy server profile using one of the following methods:

   Note: The policy server is removed immediately; there is no confirmation message.

   - Click the Remove (scissors) 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. Click the System tab; the System tab opens. 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 desired 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**.
   You are prompted, "Are you sure you want to delete this Group?"
4. Click **OK** to delete the group (or **Cancel** to cancel the request).

The policy group is deleted.

Reapplying the Configuration to a Policy Server

The CMP system lets you reapply the configuration to each MPE device. When you reapply the configuration, the CMP system completely reconfigures the MPE device with topology information (such as network elements), ensuring that the MPE device configuration matches the data in the CMP database. 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 MPE device 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 the failed servers' status.
- After upgrading the software on an MPE device, Tekelec recommends that you reapply the configuration from the CMP system to ensure that the upgraded MPE device and the CMP database are synchronized.
- There are situations in which it is possible for an MPE device configuration to 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 MPE device. If such a condition occurs, the CMP system displays the MPE device status on its System tab with the notation "Config Mismatch." You can click the notice to display a report comparing the MPE device configuration with the CMP database information. Reapplying the configuration brings the MPE device back into synchronization with the CMP database.

To reapply the configuration associated 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 **ALL** group.
   The Policy Server Administration page opens in the work area.
3. From the group **ALL**, select the desired MPE device.
   The Policy Server Administration page opens to the System tab, displaying information for that device.
4. Click **Reapply Configuration**.
   The profile information is saved to the MPE device.

The MPE device is synchronized with the CMP database.

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

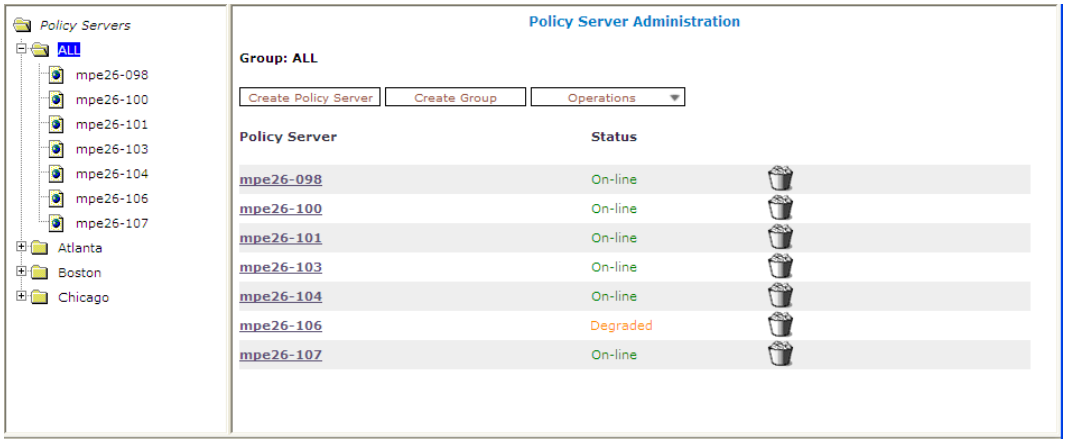

Figure 12: Group View

- **Trash can icon** — Click on the 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 the (details...) hyperlink. 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 on 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** — 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 *Configuring MPE Advanced Settings*.

The following statistics are included:

- **Ready for Cleanup** — Number of sessions that are stale (created at least 24 hours ago).
- **Removed on unknown session id** — Number of sessions removed because the session ID is no longer valid.
- **Reauthorized** — Number of sessions reauthorized.
- **Reauthorization Timeout** — Number of sessions for which the reauthorization request timed out.
- **Removed for Expiration** — Number of sessions removed.

Protocol Statistics

The Protocol Statistics section summarizes the protocol activity within the MPE device. This information is presented as a table of summary statistics for each protocol. Some protocols are broken down into sub-entries to distinguish between the different types of protocol activity.

The summary protocol statistics are the following:

- **Connections** — If the protocol is connection oriented, 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 13: Sample Protocol Statistics shows a sample.

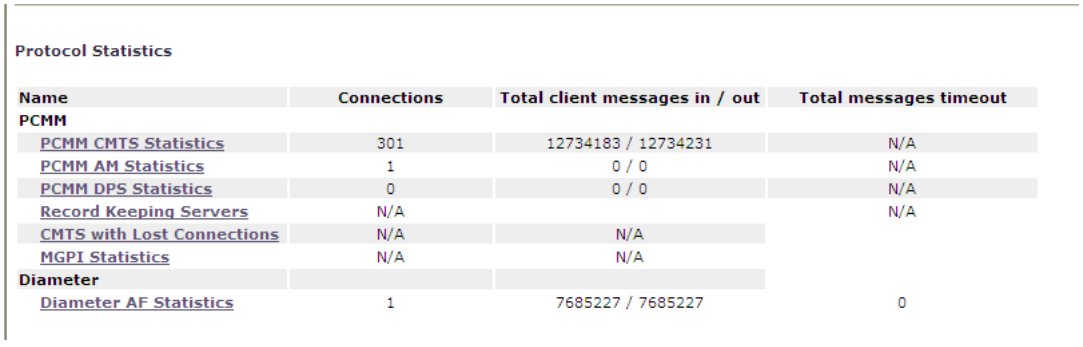

Protocol Statistics

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 13: 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.

Since many of these statistics contain detailed protocol-specific summaries of information, the specific definitions of the information that is displayed are not included here. For more specific information, see the appropriate technical specification that describes the protocol in which you are interested (see *Related Publications*).

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 resampling 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 on its 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 14: Sample Error Statistics* shows a sample.

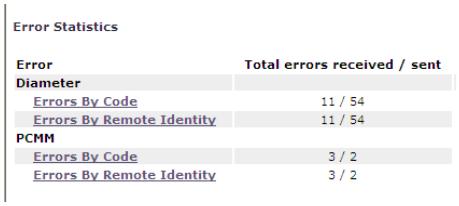

Error Statistics

Error	Total errors received / sent
Diameter	
Errors By Code	11 / 54
Errors By Remote Identity	11 / 54
PCMM	
Errors By Code	3 / 2
Errors By Remote Identity	3 / 2

Figure 14: 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.

DHCP 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 **Yes** or **No**, where **Yes** indicates that data was collected for a full interval.
- **Interval MaxSessionCount** — The highest value of the counter MaxSessionCount during the previous interval for AF and PCMM sessions.
- **Interval PCMM MaxTransactionsPerSecond** — The highest value of the PCMM transaction rate.
- **Interval Rx MaxTransactionsPerSecond** — The highest value of the Rx transaction rate.

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 5: PCMM (PacketCable MultiMedia) Protocol Statistics shows information for these protocols:

- PCMM CMTS (Cable Modem Termination System)
- PCMM AM (Application Manager)
- PCMM DPS

Table 5: 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	

Reports Display Name	OSSI XML Name
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 6: Record Keeping Servers Protocol Statistics shows information for Record Keeping Servers (RKSs).

Table 6: 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	

Reports Display Name	OSSI XML Name
Messages sent to secondary	
Ack messages received from the secondary	

Table 7: CMTS with Lost Connections Statistics shows information for individual CMTS systems with lost connections.

Table 7: CMTS with Lost Connections Statistics

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

Table 8: MGPI Statistics shows information for the MGPI protocol.

Table 8: MGPI Statistics

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

Table 9: Diameter AF (Application Function) Statistics shows information for the Diameter AF protocol.

Table 9: 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

Reports Display Name	OSSI XML Name
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
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

Reports Display Name	OSSI XML Name
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 10: Latency Statistics shows information for these statistics:

- Diameter AF
- PCMM AM
- PCMM CMTS
- PCMM DPS

Table 10: Latency Statistics

Reports Display Name	OSSI XML Name
Connections	Active Connection Count
Maximum Processing Time received / sent (ms)	Max Trans In Time\ Max Trans Out Time
Average Processing Time received / sent (ms)	Avg Trans In Time\ Avg Trans Out Time
Transactions Processed received / sent [*timeframe*] (ms)	Processing Time [0-20] ms Processing Time [20-40] ms Processing Time [40-60] ms Processing Time [60-80] ms Processing Time [80-100] ms Processing Time [100-120] ms Processing Time [120-140] ms Processing Time [140-160] ms

Reports Display Name	OSSI XML Name
	Processing Time [160-180] ms Processing Time [180-200] ms Processing Time [>200] ms

Table 11: Protocol Error Statistics shows information for these statistics:

- Diameter
- PCMM

Table 11: 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 12: Connection Error Statistics shows information for these statistics:

- Diameter
- PCMM

Table 12: Connection Error Statistics

Reports Display Name	OSSI XML Name
Total errors received	In Error Count
Total errors sent	Out Error Count
Last time for total error received	Last Error In Time
Last time for total error sent	Last Error Out Time
Last stats reset time	
Protocol Errors on each error codes	(see specific errors listed in GUI)

Table 13: KPI Interval Statistics shows information for the KPI collection interval.

Table 13: KPI Interval Statistics

Reports Display Name	OSSI XML Name
Interval StartTime	Interval Start Time
Configured Length (Seconds)	Configured Length (Seconds)

Reports Display Name	OSSI XML Name
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 14: Policy Statistics shows information for policy execution.

Table 14: 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)	
Avg Execution Time (ms)	
Processing Time Stats	
Policy write state on session termination	
Name	

Reports Display Name	OSSI XML Name
Evaluated	
Executed	
Ignored	
Total Execution Time (ms)	
Max Execution Time (ms)	
Avg Execution Time (ms)	
Processing Time Stats	

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 desired Policy Management device.
   The Policy Server Administration page opens in the work area.
3. On the Policy Server Administration page, select the **Logs** tab.

   Log information, including the log levels, is displayed. *Figure 15: Policy Server Logs Tab* shows an example. You can configure the following logs:

   - **Trace log** — Records application-level notifications.
   - **Policy Syslog** — Records policy-processing activity. Supports the standard UNIX logging system, in conformance with RFC 3164.

Figure 15: Policy Server Logs Tab

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 V7.5, 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 V7.5, 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. Select the device to view:
   - To view an MPE device, 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 appropriate Administration page opens in the work area.
3. On the Administration page, 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 the calendar icon, select the desired starting date and time, then click **Enter**.
   - **End Date/Time** — Click the calendar icon, select the desired ending date and time, then click **Enter**.
   - **Trace Code(s)** — 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 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.

   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 *Optional Policy-Processing Actions*.

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

From the Logs tab you can configure the log settings for the servers in a cluster. To configure log settings:

1. From the Logs tab, click **Modify**.
   The Modify Settings fields open in the work area.
2. In the **Modify Trace Log Settings** section of the page, configure the Trace Log Level.
   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** — Designates potentially harmful situations.
   - **Notice** — Provides messages that may be of significant interest that occur during normal operation.
   - **Info** — Designates informational messages highlighting overall progress of the application.
   - **Debug** — Designates information events of lower importance.

   

   Caution: Before changing the default logging level, consider the implications. Lowering the trace log level setting from its default value causes more notifications to be recorded in the trace log and can adversely affect performance. On the other hand, raising the log level setting causes fewer notifications to be recorded in the trace log, and could cause you to miss important notifications.

3. In the **Modify Policy Syslog Forwarding Settings** section of the page, configure the syslog forwarding settings. You can direct notifications to up to five remote systems. For each system, enter the following:
   a) **Hostname/IP Addresses** — Remote system hostname or address.

Caution: Forwarding addresses are not checked for loops. If you forward events on System A to System B, and then forward events on System B back to System A, a message flood can result, causing dropped packets.

b) **Facility** — Select from Local0 (the default) to Local7.

c) **Severity** — Filters the severity of notifications that are written to syslog:

- **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.

4. When you finish, click **Save** (or **Cancel** to discard your changes).

The log configurations are changed.

Chapter 5

Configuring Protocol Routing

Topics:

Routing enables a Policy Management device to forward requests to other Policy Management devices for further processing. The following routing messages and protocols are supported:

- PacketCable MultiMedia (PCMM) messages
- Diameter Rx messages

PCMM Routing Architectures

There are two architectures you can employ with PCMM routing: Hierarchical and Mesh.

- **Hierarchical** — In a hierarchical architecure, there is a top-level MPE cluster (an MPE-R cluster) and one or more bottom-level MPE clusters (MPE-S clusters). A PCMM message is directed to the top-level MPE cluster, which then routes the message to the appropriate MPE cluster below based on the subscriber IP address in the message.
- **Mesh** — In a mesh architecture, there is a set of two or more MPE clusters, but there is no top-level cluster. If you imagine three MPE clusters arranged in a triangle, a PCMM message coming into any one of these clusters can be forwarded out to any of the other two MPE clusters. Each cluster points to the other clusters.

In either architecture, a PCMM message is handled by the MPE cluster to which it is sent, and does not have to be forwarded. For example, in a hierarchical architecture, if a PCMM message comes into the top-level MPE cluster, and the appropriate CMTS is associated with that cluster, then the cluster handles the message itself.

Configuring PCMM Routing

Configuring PCMM routine establishes a hierarchical network of MPE-R (routing) and MPE-S systems.

To configure PCMM routing:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.
   The content tree displays a list of policy server groups; the initial group is **ALL**.
2. From the content tree, select the **ALL** group.
   The Policy Server Administration page opens in the work area.
3. From the **ALL** group, select the desired 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 16: 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 enabled, Rx traffic is routed statelessly (without translation) to other MPE devices.

c) **Downstream Policy Servers** — A list of MPE-S devices to which 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** (or **Cancel** to discard your changes).

PCMM routing is configured.

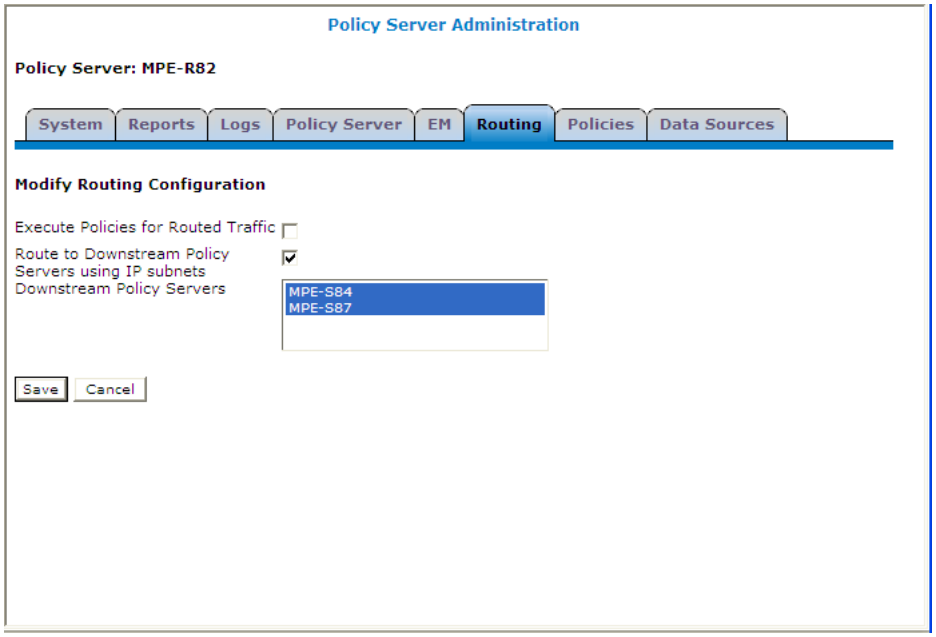

Figure 16: Modify Routing Configuration Page

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. When Rx-to-PCMM routing is desired, 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 desired 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.
6. Select **Route to Downstream Policy Servers using IP subnets**.
7. Deselect **Execute policies for Routed Traffic**.
8. When you finish, click **Save** (or **Cancel** to discard your changes).

Rx-to-PCMM routing is configured.

Chapter 6

Managing Network Elements

Topics:

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.

Understanding 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 a cable modem termination system (CMTS), a packet-switched data network (PSDN), a gateway GPRS support node (GGSN), a broadband remote access server (B-RAS), a router, a server, or a 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 **Network** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**.
2. On the Network Element Administration page, click **Create Network Element**.
   The New Network Element page opens.
3. Enter information as appropriate 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 network element's primary address 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** (the default) — Cable Modem Termination System

f) **SNMP Read Community String** — A password-like field that allows read-only access to the network element's MIBs used for SNMP polling.

If a value is not entered, SNMP data is not collected from this network element.

g) **Capacity** — The bandwidth allocated to this network element.

4. Select one or more policy servers (MPE devices) to associate with this network element.
5. To add a network element to a network element group, select the desired group (see *Adding a Network Element to a Network Element Group*).
6. When you finish, click **Save** (or **Cancel** to discard your changes).
   The network element is displayed in the Network Element Administration page.

You have created the definition for a network element.

Modifying a Network Element

To modify a network element:

1. From the **Network** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**.
2. From the content tree, select the desired network element.
   The Network Element Administration page opens in the work area.
3. On the Network Element Administration page, click **Modify**.
   The Modify Network Element page opens.
4. Modify network element information as required.
   For a description of the fields contained on this page, see *Defining a Network Element*.
5. When you finish, click **Save** (or **Cancel** to discard your changes).

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 **Network** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**.
2. From the content tree, select the **ALL** group.

   The Network Element Administration page opens in the work area, displaying all defined network elements.
3. From the work area, click the **Delete** icon, located to the right of the network element you want to delete:

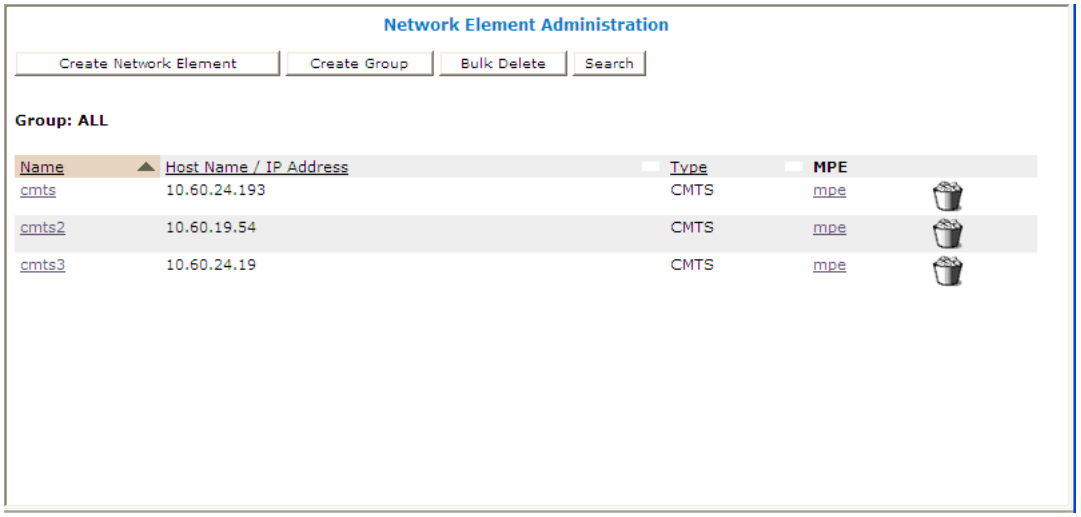

You are prompted: "Are you sure you want to delete this Network Element?"

4. Click **OK** to delete the network element (or **Cancel** to cancel the request).
   The network element is removed from the list.

You have deleted the definition of the network element.

Bulk Delete

A large network can contain a great many network elements. To perform a bulk delete of network element definitions:

1. From the **Network** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**.
2. From the content tree, select **ALL**.
   The Network Element Administration page opens in the work area.
3. On the Network Element Administration page, click **Bulk Delete**.
   The Bulk Delete Network Elements page opens.
4. Select the network elements or network element groups to delete.

   By default, the Search Pattern entry box contains an asterisk (*) to match all network elements. To search for a subset of network elements, enter a search pattern (for example, **`cmts*`**) and click **Filter**.
5. Click **Bulk Delete** (or **Cancel** to cancel the request).

The selected network element or group definition(s) are deleted from the CMP database and all associated MPE devices.

The Network Element Search Function

The Search function lets you find a specific network element within a large configuration. You can also use the function to locate all of the Cable Modem Termination Systems (CMTS) and MPE devices associated with a specified subscriber IP address or subnets. To use the network element search function:

1. From the **Network** section of the navigation pane, select **Network Elements**.

The content tree displays a list of network element groups; the initial group is **ALL**.

2. From the content tree, select **ALL**.
   The Network Element Administration page opens in the work area.
3. On the Network Element Administration page, click **Search**.
   The Network Element Search Criteria window opens.
4. Enter the desired 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** (or **Cancel** to cancel the request).

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 *Managing Scheduled Tasks* for more information.

CMTS

To configure options for a CMTS network element:

1. From 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. On the Network Element Administration page, 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 or not. If this feature is enabled, the MPE device establishes a PCMM connection to the CMTS.

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.
- **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.

5. When you finish, click **Save** (or **Cancel** to discard your changes).
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 desired MPE device.
The Policy Server Administration page opens in the work area.
3. On the Policy Server Administration page, select the **Policy Server** tab.
In the Associations section of the page, the network elements associated with this MPE device are displayed.
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:

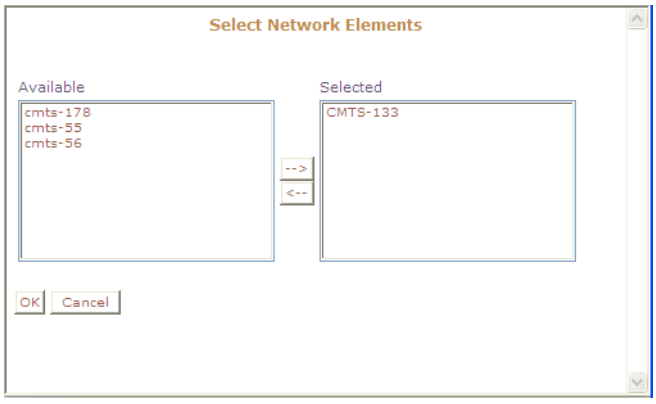

6. Select the desired network elements from the **Available** list and click **-->**.

   To disassociate a network element from the MPE device, select the network element from the **Selected** list and click **<--**. To select multiple entries, use the Ctrl and Shift keys.

7. When you finish, click **OK** (or **Cancel** to discard your changes).
   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**, located at the bottom of the page (or **Cancel** to discard your changes).

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 **Network** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**.
2. From the content tree, select the **ALL** group.
   The Network Element Administration page opens in the work area.
3. On the Network Element Administration page, 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 of the network group.
6. When you finish, click **Save** (or **Cancel** to discard your changes).
   The new group appears in the content tree.

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 **Network** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**.
2. From the content tree, select the desired network element group.

The Network Element Administration page opens in the work area, displaying the contents of the selected network element group.

3. On the Network Element Administration page, 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 17: 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** (or **Cancel** to cancel the request).

The network element is added to the selected group, and a message indicates the change; for example, "2 Network Elements were added to this group."

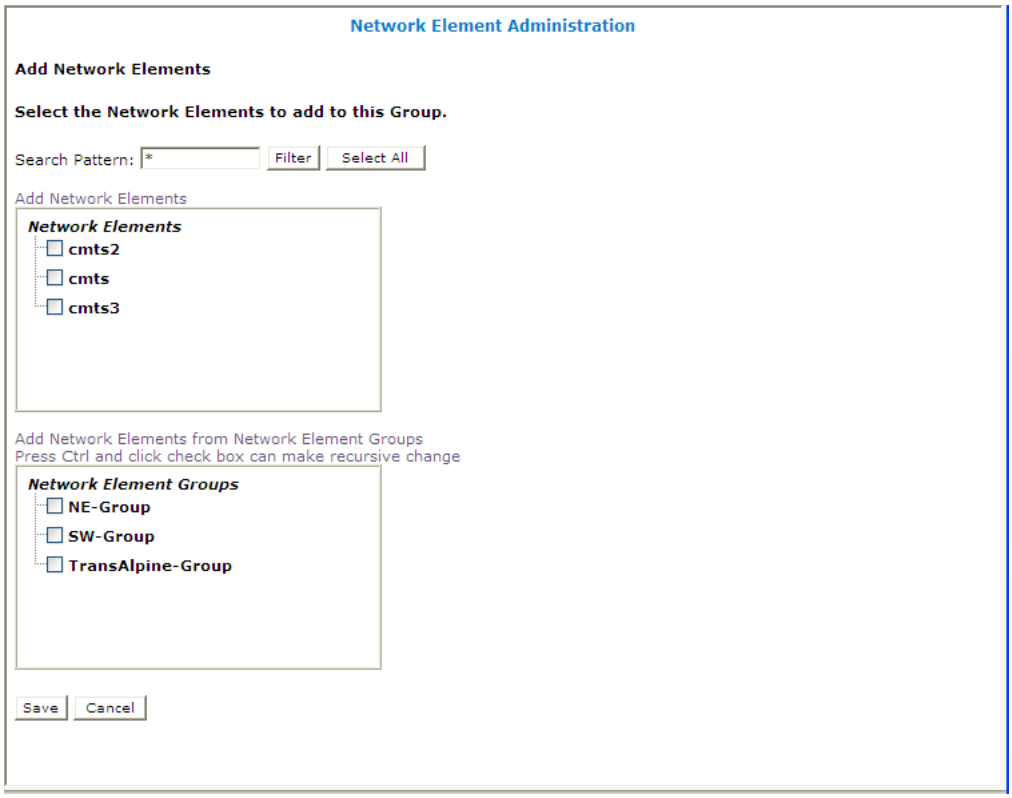

Figure 17: Add Network Element Page

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 **Network** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**.
2. From the content tree, select the desired network element group.
   The Network Element Administration page opens in the work area, displaying the contents of the selected network element group.
3. On the Network Element Administration page, 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** (or **Cancel** to discard your changes).

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 **Network** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**.
2. From the content tree, select the desired 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 Delete icon, located to the right to the network element you want to remove. You are prompted, "Are you sure you want to delete this Network Element from the group?" Click **OK** (or **Cancel** to cancel your request). The network element is removed from the group or sub-group, and a message indicates the change; for example, "Network Element deleted successfully."
   - From the content tree, select the network element; the Network Element Administration page opens. Click the System tab; the System tab opens. Click **Remove**.

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 **Network** section of the navigation pane, select **Network Elements**.

The content tree displays a list of network element groups; the initial group is **ALL**.

2. From the content tree, select the network element group or sub-group.
   The Network Element Administration page opens in the work area.
3. On the Network Element Administration page, click **Modify**.
   The Modify Group page opens.
4. Modify the name or description as desired.
5. When you finish, click **Save** (or **Cancel** to cancel the request).

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 **Network** 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. On the Network Element Administration page, click **Delete**.
   You are prompted, "Are you sure you want to delete this Group?"
4. Click **OK** to delete the group (or **Cancel** to cancel the request).

The network element group or sub-group is deleted.

Chapter 7

Managing Application Profiles

Topics:

Managing Application Profiles describes how to create and manage application profiles within the CMP system.

An application is a service provided to network subscribers for which you want to manage Quality of Service (QoS).

About Application Profiles

An application is a service provided to users of your network for which you want to manage quality of service (QoS). Examples include voice over IP (VoIP) telephony, video on demand (VoD), and gaming. Once you have defined an application profile in the CMP database, you can associate it with the MPE devices that will manage that application.

When you offer application services in your network, there are typically many servers in your network that provide that service. These servers are referred to as Application Managers or Application Servers. When these servers are establishing a session that requires quality of service they issue a request to a policy charging and rules function (PCRF).

When defining an application profile in the CMP database, you specify protocol information that is used by MPE devices to identify Application Managers and thus associate each request with its associated application. This lets the MPE device apply policy rules to the request that you have defined for the associated application.

Creating an Application Profile

To create an application profile:

1. From the **Policy Server** section of the navigation pane, select **Applications**.
   The content tree displays the **Applications** group.
2. Select the **Applications** group.
   The Application Administration page opens in the work area.
3. On the Application Administration page, click **Create Application**.
   The New Application page opens.
4. Enter the following application profile information:
   a) **General Configuration**:
      - **Name** — Name assigned to the application. The name can be up to 250 characters long and must not contain quotation marks (") or commas (,).
      - **Description/Location** (optional) — Free-form text.
      - **Connection IP Address(s)** — Enter the IP address(es), in IPv4 or IPv6 format, that are used by Application Managers for this application. To include an address in the connection list, type it and click **Add**; to remove an address from the list, select it and click **Delete**.
      - **Latency Sensitive** — Select this option if the application is latency sensitive.

   b) **Policy Servers associated with this Application**: select a policy server (MPE device) to associate it with this network element.
   c) **License Tracking**:
      - **Tracked** — Select **Yes** (the default) if the application's sessions are associated with a license. Otherwise, select **No**.
      - **Flows per Session** — Specifies the number of upstream and downstream sessions (1 to 4) that are allocated to this application.

- **License Timeout** — Specifies the duration of time for which this application requires a license. The default is 240 minutes (4 hours).

d) **PCMM**:

- **Application Manager IDs** — Enter the PCMM AMIDs that are used by Application Managers for this application. Click **Add** to define multiple values. To delete an existing value, select it from the list and click **Delete**.
- **Session Class IDs** — Enter the Session Class IDs that are used by each AM for this application. Click **Add** to define multiple values. To delete an existing value, select it from the list and click **Delete**.

e) **Diameter**:

- **Diameter Identity** — Enter the Diameter identity (typically a fully qualified domain name) or identities used by application functions for this application. Click **Add** to define multiple values. To delete an existing value, select it from the list and click **Delete**.

5. When you finish, click **Save** (or **Cancel** to discard your changes).
   The application profile is created and stored in the **Applications** group.

The application profile is created.

Modifying an Application Profile

To modify an application profile:

1. From the **Policy Server** section of the navigation pane, select **Applications**.
   The content tree displays the **Applications** group.
2. Select the **Applications** group.
   The Application Administration page opens in the work area, listing the application profiles.
3. On the Application Administration page, select the application profile you want to modify.
   The profile is displayed.
4. Click **Modify**.
   The Modify Application page opens.
5. Modify the application profile information as necessary.
   See *Creating an Application Profile* for a description of the fields on this page.
6. When you finish, click **Save** (or **Cancel** to discard your changes).

The application profile is modified.

Deleting an Application Profile

To delete an application profile:

1. From the **Policy Server** section of the navigation pane, select **Applications**.
   The content tree displays the **Applications** group.

2. Select the **Applications** group.
   The Application Administration page opens in the work area.
3. Delete the application profile using one of the following methods:
   - From the work area, click the Delete icon, located to the right of the profile you wish to delete.
   - From the content tree, select the application and click **Delete**. You are prompted, "Are you sure you want to delete this Application?"
4. Click **OK** (or **Cancel** to cancel the request).

The application profile is deleted from the CMP database and all MPE devices.

Chapter 8

Managing Traffic Profiles

Topics:

Managing Traffic Profiles defines how to create and manage traffic profiles in the CMP system.

About Traffic Profiles

A traffic profile is a set of values defined for parameters that are used in protocol messages within the MPE device. Typically, these traffic profile values are used to define the Quality of Service (QoS) for sessions that are managed by those protocol messages. You can use traffic profiles to implement policy and charging control (PCC) rules.

Traffic profiles are used in the MPE device under several situations; for example:

- They define default settings for protocol messages (see *Configuring Protocol Options on the Policy Server*)
- They modify protocol messages, thus modifying the QoS for sessions managed by those messages (see *Creating a New Policy*)

A traffic profile can be applied by a policy rule trigger, or by default if no policy rule is triggered.

Each traffic profile has a type associated with it. Since each protocol supports different parameters for controlling QoS settings, the available MPE parameters depend on the underlying protocol. Therefore, each profile type is associated with a single protocol, but a single protocol can support multiple profile types.

You can create multiple traffic profiles of the same type, as the values of the parameters for each profile determine the actual QoS that is associated with that profile. For example, one possible set of traffic profiles is as follows:

- **Default** — default predefined profile
- **P2P** — profile for peer-to-peer traffic
- **RATE_LIMIT_128K** — profile to limit download rate to 128 Kbps
- **RATE_LIMIT_64K** — profile to limit download rate to 64 Kbps

Creating a Traffic Profile

To create a traffic profile:

1. From the **Policy Server** section of the navigation pane, select **Traffic Profiles**.
   The content tree displays the **Traffic Profiles** group. The default group is **ALL**.
2. Select the **ALL** group.
   The Traffic Profile Administration page opens in the work area, listing available traffic profiles.
3. On the Traffic Profile Administration page, click **Create Traffic Profile**.
   The New Traffic Profile page opens.
4. Enter the following information:
   a) **Name** — The name assigned to the profile.
   b) **Traffic Profile Type** — Select from the following:
      - **Best Effort** (the default) — Transmission opportunities are granted on a first-come, first-served basis. Appropriate for upstream service flows such as Web browsing, e-mail, or instant messaging.
      - **Diameter QoS**

- **Downstream** — Defined through a similar set of QoS parameters that are associated with the best-effort scheduling type on upstream service flows. Appropriate for all downstream service flows.
- **Non-Real-Time Polling** — Cable modems are polled at a fixed interval for queued data. Appropriate for upstream service flows that require high throughput, and traffic that requires variable-sized data grants on a regular basis, such as high-bandwidth FTP.
- **RSVP Flow Spec** — Receivers initiate reservation requests for unidirectional data flows, and senders respond with path information.
- **Real-Time Polling** — Cable modems are polled at a fixed but short interval for queued data. Appropriate for upstream service flows of real-time traffic that generate variable-sized data packets on a periodic basis and have inflexible latency and throughput requirements, such as MPEG video.
- **Service Class** — The profile will use a service class that is configured on the CMTS.
- **Unsolicited Grant** — A fixed-size grant is offered to service flows at fixed intervals without additional polling or interaction. Appropriate for upstream service flows of real-time traffic that generate fixed-size data packets on a periodic basis, such as VoIP.
- **Unsolicited Grant with Activity Detection** — When there is activity, the CMTS sends unsolicited fixed grants at fixed intervals to the cable modem. When there is no activity, the CMTS sends unicast poll requests to the cable modem to conserve unused bandwidth. Appropriate for upstream service flows that include silence suppression.

c) Protocol Fields — The set of protocol fields displayed on the Traffic Profile page varies depending on the Traffic Profile Type selected. *Table 15: Traffic Profile Type Configuration Parameters* describes the protocol fields for each traffic profile type.

5. When you finish, click **Save** (or **Cancel** to discard your changes).

The traffic profile is defined.

Table 15: Traffic Profile Type Configuration Parameters

Traffic Profile Type	Configuration Parameter	Description
Best Effort	Traffic Priority	Priority for the service flow. Higher-priority service flows are given preference over lower-priority service flows.
	Request Transmission Policy	The interval usage code that the cable modem uses for upstream transmission requests and packet transmissions for this service flow. It also specifies whether requests can be piggybacked with data.
	Max Sustained Traffic Rate (bps)	The maximum sustained rate, in bits per second, at which traffic can operate over the service flow.
	Max Traffic Burst	The maximum burst size for the service flow.
	Min Reserved Traffic Rate (bps)	The guaranteed minimum rate, in bits per second, that is reserved for the service flow.

Traffic Profile Type	Configuration Parameter	Description
	Assumed Min Packet Size (bytes)	The assumed minimum packet size, in bytes, for which the minimum reserved traffic rate is provided.
	Maximum Concatenated Bursts (bytes)	
	Upstream Peak Traffic Rate	A four-byte unsigned integer field that specifies the peak traffic rate, in bits per second, that is allowed for a service flow. The range is 0 – 4,294,967,295 bps (4 Gbps–1).
	Required Attribute Mask	
	Forbidden Attribute Mask	
	Attribute Aggregation Rule Mask	
	Minimum Buffer	The lower limit for the size, in bits, of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). The default is a value of 0, which indicates that there is no lower limit.
	Target Buffer	The desired value for the size, in bits, of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). If the parameter is omitted or set to a value of 0, then the device selects any buffer size within the range of the minimum and maximum buffers, using a vendor-specific algorithm.
	Maximum Buffer	The upper limit for the size of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). The default is no limit.
Diameter QoS	QoS Class Identifier	Identifies the QoS class. Select from the following: • **1 = Conversational speech** • **2 = Conversational** • **3 = Streaming speech** • **4 = Streaming** • **5 = Interactive with priority 1 signalling** • **6 = Interactive with priority 1** • **7 = Interactive with priority 2** • **8 = Interactive with priority 3** • **9 = Background**

Traffic Profile Type	Configuration Parameter	Description
	Uplink Max Authorized Rate (bps)	Maximum authorized bandwidth in bits per second for uplinks (user equipment to network).
	Downlink Max Authorized Rate (bps)	Maximum authorized bandwidth in bits per second for downlinks (network to user equipment).
	Uplink Min Guaranteed Rate (bps)	Minimum guaranteed bandwidth in bits per second for uplinks (user equipment to network). Only applicable if the QoS class identifier is between 1 and 4.
	Downlink Min Guaranteed Rate (bps)	Minimum guaranteed bandwidth in bits per second for downlinks (network to user equipment). Only applicable if the QoS class identifier is between 1 and 4.
Downstream	Traffic Priority	Priority for the service flow. Higher-priority service flows are given preference over lower-priority service flows.
	Downstream Resequencing	
	Max Sustained Traffic Rate (bps)	The maximum sustained rate, in bits per second, at which traffic can operate over the service flow.
	Max Traffic Burst	The maximum burst size for the service flow.
	Min Reserved Traffic Rate (bps)	The guaranteed minimum rate, in bits per seond, that is reserved for the service flow.
	Assumed Min Packet Size (bytes)	The assumed minimum packet size, in bytes, for which the minimum reserved traffic rate is provided.
	Max Downstream Latency	The maximum latency for downstream service flows.
	Downstream Peak Traffic Rate	A four-byte unsigned integer field, specifying the rate parameter P of a token-bucket based peak rate limiter for packets of a downstream service flow. This lets you define a Max Traffic Burst value for the Max Sustained Traffic Rate much larger than a maximum packet size, but still limit the burst of packets consecutively transmitted for a service flow.
	Required Attribute Mask	
	Forbidden Attribute Mask	
	Attribute Aggregation Rule Mask	

Traffic Profile Type	Configuration Parameter	Description
	Minimum Buffer	The lower limit for the size, in bits, of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). The default is a value of 0, which indicates that there is no lower limit.
	Target Buffer	The desired value for the size, in bits, of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). If the parameter is omitted or set to a value of 0, then the device selects any buffer size within the range of the minimum and maximum buffers, using a vendor-specific algorithm.
	Maximum Buffer	The upper limit for the size of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). The default is no limit.
Non-Real-Time Polling	Traffic Priority	Priority for the service flow. Higher-priority service flows are given preference over lower-priority service flows.
	Request Transmission Policy	The interval usage code that the cable modem uses for upstream transmission requests and packet transmissions for this service flow. It also specifies whether requests can be piggybacked with data.
	Max Sustained Traffic Rate (bps)	The maximum sustained rate, in bits per second, at which traffic can operate over the service flow.
	Max Traffic Burst	The maximum burst size for the service flow.
	Min Reserved Traffic Rate (bps)	The guaranteed minimum rate, in bits per second, that is reserved for the service flow.
	Assumed Min Packet Size (bytes)	The assumed minimum packet size, in bytes, for which the minimum reserved traffic rate is provided.
	Nominal Polling Interval (microsec)	The nominal interval, in microseconds, between successive unicast request opportunities for this service flow.
	Maximum Concatenated Bursts (bytes)	The largest transmission of concatenated frames, in bytes, that a modem can make on behalf of the service flow.
	Upstream Peak Traffic Rate	A four-byte unsigned integer field that specifies the peak traffic rate, in bits per second, that is allowed for a service flow. The range is 0 – 4,294,967,295 bps (4 Gbps–1).

Traffic Profile Type	Configuration Parameter	Description
	Required Attribute Mask	
	Forbidden Attribute Mask	
	Attribute Aggregation Rule Mask	
	Minimum Buffer	The lower limit for the size, in bits, of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). The default is a value of 0, which indicates that there is no lower limit.
	Target Buffer	The desired value for the size, in bits, of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). If the parameter is omitted or set to a value of 0, then the device selects any buffer size within the range of the minimum and maximum buffers, using a vendor-specific algorithm.
	Maximum Buffer	The upper limit for the size of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). The default is no limit.
RSVP Flow Spec	Service Number	Select from the following: • **N/A** (the default) • **2 = Guaranteed Service** — controls the maximum delay and ensures no packet loss • **5 = Controlled Load Service** — appropriate for soft QoS applications
	Token Bucket Rate (bytes/sec)	The rate, in bytes, at which data arrives.
	Token Bucket Size (bytes)	The size, in bytes, of the token bucket. This dictates how "bursty" the traffic can be.
	Peak Data Rate (bytes/sec)	
	Minimum Policed Unit (bytes)	
	Maximum Packet Size (bytes)	
	Rate (bytes/sec)	
	Slack Term (microsec)	
Real-Time Polling	Request Transmission Policy	The interval usage code that the cable modem uses for upstream transmission requests and packet transmissions for this service flow. It also specifies whether requests can be piggybacked with data.

Traffic Profile Type	Configuration Parameter	Description
	Max Sustained Traffic Rate (bps)	The maximum sustained rate, in bits per second, at which traffic can operate over the service flow.
	Max Traffic Burst	The maximum burst size for the service flow.
	Min Reserved Traffic Rate (bps)	The guaranteed minimum rate, in bits per second, that is reserved for the service flow.
	Assumed Min Packet Size (bytes)	The assumed minimum packet size, in bytes, for which the minimum reserved traffic rate is provided.
	Nominal Polling Interval (microsec)	The nominal interval, in microseconds, between successive unicast request opportunities for this service flow.
	Tolerated Poll Jitter (microsec)	The maximum amount of time, in microseconds, that unicast request intervals can be delayed beyond the nominal polling interval.
	Maximum Concatenated Bursts (bytes)	The maximum size, in bytes, of a concatenated frame (a group of frames) that a service flow can transmit.
	Upstream Peak Traffic Rate	A four-byte unsigned integer field that specifies the peak traffic rate, in bits per second, that is allowed for a service flow. The range is 0 – 4,294,967,295 bps (4 Gbps–1).
	Required Attribute Mask	
	Forbidden Attribute Mask	
	Attribute Aggregation Rule Mask	
	Minimum Buffer	The lower limit for the size, in bits, of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). The default is a value of 0, which indicates that there is no lower limit.
	Target Buffer	The desired value for the size, in bits, of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). If the parameter is omitted or set to a value of 0, then the device selects any buffer size within the range of the minimum and maximum buffers, using a vendor-specific algorithm.
	Maximum Buffer	The upper limit for the size of the buffer to be provided for a service flow. The range is 0 –

Traffic Profile Type	Configuration Parameter	Description
		4,294,967,295 bits (4 Gb–1). The default is no limit.
Service Class	Service Class Name	The name of a service class.
Unsolicited Grant	Request Transmission Policy	The interval usage code that the cable modem uses for upstream transmission requests and packet transmissions for this service flow. It also specifies whether requests can be piggybacked with data.
	Unsolicited Grant Size (bytes)	The size, in bytes, of the individual data grants provided to the service flow
	Grants Per Interval	The actual number of data grants given to the service flow during each nominal grant interval.
	Nominal Grant Interval	The nominal interval between successive unsolicited data grant opportunities for this service flow.
	Tolerated Grant Jitter (microsec)	The maximum amount of time, in microseconds, that the transmission opportunities can be delayed beyond the nominal grant interval.
	Upstream Peak Traffic Rate	A four-byte unsigned integer field that specifies the peak traffic rate, in bits per second, that is allowed for a service flow. The range is 0 – 4,294,967,295 bps (4 Gbps–1).
	Required Attribute Mask	
	Forbidden Attribute Mask	
	Attribute Aggregation Rule Mask	
	Minimum Buffer	The lower limit for the size, in bits, of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). The default is a value of 0, which indicates that there is no lower limit.
	Target Buffer	The desired value for the size, in bits, of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). If the parameter is omitted or set to a value of 0, then the device selects any buffer size within the range of the minimum and maximum buffers, using a vendor-specific algorithm.
	Maximum Buffer	The upper limit for the size of the buffer to be provided for a service flow. The range is 0 –

Traffic Profile Type	Configuration Parameter	Description
		4,294,967,295 bits (4 Gb–1). The default is no limit.
Unsolicited Grant with Activity Detection	Request Transmission Policy	The interval usage code that the cable modem uses for upstream transmission requests and packet transmissions for this service flow. It also specifies whether requests can be piggybacked with data.
	Unsolicited Grant Size (bytes)	The size, in bytes, of the individual data grants provided to the service flow
	Grants Per Interval	The actual number of data grants given to the service flow during each nominal grant interval.
	Nominal Grant Interval	The nominal interval between successive unsolicited data grant opportunities for this service flow.
	Tolerated Grant Jitter (microsec)	The maximum amount of time, in microseconds, that the transmission opportunities can be delayed beyond the nominal grant interval.
	Nominal Polling Interval (microsec)	The nominal interval, in microseconds, between successive unicast request opportunities for this service flow.
	Tolerated Poll Jitter (microsec)	The maximum amount of time, in microseconds, that unicast request intervals can be delayed beyond the nominal polling interval.
	Upstream Peak Traffic Rate	A four-byte unsigned integer field that specifies the peak traffic rate, in bits per second, that is allowed for a service flow. The range is 0 – 4,294,967,295 bps (4 Gbps–1).
	Required Attribute Mask	
	Forbidden Attribute Mask	
	Attribute Aggregation Rule Mask	
	Minimum Buffer	The lower limit for the size, in bits, of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). The default is a value of 0, which indicates that there is no lower limit.
	Target Buffer	The desired value for the size, in bits, of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). If the

Traffic Profile Type	Configuration Parameter	Description
		parameter is omitted or set to a value of 0, then the device selects any buffer size within the range of the minimum and maximum buffers, using a vendor-specific algorithm.
	Maximum Buffer	The upper limit for the size of the buffer to be provided for a service flow. The range is 0 – 4,294,967,295 bits (4 Gb–1). The default is no limit.

Modifying a Traffic Profile

To modify a traffic profile:

1. From the **Policy Server** section of the navigation pane, select **Traffic Profiles**.
   The content tree opens.
2. From the content tree, select the **ALL** group.
   The Traffic Profile Administration page opens, displaying the list of defined traffic profiles.
3. Select the profile you want to modify.
   Profile information is displayed.
4. Click **Modify**.
   The Modify Traffic Profile page opens.
5. Modify profile information as required.
   For a description of the fields contained on this page, see *Creating a Traffic Profile*.
6. When you finish, click **Save** (or **Cancel** to abandon your changes).

The traffic profile is modified.

Deleting a Traffic Profile

You cannot delete a traffic profile that is deployed on an MPE device.

To delete a traffic profile:

1. From the **Policy Server** section of the navigation pane, select **Traffic Profiles**.
   The content tree opens.
2. From the content tree, select the **Traffic Profiles** group.
   The Traffic Profile Administration page opens, displaying the list of defined traffic profiles.
3. Delete the traffic profile using one of the following methods:
   - From the work area, click the Delete icon, located to the right of the traffic profile you want to delete.
   - From the content tree, select the traffic profile and click **Delete**.

   You are prompted, "Are you sure you want to delete this Traffic Profile?"

4. Click **OK** to delete the traffic profile (or **Cancel** to cancel the request).

The traffic profile is deleted.

Traffic Profile Groups

For organizational purposes, you can aggregate traffic profiles into groups. Once a traffic profile group is created, it can be populated with individual traffic profiles. The following subsections describe how to manage traffic profile groups.

Creating a Traffic Profile Group

To create a traffic profile group:

1. From the **Policy Server** section of the navigation pane, select **Traffic Profiles**.
   The content tree displays a list of traffic profile groups; the initial group is **ALL**.
2. From the content tree, select the **ALL** group.
   The Traffic Profile Administration page opens in the work area, listing all defined traffic profiles.
3. On the Traffic Profile Administration page, click **Create Group**.
   The Create Group editor page opens.
4. Enter the name of the new traffic profile group.

   The name can be up to 250 characters long and must not contain quotation marks (") or commas (,).
5. Optionally, enter a description of the traffic profile group; for example:

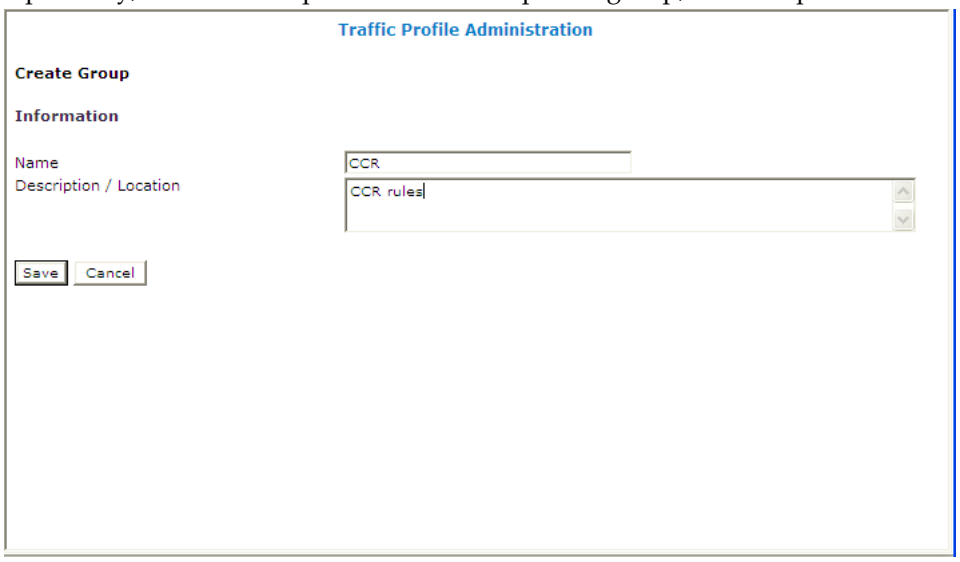

6. When you finish, click **Save** (or **Cancel** to discard your changes).
   The new group appears in the content tree.

The traffic profile group is created.

Adding a Traffic Profile to a Traffic Profile Group

To add a traffic profile to a traffic profile group:

1. From the **Policy Server** section of the navigation pane, select **Traffic Profiles**.
   The content tree displays a list of traffic profile groups; the initial group is **ALL**.
2. From the content tree, select the desired traffic profile group.
   The Traffic Profile Administration page opens in the work area, displaying the contents of the selected traffic profile group.
3. On the Traffic Profile Administration page, click **Add Traffic Profile**.
   The Add Traffic Profile page opens, displaying the traffic profiles not already part of the group.
4. Click on the traffic profile you want to add; use the Ctrl or Shift keys to select multiple traffic profiles.
5. When you finish, click **Save** to add the traffic profile to the selected group (or **Cancel** to cancel the request).

The traffic profile is added to the traffic profile group.

Modifying a Traffic Profile Group

To modify a traffic profile group:

1. From the **Policy Server** section of the navigation pane, select **Traffic Profiles**.
   The content tree displays a list of traffic profile groups; the initial group is **ALL**.
2. From the content tree, select the traffic profile group you want to modify.
   The Traffic Profile Administration page opens in the work area.
3. On the Traffic Profile Administration page, click **Modify**.
   The Modify Group page opens.
4. Edit the information in the fields.
   The name cannot contain quotation marks (") or commas (,).
5. When you finish, click **Save** (or **Cancel** to cancel the request).

The group is modified.

Removing a Traffic Profile from a Traffic Profile Group

Removing a traffic profile from a traffic profile group does not delete the profile. To delete a traffic profile, see *Deleting a Traffic Profile*.

To remove a traffic profile from a traffic profile group:

1. From the **Policy Server** section of the navigation pane, select **Traffic Profiles**.
   The content tree displays the list of traffic profile groups.
2. From the content tree, select the desired traffic profile group.
   The Traffic Profile Administration page opens in the work area, displaying the contents of the selected traffic profile group; for example:

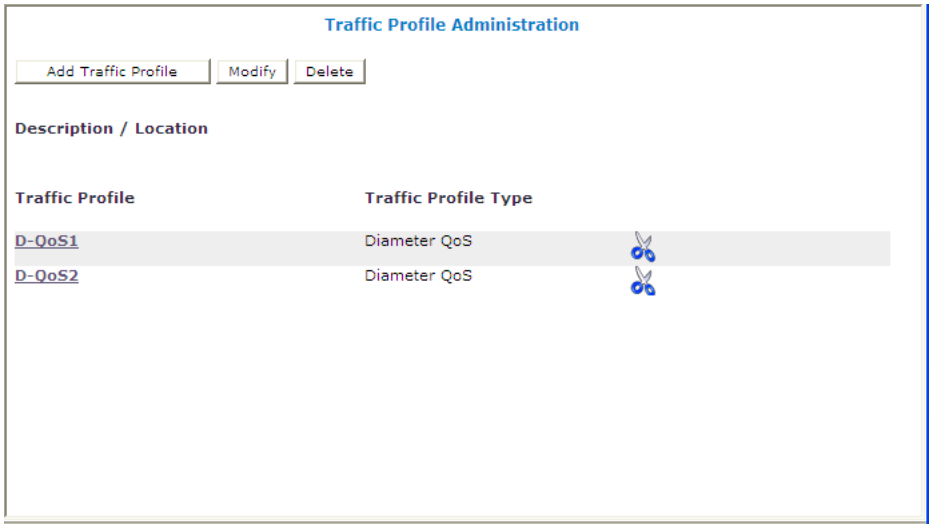

3. Remove the traffic profile using one of the following methods:
   - Click the Delete icon, located to the right of the traffic profile you want to remove.
   - From the traffic profile group in the content tree, select the traffic profile and click **Remove**.

The traffic profile is removed from the group immediately; there is no confirmation message.

Deleting a Traffic Profile Group

Deleting a traffic profile group does not delete any traffic profiles associated with the deleted group; profiles remain in the ALL group. You cannot delete the ALL group.

To delete a traffic profile group:

1. From the **Policy Server** section of the navigation pane, select **Traffic Profiles**.
   The content tree displays the list of traffic profile groups.
2. From the content tree, select the traffic profile group you want to delete.
   The Traffic Profile Administration page opens in the work area, displaying the contents of the selected traffic profile group; for example:

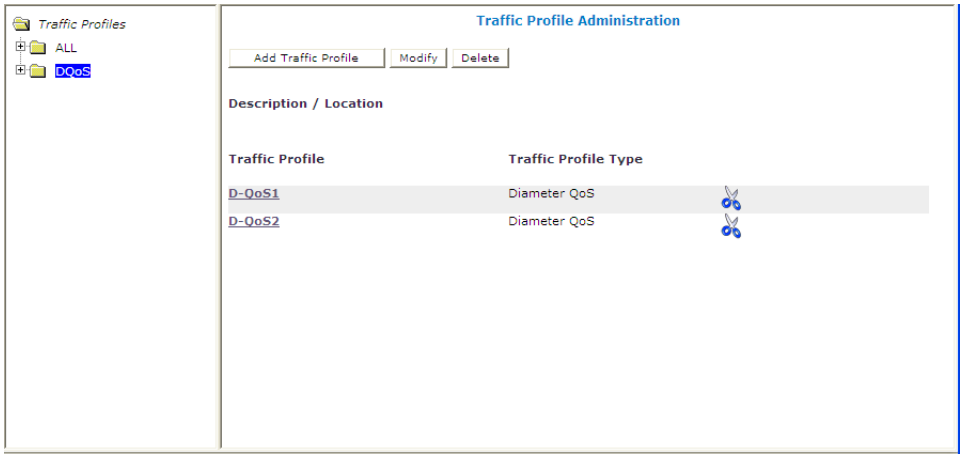

3. On the Traffic Profile Administration page, click **Delete**.
   You are prompted, "Are you sure you want to delete this Group?"
4. Click **OK** to delete the group (or **Cancel** to cancel the request).

The traffic profile group is deleted.

Chapter 9

Managing Media Profiles

Topics:

This chapter defines how to manage media profiles in the CMP system.

A media profile describes a CODEC supported for Rx-to-PCMM translation.

About Media Profiles

A media profile describes a CODEC supported for Rx-to-PCMM translation. The MPE device includes a predefined set of media profiles, and you can create new RTP (real-time transport protocol) profiles as necessary. Once you have defined a media profile in the CMP system, it is automatically deployed to MPE devices.

Media profiles are named *codec_name-transport_type-sample_rate*. Media profiles are mapped to CODECs based on the information received in a session description protocol (SDP) message.

In defining a media profile in the CMP system, you specify its name, transport type, sample rate, frame size (in both milliseconds and bytes), and packetization time.

Note: You cannot create media profiles for the UDPTL or UDP transport types.

Table 16: Predefined Media Profiles describes the predefined media profiles.

Table 16: Predefined Media Profiles

CODEC Name	AVT Profile	Frame Length (ms)	Frame Size (bytes)	Bit Rate (kbps)	Sample Rate (kHz)
PCMU	0	0.125	1	64	8
G721	2	0.125	1	64	8
GSM	3	20	33	13.2	8
G723	4	30	24	5.3, 6.3	8
PCMA	8	0.125	1	64	8
G722	9	0.125	1	64	8
G722-48	dynamic	1	6	48	8
G722-56	dynamic	1	7	56	8
G722-64	dynamic	1	8	64	8
G728	15	2.5	5	16	8
G729	18	10	10	8	8
G726-16	dynamic	0.5	1	16	8
G726-24	dynamic	1	3	24	8
G726-32	dynamic	0.25	1	32	8
G726-40	dynamic	1	5	40	8
G729D	dynamic	10	8	6.4	8
G729E	dynamic	10	15	11.8	8
GSM-EFR	dynamic	20	31	12.2	8
iLBC	dynamic	20	38	13.33	8

CODEC Name	AVT Profile	Frame Length (ms)	Frame Size (bytes)	Bit Rate (kbps)	Sample Rate (kHz)
iLBC	dynamic	30	50	15.2	8
BV16	dynamic	5	10	16	8
BV32	dynamic	5	20	32	16
RED	dynamic	10	160	128	8
VMR-WB	dynamic	20	34	13.6	8
SMV0	dynamic	20	22	8.8	8
evrc0	dynamic	20	22		8
evrcb0	dynamic	20	22		8
evrcwb0	dynamic	20	22		8
evrcwb0	dynamic	20	22		16
amr	dynamic	20	32		8
AMR/8000	dynamic	20	14	4.75	8
AMR/8000	dynamic	20	15	5.15	8
AMR/8000	dynamic	20	16	5.9	8
AMR/8000	dynamic	20	18	6.7	8
AMR/8000	dynamic	20	20	7.4	8
AMR/8000	dynamic	20	22	7.95	8
AMR/8000	dynamic	20	27	10.2	8
AMR/8000	dynamic	20	32	12.2	8
amr-wb	dynamic	20	61		16
amr-wb/16000	dynamic	20	18		16
amr-wb/16000	dynamic	20	24		16
amr-wb/16000	dynamic	20	33		16
amr-wb/16000	dynamic	20	37		16
amr-wb/16000	dynamic	20	41		16
amr-wb/16000	dynamic	20	47		16
amr-wb/16000	dynamic	20	51		16
amr-wb/16000	dynamic	20	59		16
amr-wb/16000	dynamic	20	61		16

Creating a Media Profile

To create a media profile:

1. From the **Policy Server** section of the navigation pane, select **Media Profiles.**
   The content tree displays the **Media Profiles** group.
2. Select the **Media Profiles** group.
   The Media Profile Administration page opens in the work area, listing available media profiles.
3. On the Media Profile Administration page, click **Create Media Profile.**
   The New Media Profile page opens.
4. Enter the following information:
   a) **Codec Name** — Unique media subtype assigned to the media profile.
      This is defined in the IANA MIME registration for the CODEC. Enter a string of up to 255 characters.
   b) **Transport Type** — Select from the following:
      - **RTP/AVP** (the default) — RTP audio-video profile.
      - **RTP/SAVP** — RTP secure audio-video profile.
      - **RTP/AVPF** — RTP extended audio-video profile with feedback.
   c) **Payload Number** — The payload number.
      Valid payload numbers range from 0 through 127. Enter -1 to indicate an unknown payload number.

      Note: You cannot add a CODEC that is predefined with a payload number in the range of 0 to 96.
   d) **Sample Rate (kHz)** — The sampling rate of the CODEC in KHz.
      The valid range is an integer from 1 through 100 KHz.
   e) **Frame Size in Milliseconds** — The size of one audio frame in milliseconds.
      This is the length of time represented by one audio frame. A single RTP packet may contain multiple audio frames. The bitrate is calculated using the frame size in milliseconds, the frame size in bytes, and the packetization time. The valid range is 0 through 100 ms.
   f) **Frame Size in Bytes** — The size of one audio frame size in bytes.
      This is the size represented by one audio frame. A single RTP packet may contain multiple audio frames. The bitrate is calculated using the frame size in milliseconds, the frame size in bytes, and the packetization time. The valid range is 1 through 1,500 bytes.
   g) **Packetization Time** — The length of time, in milliseconds, represented by the media in a packet.
      The bitrate is calculated using the frame size in milliseconds, the frame size in bytes, and the packetization time. The valid range is 1 through 100.
   h) **Always Use Default Ptime** — Select to always use the default packetization time, ignoring the value received in the SDP message.
      The default is unchecked.
5. When you finish, click **Save** to define the media profile (or **Cancel** to discard your changes).

The media profile is created.

Modifying a Media Profile

To modify a media profile:

1. From the **Policy Server** section of the navigation pane, select **Media Profiles**.
   The content tree opens.
2. From the content tree, select the **Media Profiles** group.
   The Media Profile Administration page opens, displaying the list of defined media profiles.
3. Select the media profile you want to modify.
   Profile information is displayed.
4. Click **Modify**.
   The Modify Media Profile page opens.
5. Modify media profile information as required.
   For a description of the fields contained on this page, see *Creating a Media Profile*.
6. When you finish, click **Save** (or **Cancel** to abandon your changes).

The media profile is modified.

Deleting a Media Profile

To delete a media profile:

1. From the **Policy Server** section of the navigation pane, select **Media Profiles**.
   The content tree opens.
2. From the content tree, select the **Media Profiles** group.
   The Media Profile Administration page opens, displaying the list of defined media profiles.
3. Delete the media profile using one of the following methods:
   a) From the work area, click the Delete icon, located to the right of the media profile you want to delete.
   b) From the content tree, select the media profile and click **Delete**.
   You are prompted, "Are you sure you want to delete this Media Profile?"
4. Click **OK** to delete the retry profile (or **Cancel** to cancel the request).

The media profile is deleted.

Chapter 10

Managing Service Classes

Topics:

This chapter defines how to create and manage service classes in the CMP system.

About Service Classes

A service class corresponds to a DOCSIS traffic description defined in a cable modem termination system (CMTS). You can define service classes using the CMP system, load them using the OSSI/XML interface, or discover them using the SNMP interface.

Creating a Service Class

To create a service class:

1. From the **Policy Server** section of the navigation pane, select **Service Classes**.
   The content tree displays the **Service Classes** group.
2. Select the **Service Classes** group.
   The Service Class Administration page opens in the work area, listing available service classes.
3. On the Service Class Administration page, click **Create Service Class**.
   The New Service Class page opens.
4. Enter the following information:
   a) **Name** — The name assigned to the service class.
   b) **Scheduling Type** — Select from the following:
      - **Downstream** (the default) — Defined through a similar set of QoS parameters that are associated with the best-effort scheduling type on upstream service flows. Appropriate for all downstream service flows.
      - **Best Effort** — Transmission opportunities are granted on a first-come, first-served basis. Appropriate for upstream service flows such as Web browsing, e-mail, or instant messaging.
      - **Non Real Time Polling** — Cable modems are polled at a fixed interval for queued data. Appropriate for upstream service flows that require high throughput, and traffic that requires variable-sized data grants on a regular basis, such as high-bandwidth FTP.
      - **Real Time Polling** — Cable modems are polled at a fixed but short interval for queued data. Appropriate for upstream service flows of real-time traffic that generate variable-sized data packets on a periodic basis and have inflexible latency and throughput requirements, such as MPEG video.
      - **Unsolicited Grant Service** — A fixed-size grant is offered to service flows at fixed intervals without additional polling or interaction. Appropriate for upstream service flows of real-time traffic that generate fixed-size data packets on a periodic basis, such as VoIP.
      - **Unsolicited Grant Service with Activity Detect** — When there is activity, the CMTS sends unsolicited fixed grants at fixed intervals to the cable modem. When there is no activity, the CMTS sends unicast poll requests to the cable modem to conserve unused bandwidth. Appropriate for upstream service flows that include silence suppression.

   c) **Maximum Traffic Rate (bps)** — The maximum sustained rate, in bits per second, at which traffic can operate over the service flow.

      Enter an integer between 0 and 4294967295.

      This field applies to the **Downstream**, **Best Effort**, **Non Real Time Polling**, and **Real Time Polling** scheduling types.

d) **Minimum Reserved Rate (bps)** — The guaranteed minimum rate, in bits per second, that is reserved for the service flow.

Enter an integer between 0 and 4294967295.

This field applies to the **Downstream**, **Best Effort**, **Non Real Time Polling**, and **Real Time Polling** scheduling types.

e) **Unsolicited Grant Size (bytes)** — The size, in bytes, of the individual data grants provided to the service flow.

Enter an integer between 0 and 65535.

This field applies to the **Unsolicited Grant Service** and **Unsolicited Grant Service with Activity Detect** scheduling types.

f) **Nominal Grant Interval (usecs)** — The nominal interval, in microseconds, between successive unsolicited data grant opportunities for this service flow.

Enter an integer between 0 and 4294967295.

This field applies to the **Unsolicited Grant Service** and **Unsolicited Grant Service with Activity Detect** scheduling types.

g) **Grants per Interval** — The actual number of data grants given to the service flow during each nominal grant interval.

Enter an integer between 0 and 127.

This field applies to the **Unsolicited Grant Service** and **Unsolicited Grant Service with Activity Detect** scheduling types.

5. When you finish, click **Save** (or **Cancel** to discard your changes).

The service class is defined.

Modifying a Service Class

To modify a service class:

1. From the **Policy Server** section of the navigation pane, select **Service Classes**.
   The content tree opens.
2. From the content tree, select the **Service Classes** group.
   The Service Class Administration page opens, displaying the list of defined service classes.
3. Select the service class you want to modify.
   Service class information is displayed.
4. Click **Modify**.
   The Modify Service Class page opens.
5. Modify service class information as required.
   For a description of the fields contained on this page, see *Creating a Service Class*.
6. When you finish, click **Save** (or **Cancel** to abandon your changes).

The service class is modified.

Deleting a Service Class

To delete a service class:

1. From the **Policy Server** section of the navigation pane, select **Service Classes**.
   The content tree opens.
2. From the content tree, select the **Service Classes** group.
   The Service Class Administration page opens, displaying the list of defined service classes.
3. Delete the service class using one of the following methods:
   - From the work area, click the Delete icon, located to the right of the service class you want to delete.
   - From the content tree, select the service class and click **Delete**.

   You are prompted, "Are you sure you want to delete this Service Class?"
4. Click **OK** to delete the service class (or **Cancel** to cancel the request).

The service class is deleted.

Chapter 11

Managing Record Keeping Servers

Topics:

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 (RKS) is a respository 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 RKSs, and then associate them with MPE devices, either by adding them to the MPE device's Record Keeping Server List, or by defining one as the default RKS.

When configuring an RKS, note that a single RKS may correspond to a single external server, but it may also correspond to a pair of external servers. This depends on how the RKS handles failover situations.

An RKS is uniquely identified by the following:

- Primary IP Address
- Primary Port
- Secondary IP Address
- Secondary Port

If you have a single server that provides both a primary and secondary address, you can configure it as a single RKS. If you have two servers, each of which only provides a single IP address/port, then you could either configure both of them as a single RKS (that acts as a backup pair) or you could configure them as two separate RKSs, each with a primary address/port and no secondary address/port. However, if an RKS does not have a secondary address/port, then that RKS will not be able to participate in the RKS failover mechanism as defined in the PCMM specification.

Creating an RKS Profile

To configure an RKS 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. On the Record Keeping Server Administration page, click **Create Record Keeping Server**.
   The New Record Keeping Server page opens.
4. Enter the following information:
   a) **Name** — The name assigned to the RKS profile.
   b) **Description/Location** (optional) — Information pertaining to the RKS that helps identify it within the network or location.
   c) **Primary Address** — IP address, in IPv4 or IPv6 format, of the primary RKS.
   d) **Primary Port** — IP port number of the primary RKS. (The port number is typically 1813.)
   e) **Secondary Address** (optional) — IP address of the secondary RKS.
   f) **Secondary Port** — IP port number of the secondary RKS. (The port number is typically 1813.)
5. When you finish, click **Save** (or **Cancel** to discard your changes).

The RKS profile is created.

Modifying an RKS Profile

To modify an RKS 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. On the Record Keeping Server Administration page, click the RKS you wish to modify.
   Configuration information for that RKS is displayed.
4. Click **Modify**.
   The Modify Record Keeping Server page opens.
5. Modify configuration information as needed.
6. When you finish, click **Save** (or **Cancel** to discard your changes).

The RKS profile is modified.

Deleting an RKS Profile

To delete an RKS 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 desired RKS profile using one of the following methods:
   - Click the **Delete** icon located to the right of the profile you want to delete.
   - From the content tree, select the profile; the Record Keeping Server Administration page opens. Click **Delete**.

   You are prompted: "Are you sure you want to delete this Record Keeping Server?"
4. Click **OK** to delete the RKS profile (or **Cancel** to cancel your request).

The RKS profile is deleted.

Chapter 12

Managing Event Messaging

Topics:

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 (EMs) 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 (via 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 (AM) 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 AM 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 AM requests event messages to be sent as part of that request, it includes a piece of information called the Billing Correlation ID (or 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 AM, 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 AM. 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. On the Event Messaging Administration page, click **Modify**.
   The Modify Event Messaging page opens. *Figure 18: 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 MSO's domain name, 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 AMs request event messages, then configure the same RKSs in both the AMs and the MPE devices.
   e) **Default Record Keeping Server** — Defines the default RKS for event messaging.
4. When you finish, click **Save** (or **Cancel** to discard your changes).

The global event message settings are defined.

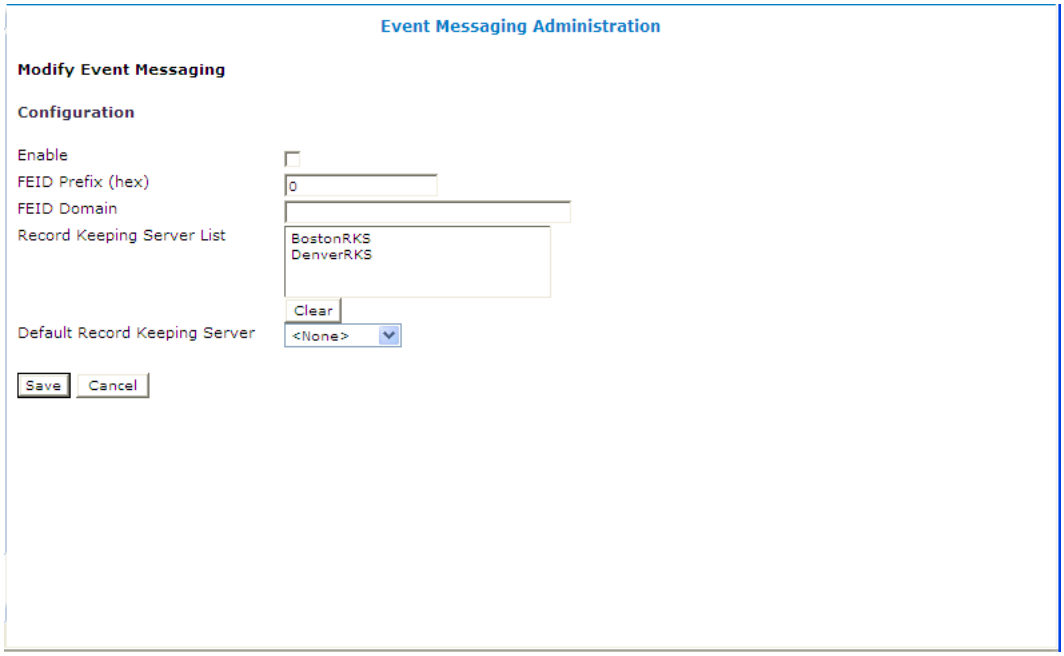

Figure 18: Modify Event Messaging Page

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 desired MPE device.
   The Policy Server Administration page opens in the work area.
3. On the Policy Server Administration page, 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 the **Overrides** radio button 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 MSO's domain name, 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 RKSs.

If you are configuring event messaging in your network so that the AMs request event messages, then configure the same RKSs in the AMs and the MPE device.

f) **Default Record Keeping Server** — Defines the default RKS for event messaging.

6. When you finish, click **Save** (or **Cancel** to discard your changes).

Local settings are defined for this MPE device.

Chapter 13

Managing Management Agent Servers

Topics:

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 (MA) server is designed specifically for network architectures that require a distributed topology and collection framework. The 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 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.

Using the MA 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 MA server is initiated by the CMP system, and optionally is performed over a secured interface.

Creating a Management Agent Profile

To create an MA profile:

1. From 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. On the Management Agent Administration page, click **Create Management Agent**.
   The New Management Agent page opens.
4. Enter the following profile information:
   a) **Associated Cluster** — Select the cluster from the pulldown list.
   b) **Name** — The name assigned to the MA.
   c) **Description/Location** — Free-form text that defines the MA's function or location.
   d) **Secure Connection** — Designates whether or not to use SSL as a secure connection for this MA.
5. When you finish, click **Save** (or **Cancel** to discard your changes).
   The MA profile is added to the list of available profiles.

The management agent profile is created.

Modifying a Management Agent Profile

To modify a management agent profile:

1. From 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 you want to modify.
   The management agent is displayed in the Management Agent Administration page.
3. On the Management Agent Administration page, click **Modify**.
   The Modify System Settings page opens.
4. Edit the profile information as desired. See *Creating a Management Agent Profile* for descriptions of these fields.
5. When you finish, click **Save** (or **Cancel** to discard your changes).

The management agent profile is modified.

Deleting a Management Agent Profile

To delete a management agent profile:

1. From 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 the **Delete** icon, located to the right of the policy you want to delete.
   - From the policy group tree, select the policy; the management agent is displayed in the Management Agent Administration page. Click **Delete**.

   You are prompted: "Are you sure you want to delete this Management Agent?"
3. Click **OK** to delete the management agent (or **Cancel** to abandon your request).

The management agent profile is deleted.

Reapplying a Management Agent Profile Configuration

To reapply a configuration to a management agent server:

1. From 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. On the Management Agent Administration page, 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 MA 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 MA task status and the current scheduled data processing:

1. From the navigation pane, select **Management Agents**.
   The Management Agent Administration page opens in the work area.
2. From the content tree, select the desired management agent.
   The management agent is displayed in the Management Agent Administration page.
3. On the Management Agent Administration page, select the **Tasks** tab.
   The various 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.

Viewing Task Status

To view the status and the current execution schedule for a specific task, click the task name. Detailed information is displayed.

The following options are available on this page:

- **Reschedule** — Reschedules when the task process starts:
  - Click on the 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 MA (or **Cancel** to discard your changes).

- **Run Now** — Runs the task process immediately.
- **Disable/Enable** — Disables or enables this feature.
- **Refresh** — Refreshes the current page.
- **Cancel** — Ignores any information added and closes this page.

Chapter 14

Understanding and Creating Policy Rules

Topics:

Policy rules dynamically control how the Multimedia Policy Engine (MPE) processes protocol messages as they pass through it. Using these rules, you can define how and when network resources are utilized by subscribers. For example, when the MPE device receives a request to establish a session with a certain Quality of Service (QoS) level, you can use a policy rule to approve the request as is, to reject the request, or to make changes in the request before it is forwarded to the intended destination network element.

Structure and Evaluation of Policy Rules

The following topics provide an overview of how policy rules are structured and evaluated.

Note: The conditions, actions, and parameters available for your use in creating policy rules depend on the mode in which the CMP system is operating.

Structure of Policy Rules

Understanding how a policy rule is structured is helpful in understanding other policy management concepts. A policy rule is defined in an if-then structure, consisting of a set of conditions that the MPE device compares to protocol messages, and a set of actions that are executed (or not executed) when the conditions match. Many conditions can be tested for existence or non-existence (by optionally selecting the operator **is** or **is not**).

Policy Parameters

When you define a policy rule, you select from a list of available conditions and actions. Most of the conditions and actions are parameterized (that is, they contain placeholders that may be replaced with specific values to allow you to customize them as needed).

For example, consider the following policy rule, which has one condition and two actions:

```
where the device will be handling greater than 100 upstream reserved flows

apply profile Default Downstream Profile to request
continue processing message
```

The condition, **where the device will be handling**..., allows the following parameters to be specified:

- An operator (greater than)
- A value (100)
- The flow direction (upstream)
- The bandwidth reservation type (reserved)

The first action, **apply profile**..., specifies a single parameter that is the name of a traffic profile to be applied to the request. The second action, **continue processing message**, instructs the MPE device to evaluate the remaining rules within the policy rules list (as opposed to immediately accepting or rejecting the request). The conditions and actions that are available for writing policies are discussed later in this section.

Policy Logical Operators

The policy wizard supports creation of rules using an explicit **AND** logical operator that contains a set of conditions. An AND operator must include at least two conditions. The actions are taken if all

conditions are evaluated as true. For example, you can use an AND operator two define two conditions as follows:

```
And
     where the request is re-authorizing an existing session
     where the enforcement session is a DPI enforcement session
.
.
.
```

The policy wizard supports creation of rules using an **OR** logical operator that contains a set of conditions. An OR operator must include at least two conditions. The actions are taken if any condition is evaluated as true. For example, you can define the following set of conditions using an OR operator:

```
Or
     where the request is creating a new session
     where the session is an enforcement session
     where the APN matches one of imode.g1t2
     where the subscriber profile data is not available
.
.
.
```

The policy wizard supports creation of rules using a **NOT** logical operator that contains a single condition. The actions are taken if the condition is evaluated as false. For example, you can define the following using a NOT operator:

```
Not
     where today is a weekend day using CONFIGURED_LOCAL_TIME
.
.
.
```

Note: Many conditions also include optional **is** and **is not** parameters. These parameters are functionally equivalent to (that is, synonymous with) using the **NOT** operator, and you are free to use or mix **NOT** with **is** and **is not** as you prefer.

Finally, the policy wizard supports creation of rules using combinations of logical operators. You can nest operators. For example, you can define the following rule:

```
Or
     And
          Not
               where the service info status is one of FINAL_SERVICE_INFORMATION

          where the session is an enforcement session
     where the session is an application session
     Not
          where the session is an application session

evaluate policy 5555
reject message
```

The policy wizard validates condition trees.

Parent and Reference Policies

As a result of evaluating conditions, a policy can execute another policy. A policy that calls another policy is called a parent policy, and a policy executed by another policy is called a reference policy. A policy can be both a parent policy and a reference policy. Additionally, you can group policies, and a parent policy can execute all the policies in the group.

Note: Do not nest policies more than five levels deep.

Evaluating Policy Rules

To write policy rules, it is important to understand how they are evaluated by the Policy Rules Engine contained within the MPE device, and how the engine fits into the protocol message processing within the MPE device.

If you look at the policy conditions that are available, you will see that many are not protocol specific. Although you can write protocol-specific policy rules, the Policy Rules Engine itself does not have any protocol knowledge. Instead, it deals with a set of abstractions that are mapped to the underlying protocol messages that are being processed. This allows the same policy rules to be used across multiple protocols.

When the MPE device receives a protocol message, it performs the initial processing of that message and then determines whether or not the message should be processed by the Policy Rules Engine. Generally, protocol messages that are either requesting bandwidth or modifying previous requests for bandwidth are processed by the Policy Rules Engine. Most other protocol messages are not. For example, a protocol message that releases bandwidth is typically not processed by the Policy Rules Engine because there is no reason to prevent or modify that action.

Once a message is identified as a candidate for the policy rules, the MPE device attempts to associate as much information with the request as possible. For example:

- Which network elements will be impacted if the request is allowed to proceed?
- Which subscriber is associated with the request? What services is that subscriber entitled to?
- Which application is associated with the message?
- What time zone is the user equipment located in?

The reason for collecting this information is to make it available to the policy rules. The information that can be associated varies and depends on a number of factors, including:

- The protocol in question and how much information is provided in the protocol message
- The amount of network topology information that has been provisioned into the MPE device
- Whether there are other protocol sessions that can be associated with this message
- Whether there are external data sources configured that the MPE device can use to associate information with the message

When the process of associating information with the request is complete, the MPE device analyzes the information and maps it into several important abstractions that are central to the functioning of the Policy Rules Engine:

1. A list of network devices that the request affects. A network device is any network element, any logical or physical sub-component of a network element, or any other network equipment.
2. A list of flows associated with the request. A flow is a logical representation of a QoS enforcement point that is used for a specific purpose (typically in a single direction, either upstream or downstream). A flow is usually characterized by a collection of bandwidth parameters. Different

protocols can have a different number of flows associated with a message. For example, PCMM messages have only one flow per request.

3. A list of policies associated with the request. This includes policy groups and reference policies called by the parent policy.

After constructing these lists, the Policy Rules Engine applies the policy rules according to the following algorithm:

```
For each network device:
   For each flow that is being created or modified:
      For each policy that is being evaluated:
         Evaluate all policy rules
      End
   End
End
```

It should be clear from this algorithm that a single message can result in multiple policies being evaluated, and a policy rule being evaluated multiple times. This is important to understand to ensure that the policy rules you write operate in the way you intended.

By using parent policies, reference policies, and policy groups, you can control the order of policy execution. For example, assume there are four policies: two parent policies, $policy_1$ and $policy_4$, and two reference policies, $policy_2$ and $policy_3$ that are in a policy group, $group_1$. The hierarchy is as follows:

```
policy1
   policy2
   policy3
policy4
```

The order of execution can vary, depending on how each policy evaluates and what actions each contains:

- The normal order of execution would be $policy_1$, $policy_2$, $policy_3$, $policy_4$.
- If the conditions in $policy_1$ evaluate to false, the order of execution would be $policy_1$, $policy_4$.
- if $policy_2$ includes the mandatory action "break from policy level," the order of execution would be $policy_1$, $policy_2$, $policy_4$.

Note: Policies created using a more recent version of the CMP software may not evaluate and execute as intended on an MPE device running an older version of the MPE software. To ensure that policies are evaluated and executed as intended, update all systems to the same version of the software.

Using Reference Policies

Multiple policies that share the same conditions can be simplified by including the common conditions in a parent policy and any unique conditions in reference policies. During execution, the common conditions are only evaluated once.

For example, consider the following policies, which apply tiers to session requests. Each policy uses the same conditions, and the Policy Rules Engine evaluates the same conditions up to three times:

```
Bronze Policy
where the request is creating a new session
   and where the flow is an application flow
   and where the AF-Application-ID matches one of voip
```

```
   and where the tier is one of Bronze
apply bronze to request
accept message
```

```
Silver Policy
where the request is creating a new session
   and where the flow is an application flow
   and where the AF-Application-ID matches one of voip
   and where the tier is one of Silver
apply silver to request
accept message
```

```
Gold Policy
where the request is creating a new session
   and where the flow is an application flow
   and where the AF-Application-ID matches one of voip
   and where the tier is one of Gold
apply gold to request
accept message
```

Using reference policies in a policy group, the same results can be obtained with the following policies:

```
where the request is creating a new session
   and where the flow is an application flow
   and where the AF-Application-ID matches one of voip
evaluate policy group Tier Policies
```

```
   Bronze Policy
   where the tier is one of Bronze
   apply bronze to request
   accept message
```

```
   Silver Policy
   where the tier is one of Silver
   apply silver to request
   accept message
```

```
   Gold Policy
   where the tier is one of Gold
   apply gold to request
   accept message
```

Creating a New Policy

Policy rules are created and modified using the policy wizard in the CMP system. Once created or modified, the rule is stored in the policy library. The policy wizard guides you step by step to creating a new policy rule. The wizard displays only the options available at each step.

The following procedure describes how to create a new policy rule, using this policy as an example:

```
And
   where the request is creating a new session
```

```
   where the session is an application session
   where the flow media type is one of Video

Advanced: set values for QoS and Charging parameters to
   PCMM Classifier - Priority6
send notification to trace log with `VideoPolicy triggered` and severity `Info`

continue processing message
```

To create a new policy rule:

1. From the **Policy Management** section of the navigation pane, select **Policy Library**.
   The content tree displays a list of policy library groups; the default is **ALL**.
2. From the content tree, select the **ALL** group.
   The Policy Administration page opens in the work area.
3. On the Policy Administration page, click **Create Policy**.
   The Create Policy page opens.
4. Select a starting point for the new policy:
   - **Blank** — The policy rule is created from the beginning, without any attributes being pre-defined.
   - **Use Template** — The policy rule is created based on a user-defined template that may have policy parameters pre-defined. This template can be modified as needed.
   - **Copy Existing Policy** — The policy rule is created based on an existing policy rule, which you modify as needed.
5. Click **Next** (or **Cancel** to close the wizard without saving the policy).
   The Tables page opens.
6. Specify the table(s) you want to use in the policy. For more details on associating a table with a policy, see *Associating Policy Tables with a Policy Rule*.

   If no tables are associated with the policy, click **Next**.
   - To specify multiple tables, click the selection icon ( ) multiple times
   - To move a table so that it is evaluated earlier in the rule, click the up icon ( )
   - To move a table so that it is evaluated later in the rule, click the down icon ( )
   - To delete a table, click the delete icon ( )
7. When you finish defining tables, click **Next** (or **Cancel** to close the wizard without saving the policy).
   The Conditions page opens.
8. Select the desired policy conditions.

   As a condition is selected, it appears in the Description area at the bottom of the page.

   You can select multiple conditions, enter multiple instances of each condition, change the order of conditions, group conditions logically, or remove conditions:
   - To enter multiple instances of a condition, click the selection icon ( ) in the Conditions window multiple times.
   - To combine a logical group of conditions, click **And** or **Or**, located in the upper right corner of the Description window, and drag the conditions into the container that appears (represented by a folder icon). You can toggle a container between **And** and **Or** by double-clicking on the folder.

- To change a condition's order of evaluation or include it within a logical container, drag and drop the condition within the Description window. You cannot drop a container onto itself or one of its sub-containers.
- To negate a condition, change the **is** parameter if present, or click **Not**, located in the upper right corner of the Description window, and drag the condition into the container that appears (represented by a folder icon).
- To delete a condition or container from the rule, select it and click **Delete**. You are prompted, "The focused item and all its children will be deleted. Continue?" Click **OK** (or **Cancel** to keep the condition or container).

Tip: To add conditions directly to an existing container, select the container first.

For example:

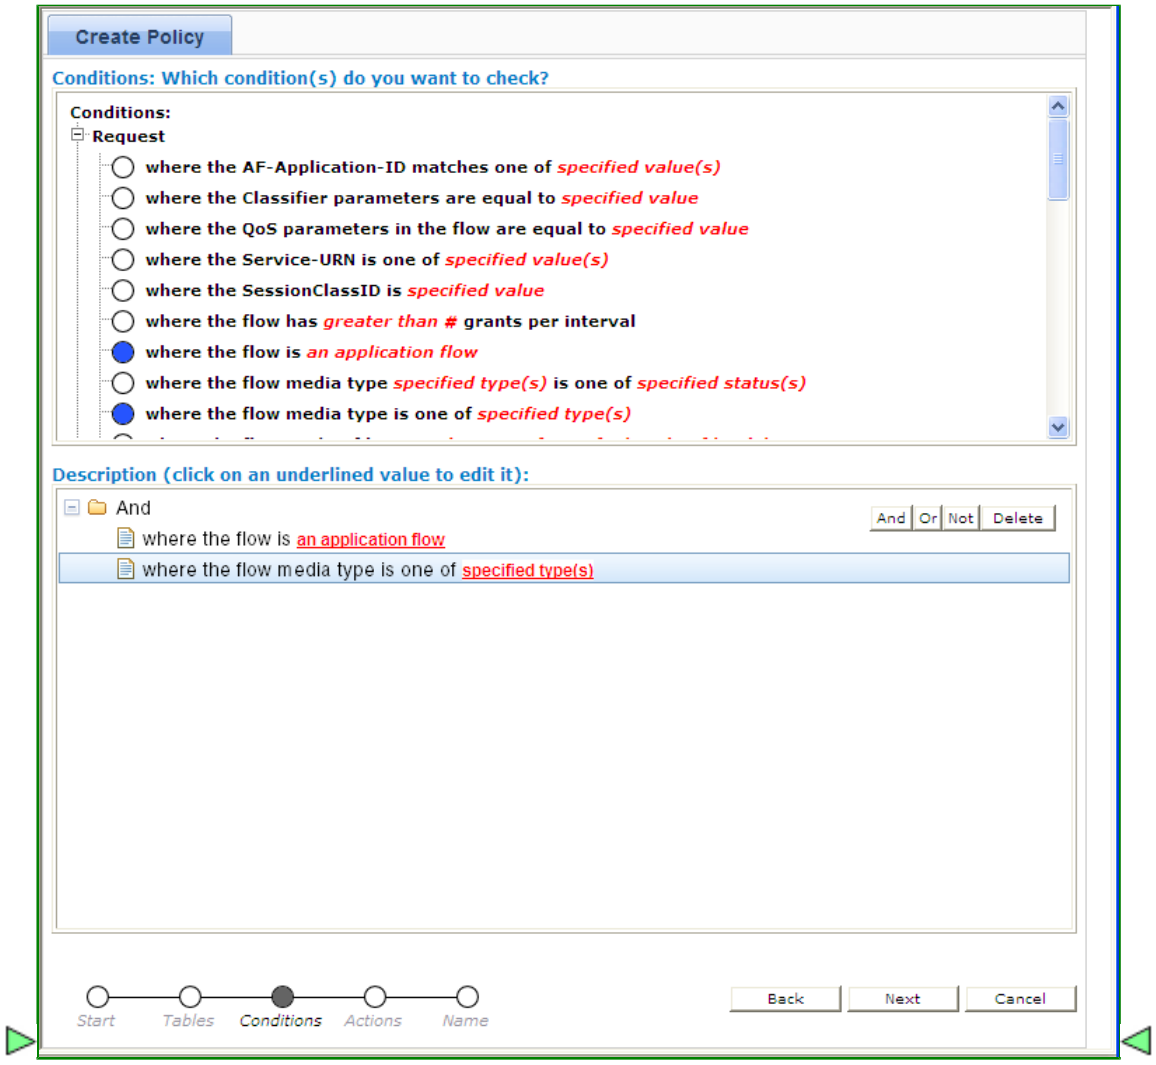

9. If a policy condition includes a parameter that requires further input, it displays red underlined text in the Description area. To provide the input, click the red underlined text; a popup window opens, from which you can do one of the following:

   - Select one or more options
   - Enter a value (such as a traffic bit rate or percentage)

When you finish, click **OK** (or **Cancel** to discard your changes). The popup window closes and the input is added to the policy condition.

10. When you finish defining policy conditions, click **Next** (or **Cancel** to close the wizard without saving the policy).
    The Actions page opens.
11. Select the required action and any optional actions that the MPE device should execute if the policy request matches the defined conditions of the policy rule.

    For example:

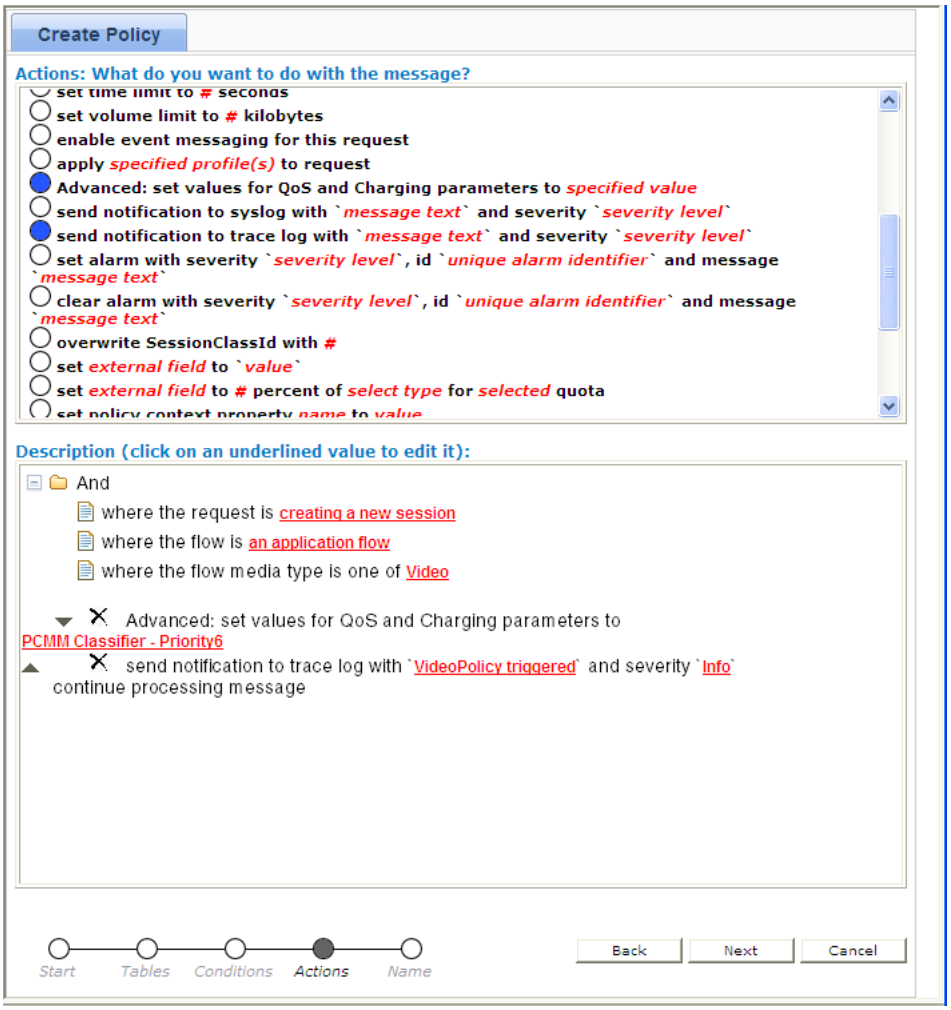

- To enter multiple instances of an action, click the selection icon ( ) multiple times
- To move an action so that it is evaluated earlier in the rule, click the up icon ( )
- To move an action so that it is evaluated later in the rule, click the down icon ( )
- To delete an action from the rule, click the delete icon ( )

12. When you finish, click **Next** (or **Cancel** to close the wizard without saving the policy).
    The Name page opens.
13. Assign a unique name (where uniqueness is not case sensitive) to the new policy rule; for example:

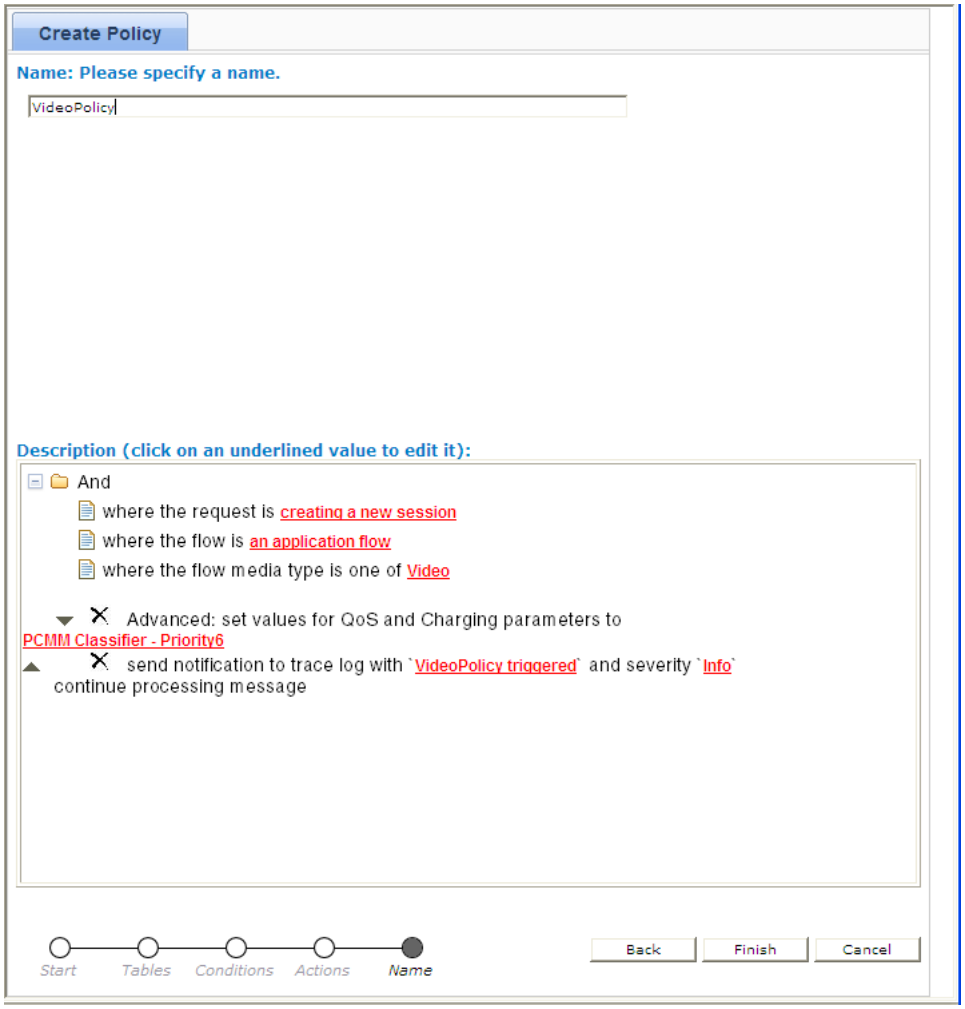

Note: The name can be up to 255 characters long and cannot contain the following characters: **< > \ ; & ' " =**

14. Click **Finish** (or **Cancel** to close the wizard without saving the policy).
    The Create Policy page closes.

The policy rule is saved to the policy library in the CMP database.

Once a policy rule is created, you must deploy it to MPE devices so it can take effect. Reference policy rules (rules called by parent policy rules) do not need to be deployed; they are deployed automatically when called by a parent rule. See *Managing Policy Rules*.

Modes Within the Policy Wizard

The behavior of the policy wizard varies depending on the mode in which your CMP is running. The mode can affect many policy wizard behaviors, including the following:

- Entire categories of conditions are enabled or disabled.

- Specific conditions and/or actions are enabled or disabled.
- Some conditions will have a slightly different appearance.
- The set of valid values for some parameters will vary.

If your policy wizard does not include a category, condition, or action documented here, it means that those categories, conditions, or actions are not relevant in your present CMP mode.

Parameters Within Policy Rules

When you are defining policy rules, both the conditions and actions may contain parameters. Parameters let you customize the specific situation in which a policy rule will be applied. Some conditions and actions may contain multiple parameters. For example, one possible condition is as follows:

```
where the device will be handling greater than 100 upstream reserved flows
```

This condition contains four different parameters. The policy wizard displays the parameters using a red font, with each parameter having a single continuous underline. In this example, greater than is a single parameter, as is 100, upstream, and reserved.

You can click on any parameter to open a pop-up window that lets you specify the value of that parameter. Each parameter has a data type associated with it that determines the values that can be specified: some may be numbers, some may be free-form text, and some may be limited to specific sets of values. For example, the following parameter is limited to a set of text values:

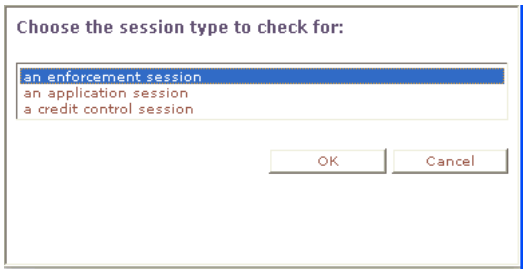

If you have many policies with similar structures, you can consolidate them using policy tables to capture the differences. To specify a parameter in a rule that uses a policy table, instead of selecting a value click **Use Policy Table**. For more information on table-driven policies see *Managing Policy Tables*.

Table 17: Common Parameters defines some common parameter types that are used in many of the policy rules. In this table, the column labeled "Default Text" shows the text value that is displayed in the condition or action text when they are initially displayed. (This may be different in some instances, but this value is the default.)

There are also many parameter types that are used in only one condition or action. These parameter types are defined in the sections where those conditions or actions are defined.

Table 17: Common Parameters

Parameter Type	Default Text	Description of Values
app-name	specified name	Names of applications that have been defined in the CMP database.
bandwidth	#	A numeric value that specifies bandwidth in bits per second (bps). You can also type "k", "K", "m","M", "g", or "G" in the value to specify the value in units of kilobits, megabits, or gigabits per second instead.
class-of-service	specified class of	One (or more) of the following: • **Best Effort** • **Non Real-time Polling** • **Real-time Polling** • **UGS** • **Background** • **Conversational** • **Streaming** • **Interactive**
flow-direction	upstream	One of the following: • **upstream** • **downstream** • **upstream or downstream**
ip-address	specified address	An IPv4 or IPv6 address.
log-message	text	Any string. This text may contain policy parameters (as described later in this section) that perform parameter substitution within the message text.
matches-op	matches one of	One of the following: • **matches one of** • **does not match any of**
match-list		A comma-separated list of values, where each value is a wildcard match pattern that uses the "*" character to match zero or more characters or the "?" character to match exactly one character.
number	#	A numeric value. In some circumstances, the numeric value may be required to fall within a certain range of valid values.
operator	greater than	One of the following: • **greater than or equal to** • **greater than** • **less than or equal to** • **less than** • **equal to**

Parameter Type	Default Text	Description of Values
		• **not equal to**
operator-binary	is	One of the following: • **is** • **is not**
operator-greater	greater than	One of the following: • **greater than or equal to** • **greater than**
operator-less	less than	One of the following: • **less than or equal to** • **less than**
percent	#	An integer value between 0 and 100; for certain values, an extended, non-integer percentage that can exceed 100 (for example, 102.4%).
qos-direction	upstream	One of the following: • **upstream** • **downstream**
qos-status	reserved	One or more of the following: • **reserved** • **committed**
seconds	#	A numeric value that specifies time in units of seconds.
string	specified	Any string.
subnet	specified subnet	An IPv4 subnet in CIDR notation (for example, 1.2.3.0/24); or an IPv6 subnet (for example, fc00::1006/64).

Conditions Available for Writing Policy Rules

The policy wizard supports a large number of conditions that can be used for constructing policy rules. To help you find the conditions you want, the conditions are organized into different categories, which are summarized in *Table 18: Policy Condition Categories*.

Table 18: Policy Condition Categories

Category	Description
Request	Conditions that are based on information that is explicitly contained within or related to the protocol message (request) that triggered the policy rule execution.

Category	Description
Application	Conditions related to the application associated with the request.
Network Device Identity	Conditions related to the specific network device for which the policy rule is being evaluated. This includes conditions based on the network device type, as well as those that refer to specific unique identifiers for network devices.
Network Device Usage	Conditions related to the calculated usage for the network device for which the policy rule is being evaluated. This usage includes device-level tracking of both bandwidth and flow/session counts.
User	Conditions related to the subscriber, or subscriber account, that is associated with the protocol message that triggered the policy rule execution. This includes subscriber-level and account-level tracking of usage.
Policy Context Properties	Conditions related to the context in which a policy is evaluated.
Time of Day	Conditions related to the time at which the policy rules are being executed.

The conditions that are included within each of these categories are described in the sections that follow. Conditions are listed in alphabetical order. The parameters that can be modified within each condition are also detailed.

Request Conditions

Request conditions are based on information that is explicitly contained within, or related to, the protocol message (request) that triggered the policy rule execution.

where the AF-Application-ID matches one of *specified value(s)*

Syntax
where the AF-Application-ID matches one of *csv*

Parameters
csv

Comma-separated list of text values.

Description

Selects protocol messages based on the Diameter AF Application Identifier field. A valid AF Application identifier is any string describing the application, for example VoIP or streaming.

where the Classifier parameters are equal to *specified value*

Syntax
where the Classifier parameters are equal to *classifier*

Parameters
classifier

One or more of the following:

- **PCMM Classifier (Extended) - Action**
- **PCMM Classifier (Extended) - Activation State**
- **PCMM Classifier (Extended) - Classifier Id**
- **PCMM Classifier (Extended) - Destination Mask**
- **PCMM Classifier (Extended) - Destination Port End**
- **PCMM Classifier (Extended) - Source Mask**
- **PCMM Classifier (Extended) - Source Port End**
- **PCMM Classifier (IPv6) - Destination Address**
- **PCMM Classifier (IPv6) - Destination Prefix Length**
- **PCMM Classifier (IPv6) - Flags**
- **PCMM Classifier (IPv6) - Flow Label**
- **PCMM Classifier (IPv6) - Next Header Type**
- **PCMM Classifier (IPv6) - Source Address**
- **PCMM Classifier (IPv6) - Source Prefix Length**
- **PCMM Classifier (IPv6) - tc-high**
- **PCMM Classifier (IPv6) - tc-low**
- **PCMM Classifier (IPv6) - tc-mask**
- **PCMM Classifier - Destination Address**
- **PCMM Classifier - Destination Port**
- **PCMM Classifier - DSCP/TOS Field**
- **PCMM Classifier - DSCP/TOS Mask**
- **PCMM Classifier - Priority**
- **PCMM Classifier - ProtocolId**
- **PCMM Classifier - Source Address**
- **PCMM Classifier - Source Port**

Description

Distinguishes between different types of PCMM classifier parameters.

where the flow has *greater than* # grants per interval

Syntax
where the flow has *operator number* grants per interval

Parameters
operator

See common parameters.

number

See common parameters.

Description

Selects protocol messages based on the number of grants per interval in the flow.

where the flow is *an application flow*

Syntax
where the flow is *flow-type*

Parameters
flow-type

One or more of the following:

- **an application flow** (the default)
- **a UE flow**
- **the default flow**

Description

Selects protocol messages based on the type of flow.

where the flow media type is one of *specified type(s)*

Syntax
where the flow(s) media type is one of *media-type*

Parameters
media-type

One or more of the following:

- **Audio**
- **Video**
- **Data**
- **Application**
- **Control**
- **Text**
- **Message**
- **Other**

Description

Selects protocol messages based on the flow or flows' media type.

where the flow media type *specified type(s)* is one of *specified status(s)*

Syntax
where the flow media type *media-type* is one of *media-status*

Parameters
media-type

One or more of the following media types:

- **Audio**
- **Video**
- **Data**
- **Application**
- **Control**
- **Text**
- **Message**
- **Other**

media-status

One or more of the following status type:

- **Enabled**
- **Enabled Uplink**
- **Enabled Downlink**
- **Disabled**
- **Removed**

Description

Selects protocol messages that matches the flow's media type and status type.

where the flow packet filter *matches one of specified packet filter(s)*

Syntax
where the flow packet filter *matches-op csv*

Parameters
matches-op

See common parameters.

csv

Comma-separated list of values.

Description

Selects protocol messages based on the packet filters. The packet filters use IPFilterRule format, as defined in the Diameter base protocol (RFC 3588). For example: `permit in ip from 10.0.0.1 to 10.0.0.2 5060`.

where the flow usage is one of *specified usage(s)*

Syntax
where the flow usage is *flow-usage-type*

Parameters
flow-usage-type

One or more of the following:

- **No Information**
- **RTCP**
- **AF Signaling**

Description

Selects protocol messages based on the flow usage.

where the flow(s) media types Matches *specified type(s)*

Syntax
where the flow(s) media type matches *media-type*

Parameters
media-type

One or more of the following media types:

- **Audio**
- **Video**
- **Data**
- **Application**
- **Control**
- **Text**
- **Message**
- **Other**

Description

Selects one or more protocol messages that match one or more flow's media types.

where the protocol being executed is *PCMM*

Syntax
where the protocol being executed is *protocol*

Parameters
protocol

One of the following:

- **PCMM** (the default)
- **Diameter AF**

Description

Distinguishes between protocols being executed.

where the QoS parameters in the flow are equal to *specified value*

Syntax
where the QoS parameters in the flow are equal to *profile-param*

Parameters
profile-param

Names of profile parameters that are derived from internal representations of protocol messages. For the specific meaning of the fields, consult the specific protocol specifications.

- **Diameter AF Flow-Description**
- **Diameter AF Flow-Status**
- **Diameter AF Flow-Usage**
- **Diameter AF Maximum-Authorized-Data-Rate**
- **Diameter AF Media-Type**
- **Diameter AF PacketTime**
- **Diameter AF QCI**
- **Diameter AF Reservation-Priority**
- **Diameter AF RTCP RR-Bandwidth**
- **Diameter AF RTCP RS-Bandwidth**
- **Diameter Flow-Status**
- **PCMM AMID**
- **PCMM Classifier (Extended) - Action**
- **PCMM Classifier (Extended) - Activation State**
- **PCMM Classifier (Extended) - Classifier Id**
- **PCMM Classifier (Extended) - Destination Mask**
- **PCMM Classifier (Extended) - Destination Port End**
- **PCMM Classifier (Extended) - Source Mask**
- **PCMM Classifier (Extended) - Source Port End**
- **PCMM Classifier (IP6v6) - Destination Address**
- **PCMM Classifier (IP6v6) - Destination Prefix Length**
- **PCMM Classifier (IP6v6) - Flags**
- **PCMM Classifier (IP6v6) - Flow Label**
- **PCMM Classifier (IP6v6) - Next Header Type**
- **PCMM Classifier (IP6v6) - Source Address**
- **PCMM Classifier (IP6v6) - Source Prefix Length**
- **PCMM Classifier (IP6v6) - tc-high**
- **PCMM Classifier (IP6v6) - tc-low**
- **PCMM Classifier (IP6v6) - tc-mask**
- **PCMM Classifier - Destination Address**
- **PCMM Classifier - Destination Port**
- **PCMM Classifier - DSCP/TOS Field**
- **PCMM Classifier - DSCP/TOS Mask**
- **PCMM Classifier - Priority**

- **PCMM Classifier - ProtocolId**
- **PCMM Classifier - Source Address**
- **PCMM Classifier - Source Port**
- **PCMM Gate Id**
- **PCMM GateSpec - DSCP/TOS Enabled**
- **PCMM GateSpec - DSCP/TOS Field**
- **PCMM GateSpec - DSCP/TOS Mask**
- **PCMM GateSpec - Session Class Id**
- **PCMM GateSpec - Timer T1 (secs)**
- **PCMM GateSpec - Timer T2 (secs)**
- **PCMM GateSpec - Timer T3 (secs)**
- **PCMM GateSpec - Timer T4 (secs)**
- **PCMM Traffic Profile - Authorized Assumed Minimum Reserved Traffic Rate Packet Size (bytes)**
- **PCMM Traffic Profile - Authorized Attribute Aggregation Rule Mask**
- **PCMM Traffic Profile - Authorized Downstream Peak Traffic Rate**
- **PCMM Traffic Profile - Authorized Downstream Resequencing**
- **PCMM Traffic Profile - Authorized Forbidden Attribute Mask**
- **PCMM Traffic Profile - Authorized Grants Per Interval**
- **PCMM Traffic Profile - Authorized Maximum Buffer**
- **PCMM Traffic Profile - Authorized Maximum Concatenated Bursts**
- **PCMM Traffic Profile - Authorized Maximum Downstream Latency**
- **PCMM Traffic Profile - Authorized Maximum Packet Size [M] (bytes)**
- **PCMM Traffic Profile - Authorized Maximum Sustained Traffic Rate (bps)**
- **PCMM Traffic Profile - Authorized Maximum Traffic Burst (bytes)**
- **PCMM Traffic Profile - Authorized Minimum Buffer**
- **PCMM Traffic Profile - Authorized Minimum Policed Unit [m] (bytes)**
- **PCMM Traffic Profile - Authorized Minimum Reserved Traffic Rate (bps)**
- **PCMM Traffic Profile - Authorized Nominal Grant Interval (microsec)**
- **PCMM Traffic Profile - Authorized Nominal Polling Interval (microsec)**
- **PCMM Traffic Profile - Authorized Peak Data Rate [p] (bytes/sec)**
- **PCMM Traffic Profile - Authorized Rate [R] (bytes/sec)**
- **PCMM Traffic Profile - Authorized Request Transmission Policy**
- **PCMM Traffic Profile - Authorized Required Attribute Mask**
- **PCMM Traffic Profile - Authorized Slack Term [S] (microsec)**
- **PCMM Traffic Profile - Authorized Target Buffer**
- **PCMM Traffic Profile - Authorized Token Bucket Rate [r] (bytes/sec)**
- **PCMM Traffic Profile - Authorized Token Bucket Size [b] (bytes)**
- **PCMM Traffic Profile - Authorized Tolerated Grant Jitter (microsec)**
- **PCMM Traffic Profile - Authorized Tolerated Poll Jitter (microsec)**
- **PCMM Traffic Profile - Authorized Traffic Priority (bytes/sec)**
- **PCMM Traffic Profile - Authorized Unsolicited Grant Size (bytes)**
- **PCMM Traffic Profile - Authorized Upstream Peak Traffic Rate**
- **PCMM Traffic Profile - Committed Assumed Minimum Reserved Traffic Rate Packet Size (bytes)**

- **PCMM Traffic Profile - Committed Attribute Aggregation Rule Mask**
- **PCMM Traffic Profile - Committed Downstream Peak Traffic Rate**
- **PCMM Traffic Profile - Committed Downstream Resequencing**
- **PCMM Traffic Profile - Committed Forbidden Attribute Mask**
- **PCMM Traffic Profile - Committed Grants Per Interval**
- **PCMM Traffic Profile - Committed Maximum Buffer**
- **PCMM Traffic Profile - Committed Maximum Concatenated Bursts**
- **PCMM Traffic Profile - Committed Maximum Downstream Latency**
- **PCMM Traffic Profile - Committed Maximum Packet Size [M] (bytes)**
- **PCMM Traffic Profile - Committed Maximum Sustained Traffic Rate (bps)**
- **PCMM Traffic Profile - Committed Maximum Traffic Burst (bytes)**
- **PCMM Traffic Profile - Committed Minimum Buffer**
- **PCMM Traffic Profile - Committed Minimum Policed Unit [m] (bytes)**
- **PCMM Traffic Profile - Committed Minimum Reserved Traffic Rate (bps)**
- **PCMM Traffic Profile - Committed Nominal Grant Interval (microsec)**
- **PCMM Traffic Profile - Committed Nominal Polling Interval (microsec)**
- **PCMM Traffic Profile - Committed Peak Data Rate [p] (bytes/sec)**
- **PCMM Traffic Profile - Committed Rate [R] (bytes/sec)**
- **PCMM Traffic Profile - Committed Request Transmission Policy**
- **PCMM Traffic Profile - Committed Required Attribute Mask**
- **PCMM Traffic Profile - Committed Slack Term [S] (microsec)**
- **PCMM Traffic Profile - Committed Target Buffer**
- **PCMM Traffic Profile - Committed Token Bucket Rate [r] (bytes/sec)**
- **PCMM Traffic Profile - Committed Token Bucket Size [b] (bytes)**
- **PCMM Traffic Profile - Committed Tolerated Grant Jitter (microsec)**
- **PCMM Traffic Profile - Committed Tolerated Poll Jitter (microsec)**
- **PCMM Traffic Profile - Committed Traffic Priority (bytes/sec)**
- **PCMM Traffic Profile - Committed Unsolicited Grant Size (bytes)**
- **PCMM Traffic Profile - Committed Upstream Peak Traffic Rate**
- **PCMM Traffic Profile - Envelope**
- **PCMM Traffic Profile - Reserved Assumed Minimum Reserved Traffic Rate Packet Size (bytes)**
- **PCMM Traffic Profile - Reserved Attribute Aggregation Rule Mask**
- **PCMM Traffic Profile - Reserved Downstream Peak Traffic Rate**
- **PCMM Traffic Profile - Reserved Downstream Resequencing**
- **PCMM Traffic Profile - Reserved Forbidden Attribute Mask**
- **PCMM Traffic Profile - Reserved Grants Per Interval**
- **PCMM Traffic Profile - Reserved Maximum Buffer**
- **PCMM Traffic Profile - Reserved Maximum Concatenated Bursts**
- **PCMM Traffic Profile - Reserved Maximum Downstream Latency**
- **PCMM Traffic Profile - Reserved Maximum Packet Size [M] (bytes)**
- **PCMM Traffic Profile - Reserved Maximum Sustained Traffic Rate (bps)**
- **PCMM Traffic Profile - Reserved Maximum Traffic Burst (bytes)**
- **PCMM Traffic Profile - Reserved Minimum Buffer**
- **PCMM Traffic Profile - Reserved Minimum Policed Unit [m] (bytes)**

- **PCMM Traffic Profile - Reserved Minimum Reserved Traffic Rate (bps)**
- **PCMM Traffic Profile - Reserved Nominal Grant Interval (microsec)**
- **PCMM Traffic Profile - Reserved Nominal Polling Interval (microsec)**
- **PCMM Traffic Profile - Reserved Peak Data Rate [p] (bytes/sec)**
- **PCMM Traffic Profile - Reserved Rate [R] (bytes/sec)**
- **PCMM Traffic Profile - Reserved Request Transmission Policy**
- **PCMM Traffic Profile - Reserved Required Attribute Mask**
- **PCMM Traffic Profile - Reserved Slack Term [S] (microsec)**
- **PCMM Traffic Profile - Reserved Target Buffer**
- **PCMM Traffic Profile - Reserved Token Bucket Rate [r] (bytes/sec)**
- **PCMM Traffic Profile - Reserved Token Bucket Size [b] (bytes)**
- **PCMM Traffic Profile - Reserved Tolerated Grant Jitter (microsec)**
- **PCMM Traffic Profile - Reserved Tolerated Poll Jitter (microsec)**
- **PCMM Traffic Profile - Reserved Traffic Priority (bytes/sec)**
- **PCMM Traffic Profile - Reserved Unsolicited Grant Size (bytes)**
- **PCMM Traffic Profile - Reserved Upstream Peak Traffic Rate**
- **PCMM Traffic Profile - Service Class Name**
- **PCMM Traffic Profile - Service Number**
- **PCMM Traffic Profile - Type**
- **PCMM Transaction Id**
- **PCMM User Id**

Description

Selects protocol messages based on values of specific parameters in the protocol message for which there may be an explicit condition already. Depending on the parameter chosen, you may be prompted to enter the value to compare against.

where the request AVP Media-Component-Description *exists*

Syntax
where the request AVP Media-Component-Description *accessibility*

Parameters
accessibility

One of the following:

- **exists** (the default)
- **does not exist**

Description

Determines whether the AVP Media-Component-Description is accessible.

where the request is *creating a new flow*

Syntax
where the request is *change-type*

Parameters
change-type

One or more of the following:

- **creating a new flow** (the default)
- **modifying an existing flow**
- **provisioning a default flow**
- **terminating an existing flow**

Description

Distinguishes between protocol messages based on the type of operation being performed on the flow.

where the request is *creating a new session*

Syntax
where the request is *request-type*

Parameters
request-type

One or more of the following:

- **creating a new session** (the default)
- **modifying an existing session**
- **re-authorizing an existing session**
- **terminating an existing session**

Description

Distinguishes between protocol messages based on the type of operation being performed on the subscriber's session.

where the request is for *reserved* bandwidth

Syntax
where the request is for *qos-status* bandwidth

Parameters
qos-status

See common parameters.

Description

Distinguishes between protocol messages based on the type of bandwidth that is being updated.

where the request *is* for *specified class of* traffic

Syntax
where the request *operator* for *class-of-service* traffic

Parameters
operator

See common parameters.

class-of-service

One or more of the following:

- **Best Effort**
- **Non Real-Time Polling**
- **Real-Time Polling**
- **UGS**
- **Background**
- **Conversational**
- **Streaming**
- **Interactive**

Description

Distinguishes between protocol messages based on the class of service for the network traffic that is being updated.

where the request is for *upstream* bandwidth

Syntax
where the request is for *qos-direction* bandwidth

Parameters
qos-direction

See common parameters.

Description

Distinguishes between protocol messages based on the direction of bandwidth that is being updated.

where the request MPS Identifier *matches one of value(s)*

Syntax
where the MPS Identifier *matches-op csv*

Parameters
matches-op

See common parameters.

csv

Comma-separated list of text values.

Description

Determines whether the MPS Identifier matches a specified value(s).

where the requested guaranteed *upstream* bandwidth is *greater than #* bps

Syntax
where the requested guaranteed *flow-direction bandwidth* is *operator bandwidth* bps

Parameters
flow-direction

See common parameters.

bandwidth

See common parameters.

operator

See common parameters.

Description

Selects protocol messages based on the amount of bandwidth being requested in a specific direction relative to a numeric value.

where the requested maximum *upstream* bandwidth is *greater than specified* bps

Syntax
where the requested maximum *flow-direction* bandwidth is *operator bandwidth* bps

Parameters
flow-direction

See common parameters.

operator

See common parameters.

bandwidth

See common parameters.

Description

Selects protocol messages based on the maximum amount of bandwidth being requested in a specific direction relative to a numeric value.

Example

```
And
    where the request is creating a new session
    where the session is an application session
   where the requested maximum upstream or downstream bandwidth is greater
 than 2400 bps
reject message
```

where the requested media component description reservation priority is one of *specified*

Syntax
where the requested media component description reservation priority is one of *priority*

Parameters
priority

One or more of the following:

- **DEFAULT**
- **PRIORITY_ONE**
- **PRIORITY_TWO**
- **PRIORITY_THREE**
- **PRIORITY_FOUR**
- **PRIORITY_FIVE**
- **PRIORITY_SIX**
- **PRIORITY_SEVEN**
- **PRIORITY_EIGHT**
- **PRIORITY_NINE**
- **PRIORITY_TEN**
- **PRIORITY_ELEVEN**
- **PRIORITY_TWELVE**
- **PRIORITY_THIRTEEN**
- **PRIORITY_FOURTEEN**
- **PRIORITY_FIFTEEN**

Description

Selects Rx protocol messages based on the requested media component description reservation priority.

where the requested QCI is one of *specified*

Syntax
where the requested QCI is one of *class-of-service*

Parameters
class-of-service

One or more of the following:

- **1** (Conversational speech)
- **2** (Conversational)
- **3** (Streaming speech)
- **4** (Streaming)
- **5** (Interactive with priority 1 signalling)
- **6** (Interactive with priority 1)
- **7** (Interactive with priority 2)
- **8** (Interactive with priority 3)
- **9** (Background)

Description

Selects protocol messages based on the QoS class identifier (QCI).

where the requested service class *matches one of specified name(s)*

Syntax
where the requested service class *matches-op service-class-name*

Parameters
matches-op

See common parameters.

service-class-name

Names of service classes that are defined in the CMP database or that have been discovered via SNMP.

Description

Selects protocol messages based on the service class name in the request. See *Managing Traffic Profiles* for information on service classes.

where the requested session reservation priority is one of *specified*

Syntax
where the requested session reservation priority is one of *priority*

Parameters
priority

One or more of the following:

- **DEFAULT**
- **PRIORITY_ONE**
- **PRIORITY_TWO**

- **PRIORITY_THREE**
- **PRIORITY_FOUR**
- **PRIORITY_FIVE**
- **PRIORITY_SIX**
- **PRIORITY_SEVEN**
- **PRIORITY_EIGHT**
- **PRIORITY_NINE**
- **PRIORITY_TEN**
- **PRIORITY_ELEVEN**
- **PRIORITY_TWELVE**
- **PRIORITY_THIRTEEN**
- **PRIORITY_FOURTEEN**
- **PRIORITY_FIFTEEN**

Description

Selects Rx protocol messages based on the requested session reservation priority.

where the requested time limit is *greater than* # seconds

Syntax
where the requested time limit is *operator seconds* seconds

Parameters
operator

See common parameters.

seconds

See common parameters.

Description

Selects protocol messages based on the specified time limit.

where the requested time limit is unlimited (or unspecified)

Description

Selects protocol messages that have no time limit.

where the requested volume limit is *greater than* # kilobytes

Syntax
where the requested volume limit is *operator bandwidth* kilobytes

Parameters
operator

See common parameters.

bandwidth

See common parameters.

Description

Selects protocol messages based on the specified volume limit.

where the requested volume limit is unlimited or unspecified

Description

Selects protocol messages that have no volume limit.

where the service info status is one of *specified*

Syntax
where the service info status is one of *status*

Parameters
status

One of the following:

- **FINAL_SERVICE_INFORMATION**
- **PRELIMINARY_SERVICE_INFORMATION**

Description

Selects Rx protocol messages based on the service information status.

where the Service-URN is one of *specified value(s)*

Syntax
where the Service-URN is one of *csv*

Parameters
csv

Comma-separated list of text values.

Description

Selects Rx protocol messages based on the value of the Service-URN field.

where the session is *an enforcement session*

Syntax
where the session is *session-type*

Parameters
session-type

One of the following:

- **an enforcement session** (the default)
- **an application session**
- **a credit control session**

Description

Distinguishes between protocol messages that are operating on different sessions.

where the SessionClassID is specified value

Syntax
where the SessionClassID is *unit*

Parameters
unit

A number between 0 and 255.

Description

Selects protocol messages based on the value of the SessionClassID field.

Application Conditions

Application conditions are related to the application associated with the request. See *Managing Application Profiles* for information on creating and managing application profiles.

where *AMID* is the application manager ID

Syntax
where *number* is the application manager ID

Parameters
number

A 32-bit numeric value that is greater than 0.

Description

Triggers a policy based on the access manager ID in the message.

where *AppType* is the application type

Syntax
where *number* is the application type

Parameters
number

A 16-bit numeric value that is greater than 0.

Description

Triggers a policy based on the application type in the message (this is a sub-field within the AMID).

where the application is latency sensitive

Description

Triggers a policy when the associated application is latency sensitive (can be set in the CMP system when applications are defined).

where the application *is* one of *specified name*

Syntax
where the application *operator-binary* one of *app-name*

Parameters
operator-binary

See common parameters.

app-name

See common parameters.

Description

Triggers a policy based on the associated application.

where the application will be using *greater than #* bps *upstream reserved* bandwidth

Syntax
where the application will be using *operator-greater bandwidth* bps *qos-direction qos-status* bandwidth

Parameters
operator-greater

See common parameters.

bandwidth

See common parameters.

qos-direction

See common parameters.

qos-status

See common parameters.

Description

Triggers a policy based on the total amount of bandwidth used by the associated application as it relates to a defined threshold. This can be further qualified by both the direction and allocation status of the bandwidth. The total represents the amount of bandwidth that is allocated if the current request is approved.

where the application will be using *greater than # upstream reserved* flows

Syntax
where the application will be using *operator-greater bandwidth qos-direction qos-status* flows

Parameters
operator-greater

See common parameters.

bandwidth

See common parameters.

qos-direction

See common parameters.

qos-status

See common parameters.

Description

Triggers a policy based on the total number of flows used by the associated application as it relates to a defined threshold. This can be further qualified by both the direction and allocation status of the flows. The total represents the number of flows that is allocated if the current request is approved.

where there is no application associated with the request

Description

Triggers a policy when there is no associated application.

Network Device Identity Conditions

Network Device Identity conditions are related to the specific network device for which the policy rule is being evaluated. This includes conditions based on the network device type, as well as those that refer to specific unique identifiers for network devices. See *Managing Network Elements* for information on defining the network elements available.

where # is the CMTS blade index

Syntax
where *number* is the CMTS blade index

Parameters
number

A numeric value between 0 and 255.

Description

Triggers a policy that is only evaluated for a specific CMTS blade (based on the index number of the blade).

where # is the CMTS channel index

Syntax
where *number* is the CMTS channel index

Parameters
number

A numeric value between 0 and 255.

Description

Triggers a policy that is only evaluated for a specific CMTS channel (based on the index number of the channel).

where the cable modem IP address is *specified address*

Syntax
where the cable modem IP address is *ip-address*

Parameters
ip-address

See common parameters.

Description

Triggers a policy that is only evaluated for a specific cable modem (based on its IP address).

where the cable modem IP address is in *specified subnet*

Syntax
where the cable modem IP address is in *subnet*

Parameters
subnet

See common parameters.

Description

Triggers a policy that is only evaluated for cable modems whose IP address falls within a specific subnet.

where the cable modem MAC address is *specified address*

Syntax
where the cable modem MAC address is *mac-address*

Parameters
mac-address

MAC address, in the format *hh:hh:hh:hh:hh:hh*.

Description

Triggers a policy that is only evaluated for protocol messages that are using the MAC address of the cable modem. To evaluate this condition, the MPE device must be configured with cable modem provisioning information.

where the device name *matches one of specified name(s)*

Syntax
where the device name *matches-op match-list*

Parameters
matches-op

See common parameters.

match-list

See common parameters.

Description

Triggers a policy based on whether the device name matches one or more wildcard match patterns.

where the device type *is specified type*

Syntax
where the device type *operator-binary device-type*

Parameters
operator-binary

See common parameters.

device-type

One or more of the following:

- **CMTS**

- **Blade**
- **Channel**
- **Cable Modem**
- **CPE**

Description

Triggers a policy based on the device type for which it is evaluated.

where the endpoint IP address is in *specified subnet*

Syntax
where the endpoint IP address is in *subnet*

Parameters
subnet

See common parameters.

Description

Triggers a policy that is only evaluated for endpoints whose IP address falls within a specific subnet.

where the endpoint IP address is *specified address*

Syntax
where the endpoint IP address is *ip-address*

Parameters
ip-address

See common parameters.

Description

Triggers a policy that is only evaluated for a specific endpoint (based on its IP address).

where the network element name *matches one of specified name(s)*

Syntax
where the network element name *matches-op csv*

Parameters
matches-op

See common parameters.

csv

Comma-separated list of values.

Description

Triggers a policy based on the name of the network element for which it is being evaluated.

where the network element type *is specified type*

Syntax
where the network element type *operator-binary element-type*

Parameters
operator-binary

See common parameters.

element-type

One or more of the following:

- **CMTS**

Description

Triggers a policy based on the type of network element for which it is being evaluated. Note that if the policy is being evaluated for a device that is not a network element but is contained within a network element (such as an interface within a router) then the network element "container" is used as the basis of comparison.

where the request is not using the cable modem IP address

Description

Triggers a policy that is only evaluated for protocol messages that are using the IP address of the cable modem. In order to know this, the MPE device must be configured with cable modem provisioning information.

where the request is using the cable modem IP address

Syntax
where *number* is the CMTS blade index

Description

Triggers a policy that is only evaluated for protocol messages that are not using the IP address of the cable modem. In order to know this, the MPE device must be configured with cable modem provisioning information.

Network Device Usage Conditions

Network Device Usage conditions are related to the calculated usage for the network device for which the policy rule is being evaluated. This usage includes device-level tracking of both bandwidth and flow/session counts.

where the device will be handling *greater than #* bps *reserved* bandwidth in total for *specified class of* traffic

Syntax
where the device will be handling *operator bandwidth* bps *qos-status* bandwidth in total for *class-of-service* traffic

Parameters
operator

See common parameters.

bandwidth

See common parameters.

qos-status

See common parameters.

class-of-service

See common parameters.

Description

Triggers a policy based on the total amount of bandwidth allocated for specific classes of service by the current device as it relates to a defined threshold. This can be further qualified by the allocation status of the bandwidth. The total represents the bandwidth that is allocated if the current request is approved.

where the device will be handling *greater than #* bps *upstream reserved* bandwidth

Syntax
where the device will be handling *operator bandwidth* bps *qos-direction qos-status* bandwidth

Parameters
operator

See common parameters.

bandwidth

See common parameters.

qos-direction

See common parameters.

qos-status

See common parameters.

Description

Triggers a policy based on the total amount of bandwidth used by the current device as it relates to a defined threshold. This can be further qualified by both the direction and allocation status of the bandwidth. The total represents the bandwidth that is allocated if the current request is approved.

where the device will be handling *greater than #* bps *upstream reserved* bandwidth in total for *specified application*

Syntax
where the device will be handling *operator bandwidth* bps bandwidth *qos-direction qos-status* bandwidth in total for *app-name*

Parameters
operator

See common parameters.

bandwidth

See common parameters.

qos-direction

See common parameters.

qos-status

See common parameters.

app-name

Names of the applications that are defined in the CMP database.

Description

Triggers a policy based on the total amount of bandwidth allocated for specific applications by the current device as it relates to a defined threshold. This can be further qualified by both the direction and allocation status of the bandwidth. The total represents the bandwidth that is allocated if the current request is approved.

where the device will be handling *greater than #* percent of *reserved* capacity for *specified class of* traffic

Syntax
where the device will be handling *operator percent* percent of *qos-status* capacity for *class-of-service* traffic

Parameters
operator

See common parameters.

percent

See common parameters.

qos-status

See common parameters.

class-of-service

See common parameters.

Description

Triggers a policy based on the percent of bandwidth capacity allocated for specific classes of service by the current device as it relates to a defined threshold. This can be further qualified by the allocation status of the bandwidth. The total represents the bandwidth that is allocated if the current request is approved.

where the device will be handling *greater than #* percent of *upstream reserved* capacity

Syntax

where the device will be handling *operator percent* percent of *qos-direction qos-status* capacity

Parameters

operator

See common parameters.

percent

See common parameters.

qos-direction

See common parameters.

qos-status

See common parameters.

Description

Triggers a policy based on the percent of bandwidth capacity used by the current device as it relates to a defined threshold. This can be further qualified by both the direction and allocation status of the bandwidth. The total represents the bandwidth that is allocated if the current request is approved.

where the device will be handling *greater than #* percent of *upstream reserved* capacity for *specified application*

Syntax

where the device will be handling *operator percent* percent of *qos-direction qos-status* capacity for *app-name*

Parameters

operator

See common parameters.

percent

See common parameters.

qos-direction

See common parameters.

qos-status

See common parameters.

app-name

Names of the applications that are defined in the CMP database.

Description

Triggers a policy based on the percent of bandwidth capacity allocated for specific applications by the current device as it relates to a defined threshold. This can be further qualified by both the direction and allocation status of the bandwidth. The total represents the bandwidth that is allocated if the current request is approved.

where the device will be handling *greater than # reserved* flows in total for *specified class of* traffic

Syntax

where the device will be handling *operator number qos-status* flows in total for *class-of-service* traffic

Parameters

operator

See common parameters.

number

See common parameters.

qos-status

See common parameters.

class-of-service

See common parameters.

Description

Triggers a policy based on the total number of flows for specific classes of service used by the current device as it relates to a defined threshold. This can be further qualified by the allocation status of the flows. The total represents the number of flows that are allocated if the current request is approved.

where the device will be handling *greater than # upstream reserved* flows

Syntax

where the device will be handling *operator number qos-direction qos-status* flows

Parameters

operator

See common parameters.

number

See common parameters.

qos-direction

See common parameters.

qos-status

See common parameters.

Description

Triggers a policy based on the total number of flows used by the current device as it relates to a defined threshold. This can be further qualified by both the direction and allocation status of the flows. The total represents the number of flows that are allocated if the current request is approved.

where the device will be handling *greater than # upstream reserved* flows in total for *specified application*

Syntax
where the device will be handling *operator number qos-direction qos-status* flows in total for *app-name*

Parameters
operator

See common parameters.

number

See common parameters.

qos-direction

See common parameters.

qos-status

See common parameters.

app-name

Names of the applications that are defined in the CMP database.

Description

Triggers a policy based on the total number of flows for specific applications used by the current device as it relates to a defined threshold. This can be further qualified by both the direction and allocation status of the flows. The total represents the number of flows that are allocated if the current request is approved.

User Conditions

User conditions are related to the subscriber or subscriber account that is associated with the protocol message that triggered the policy rule execution. This includes subscriber-level and account-level tracking of usage. The following conditions are available.

where the account id matches one of *specified id(s)*

Syntax
where the account id matches one of *match-list*

Parameters
match-list

See common parameters.

Description

Triggers a policy that is only evaluated for one or more specific user ID values (based on matching wildcard patterns).

where the account will be handling *greater than #* percent of *upstream reserved* limit

Syntax
where the account will be handling *operator percent* percent of *qos-direction qos-status* limit

Parameters
operator

See common parameters.

percent

See common parameters.

qos-direction

See common parameters.

quo-status

See common parameters.

Description

Triggers a policy based on the percent of the bandwidth limit used by the account as it relates to a defined threshold. This can be further qualified by both the direction and allocation status of the bandwidth. The total is the bandwidth allocated if the request is approved.

where the account will be using *greater than #* bps upstream bandwidth in total for *specified application*

Syntax
where the account will be using *operator bandwidth* bps upstream bandwidth in total for *app-name*

Parameters
operator

See common parameters.

bandwidth

See common parameters.

app-name

Names of applications that are defined in the CMP database.

Description

Triggers a policy based on the total amount of bandwidth allocated for specific applications by the associated account as it relates to a defined threshold. The total represents the bandwidth that is allocated if the current request is approved. See *Managing Application Profiles* for information on applications.

where the account will be using *greater than #* bps *reserved* bandwidth in total for *specified class of* traffic

Syntax

where the account will be using *operator number* bps *qos-status* bandwidth in total for *class-of-service* traffic

Parameters

operator

See common parameters.

number

See common parameters.

quo-status

See common parameters.

class-of-service

See common parameters.

Description

Triggers a policy based on the total amount of bandwidth for specific classes of service used by the associated accounts as it relates to a defined threshold. This can be further qualified by the allocation status of the bandwidth. The total represents the amount of bandwidth that are allocated if the current request is approved.

where the account will be using *greater than #* bps *upstream reserved* bandwidth

Syntax

where the account will be using *operator bandwidth* bps *qos-direction qos-status* bandwidth

Parameters

operator

See common parameters.

bandwidth

See common parameters.

qos-direction

See common parameters.

qos-status

See common parameters.

Description

Triggers a policy based on the total amount of bandwidth used by the account as it relates to a defined threshold. This can be further qualified by both the direction and allocation status of the bandwidth. The total is the bandwidth allocated if the request is approved.

where the account will be using *greater than # reserved* flows in total for *specified class of* traffic

Syntax

where the account will be using *operator number qos-status* flows in total for *class-of-service* traffic

Parameters

operator

See common parameters.

number

See common parameters.

qos-status

See common parameters.

class-of-service

See common parameters.

Description

Triggers a policy based on the total number of flows for specific classes of service used by the associated accounts as it relates to a defined threshold. This can be further qualified by the allocation status of the flows. The total represents the number of flows that are allocated if the current request is approved.

where the account will be using *greater than #* upstream flows in total for *specified application*

Syntax

where the account will be using *operator number* upstream flows in total for *app-name*

Parameters

operator

See common parameters.

number

See common parameters.

app-name

Names of applications that are defined in the CMP database.

Description

Triggers a policy based on the total number of flows for specific applications used by the associated accounts as it relates to a defined threshold. The total represents the number of flows that are allocated if the current request is approved. See *Managing Application Profiles* for information on applications.

where the account will be using *greater than # upstream reserved* flows

Syntax
where the account will be using *operator number qos-direction qos-status* flows

Parameters
operator

See common parameters.

number

See common parameters.

qos-direction

See common parameters.

qos-status

See common parameters.

Description

Triggers a policy based on the total number of flows used by the associated account as it relates to a defined threshold. This can be further qualified by both the direction and allocation status of the flows. The total represents the number of flows that are allocated if the current request is approved.

where the tier *is* one of *specified tier(s)*

Syntax
where the tier *operator* one of *tiers*

Parameters
operator

See common parameters.

tiers

A comma-separated list of names of one more tiers defined in the CMP database.

Description

Triggers a policy that is or is not evaluated for one or more specific tiers.

where the User's Tier *downstream* bandwidth limit is between # bps and # bps

Syntax
where the User's Tier *qos-direction* bandwidth limit is between *bandwidth* bps and *bandwidth* bps

Parameters
qos-direction

See common parameters.

bandwidth

See common parameters.

Description

Triggers a policy that is evaluated for a user tier based on the bandwidth limit. This can be further qualified by the direction of the bandwidth.

Example

```
where the User's Tier downstream bandwidth limit is between 2M bps and 25M
 bps
```

where the User's Tier *downstream* bandwidth limit is *greater than* # bps

Syntax
where the User's Tier *qos-direction* bandwidth limit is *operator bandwidth* bps

Parameters
qos-direction

See common parameters.

operator

See common parameters.

bandwidth

See common parameters.

Description

Triggers a policy that is evaluated for a user tier based on the comparison between the bandwidth limit and a numerical value. This can be further qualified by the direction of the bandwidth.

Example

```
where the User's Tier downstream bandwidth limit is less than or equal to
 25M bps
```

Policy Context Property Conditions

Policy Context Properties are related to policy contexts.

where the policy context property *name exists*

Syntax
where the policy context property *property-name accessibility*

Parameters
property-name

String.

accessibility

One of the following:

- **exists** (the default)
- **does not exist**

Description

Triggers a policy based on whether or not the specified policy context property exists.

where the policy context property *name* is numerically *equal to value*

Syntax
where the policy context property *property-name* is numerically *operator value*

Parameters
property-name

String.

operator

See common parameters.

value

Integer value in the inclusive range of -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807 (that is, -2^{63} to $2^{63} - 1$).

Description

Triggers a policy based on a numerical comparison between the specified policy context property value and a specified value.

where the policy context property *name matches one of `value(s)`*

Syntax
where the policy context property *property-name matches-op `match-list`*

Parameters

property-name

String.

matches-op

See common parameters.

match-list

See common parameters.

Description

Triggers a policy based on whether the specified policy context property value matches a list of specified values (based on matching wildcard patterns).

Time-of-Day Conditions

Time-of-Day conditions are related to the time at which the policy rules are being executed.

where the current time *is* between *start time* and *end time* using *configured local time*

Syntax

where the current time *operator-binary* between *time-of-day* and *time-of-day* using *time-zone*

Parameters

operator-binary

See common parameters.

time-of-day

A time, in the format of *hh*:*mm*, where *hh* is a number in the range from 0 to 23.

time-zone

One of the following:

- **CONFIGURED LOCAL TIME** (the default) — Calculate the time from the location configured for this MPE device
- **SYSTEM LOCAL TIME** — Calculate the time from the location of this MPE device
- **USER LOCAL TIME** — Calculate the time from the location configured for the user equipment's location

Description

Triggers a policy based on time. If time-zone information is available from the user equipment, time can be calculated from either the MPE device or the user equipment's location.

where today is a week day using *configured local time*

Syntax
where today is a week day using *time-zone*

Parameters
time-zone

One of the following:

- **CONFIGURED LOCAL TIME** (the default) — Calculate the time from the location configured for this MPE device
- **SYSTEM LOCAL TIME** — Calculate the time from the location of this MPE device
- **USER LOCAL TIME** — Calculate the time from the location configured for the user equipment's location

Description

Triggers a policy based on the day of the week. If time-zone information is available from the user equipment, time can be calculated from either the MPE device or the user equipment's location.

where today is a weekend day using *configured local time*

Syntax
where today is a weekend day using *time-zone*

Parameters
time-zone

One of the following:

- **CONFIGURED LOCAL TIME** (the default) — Calculate the time from the location configured for this MPE device
- **SYSTEM LOCAL TIME** — Calculate the time from the location of this MPE device
- **USER LOCAL TIME** — Calculate the time from the location configured for the user equipment's location

Description

Triggers a policy based on the day of the week. If time-zone information is available from the user equipment, time can be calculated from either the MPE device or the user equipment's location.

where today *is day* using *configured local time*

Syntax
where today *operator-binary day-of-week* using *time-zone*

Parameters
operator-binary

See common parameters.

day-of-week

One of the following:

- **Sunday**
- **Monday**
- **Tuesday**
- **Wednesday**
- **Thursday**
- **Friday**
- **Saturday**

time-zone

One of the following:

- **CONFIGURED LOCAL TIME** (the default) — Calculate the time from the location configured for this MPE device
- **SYSTEM LOCAL TIME** — Calculate the time from the location of this MPE device
- **USER LOCAL TIME** — Calculate the time from the location configured for the user equipment's location

Description

Triggers a policy based on the day of the week. If time-zone information is available from the user equipment, time can be calculated from either the MPE device or the user equipment's location.

Actions Available for Writing Policy Rules

The policy wizard supports a large number of actions that can be used for constructing policy rules. There are two types of actions:

- **Mandatory policy-processing actions** — This action defines what should happen when the current policy is through executing. When you are creating a policy rule in the policy wizard, these actions are displayed at the top of the list of available actions with a radio button that forces you to select only one of these actions.
- **Optional actions** — This action contains a list of optional actions that you can add to your policy rule. These actions are then executed when the policy rule's conditions have been met. You can select none, one, several, or all of these optional actions. However, each action is limited, so that it can be executed only once per policy rule.

In the same way that you can customize the conditions by editing parameters, many of these actions can be customized by specifying parameter values as well. Actions are listed in alphabetical order. Actions also may be affected by the current mode; hence, some of the actions documented here may not be available in your policy wizard.

Mandatory Policy-Processing Actions

Policy-processing actions define what the Policy Engine should do when the current policy is through executing. The following are the mandatory policy-processing actions; one of these actions must be selected in each policy.

accept message

Description

After executing the current policy rule, the Policy Engine continues with the normal processing of the protocol message but no further policy rules are evaluated.

break from policy level

Description

Stop evaluating the current policy and continue policy evaluation with the next policy at the parent's level. You should use this action only in reference policies.

continue processing message

Description

After executing the current policy rule, the Policy Engine continues with the next policy rule.

reject message

Description

After executing the current policy rule, the Policy Engine terminates all policy-rule processing and rejects the current protocol message. The specific interpretation of “rejecting” the message varies depending on the associated protocol. For most application-level requests this translates into some type of error being sent back to the application.

skip to next device

Description

Stop evaluating policies for the current device and continue policy evaluation with the next device. If there is no next device, policy execution ends.

skip to next flow

Description

Stop evaluating policies for the current flow and continue policy evaluation with the next flow. If there is no next flow, evaluation continues with the next device; if there is no next device, policy execution ends.

Optional Policy-Processing Actions

The following optional policy-processing actions are available.

Advanced: set values for QoS and Charging parameters to *specified value*

Syntax

Advanced: set values for QoS and Charging parameters to *profile-param*

Parameters

profile-param

Names of profile parameters that are derived from internal representations of protocol messages. This list is lengthy and subject to change as new protocols are supported, and therefore is not given here. The policy wizard includes a customized dialog to help you in the selection of valid values. For the specific meaning of the fields it may be necessary to consult protocol specifications.

Description

Overwrites the corresponding settings in the current protocol message. If you specify settings that are not relevant in the current protocol message, they are ignored. If you select Diameter Enforcement Session Event Triggers, you are presented with another dialog where you can select ECGI_CHANGE and TAI_CHANGE, in addition to the list of previous triggers.

apply *specified profile(s)* to all flows in the request

Syntax

apply *traffic-profile* to all flows in the request

Parameters

traffic-profile

One or more traffic profiles. For more information on traffic profiles, see *Managing Traffic Profiles*.

Description

This parameter allows you to choose different traffic profiles to apply to different types of calls.

apply *specified profile(s)* to request

Syntax
apply *traffic-profile* to request

Parameters
traffic-profile

One or more traffic profiles. For more information on traffic profiles, see *Managing Traffic Profiles*.

Description

Overwrites the corresponding settings in the current protocol message. If multiple traffic profiles are selected they are applied in the order in which they are specified. If the traffic profile contains settings that are not relevant in the current protocol message, they are ignored.

apply *specified profile(s)* to selected *specified type(s)* flows in the request

Syntax
apply *traffic-profile* to selected *media-type* flows in the request

Parameters
traffic-profile

One or more traffic profiles. For more information on traffic profiles, see *Managing Traffic Profiles*.

media-type

One or more of the following, used to determine the type of media:

- **Audio**
- **Video**
- **Data**
- **Application**
- **Control**
- **Text**
- **Message**
- **Other**

Description

Overwrites the corresponding settings in the protocol messages of the specified type. If multiple traffic profiles are selected, they are applied in the order in which they are specified. If the traffic profile contains settings that are not relevant in the current protocol message, they are ignored. The second parameter lets you choose different traffic profiles to apply to different types of calls.

clear alarm with severity `*severity level*`, id `*unique alarm identifier*` and message `*message text*`

Syntax
clear alarm with severity `*level*`, id `*alarm-id*` and message `*message*`

Parameters
level

One of the following, used to determine which alarm ID is cleared:

- **Critical** (ID 74000)
- **Major** (ID 74001)
- **Minor** (ID 74002)

alarm-id

The alarm ID. If you select **Evaluate as expression**, the text in the field is evaluated as an arithmetic expression, and the result is used.

message

String. This text may contain policy parameters (described later in this section) to perform parameter substitution within the message text. If you select **Evaluate as expression**, the text in the field is evaluated as an arithmetic expression, and the result is used.

Description

Clears an alarm on the CMP Active Alarms display containing the specified severity level and message text. This notification is written to the Alarm History Report with severity Clear. To be cleared, a notification must be uniquely identified by severity and alarm ID. For more information, see *Viewing Active Alarms*.

enable event messaging for this request

Description

Enables event messaging for the current message, using the default Event Messaging parameters for this MPE device. If there is no EventGenerationInfo object in the current message, a new one is added.

evaluate policy group *select policy group*

Syntax
evaluate policy group *group-name*

Parameters
group-name

Name of a policy group defined in the CMP database.

Description

If the conditions evaluates to true, evaluate the rules in a policy group. When you click on the **select policy group** parameter, a pop-up window opens so you can select an existing policy group.

evaluate policy *select policy*

Syntax
evaluate policy *policy-name*

Parameters
policy-name

Name of a policy defined in the CMP database.

Description

If the conditions evaluate to true, evaluate a policy. When you click on the **select policy parameter**, a pop-up window opens, giving you the choice of selecting an existing policy or creating a new policy. If you click **Create**, a new Policy Wizard tab opens so you can create the new policy. When you save the new policy, it is added to the list of policies available for selection at this point.

overwrite DSCP/TOS field with

Syntax
overwrite DSCP/TOS field with *dscp*

Parameters
dscp

A numeric representation of DSCP bits to be inserted into the message.

Description

Overwrites the DSCP/TOS field with a value. Although this is a number, the policy wizard includes a customized dialog to help you construct the value.

overwrite SessionClassId with

Syntax
overwrite SessionClassId with *number*

Parameters
number

See common parameters.

Description

Overwrites the SessionClassId field in the message with the specified value.

remove all policy context properties

Description

Removes all policy context properties.

remove policy context property *name*

Syntax

remove policy context property *property-name*

Parameters

property-name

String. May contain policy rule variables (see *Policy Rule Variables*) to perform parameter substitution within the property name.

Description

Removes a subscriber property in the SPR.

send notification to syslog with `*message text*` and severity `*severity level*`

Syntax

send notification to syslog with `*message*` and severity `*level*`

Parameters

message

String. This text may contain policy parameters (described later in this section) to perform parameter substitution within the message text. If you select **Evaluate as expression**, the text in the field is evaluated as an arithmetic expression, and the result is used.

level

The sevlog severity. One of the following:

- **Emergency**
- **Alert**
- **Critical**
- **Error**
- **Warning**
- **Notice**
- **Info**
- **Debug**

Description

Sends a message to the syslog service containing the specified message text and at the specified severity level.

Note: Policies written before V7.5 that used the action **send alert with `*text*` and severity `*severity level*`** will be converted to use this action instead, which will send a notification to syslog instead of an alarm to the CMP system.

send notification to trace log with `*message text*` and severity `*severity level*`

Syntax
send notification to trace log with `*message*` and severity `*level*`

Parameters
message

String. This text may contain policy parameters (described later in this section) to perform parameter substitution within the message text. If you select **Evaluate as expression**, the text in the field is evaluated as an arithmetic expression, and the result is used.

level

One of the following:

- **Emergency** (ID 4560)
- **Alert** (ID 4561)
- **Critical** (ID 4562)
- **Error** (ID 4563)
- **Warning** (ID 4564)
- **Notice** (ID 4565)
- **Info** (ID 4566)
- **Debug** (ID 4567)

Description

Sends a message to the trace log containing the specified message text and at the specified severity level. If the configured minimum notification severity level is higher than that specified in the policy action, then the policy action does not generate the notification.

Note: Policies written before V7.5 that used the action **write `text` to the log file** will be converted to use this action instead, with the severity Info.

set alarm with severity `*severity level*`, id `*unique alarm identifier*` and message `*message text*`

Syntax
set alarm with severity `*level*`, id `*alarm-id*` and message `*message*`

Parameters
level

One of the following:

- **Critical** (ID 74000)
- **Major** (ID 74001)
- **Minor** (ID 74002)

alarm-id

The alarm ID. If you select **Evaluate as expression**, the text in the field is evaluated as an arithmetic expression, and the result is used.

message

String. This text may contain policy parameters (described later in this section) to perform parameter substitution within the message text. If you select **Evaluate as expression**, the text in the field is evaluated as an arithmetic expression, and the result is used.

Description

Sends an alarm to the CMP system containing the specified severity level and message text. This alarm is written to the Alarm History Report, and will appear in the Active Alarms display for one hour, until cleared, or unless the server fails over, whichever comes first. Alarms generated by policy actions do not affect the HA score of a server, and will not cause a failover. For more information, see *Viewing Active Alarms*.

set *external field* to # percent of *select type* for *selected* quota

Syntax

set *field* to *value* percent of *type* for *quota-name* quota

Parameters

field

String name of field in external database.

value

String name of field in external database.

type

One of the following:

- **service-specific**
- **time**
- **total volume**
- **uplink volume**
- **downlink volume**

quota-name

Name(s) of quotas defined in the CMP database.

Description

Sets a field in an external database to a percentage of the time, total volume, or service-specific quota of one or more selected quotas. This can be an LDAP server. The MPE device on which this policy is

executed must have write access to the database, and the external field must be defined on the MPE device. For more information, see *Configuring Data Source Interfaces*.

set *external field* to `*value*`

Syntax
set *field* to `*value*`

Parameters
field

String name of field in external database.

value

String value of field in external database. If you select **Evaluate as expression**, the text in the field is evaluated as an arithmetic expression, and the result is used.

Description

Sets the value of a field in an external database. This can be an LDAP server. The MPE device on which this policy is executed must have write access to the database, and the external field must be defined on the MPE device. For more information, see *Configuring Data Source Interfaces*.

set policy context property *name* to *value*

Syntax
set policy context property *property-name* to *value*

Parameters
property-name

String. May contain policy rule variables (see *Policy Rule Variables*) to perform parameter substitution within the property name.

value

String.

Description

Sets and saves a subscriber property in the SPR. You can specify that the property is not saved if the policy rejects the message.

set time limit to # seconds

Syntax
set time limit to *seconds* seconds

Parameters
seconds

See common parameters.

Description

Overwrites the time limit in the current message. If there is no TimeLimit object in the current message, a new one is added with the specified value.

set volume limit to # kilobytes

Syntax
set volume limit to *bandwidth* kilobytes

Parameters
bandwidth

See common parameters.

Description

Overwrites the volume limit in the current message. If there is no VolumeLimit object in the current message, a new one is added with the specified value.

Policy Rule Variables

During policy rule execution within the MPE device, some actions (for example, **`send notification`**) allow for substitution of policy rule variables with contextual information. Each time the policy rules are evaluated, the unique set of policy rule variables is referred to as the *policy context*. This section summarizes these policy rule variables.

Using Policy Rule Variables

Typically, policy rule variables are used to perform substitution of textual information into a text message that is being used for some type of logging. This is typically done in an action. To use a policy rule variable, insert the variable into the text message when you define the action.

The format of a policy rule variable is as follows:

`“{” name [ “:” default-value ] “}”`

The name can contain the characters A–Z, a–z, 0–9, underscore (_), period (.), and backslash (\).

The following are examples of policy rule variables:

```
{Bandwidth}
{Device.Name}
{Device.Name:UNKNOWN}
```

Basic Policy Rule Variables

Table 19: Basic Policy Rule Variables displays some of the basic policy rule variables that are available.

Under certain circumstances the MPE device can associate additional context information with a request. This information may be used during the policy rule execution. The availability of this information depends on:

- The mode (for example, cable) in which the MPE device is executing
- Whether the information is provisioned on the MPE device or, if present, a Subscriber Profile Repository (SPR)
- The protocol in use and how much information is available in the request (some protocols have optional information which, if specified, can be used to associate additional information)

There are a number of policy rule variables that can be used to provide information about the device for which a policy rule is being executed. Some of these variables are only available for certain device types, while others are available for all devices.

Table 19: Basic Policy Rule Variables

Variable Name	Description	Modes, Protocols, Device Type
{Policy}	The name of the policy rule that is being executed.	--
{Date}	The date when the policy rule is executed,	--
{Time}	The time when the policy rule is executed, in the format *hh*:*mm*:*ss*.*SSS*.	--
{Conditions}	A list of (variable, value) tuples that lists the variables whose values were referenced in the conditions of the policy rule. The list is inserted with one variable per line in the format *variable=value*.	--
{Device}	The name of the device for which the policy rule is being evaluated.	--
{DeviceId}	ID of the device for which the policy rule is being evaluated.	--
{QosDir}	The direction of the flow for which the policy rule is being evaluated, either “Up” or “Down.”	--
{Bandwidth}	The DOCSIS type of the flow for which the policy rule is being evaluated: “BES,” “NRTP,” “RTP,” “UGS,” or “UGSAD.”	--
{Device.Name}	The name (as defined in the CMP database) of the device.	Any
{Device.UpstreamCapacity}	The upstream bandwidth capacity of the device.	Any

Variable Name	Description	Modes, Protocols, Device Type
{Device.DownstreamCapacity}	The downstream bandwidth capacity of the device.	Any
{Device.FlowCount}	The number of active flows for the device.	Any
{Element.Name}	The name (as defined in the CMP database) of the network element associated with the current device. If the device is a network element, then this is the same as the {Device.Name}. However, if the device is contained within a network element (as is the case with Interfaces, Channels, and so forth), then this will have a different value.	Any
{Element.Hostname}	The hostname (or IP address) of the network element associated with the current device. If the device is a network element, then this is the same as the {Device.Name}. However, if the device is contained within a network element (as is the case with Interfaces, Channels, and so forth), then this will have a different value.	Any
{Element.BackupHostname}	The hostname (or IP address) of the backup network element associated with the current device. If the device is a network element, then this is the same as the {Device.Name}. However, if the device is contained within a network element (as is the case with Interfaces, Channels, and so forth), then this will have a different value.	Any
{Element.UpstreamCapacity}	The upstream bandwidth capacity of the network element associated with the current device. If the device is a network element, then this is the same as the {Device.Name}. However, if the device is contained within a network element (as is the case with Interfaces, Channels, and so forth), then this will have a different value.	Any

Variable Name	Description	Modes, Protocols, Device Type
{Element.DownstreamCapacity}	The downstream bandwidth capacity of the network element associated with the current device. If the device is a network element, then this is the same as the {Device.Name}. However, if the device is contained within a network element (as is the case with Interfaces, Channels, and so forth), then this will have a different value.	Any
{Session.IMEI}	This variable expands to the IMEI of the subscriber's phone or equipment associated with the request.	Any
{Session.IMEISV}	This variable expands to the IMEISV of the subscriber's phone or equipment associated with the request.	Any

Chapter 15

Managing Policy Rules

Topics:

Policy rules are created and saved within the CMP database and then deployed to MPE devices. The CMP system lets you create and modify the details within policy rules, as well as edit the order in which policy rules are applied to a protocol message.

To create policy rules, see *Understanding and Creating Policy Rules*. *Managing Policy Rules* describes how to manage your library of policy rules and policy groups.

Displaying a Policy

To display a policy:

1. From the **Policy Management** section of the navigation pane, select **Policy Library.**
   The content tree displays a list of policy library groups; the initial group is **ALL**. If a policy references another policy or policy group, a gear icon (⚙) appears next to the policy name in the content tree.
2. From the content tree, select the desired policy.
   The policy is displayed. *Figure 19: Sample Policy Description* shows an example.

Figure 19: Sample Policy Description

You can choose from two logical views of policy conditions:

- A tree format (the default, shown)
- A Boolean expression format similar to SQL

To switch between one view and the other, click **Toggle View**.

If the policy evaluates a policy group, the policies in the group (which are referenced policies) are displayed. Click on a policy name to see details of that policy. If a referenced policy itself refers to other policies or groups, those policies or groups are also displayed.

Deploying Policy Rules

Deploying a policy (or policy group) is the act of transferring the policy from the CMP policy database to an MPE device. Once deployed, the policy rules defined within the policy or policy group are used as decision-making criteria by the MPE device.

Figure 20: Policy Deployment shows how policies P1 through P7 are created in the CMP database and then deployed individually to different MPE devices within the network. Each of the policies is associated individually with the MPE device where it is deployed. In the example, each policy server (MPE device) displays the policies that have been deployed to it and the order in which they are applied to policy requests, from top to bottom.

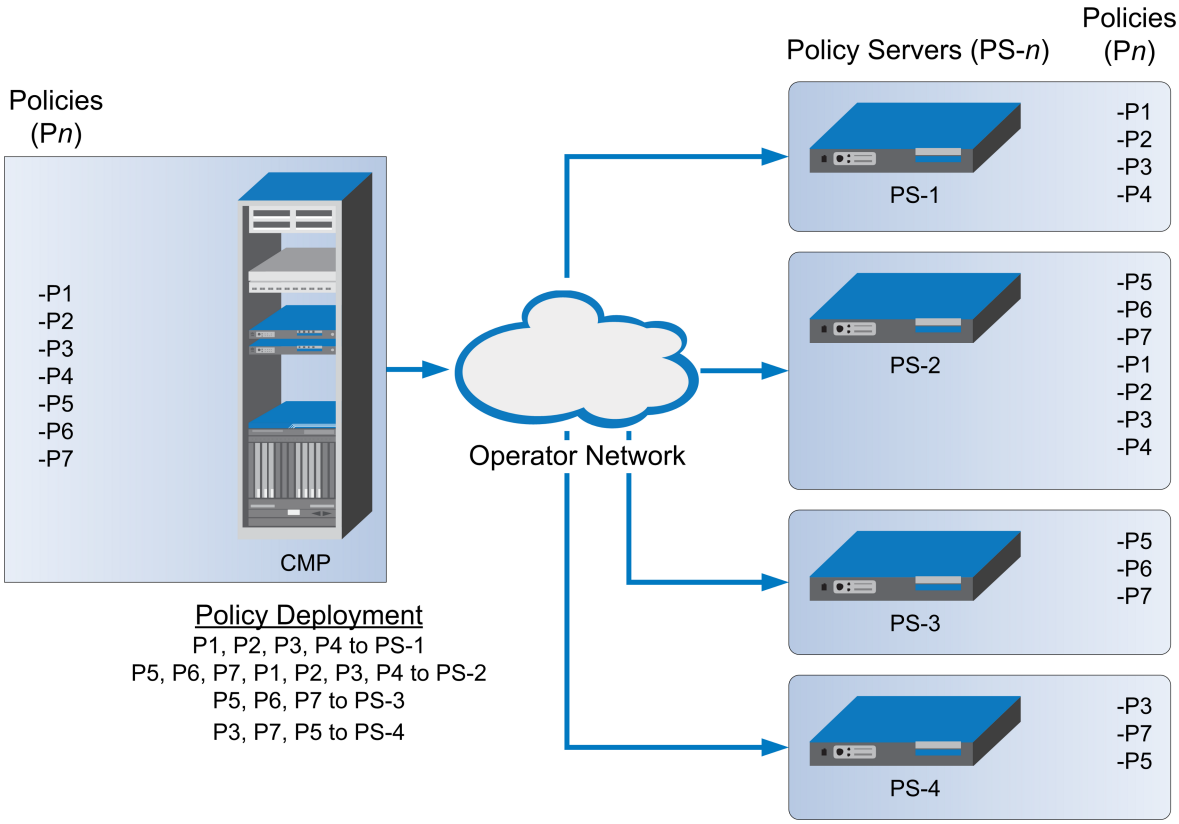

Figure 20: Policy Deployment

Figure 21: Policy Group Deployment shows how the same library of policies can be grouped first and then deployed as policy groups. When a policy group is created, the policies are arranged in the order in which they are to be evaluated. Grouping policies makes deployment of multiple policies easier and helps to ensure consistency in how policies are applied to policy requests on different MPE devices.

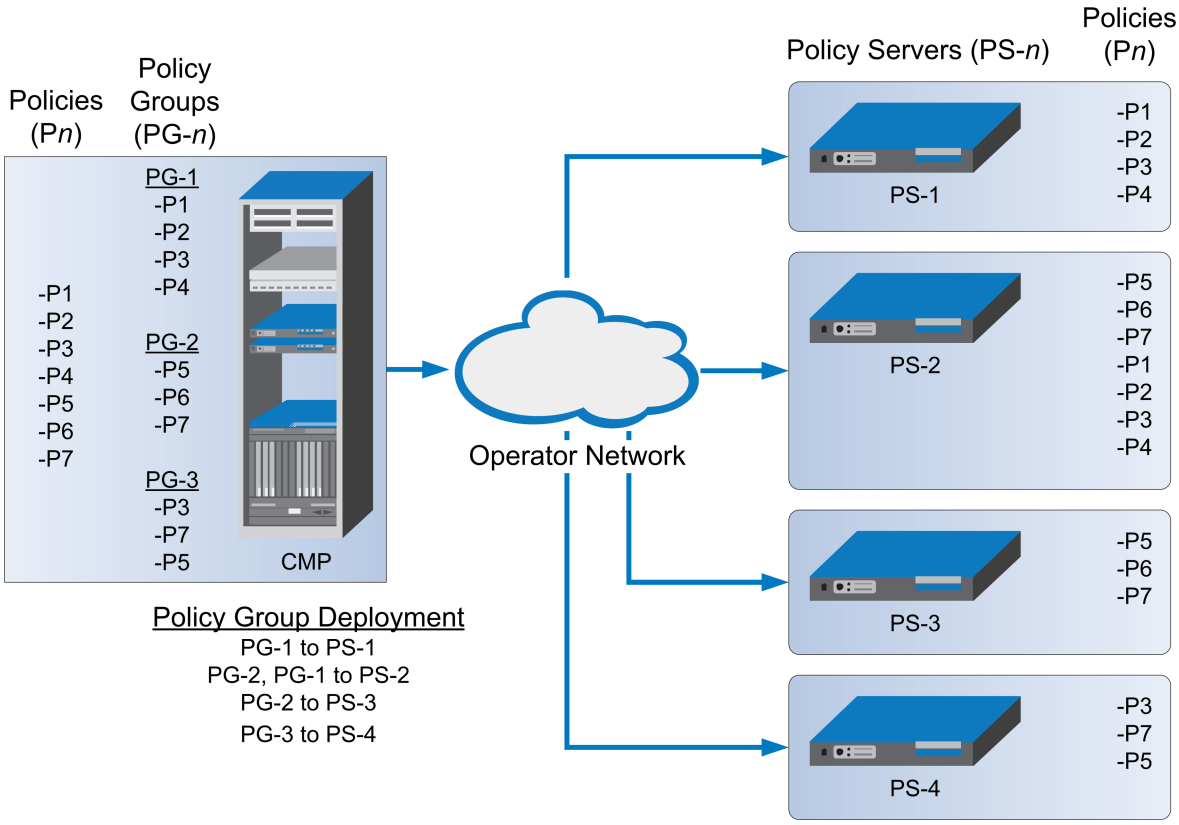

Figure 21: Policy Group Deployment

When you first create a policy rule, that rule exists only within the CMP database. Once the policy rule is deployed, any change to the policy rule is automatically redeployed when you complete your changes. Automatic redeployment also applies to policy groups as well: any change to a policy group triggers automatic redeployment. If you add a policy rule that was not previously deployed to a policy group that is deployed to one or more MPE devices, then the rule is deployed automatically to those MPE devices.

Figure 22: Policy Redeployment shows that when a policy (P3) is modified, its associated groups (PG-1 and PG-3) are redeployed automatically.

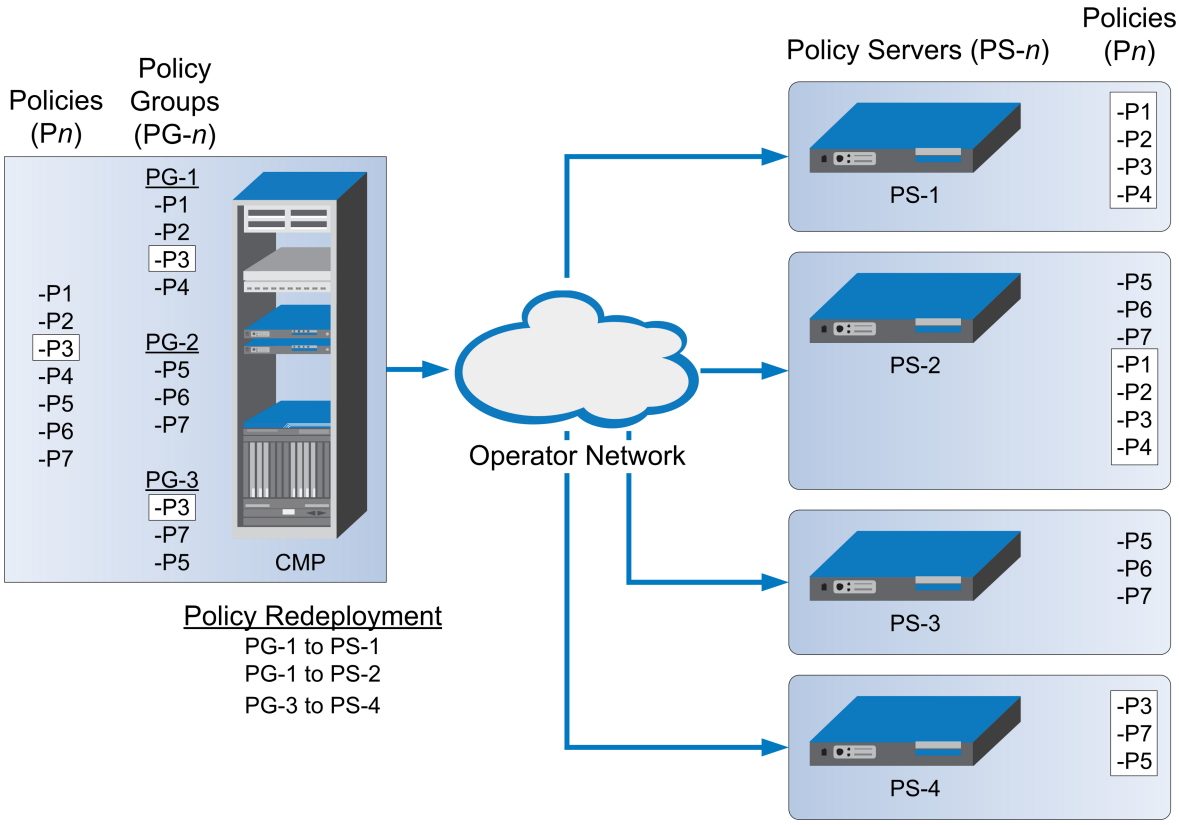

Figure 22: Policy Redeployment

When a policy rule is used as a reference policy, you do not need to deploy it; it is deployed automatically when called by a parent, or top-level, policy.

Modifying and Deleting a Policy

Policies can be modified and then redeployed to MPE devices. When a policy that resides in multiple policy groups is modified, the changes are propagated to the various groups.

Modifying a Policy

To modify an existing policy:

1. From the **Policy Management** section of the navigation pane, select **Policy Library.**
   The content tree displays a list of policy library groups; the initial group is **ALL**.
2. From the content tree, select the **ALL** group.
   The Policy Administration page opens in the work area, listing the available policies.
3. Select the policy you want to edit.
   The Policy Administration page displays information about the policy.
4. Click **Modify**.

The policy wizard opens in a Modify Policy tab.

5. Edit the desired policy information.
   See *Creating a New Policy* for details on the fields within the policy wizard.
6. When you finish, click **Finish** (or **Cancel** to discard your changes).

The policy is modified. The modified policy is now ready to be added to a policy group (see *Adding a Policy or a Policy Group to a Policy Group*), or deployed to one or more MPE devices (see *Deploying a Policy or Policy Group to MPE Devices*).

Note: Redeployment of a policy is automatically performed to those MPE devices where the policy was initially deployed.

Deleting a Policy

Policies, policies within a policy group, and entire policy groups can be removed from an MPE device when they are no longer needed. Because the policy still resides in the CMP database, it can be redeployed at a later date if needed. If a policy is no longer needed, it can be deleted from the CMP database as well.

Note: Deleting a policy from the CMP database automatically removes the policy from all associated MPE devices.

To delete a policy:

1. From the **Policy Management** section of the navigation pane, select **Policy Library**.
   The content tree displays a list of policy groups; the initial group is **ALL**.
2. From the content tree, select the **ALL** group.
   The Policy Administration page opens in the work area, displaying all defined policies.
3. Use one of the following methods to select the policy to delete:
   - From the work area, click the **Delete** icon located to the right of the policy you want to delete.
   - From the policy group tree, select the policy; the Policy Administration page opens. Click **Delete**.

   You are prompted, "Are you sure you want to delete this Policy?"
4. Click **OK** to delete the policy (or **Cancel** to cancel the request).

The policy is deleted.

To remove a deployed policy from an MPE device, see *Removing a Policy or Policy Group from an MPE Device*.

Policy Templates

The CMP system lets you create policy templates to simplify the creation of multiple policies with similar conditions and actions. A policy template is similar to a policy, except that some (or all) of the parameters in the conditions and actions are not completely defined. Those parameters are defined later, when you use the policy template to create policy rules.

The policy template wizard is used to create or modify a policy template. This wizard is similar to the policy wizard; however, the policy template wizard allows parameters to be only partially defined. For example, a template may only be configured for policy requests requiring bandwidth above a

certain value, but not define the exact bandwidth value. You can then specify a specific bandwidth value when you use the template to create the new policy rule.

Creating a Policy Template

To create a new policy template:

1. From the **Policy Management** section of the navigation pane, select **Template Library**.
   The content tree displays the Template Library group.
2. Select the **Template Library** group.
   The Template Administration page opens in the work area.
3. On the Template Administration page, click **Create Template**.
   The Create New Policy Template window opens.
4. Select the base policy or policy template with which to begin:
   - **Blank** — No policy template attributes are pre-defined.
   - **Use Template** — Select an existing template with pre-defined attributes. Modify the template as needed, then save the template with a new template name.
   - **Copy Existing Policy** — Select an existing policy. Modify the policy as needed, then save the policy as a policy template.
5. Edit the desired policy information from one or more of the policy wizard pages.
   See *Creating a New Policy* for details on the fields within the policy wizard.
6. When you finish, click **Finish** to save the policy template (or **Cancel** to discard your changes).
   The window closes.

The policy template is created.

Modifying a Policy Template

You can edit a policy template to make changes. Modifying a policy template does not modify previously configured policies.

To modify an existing policy template:

1. From the **Policy Management** section of the navigation pane, select **Template Library**.
   The content tree displays the **Template Library** group.
2. Select the **Template Library** group.
   The Template Administration page opens in the work area.
3. Select the template you want to modify.
   The Template Administration page displays a description of the template.
4. Click **Modify**.
   The Modify Policy tab opens with the last step of the template creation process.
5. The wizard begins at the last step of the template creation process. Click **Back** to return to where you want to edit the template and modify the desired information.
6. When you finish, click **Finish** to save the modified template (or **Cancel** to discard your changes).
   The window closes.

The template is modified.

Deleting a Policy Template

To delete a policy template:

1. From the **Policy Management** section of the navigation pane, select **Template Library**.
   The Template Administration page opens in the work area, displaying all defined policy templates.
2. Use one of the following methods to select the policy template to delete:
   - From the work area, click the **Delete** icon, located to the right of the policy template you want to delete.
   - From the template library, select the template; the Template Administration page displays the template. Click **Delete**.

   You are prompted, "Are you sure you want to delete this template?"
3. Click **OK** to delete the policy template (or **Cancel** to abandon the request).

The policy template is deleted.

Managing a Policy Group

The CMP lets you create policy groups. Policy groups are an organizational aid that provide for flexible policy management, deployment, and execution. Policies are saved to a group in the order in which the MPE device applies them to a policy request. If needed, you can change that order. You can save a policy to multiple policy groups and add a policy to, or remove it from, a policy group at any time.

Creating a Policy Group

To create a new policy group:

1. From the **Policy Management** section of the navigation pane, select **Policy Library.**
   The content tree displays a list of policy library groups; the initial group is **ALL**.
2. From the content tree, select the **ALL** group.
   The Policy Administration page opens in the work area, listing available policies.
3. On the Policy Administration page, click **Create Group**.
   The group naming field opens in the work area; for example:

4. Enter the name to assign to the new group.

   The name can be up to 64 characters long and must not contain quotation marks (") or commas (,).

5. Click **Save** (or **Cancel** to discard your changes).

The new group information is saved to the CMP database and displayed in the content tree.

Adding a Policy or a Policy Group to a Policy Group

Once you create a policy group, you can add policies to it. You can also add policy groups to a policy group.

Note: Tekelec recommends that you only nest policy groups two levels deep.

To add one or more policies or policy groups to a policy group:

1. From the **Policy Management** section of the navigation pane, select **Policy Library.**
   The content tree displays a list of policy library groups; the initial group is **ALL**.
2. From the content tree, select the policy group to which you want to add the policy or policy group.
   The Policy Administration page opens in the work area, listing the policies and policy groups currently in the group.
3. On the Policy Administration page, click **Modify**.
   The Policy Administration page opens in the work area.
4. Click **Add**.
   A window opens, displaying the policies and policy groups available.
5. You can optionally filter the list by policies or policy groups. From the pulldown list, select **Policy** to display policies, **Group** to display policy groups, or **All** (the default) to list both policies and policy groups.
6. Select the desired policy or group to add to this group and click **Add** (or **Cancel** to cancel the request). Use Shift/click to select multiple policies or policy groups.
   The policies or policy groups are added to the policy group and the window closes.

Note: Policies or policy groups are applied to messages in the order in which they appear in the policy group. You can change the sequential order as desired (see *Changing the Sequence of Deployed Policies or Policy Groups*).

7. When you finish, click **Save** (or **Cancel** to discard your changes).
   The added policies and policy groups are displayed in the policy group tree.

Now you can deploy the policy group to the policy servers (see *Deploying a Policy or Policy Group to MPE Devices*).

Note: If this group had been deployed previously, it is automatically redeployed at this time, ensuring the MPE devices are resynchronized with the CMP database.

Removing a Policy from a Policy Group

Removing a policy from a policy group that has been saved to the CMP database only removes the policy from the selected policy group. The policy itself remains in the **ALL** group, as well as any other group to which it had been added. (To remove a policy from all groups in the Policy Library, see *Removing a Policy or Policy Group from an MPE Device*.)

To remove a policy from a policy group:

1. From the **Policy Management** section of the navigation pane, select **Policy Library.**
   The content tree displays a list of policy library groups; the initial group is **ALL**.
2. From the content tree, select the desired policy group.
   The Policy Administration page opens in the work area, listing the policies it contains.
3. Remove the desired policy using one of the following methods:
   - From the content tree, select the desired policy within the policy group; its profile information is displayed. Click **Remove**.
   - From the content tree, select the desired policy group and click **Modify**. Select the remove icon, located to the right of the policy you want to remove.

The modified policy group is redeployed, ensuring that the MPE devices are resynchronized with the CMP database.

Note: If the policy group has never been deployed, you can now deploy it to MPE devices (see *Deploying a Policy or Policy Group to MPE Devices*).

Removing a Policy Group

Removing a policy group removes the policy group from all policy groups to which it has been added.

To remove a policy group:

1. From the **Policy Management** section of the navigation pane, select **Policy Library.**
   The content tree displays a list of policy library groups; the initial group is **ALL**.
2. From the content tree, select the desired policy group.
   The Policy Administration page opens in the work area, listing policies and policy groups.
3. From the content tree, select the desired policy group; its profile information is displayed. Click **Delete**.
   You are prompted, "Are you sure you want to delete this Group?"

4. Click **OK** to delete the policy group (or **Cancel** to abandon the request).
   The policy group is removed from the CMP database.

Any policy groups that contained the deleted policy group are redeployed, ensuring that the MPE devices are resynchronized with the CMP database.

Changing the Sequence of Policies or Policy Groups Within a Policy Group

The order in which policies or policy groups appear in a policy group is the order in which they are deployed and applied to policy requests. You can modify the order of policies or policy groups, both inside and outside of a policy group.

To change the order of the policies or policy groups within a group:

1. From the **Policy Management** section of the navigation pane, select **Policy Library.**
   The content tree displays a list of policy library groups; the initial group is **ALL**.
2. From the content tree, select the desired policy group.
   The Policy Administration page opens in the work area, displaying policies or policy groups in their current sequential order.
3. On the Policy Administration page, click **Modify**.
   The Manage Policies page opens.
4. Use any of the following options to change the sequence of policies or policy groups within the group:
   - Use the up and down arrow icons
5. When you finish, click **Save** (or **Cancel** to discard your changes).

The modified policy group is redeployed, ensuring that the MPE devices are resynchronized with the CMP database.

Note: If the policy group has never been deployed, you can now deploy it to MPE devices (see *Deploying a Policy or Policy Group to MPE Devices*).

Displaying Policy Details Contained Within a Policy Group

To display the policies within a policy group:

1. From the **Policy Management** section of the navigation pane, select **Policy Library.**
   The content tree displays a list of policy library groups; the initial group is **ALL**.
2. From the content tree, select the desired policy group.
   The Policy Administration page opens in the work area, listing the policies it contains.
3. Click **Show Details**.
   The configured policies, including the configured parameters for the policies, are displayed. To switch between logical views of policy conditions, click **Toggle View**.
4. When you finish, click **Cancel**.

Deploying a Policy or Policy Group to MPE Devices

The basic procedure for deploying either a policy or a policy group to MPE devices is the same. The following procedure uses the example of deploying a policy group:

1. From the **Policy Management** section of the navigation pane, select **Policy Library.**
   The content tree displays a list of policy library groups; the initial group is **ALL**.
2. From the content tree, select the policy or policy group to deploy.
   The Policy Administration page opens in the work area, listing the policies it contains.
3. On the Policy Administration page, click **Deploy**.
   The policy server tree is displayed, listing all possible target policy servers (MPE devices) and server groups. You can expand the tree view if necessary.
4. Select the desired target MPE devices or policy server groups.

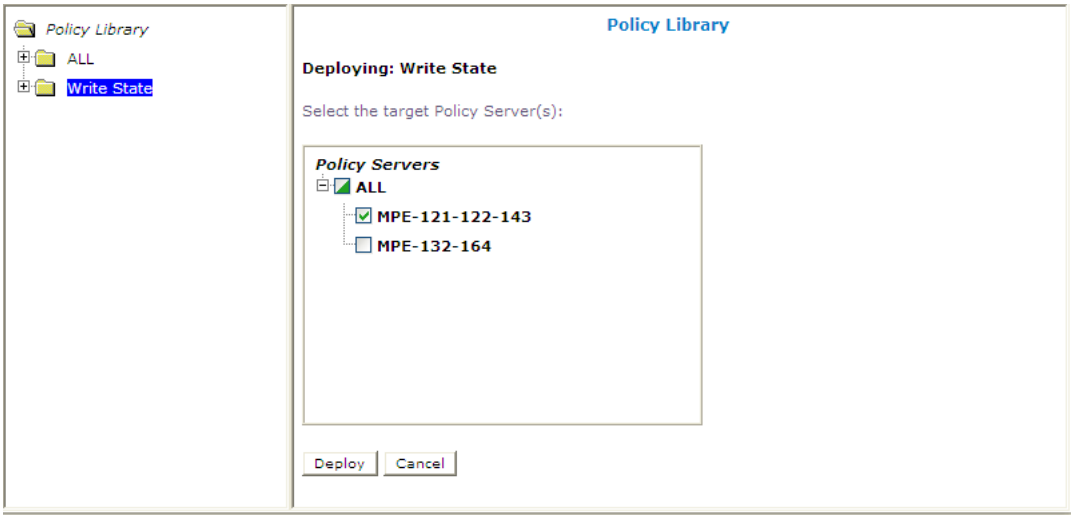

5. Click **Deploy** (or **Cancel** to cancel the request).
   You are prompted, "Policy Servers - Deployment Succeeded" followed by a list of MPE devices to which the policy or policy group was deployed.

The policy information is saved to each selected MPE device.

Removing a Policy or Policy Group from an MPE Device

Removing a deployed policy or policy group from an MPE device is performed from the Policy Server Administration page.

To remove a policy or policy group from 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, displaying information about the MPE device.
3. On the Policy Server Administration page, select the **Policies** tab.
4. Click **Modify**.
   The Manage Policies page opens.
5. Click the Remove icon, located to the right of the policy or policy group that you want to remove.
   The policy or policy group is removed from the list.
6. Repeat step 5 as required.
7. When you finish, click **Save** (or **Cancel** to abandon the request).

You are prompted, "The policies were redeployed successfully to Policy Server '*mpe*'."

The policy or policy group is redeployed to the MPE device, minus the removed policy or policy group.

Removing a Policy or Policy Group from an MPE Device

Removing a deployed policy or policy group from an MPE device is performed from the Policy Server Administration page.

To remove a policy/policy group from 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, displaying information about the MPE device.
3. On the Policy Server Administration page, select the **Policies** tab.
4. Click **Modify**.
   The Manage Policies page opens.
5. Click the Remove icon, located to the right of the policy or policy group that you want to remove.
6. When you finish, click **Save** (or **Cancel** to cancel the request).

The policy or policy group is redeployed to the MPE device, minus the removed policy or policy group.

Changing the Sequence of Deployed Policies or Policy Groups

Changing the sequential order of deployed policies or policy groups is performed directly on an MPE device using the Policy Server Administration page.

To change the sequential order of policies or policy groups:

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, displaying information about the MPE device.
3. On the Policy Server Administration page, select the **Policies** tab.
4. Click **Modify**.
   The Manage Policies page opens in the work area.
5. Use any of the following options to change the sequential positioning of the policies or policy groups:
   - Use the up and down arrow icons
6. When you finish, click **Save** (or **Cancel** to cancel the request).

The policies or policy groups are redeployed to the MPE device in their new sequential order. A confirmation message displays in the work area.

Importing and Exporting Policies, Policy Groups, and Templates

Policies, policy groups, and templates can be exported from the CMP database for inspection or backup purposes. These items are exported as a whole and cannot be exported individually, as every policy, policy group, and policy template in the database is saved to a single file when performing the export function.

For information only, exported policies are marked with policy version numbers as well as the version number of the CMP software under which they were created. This does not affect importation of policies created under different versions of the CMP software.

Importing Policies

To import a policy file into the policy library:

1. From the **Policy Management** section of the navigation pane, select **Policy Import / Export**. The Import/Export page opens.
2. On the Import/Export page, click **Browse** to locate the policy file to import.
3. Select the desired collision handling option:
   - **Delete all before importing** — All policies, policy groups, and templates currently in the CMP database are deleted first; then the imported versions are saved to the MPE device.
   - **Overwrite with imported version** — All items are imported. If the CMP database currently contains any policies, policy groups, or templates using the same names as the ones being imported, they are overwritten with the imported versions.
   - **Reject any that already exist** — All items are imported except for imported versions with the same name as any policy, policy group, or template currently in the CMP database.
   - **Any collisions prevent all importing** (the default) — No items are imported if any of the imported versions has the same name as any policy, policy group, or template currently in the CMP database.
4. Click **Import**.

The policies are imported.

If you try to import an invalid file you receive a validation error: "You must correct the following error(s) before proceeding: There is a problem with the import file. The name is required, the file must be present, and the file must be in the correct format."

Exporting Policies

To export the policies or policy templates that reside in the policy library:

1. From the **Policy Management** section of the navigation pane, select **Policy Import / Export**. The Import/Export page opens.
2. Select the type of export: **Policies** (the default) or **Templates.**
3. Select the policy group to export: **All** (the default) or a named group.
4. Click **Export** to export the policy group in XML format, or **Text** to export the policy group in descriptive format. Policies exported in text format cannot be reimported.

A standard File Download window opens.

5. Click **Save** (or **Cancel** to close the window and cancel the request).
   A standard Save As window opens.
6. Assign a name to the policy file (the default is PolicyExport.xml), use the browse function to map to the desired location, and click **Save**.
   When the policies are successfully exported, a standard Download Complete window opens.
7. Select **Close** to close the Download Complete window.

The policies or templates are exported to a file.

Chapter 16

Managing Policy Tables

Topics:

Managing Policy Tables describes how to create, modify, delete, and view policy tables, which are independent objects that you can use to capture differences in policy structures.

You can manage multiple policies with small differences by abstracting the differences into tables. The process of modifying the policies, or creating new, similar policies then becomes a matter of modifying the policy table, which is simpler and less prone to error.

About Policy Tables

In practical use, many policies are very similar, having only small differences between them. Policy tables are an available option in the policy wizard. A policy table abstracts the differences between related policies.

Using a policy table instead of creating many similar policies makes the tasks of adding new policies, modifying existing sets of policies, and checking consistency among related policies simpler and less prone to error.

Policy tables resemble database tables, and contain the following elements:

- Table name
- Table description
- Column definitions — Every column has a definition that contains a name, data type, and indication if the column is a key column. Every entry in the column must have the same data type. Any data associated with a message, including fields (such as a quota or RAT type) and sub-fields (such as a user account ID or tier name), can be used as a key.
- Policy variable (for key columns only) — Used to obtain the value from the policy context when using the policy table to look up a row.
- Data — The contents of the table cells. (Blank cells are not allowed in a policy table.)

Each row in a policy table can be thought of as a scenario, and each row can replace a policy. Substitutions in policy condition and action parameters can include the values in a specified policy table.

Table 20: Example of a Policy Table shows an example of a simple policy table. The first column lists one or more access point names (APN), and is the key column. The second column contains a PCC rule that will be installed as part of the execution of a policy. The third column contains one or more PCC rules that will be removed as part of the execution of a policy. The second and third columns must contain names of PCC rules defined as traffic profiles in the CMP database.

Table 20: Example of a Policy Table

APN	Install	Remove
apn1.com	pcc_rule_1	pcc_default_1, pcc_basic
apn2.com	pcc_rule_2	pcc_default_2, pcc_basic
apn3.com, apn4.com	pcc_rule_1	pcc_default_1
apn5.com, apn6.com	pcc_rule_2	pcc_default_2

Each policy can have zero or more policy tables. To support the use of multiple policy tables, policies refer to a policy table using an alias. Each policy can use a different alias for the same policy table. For example, a policy table named “PCC rules to install and remove, based on APN” can be referred to in a policy as “pcc_rules.” Policies can use table cells addressed as *table_name.column_name*.

The following policy rule uses the defined policy table. The italicized text represent substitutions. The table references begin with "pcc_rules."

```
use table 'PCC rules to install and remove, based on APN' called 'pcc_rules'
where the request is modifying an existing session
     and where the session is a credit control session
     and where the requested quota is one of Bucket Exceeded,OS_no_TV_volume
     and where the quota usage reporting reason is one of validity time expired
     and where the APN matches one of pcc_rules.apn
     and where the user Custom1 matches one of 101
install pcc_rules.install PCC rule(s) for flow
remove pcc_rules.remove PCC rule(s)
send notification to syslog with
`100;{User.MSISDN};{User.AccountId};{User.IMSI};{Session.IMEI};{Date} {Time};
Info GalacTel : You have a new 500 minutes to enjoy your mobile Internet offer.
Beyond that the flow will be reduced.; {Date} {Time};{Date}
{Time};{User.Custom1};{User.BillingDay}` and severity `Emergency`

accept message
```

The use of policy tables is not required. The decision to use a policy table may arise after you have created a series of production policy rules, if you notice that the policies differ only in a few small ways.

Creating Policy Tables

When you define a policy table, it must contain at least one key column and one row, and you must populate every cell in the table.

To create a policy table:

1. From the **Policy Management** section of the navigation pane, select **Policy Table Library**.
   The content tree displays the Policy Table Library group.
2. Select the **Policy Table Library** group.
   The Policy Table Administration page opens in the work area.
3. On the Policy Table Administration page click **Create Policy Table**.
   The New Policy Table Administration page opens.
4. Enter information as appropriate:
   a) **Name** (required) — The name you assign to the policy table.
      The name can be up to 255 characters long and must not contain quotation marks (") or commas (,).
   b) **Description/Location** (required) — Free-form text that identifies the policy table.
5. Click **Add Row** or **Add Column** (required) — You must define at least one key column.
   If you click **Add Column**, a Policy Table Column window opens. Enter the following information:
   - **Column Name** (required) — Policies will use this name as part of the address of cells in this column.
   - **Key** — If this is a key column, check the box and either select a policy variable from the pulldown list or type the name of the variable you want to use. The policy variable is used to obtain the value from the policy context when using the table to look up a row.

- **Column Type** (required) — The datatype of cells in the column. Click the folder icon; a selection window opens, displaying the Policy Wizard actions and conditions. Locate the condition or action you wish to abstract and select the variable you wish to use (displayed in red text); the datatype is taken from the variable.
- When you finish, click **Save** (or **Cancel** to abandon your changes).

If you click **Add Row**, a row is added below the current row in the table. Select each cell in the row; a window opens so you can enter the value of that cell. The data in cells must match the datatype of the column. Enter the value and click **OK** (or **Cancel** to abandon your changes). You can also enter a comma-separated list of values.

The column or row is displayed.

6. To manage a row or column, select it and click **Operations**, then select from the pulldown list:
   - **Delete Row** — Deletes the table row.
   - **Move Row Up** — Moves the table row up.
   - **Move Row Down** — Moves the table row down.
   - **Delete Column** — Deletes the column in the table.
   - **Move Column Left** — Moves the column left in the table.
   - **Move Column Right** — Moves the column right in the table.
   - **Sort Column** — Sorts the column in the table.
   - **UnSort Column** — Reverts the column to its original order.
7. When you finish defining the table, click **Validate**; the table definition is validated. Validation ensures that tables contain a key column, at least one row, and no empty cells. If the table is invalid, a diagnostic message appears.
8. When you finish, click **Save** (or **Cancel** to discard your changes).
   The policy table is validated, and if valid is displayed on the Policy Table Administration page.

You have defined a policy table. You can now use the table in a policy.

Associating Policy Tables with a Policy Rule

To associate a policy table with a new or existing policy rule, the policy table must already be defined. See *About Policy Tables* for more information on what a policy table is. See *Creating Policy Tables* for more information on how to define a policy table. See *Creating a New Policy* for more information on creating and modifying a policy definition.

One or more policy tables can be associated with a new or existing policy rule from the **Table Associations** page of the policy wizard using this procedure:

1. Start the Policy Wizard.
2. On the **Table Associations** page, click the selection icon next to **Use table *policy table* called *specified alias name***.
   The policy table option is added to the **Description** section of the page, where you select an existing policy table to use, and define an alias name for this policy table, if needed.
3. In the **Description** section of the page, click *policy table* to select an existing policy table.
   The **Policy Table Data** window appears.
4. Click to highlight the existing table to use, and click **OK**.

5. Click *specified alias name* to associate a unique name with this table. An alias name is required; enter a name here to specify the purpose of this policy table in this policy. You can then use the same policy table in multiple policies but define a different purpose each time with the alias name field. An **Input a Value** window opens.
6. Enter an alias name following the format specified in the window, and click **OK**.
7. Repeat these steps to associate another policy table with this policy rule, if needed.
8. If multiple policy tables are associated with this policy rule, use the up or down icon to move a table up or down to change the order in which it is evaluated in the rule.
9. Click **Next** to continue to the **Conditions** page.

The selected policy table(s) are associated with this policy definition.

Modifying Policy Tables

1. From the **Policy Management** section of the navigation pane, select **Policy Table Library**. The Policy Table Administration page opens in the work area.
2. On the Policy Table Administration page, select the policy table you want to modify. The Policy Table Administration page displays information about the policy table.
3. Click **Validate**. If selected, the data modified is validated. If invalid, a diagnostic message appears.
4. Click **Modify**. The table fields become editable. See *Creating Policy Tables* for information about the table fields.
5. When you finish, click **Save** (or **Cancel** to discard your changes).

The policy table content is modified.

Deleting Policy Tables

1. From the **Policy Management** section of the navigation pane, select **Policy Table Library**. The Policy Table Administration page opens in the work area.
2. Delete the Policy Table using one of the following methods:
   - From the work area, click the **Delete** icon located to the left of the policy table you wish to delete.
   - Open the policy and click **Delete**.

   You are prompted, "Are you sure you want to delete this policy table?"
3. Click **OK** (or **Cancel** to abandon the request).

The policy table is deleted.

Viewing Policy Tables

From the **Policy Management** section of the navigation pane, select **Policy Table Library**.

A tree frame view displays all existing policy tables. You will see all of the existing policy tables in the main frame when you click **ALL**.

Note: The policy table details are viewed by clicking the actual policy table name in the tree frame.

Chapter 17

System-Wide Reports

Topics:

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.

Viewing Active Alarms

The Active Alarms report provides an aggregate view of timestamped 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 Active Alarms report provides details about active alarms. To view the Active Alarms report, from the **System Wide Reports** section of the navigation pane, select **Active Alarms**.

Figure 23: Sample Active Alarms Report shows a sample active alarms report.

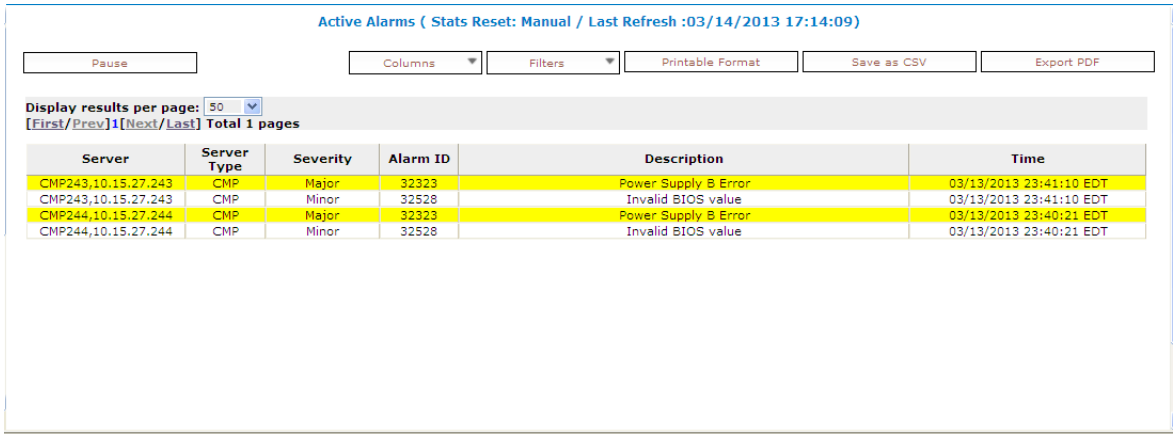

Figure 23: Sample Active Alarms Report

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 yellow.)
- **Minor** — Non-service affecting fault.

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*.

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.

- **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 Alarm History Report

The Alarm History Report displays historical alarm information.

To view the alarm history report, from the **System Wide Reports** section of the navigation pane, select **Alarm History Report**.

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. 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 (the default is **All**) from the list.
- **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.)

After entering filtering information, click **Filter** to refresh the display with the filtering applied.

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.
  - **Major** — Service may be interrupted if the issue is not corrected.
  - **Minor** — Non service affecting fault.
  - **Info** — Informational message only.
  - **Clear** — Alarm has been cleared.

  Note: Alarms generated by Policy Management systems running software before V7.5 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 the binoculars icon, located to the right of the alert. A window displays additional information; for example:

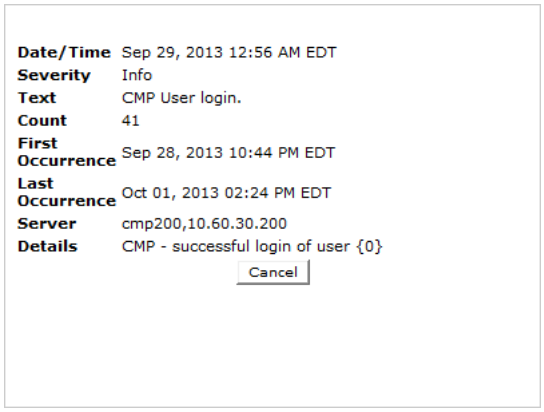

Click **Cancel** to close the window.

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

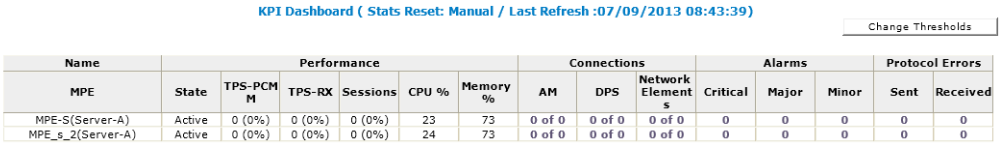

Figure 24: Example of KPI Dashboard

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

In the top right corner there is a Change Thresholds button that allows you to change threshold settings used to determine cell coloring (discussed below).

Each MPE cluster has one row in the table per server. The first row displays information for the first server that was configured (Server-A) in the cluster. The second row displays information for the second server that was configured (Server-B) in the cluster, if present. The third row displays information for the third server that was configured (Server-C) in the cluster, 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 is 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 *Policy Management Troubleshooting Guide* for this release.

Mapping Display to KPIs

Table 21: 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 21: 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

KPI Dashboard Column	Mapping to Statistics	
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

Color Threshold Configuration

The Color Threshold Configuration popup window is brought up when you click the **Change Thresholds** button, located in the top right corner of the KPI Dashboard.

The values displayed in the dialog boxes are the current settings. The user can modify the values and click **Save** to put the new values into effect. The values is saved so the next time the dashboard is opened it uses the same values.

Note: Saving the thresholds affects other users that may be viewing the dashboard at the same time.

The **Cancel** button closes the popup dialog without any changes to the KPI dashboard display. The **Reset** button restores the values to their defaults. The TPS and session limits for the Policy Management device will be set to the officially supported rates for the current software release.

Viewing the Trending Reports

To view the trending reports, from the **System Wide Reports** section of the navigation pane, select **Trending Reports**.

The navigation pane displays the trending reports. The reports display aggregated or individual MPE statistics in graph tables.

The following trending reports are available:

- **Session Count** — The maximum number of sessions per interval which were maintained over a period of time in selected or all MPE devices.
- **PCMM Transaction Per Second** — The number of PCMM requests and answer pairs processed in a second.
- **Rx Transaction Per Second** — The number of Rx requests and answer pairs processed in a second.

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. 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 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.

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.

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 that allow you to configure a customized trending report (*Figure 25: Trending Report Definition Configuration Page* shows a sample).

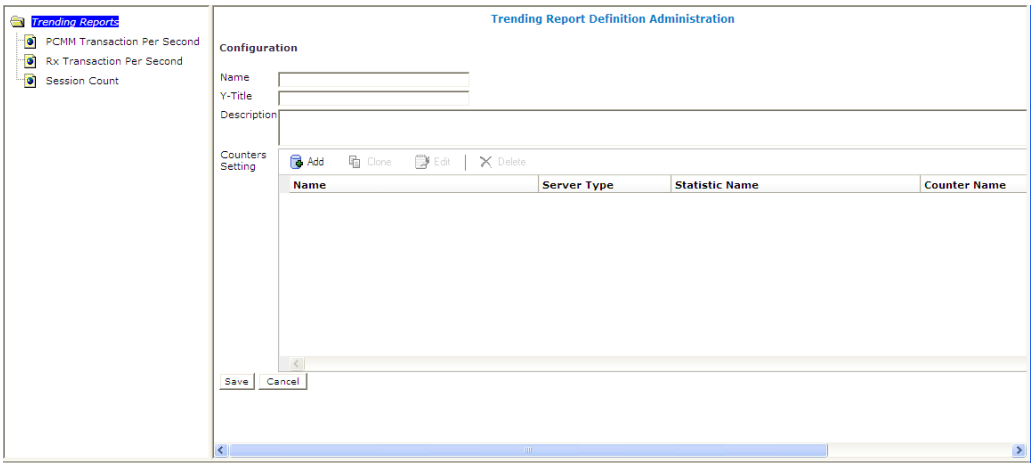

Figure 25: Trending Report Definition Configuration Page

3. Enter the following information for the new trending report:
   1. **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.

   2. **Y-title** — The title of the Y series.

      The title can contain up to 40 characters and cannot begin or end with a space.

   3. **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 the desired 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 desired 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. Click **Cancel** to exit the popup without adding a counter.

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 (see below).

5. After adding the first counter to the trending report, you can edit the counter information, clone the counter to create a new counter, or delete the counter.
   a) To edit a counter, select the counter, and click **Edit**.
   The Edit Stats Definition popup appears. Edit the information as desired. Click **Save** to save the edits. Click **Cancel** to exit the popup without saving the information.
   b) To add a new counter by cloning an existing counter, select the counter and click **Clone**.
   The Clone Stats Definition popup appears, containing the information that was used to create the selected counter. Edit the information as desired to create the new counter. Click **Save** to create a new counter. Click **Cancel** to exit the popup without creating a new counter.
   c) To delete an existing counter, select the counter and click **Delete**. You will be asked if you want to delete the counter. Click **Yes** to delete the counter. Click **No** to exit the popup without deleting the counter.
6. Click **Save** at the bottom of the Trending Report Definition page to save the report. Click **Cancel** to exit the Trending Report Definition page without saving the report.
   The custom trending report appears, in alphabetical order by name, in the list of custom trending reports.

You have defined and saved a custom trending report.

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 desired 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 desired custom trending report.
   The report opens.

3. Click **Settings.**
   The Trending Report Definition Administration page displays for the report.
4. Click **Delete**.
   The report is deleted.

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, from the **System Wide Reports** section of the navigation pane, select **Connection Status**.

Figure 26: Sample Connection Status Report shows a sample connection status report.

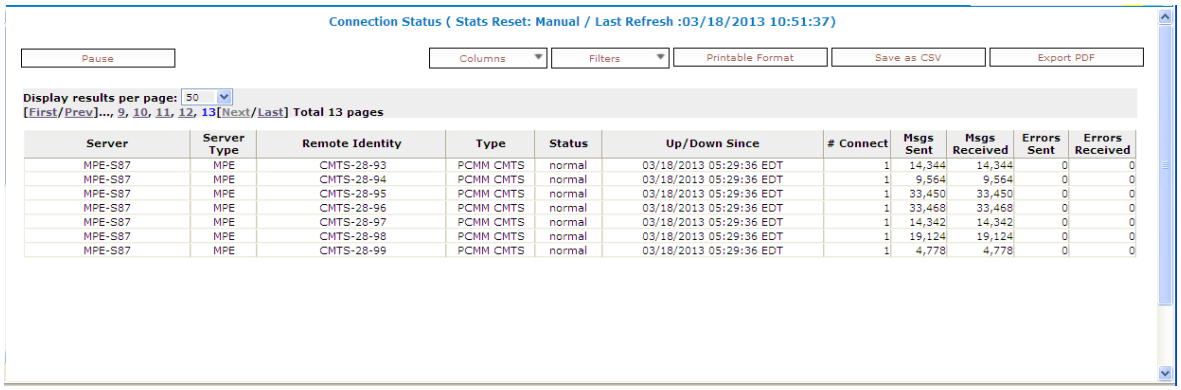

Figure 26: Sample Connection Status Report

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)
- **# Connect** — the number of times that the connection has been re-established

  Note: This counter is reset if the cluster is restarted.

- **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.
- **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, from the System Wide Reports section of the navigation pane, select **Protocol Errors**.

Figure 27: Sample Protocol Errors Report shows a sample protocol errors report.

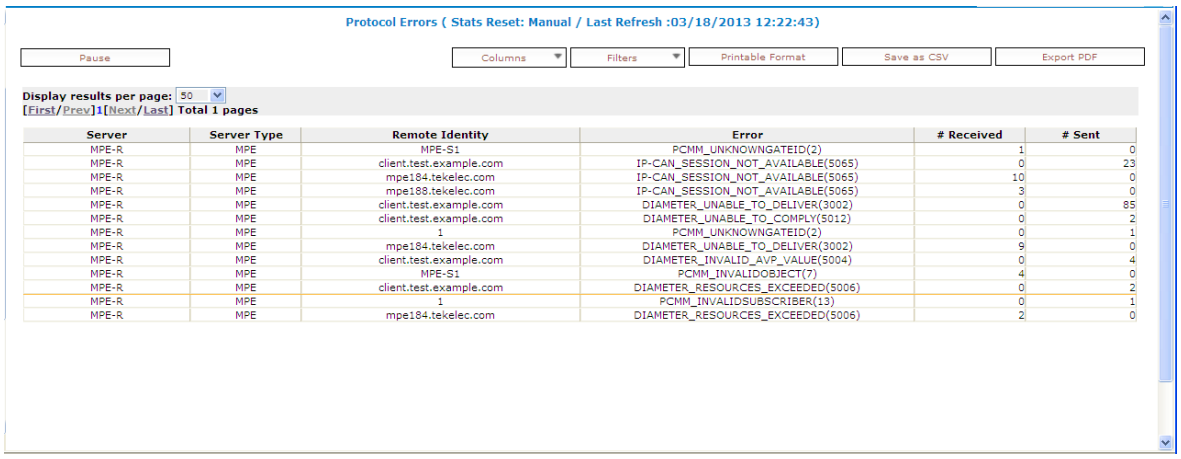

Server	Server Type	Remote Identity	Error	# Received	# Sent
MPE-R	MPE	MPE-S1	PCMM_UNKNOWNGATEID(2)	1	0
MPE-R	MPE	client.test.example.com	IP-CAN_SESSION_NOT_AVAILABLE(5065)	0	23
MPE-R	MPE	mpe184.tekelec.com	IP-CAN_SESSION_NOT_AVAILABLE(5065)	10	0
MPE-R	MPE	mpe188.tekelec.com	IP-CAN_SESSION_NOT_AVAILABLE(5065)	3	0
MPE-R	MPE	client.test.example.com	DIAMETER_UNABLE_TO_DELIVER(3002)	0	85
MPE-R	MPE	client.test.example.com	DIAMETER_UNABLE_TO_COMPLY(5012)	0	2
MPE-R	MPE	1	PCMM_UNKNOWNGATEID(2)	0	1
MPE-R	MPE	mpe184.tekelec.com	DIAMETER_UNABLE_TO_DELIVER(3002)	9	0
MPE-R	MPE	client.test.example.com	DIAMETER_INVALID_AVP_VALUE(5004)	0	4
MPE-R	MPE	MPE-S1	PCMM_INVALIDOBJECT(7)	4	0
MPE-R	MPE	client.test.example.com	DIAMETER_RESOURCES_EXCEEDED(5006)	0	2
MPE-R	MPE	1	PCMM_INVALIDSUBSCRIBER(13)	0	1
MPE-R	MPE	mpe184.tekelec.com	DIAMETER_RESOURCES_EXCEEDED(5006)	2	0

Figure 27: Sample Protocol Errors Report

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.
- **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, from the System Wide Reports section of the navigation pane, select **Policy Statistics Report**.

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 on 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:

- **Server Name** — name of the associated system
- **Server Type — 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** pulldown menu. You can define filtering criteria using the following fields:

- **Server Name** — Filter in all servers (the default) or one specific server.
- **Policy Name** — Filter in all policies (the 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 **Connection Status Report** window opens.

You can save the report in comma-separated value (CSV) format, suitable for importing into a spreadsheet application. Click **Save as CSV**. A file named `report.csv` is generated, and a standard **File Download** window opens, so you can save or open the file.

You can save the report as a Portable Document Format (PDF) file, suitable for storage or online display. Click **Export PDF**. A file named `report.pdf` is generated, and a standard **File Download** window opens, so you can save or open the file.

Chapter 18

Upgrade Manager

Topics:

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. Note that access to these GUI options can be affected by settings on the role setting page.

For specific steps on performing an upgrade, contact the Tekelec *Customer Care Center*.

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 monitor the current upgrade status for all servers on the system, monitor the ISO download process, and perform upgrade-related operations. 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 Elements

On the **Upgrade Manager** menu, **ISO 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 **ISO Maintenance** GUI page: Checkboxes to select servers on which to perform operations, the table of filtered servers, and pulldown menus (**Columns**, **Filters**, and **Operations**) for changing what displays in the table and for performing operations. The following list describes all of these elements.

Table 22: ISO Maintenance Elements

Element	Description
<Checkbox>	Use the checkbox column to check mark the servers on which an operation is to be performed. If you check mark a main cluster server, all servers in that cluster are check marked. Note that at least one server must be check marked before you can select an operation from the **Operations** pulldown menu.
Name	Displays the server names of all filtered servers. When a server is downloading an ISO file, a special download icon appears next to the name.
Appl Type	Displays the type of application running on each server. The **Filters** pulldown menu lets you select **CMP Site1 Cluster**, **CMP Site2 Cluster**, **MPE**, **BoD**, **MA**, or **All** servers.
IP	Displays the OAM server IP address of each server. The **Filters** pulldown menu lets you select only a server with a specific IP address or **All** servers.
Running Release	Displays the current Policy Management software release of each server. The **Filters** pulldown menu lets you display a specific release only or **All** releases.
ISO	Displays the ISOs or CD-ROM on each server. Use the checkbox to select the ISO to delete during the Delete ISO operation.
Columns	Use the **Columns** pulldown menu to change the columns that appear in this table. 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** pulldown menu to select a subset of servers to appear on this page. On this menu are the following pulldown filter submenus: **Appl Type**, **IP**, 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	Use the **Operations** pulldown menu to select an ISO operation to perform. **Note:** The servers on which the operation is being performed must be check marked (in the first column of the table) before that or any operation can be selected. The operations that appear in the pulldown menu depend on the state of the servers that are selected; that is, when more than one server is selected, only the operations that are available on all of these servers appear. Possible operations are **Push Script**, **Upload ISO**, and **Delete ISO**. As a protective feature, when a command is executed, a warning message pops up, asking if you are sure you want to execute this operation (click **OK** or **Cancel**). When **OK** is clicked, a progress bar displays the status of the command completion in a pop-up window. Note that once the operation is confirmed, it cannot be cancelled.

Viewing 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 **Upgrade Manager** section of the navigation pane, select **ISO Maintenance**.

   The **ISO Maintenance** page appears.

2. (Optional) Use the filter criteria as needed, accessed from the Filters pulldown menu, to customize the list of servers that display in the table.
3. (Optional) Use the Columns pulldown menu as needed, to check and uncheck columns, to customize the data that displays in the table.

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

Pushing a Script to a Server

Use this procedure to push the upgrade script to the remote servers receiving the software upgrade. This step is required before a software upgrade can occur on a server. An error message displays in the Upgrade Status column until Push Script has been run.

1. From the **Upgrade Manager** section of the navigation pane, select **ISO Maintenance**.
   The **ISO Maintenance** page appears.
2. Select the server(s) receiving the upgrade script.
3. Click on the Operations pulldown menu and select **Push Script**.
   You are prompted, "Are you sure you want to execute Push Script?"

4. Click **OK** (or **Cancel** to abandon your request).
   A progress bar displays the progress of the operation.

The upgrade script is downloaded to the selected servers.

Adding an ISO File to a Server

Use this procedure to download an upgrade ISO file to a remote server in preparation for a software upgrade.

1. From the **Upgrade Manager** section of the navigation pane, select **ISO Maintenance**.
   The **ISO Maintenance** page appears.
2. Select the server(s) to receive the ISO file.
3. Click the Operations pulldown menu and select **Upload ISO**.
   An Upload/Add ISO window appears.
4. Enter the ISO Server Hostname or IP address, User, Password, and ISO file full path for the ISO file being added.

Option	Description
Mode	Mode used to transfer file to remote servers. Currently, SCP is available.
ISO Server Hostname/IP	Enter the name or address of the server receiving the ISO file. This field is required.
User	Enter your user name. This field is required.
Password	Enter your password. This field is required.
ISO file full path	Enter the location where the ISO file is to be stored on the remote server. This field is required.

5. Click **Add** (or **Back** to abandon your request).
   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 server(s).

Deleting an ISO File from a Server

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

1. From the **Upgrade Manager** section of the navigation pane, select **ISO Maintenance**.
   The **ISO Maintenance** page appears.
2. Select the server(s) from which the ISO file is being removed.
3. Select the ISO file on the server that is being removed.
4. Click the Operations pulldown menu and select **Delete ISO**.
   You are prompted, "Are you sure you want to execute Delete ISO?"
5. Click **OK** (or **Cancel** to cancel the request).
   A progress bar displays the progress of this operation.

The selected ISO file(s) are deleted from the selected remote server(s).

About Performing an Upgrade

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.

The information in this section is a general overview of what happens during the upgrade process. Steps for performing an upgrade are provided by the Tekelec *Customer Care Center*.

Caution: Use only the upgrade procedure provided by the Tekelec Customer Care Center. Before upgrading any system, please go to the Tekelec Customer Support website and review any Technical Service Bulletins (TSBs) that relate to this upgrade. 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.

Warning: Contact the Tekelec Customer Care Center and inform them of your upgrade plans prior to beginning this or any upgrade procedure.

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, the table of filtered servers, and pulldown menus (**Columns**, **Filters**, and **Operations**) for changing what displays in the table and for performing operations. The following list describes all of these elements.

Table 23: System Maintenance Elements

Element	Description
<Checkbox>	Use the checkbox column to check mark the servers on which an operation is to be performed. If you check mark a main cluster server, all servers in that cluster are check marked. Note that at least one server must be check marked before you can select an operation from the **Operations** pulldown menu.
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.

Appl Type	Displays the type of Policy Management application running on each server. The **Filters** pulldown menu allows you to display **CMP Site1 Cluster only**, **CMP Site2 Cluster only**, **MPE**, **BoD**, **MA**, or **All**.
IP	Displays the IP address of each server. The **Filters** pulldown menu allows you to display only the server with a specific IP address or All servers.
Server State	Displays the state of each server. The server state can appear in different colors, depending on the state displayed. The **Filters** pulldown menu allows you to display **Active** only, **Standby** only, **Out-Of-Service** only, **Force Standby** only, or **All** servers (the default).
ISO	Displays the ISOs or CD-ROM on each server. Use the checkbox to select an ISO to use during an upgrade on that server.
Prev Release	Displays the previous Policy software release of each server, if known. The **Filters** pulldown menu lets you display a specific release only or All releases.
Running Release	Displays the current Policy software release of each server. The **Filters** pulldown menu allows you to display a specific release only or All releases.
Replication	Displays whether replication is On or Off.
Compatible Replication	Displays whether compatible replication is On or Off.
Legacy Sync	Displays whether legacy synchronization is On or Off.
Upgrade Status	Displays details of last upgrade performed on each server.
Columns	Use the **Columns** pulldown menu 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** pulldown menu 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	Use the **Operations** pulldown menu to select an upgrade operation to perform. **Note:** The servers on which the operation is being performed must be check marked (in first column of table) before that or any operation can be selected. The operations that appear in the pulldown menu depend on the state of the servers that are selected; that is, when more than one server is selected, only the operations that are available on all of these servers appear. Possible operations are listed below. As a protective feature, when a command is executed, a warning message pops up, asking if you are sure you want to execute this operation (click **OK** or **Cancel**). When you click **OK**, a progress bar displays the status of the command completion in a pop-up window. **Note:** Once the operation is confirmed, it cannot be cancelled.

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 (CMP/MPE/BoD/MA).
Force Standby	Forces the selected server(s) into standby status.
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).

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. Use the filter criteria as needed, accessed from the Filters pulldown menu, to customize the list of servers that display in the table.
3. Use the Columns pulldown menu as needed, to check and uncheck columns, to customize the data that displays in the table.

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

About Preparing for an Upgrade

Upgrading a server requires a large amount of preparation. For detailed information about preparing for an upgrade, please access Tekelec's Customer Support site.

Caution: Use only the upgrade procedure provided by the Tekelec Customer Care Center. Before upgrading any system, please go to the Tekelec Customer Support website and review any Technical Service Bulletins (TSBs) that relate to this upgrade. 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.

Warning: Contact the Tekelec Customer Care Center and inform them of your upgrade plans prior to beginning this or any upgrade procedure.

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. Procedures and scripts are available to preserve the current state of subscriber data, such as MPE sessions. Before beginning a rollback, contact the Tekelec Customer Care Center and inform them of your plans.

Chapter 19

System Administration

Topics:

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. On the System Settings page, 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.
   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 select a new password.)
   h) **Minimum Password Length** — The minimum allowable length in characters for a password, from 6 to 64 characters.
   The default is six characters.
   i) **Login Banner Title** — The title that displays at the top of the login page. The default is "Welcome." You can enter up to ten 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 GUI menu. 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; clear the check box to disable this feature.
   b) **Invalid Login Threshold Value (number of failed logins)** — Defines the maximum number of consecutive failed logins after which action is taken.
      Enter a value from 1 through 500; the default is 3 attempts.
   c) **Action(s) upon Crossing Threshold** — The system action to take if a user reaches the invalid login threshold:
      - **Lock user** — 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 letter(s) (1-64)** — By default, this setting is 0 (disabled). Select it to require users to include from 1 to 64 lowercase letters in their passwords.
   - **Require at least upper-case letter(s) (1-64)** — By default, this setting is 0 (disabled). Select it to require users to include from 1 to 64 uppercase letters in their passwords.
   - **Require at least numeral(s) (1-64)** — By default, this setting is 0 (disabled). Select it to require users to include from 1 to 64 numerals in their passwords.
   - **Require at least non-alphabetic character(s) (1-64)** — By default, this setting is 0 (disabled). Select it to require users to include from 1 to 64 nonalphabetic 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** (or **Cancel** to discard your changes).

The system settings are configured.

Figure 28: 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.

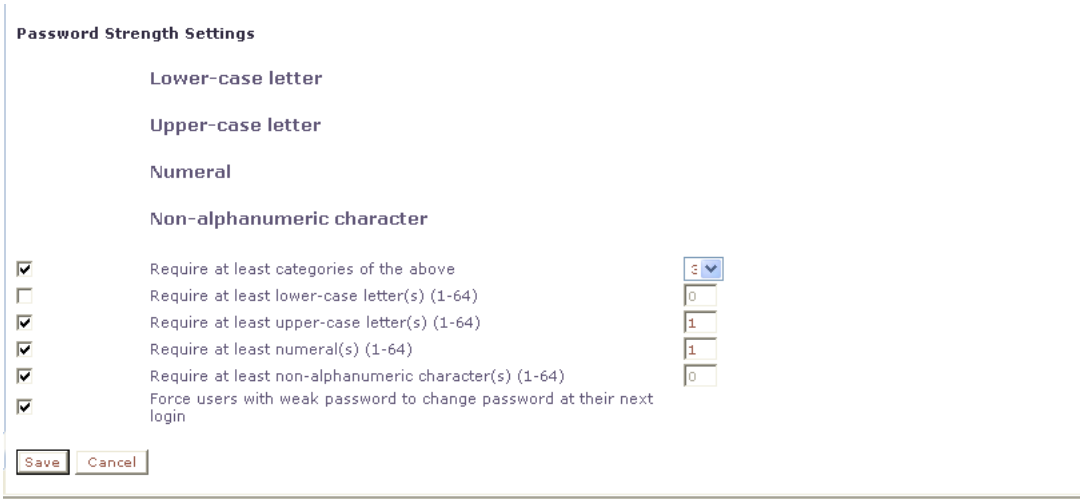

Figure 28: Sample Password Strength Policy

Importing to and Exporting from the CMP Database

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.

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).

Tekelec recommends that you export the existing database of objects before starting an importation 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 the type of import: * (the default, to import all types), **Network Elements**, **Accounts**, **Tiers**, **Traffic Profiles**, **Applications**, **Policy Table**, **Roles**, **Scopes**, or **Users**.
   If you select **Network Elements**, additional filtering 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 **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. Tekelec recommends that the 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** (the default), **Accounts**, **Tiers**, **Traffic Profiles**, **Applications**, **Policy Table**, **Roles**, **Scopes**, or **Users**.
   If you select **Network Elements**, additional filtering fields are available to help you manage the potential volume of data being exported. 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.

3. Click **Export**.
   A standard file download window opens, and you are prompted, "Do you want to open or save this file?"
4. Click **Save** to save the file (or **Cancel** to abandon the request).
   Data is exported to an XML file. If the operation takes more than five seconds, a progress bar appears.

The user accounts datacollector and _policy_server cannot be exported.

User passwords are exported in encrypted text. Password expiration timestamps are retained, so the passwords will expire on the schedule of the CMP system from which they were 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 are displayed in the Manager Report 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.
- **Cluster 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 pause refreshing the display, click **Pause**. To resume refreshing, click **Resume**. To reset the display counters, click **Reset All Counters**.

- **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** buttons let you restart the CMP software on the server or restart the server itself.

The Trace Log

The Trace Log is part of system administration records notifications for management activity on the CMP system. You can configure the severity level of messages written to the Trace Log; for information, see *Configuring Log Settings*.

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 in-progress message "Scanning Trace Logs" appears.
3. When you finish, click **Close**.
   The Trace Log Viewer window closes.

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, affording 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. On the Audit Log page, click **Show All**.
   The Audit Log opens. (*Figure 29: Audit Log* shows an example.)

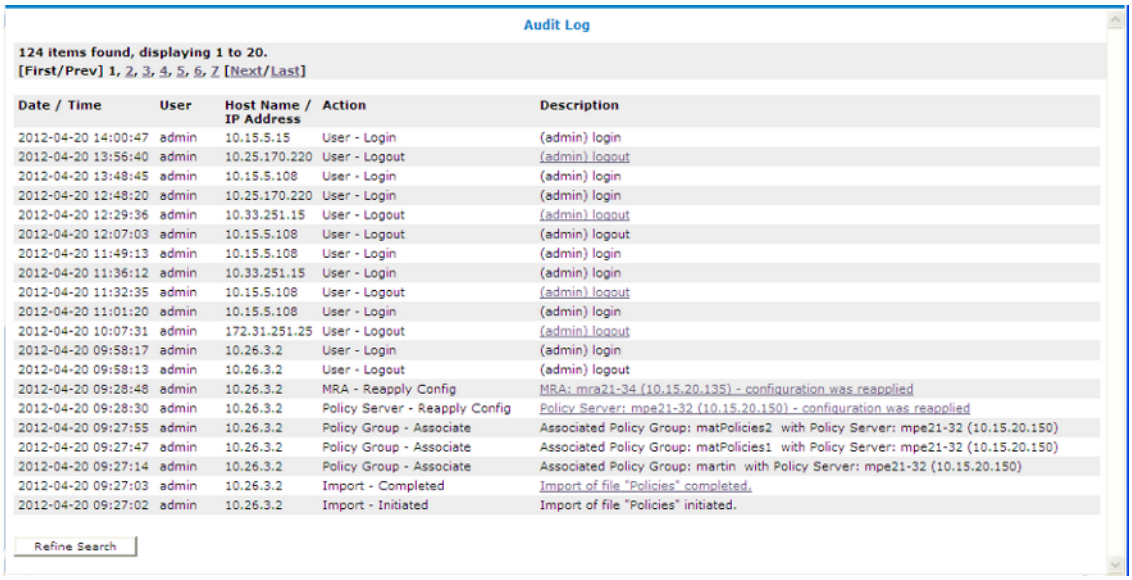

Audit Log

124 items found, displaying 1 to 20.
[First/Prev] 1, 2, 3, 4, 5, 6, 7 [Next/Last]

Date / Time	User	Host Name / IP Address	Action	Description
2012-04-20 14:00:47	admin	10.15.5.15	User - Login	(admin) login
2012-04-20 13:56:40	admin	10.25.170.220	User - Logout	(admin) logout
2012-04-20 13:48:45	admin	10.15.5.108	User - Login	(admin) login
2012-04-20 12:48:20	admin	10.25.170.220	User - Login	(admin) login
2012-04-20 12:29:36	admin	10.33.251.15	User - Logout	(admin) logout
2012-04-20 12:07:03	admin	10.15.5.108	User - Logout	(admin) logout
2012-04-20 11:49:13	admin	10.15.5.108	User - Login	(admin) login
2012-04-20 11:36:12	admin	10.33.251.15	User - Login	(admin) login
2012-04-20 11:32:35	admin	10.15.5.108	User - Logout	(admin) logout
2012-04-20 11:01:20	admin	10.15.5.108	User - Login	(admin) login
2012-04-20 10:07:31	admin	172.31.251.25	User - Logout	(admin) logout
2012-04-20 09:58:17	admin	10.26.3.2	User - Login	(admin) login
2012-04-20 09:58:13	admin	10.26.3.2	User - Logout	(admin) logout
2012-04-20 09:28:48	admin	10.26.3.2	MRA - Reapply Config	MRA: mra21-34 (10.15.20.135) - configuration was reapplied
2012-04-20 09:28:30	admin	10.26.3.2	Policy Server - Reapply Config	Policy Server: mpe21-32 (10.15.20.150) - configuration was reapplied
2012-04-20 09:27:55	admin	10.26.3.2	Policy Group - Associate	Associated Policy Group: matPolicies2 with Policy Server: mpe21-32 (10.15.20.150)
2012-04-20 09:27:47	admin	10.26.3.2	Policy Group - Associate	Associated Policy Group: matPolicies1 with Policy Server: mpe21-32 (10.15.20.150)
2012-04-20 09:27:14	admin	10.26.3.2	Policy Group - Associate	Associated Policy Group: martin with Policy Server: mpe21-32 (10.15.20.150)
2012-04-20 09:27:03	admin	10.26.3.2	Import - Completed	Import of file "Policies" completed.
2012-04-20 09:27:02	admin	10.26.3.2	Import - Initiated	Import of file "Policies" initiated.

Refine Search

Figure 29: Audit Log

For a detailed description of an item, click the underlined description. The details of the event display. (*Figure 30: Audit Log Details* shows an example.)

To filter search results, click **Refine Search**, located at the bottom of the page. (See *Searching for Audit Log Entries*.)

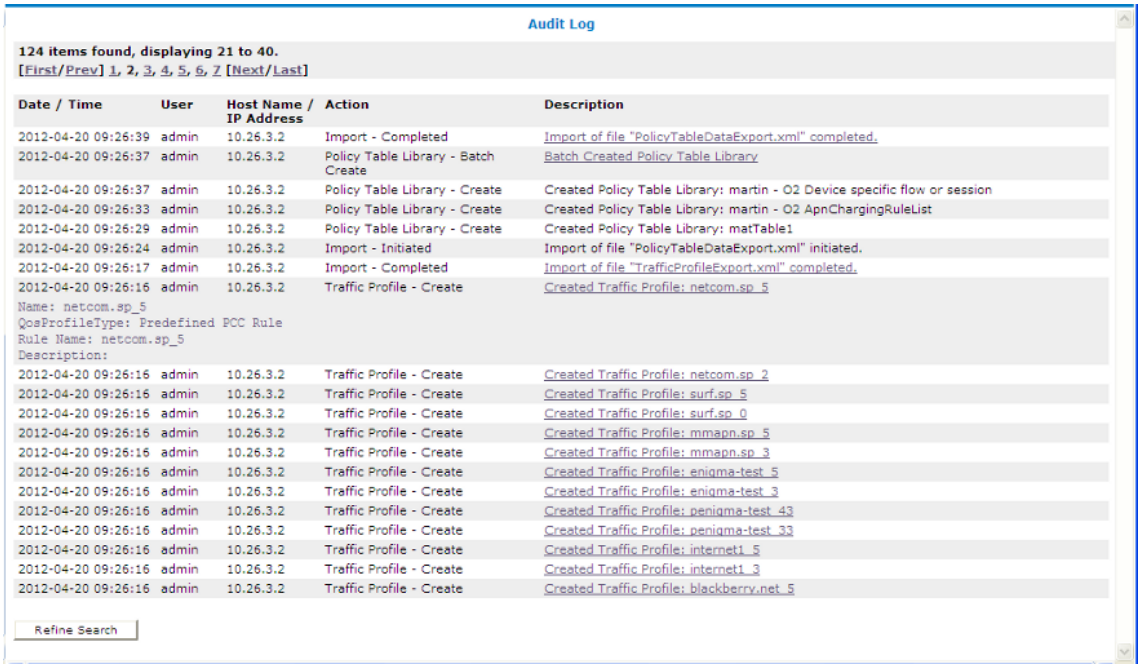

Figure 30: Audit Log Details

Searching for Audit Log Entries

To search for entries in 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. On the Audit Log page, click **Search**.
   The Audit Log Search Restrictions Page opens.
3. Define the following items, depending on how restrictive you want the audit log search to be:
   - **From/To** — Enter the start and end dates and times for this search.
   - **Action by User Name(s)** — Enter the name of the user or users to audit.
   - **Action on Policy Server(s)** — Enter the name of the Policy Management device to audit.
   - **Audit Log Items to Show** — Specifies an item to audit for display: **Policy Server**, **Management Agent**, **Network Element**, **Network Element Group**, **Application**, **Policy**, **Policy Group**, **Account**, **User**, **Audit**, **OM Statistics**, **MPE Manager**, **Upgrade Manager**, **Topology Setting**, **Global Configuration Settings**, **Trending Report**, **BoD**, **BoD Services**, or **BoD Group**. By default you can specify three items; click **More Lines** to add an additional item.
   - **Results Forms** — Specifies the number of items per page to display, along with which data to display (most recent or oldest items).
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; the default filename is `AuditLogExport.txt`.

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. On the Audit Log page, 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 the calendar icon, located to the right of the field, to select a date from the pop-up window.
4. When you finish, click **Export**.
   A standard File Download window opens; you can open or save the export file.

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. On the Audit Log page, click **Export/Purge**.
   The Export and Purge Audit Log Items page opens.
3. In the **Items to Purge** section, enter a date in the format *mm/dd/yyyy*, or click the 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.

Managing 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: Tekelec strongly recommends that you perform these tasks in the order in which they are listed, or serious system problems can occur. Consult Tekelec Technical Support before changing the order of any task.

The tasks include:

- **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.

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 on its name; the current settings and status are displayed; for example:

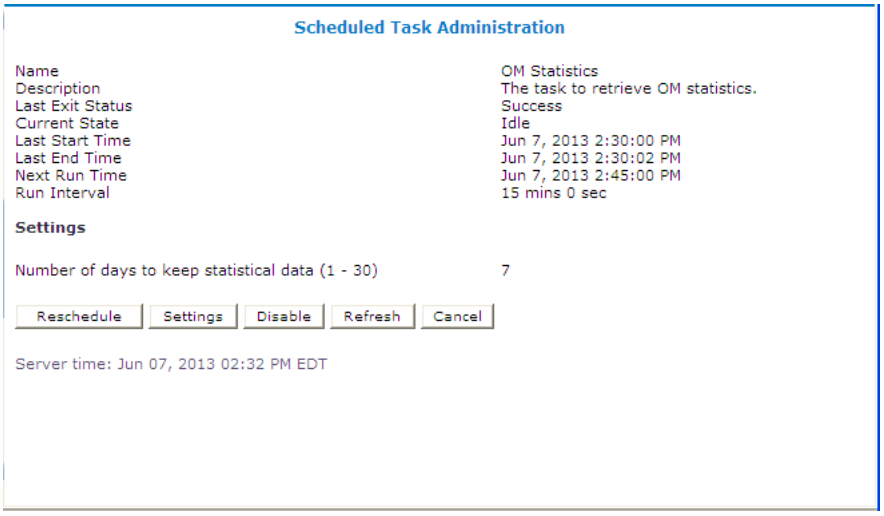

3. The options for this task are as follows:

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

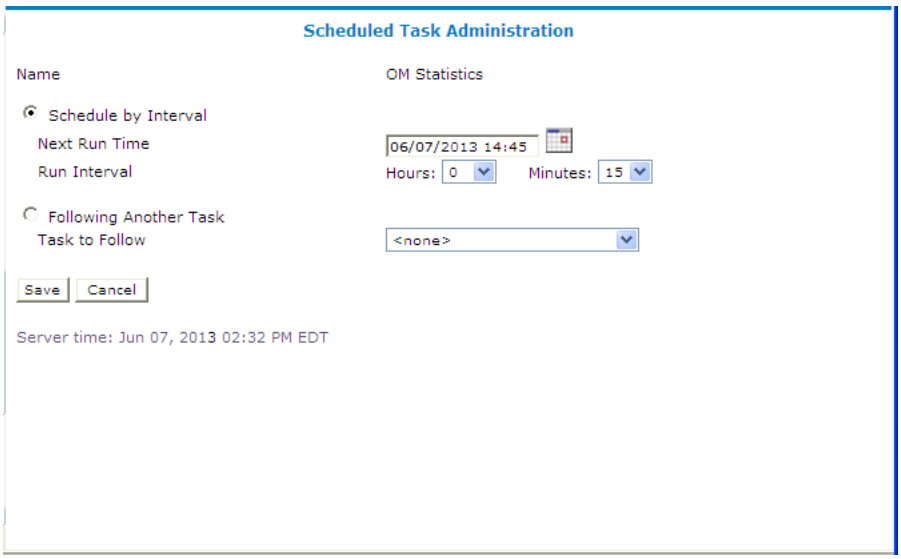

  - **Schedule by Interval** (**Next Run Time** or **Run Interval**) — Defines the run interval for the task to follow.

    Valid run intervals are from 0 to 24 hours in 5-minute increments.

  - **Following Another Task** — Defines the run time as following the completion of another scheduled task that you select from the list.

- **Settings** — Number of days to keep data; the default is seven days.
- **Run Now** — Runs the process immediately.

You are prompted, "Click 'OK' to run this task now." Click **OK** to run the task (or **Cancel** to cancel the request).

- **Disable** or **Enable** — Disables or enables the next scheduled execution of this process.

  If you click **Disable**, you are prompted, "Click 'OK' to disable this task." Click **OK** (or **Cancel** to cancel the request); the task is disabled and will not run at its 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.

Configuring Roles

Assigning roles to the various users that access the CMP system lets you control who can configure and access what within the CMP system. The default roles are:

- **Viewer** — Permits read-only access to functions associated with Policy Management device management and configuration. Access is also permitted to limited system administration functions, such as Change Password.
- **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.
- **Administrator** — Permits full read/write access to all functions. You cannot delete the Administrator role.

Creating a New Role

To create a new 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** (that is, the functions do not appear to users of the role, so access must be explicitly granted) or **Read-Only**.
4. Enter the following information:

a) **Name** — The desired name for the new role
b) **Description/Location** (optional) — Free-form text
c) **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**
- **AVP Definition**
- **Global Configuration Settings**

d) **Network Privileges** — Defines access to the network management Paths function (assigning the privilege **Hide**, **Read-Only**, or **Read-Write**):

Topology

e) **BoD Privileges** — Defines access to the Bandwidth on Demand Application Manager (with the privileges **Hide**, **Read-Only**, or **Read-Write**.

- **Configuration**
- **Services**
- **Service Import/Export**

f) **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**)

g) **System Wide Reports Privileges** — Defines access to the system-wide reports functions:

Trending Reports Configuration (with the privileges **Hide**, **Read-Only**, or **Read-Write**)

h) **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**)

i) **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**)

j) **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 of Policy Server** (with the privileges **Hide**, **Read-Only**, or **Read-Write**)
- **Trace Log** (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**)

5. When you finish, click **Save** (or **Cancel** to discard your changes).

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. On the Role page, click **Modify**.
   The Modify Role page opens.
5. Modify role information as necessary.
   See *Creating a New Role* for a description of the fields contained within this page.
6. When you finish, click **Save** (or **Cancel** to discard your changes).

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 the Delete icon located next to the role to delete.
   - From the content tree, select the role to delete (role information displays in the work area), then click **Delete**.

   You are prompted: "Are you sure you want to delete this Role?"
4. Click **OK** or **Cancel** to cancel the request.

The role's information is deleted from the CMP database.

Creating a New Scope

The CMP lets you configure scopes that 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. On the Scope Administration page, click **Create Scope**.
   The New Scope page opens.
4. Enter the following information:
   a) **Name** — The desired 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 (or **Cancel** to discard your changes).

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. On the Scope Administration page, select the scope you want to modify.
   The scope description opens.
4. Click **Modify**.
   The Modify Scope page opens. *Creating a New Scope* describes the fields on this page.
5. Modify scope information as necessary.
6. When you finish, click **Save** (or **Cancel** to discard the request).

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. (*Figure 31: Deleting a Scope* shows an example.)
3. Delete the role using one of the following methods:
   - From the work area, click the Delete icon, located to the right of the role to delete.
   - From the content tree, select the role to delete (role information displays in the work area), then click **Delete**.

   You are prompted: "Are you sure you want to delete this Scope?"
4. Click **OK** (or **Cancel** to cancel the request).

The scope is deleted.

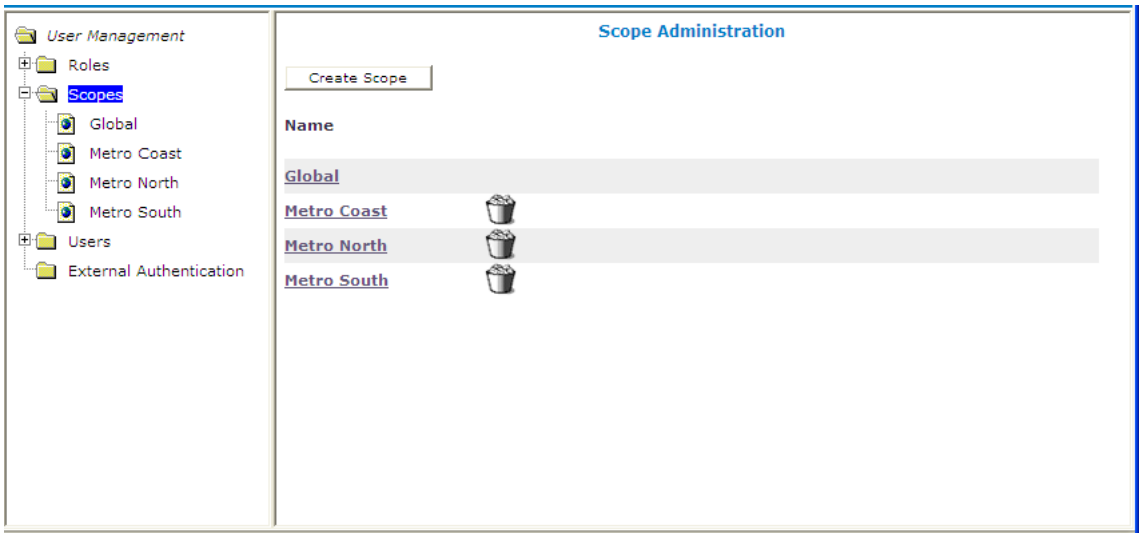

Figure 31: Deleting a Scope

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, Tekelec recommends changing this value from its 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 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.

   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 attributes:
   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 pulldown 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** (or **Cancel** to discard your changes).

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 desired user profile from the content tree.
   The profile information page opens.

5. Click **Modify**.
   The Modify User page opens. (*Figure 32: Modify User Page* shows an example.)
6. Modify the user profile as desired.
   (For field descriptions, see *Creating a User Profile*.)
7. When you finish, click **Save** (or **Cancel** to discard your changes).

The user profile is modified.

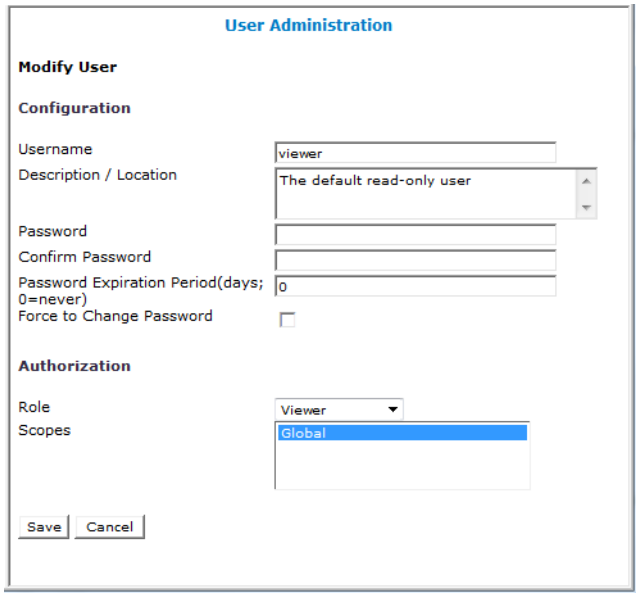

Figure 32: Modify User Page

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; for example:

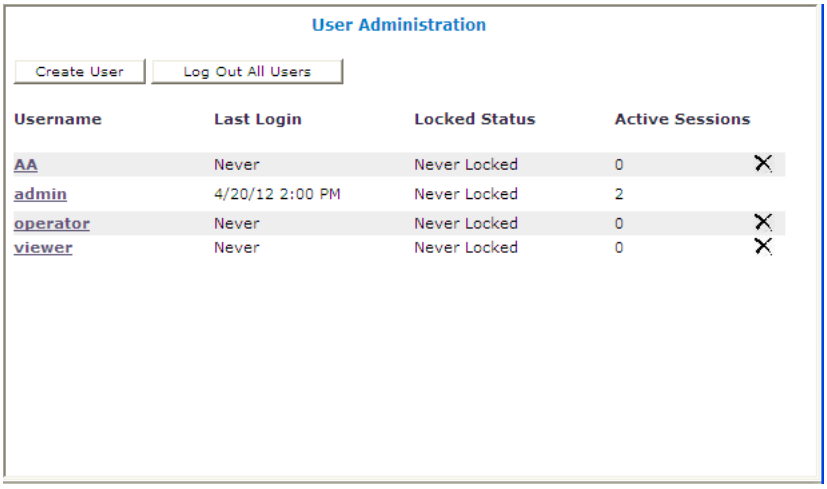

4. Delete the desired user profile using one of the following methods:
   - From the work area, select the delete icon, located to the right of the profile you want to delete.
   - From the content tree, select the user profile that you want to delete (profile information displays in the work area), then click **Delete**.

   You are prompted: "Are you sure you want to delete this user?"
5. Click **OK** to delete the user profile (or **Cancel** to abandon the request).

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 itself.

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 desired user profile from the content tree.
   The User Administration page opens.
5. Click **Lock**.
   You are prompted: "Are you sure you want to lock out this user?"
6. Click **OK** (or **Cancel** to cancel the request).

The account is locked. The page displays: "User account locked successfully." The **Lock** button becomes an **Unlock** button. On the User Administration page, the user's Locked Status 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. Select the desired user profile from the content tree.
   The User Administration page opens.
4. Click **Unlock**.
   You are prompted: "Are you sure you want to unlock this user?"
5. Click **OK** (or **Cancel** to cancel the request).
   The account is unlocked. The page displays: "User account unlocked successfully." The **Unlock** button becomes a **Lock** button. On the User Administration page, the user's Locked Status changes to "Unlocked by Admin."

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 any user's password.

If a system administrator has configured your account for password expiration, you will receive a warning when you log in that you will need to change your password.

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** — Retype the new password.

      If your new password does not conform to the password strength rules, a validation error message appears; for example:

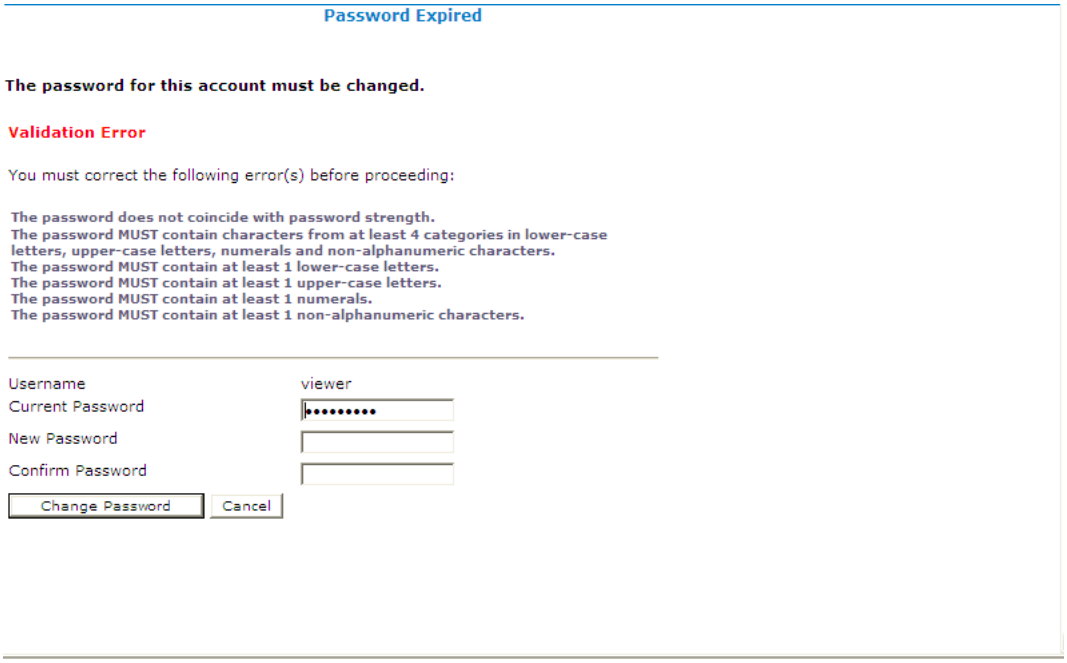

3. When you finish, click **Change Password**.

Your password is changed.

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. Associating user roles and scopes on the CMP system

3. Configuring the CMP system 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 will be aware of the names and locations of the configuration files.)

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 = camiant
        shortname = MPE5
}
client 10.0.10.23 {
        secret = camiant
        shortname = CMP56
}
```

The best practice is to define IP addresses rather than FQDNs. If no netmask is given, the default is /32. The shared secret (in this example, "**camiant**") must be both defined on 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 to define what attributes to send back to the client. Tekelec provides a VSA dictionary file, dictionary.camiant, in the directory **/opt/camiant/install/radius**. *Figure 33: Tekelec VSA Dictionary For RADIUS* shows the contents of this file. The local RADIUS administrator is responsible for incorporating the Tekelec VSA dictionary into the RADIUS server.

```
========== dictionary.camiant ==================
# Camiant Inc VSA's, from RFC 2548
# The filename given here should be an absolute path.
#
# Place additional attributes or $INCLUDEs here.

VENDOR Camiant 21274
BEGIN-VENDOR Camiant
ATTRIBUTE Camiant-MI-role 1 string
```

```
ATTRIBUTE Camiant-MI-scope 3 string
END-VENDOR Camiant
======================
```

Figure 33: Tekelec VSA Dictionary For RADIUS

The attributes **Camiant-MI-role** and **Camiant-MI-scope** are for access to the GUI. The GUI has both a scope and a role 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 GUI 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 GUI scope is **Global**, and the administrator can create other scopes within the GUI.

Associating Roles and Scopes

The GUI of 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 Tekelec VSA dictionary, by defining the attributes **Camiant-MI-role** and **Camiant-MI-scope**. The flexibility of roles and scopes can be supported by RADIUS if the Tekelec dictionary is integrated.

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

```
Jeff      Cleartext-Password := "garbage"
          Camiant-MI-role = "Administrator",
          Camiant-MI-scope = "Global"

view      Cleartext-Password := "camiant"
          Class = "Viewer",
          Camiant-MI-role = "Viewer",
          Camiant-MI-scope = "Global"
```

In this example, the user Jeff has access to the GUI as an administrator, and the user view has access to the GUI as a viewer (read-only access).

However, if Tekelec 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 GUI can use the Class attribute for RADIUS authentication.

To accept the Class attribute for GUI login, define a scope and a role that matches what the RADIUS server returns as the Class attribute. The GUI uses the Class attribute for both required credentials. For example, consider this user defined in RADIUS:

```
Dawn      Cleartext-Password := "camiant"
          Class = "Viewer"
```

Dawn can get access to the GUI 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 **RADIUS Authentication**.
   The RADIUS 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 RADIUS Authentication**.
   Additional fields appear (*Figure 34: 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** (the 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** (the 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** (the 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** (or **Cancel** to discard your changes).
   The window closes.

RADIUS Authentication and Accounting is configured.

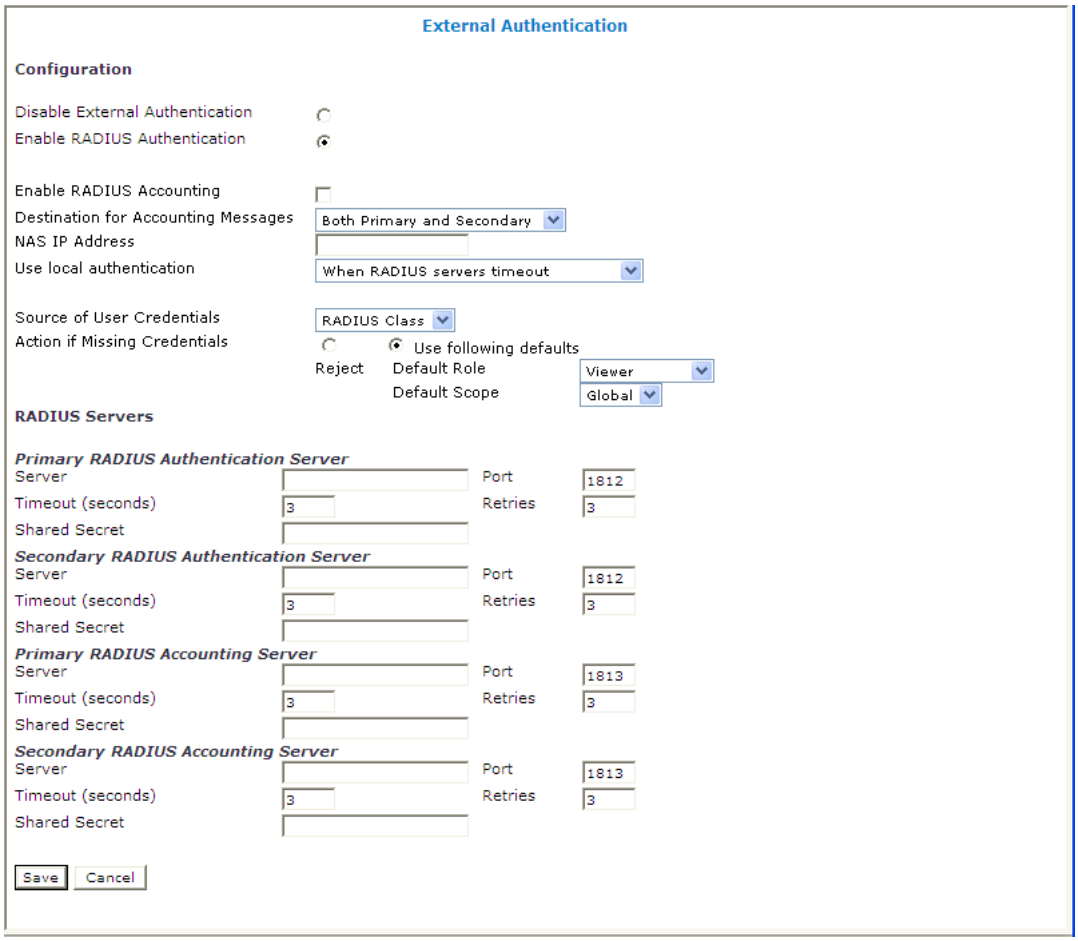

Figure 34: RADIUS External Authentication Configuration Page

Appendix A

CMP Modes

Topics:

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, Tekelec pre-configures servers delivered to customers. 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 Tekelec Technical 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 35: Mode Settings Page*). Select modes, sub-modes, 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.

Table 24: CMP Modes and Sub-Modes briefly describes each mode and sub-mode.

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
- **Manage Direct Link** — If enabled, all replication and HA traffic goes through the backplane interface; if disabled, all replication and HA traffic goes through the OAM interface

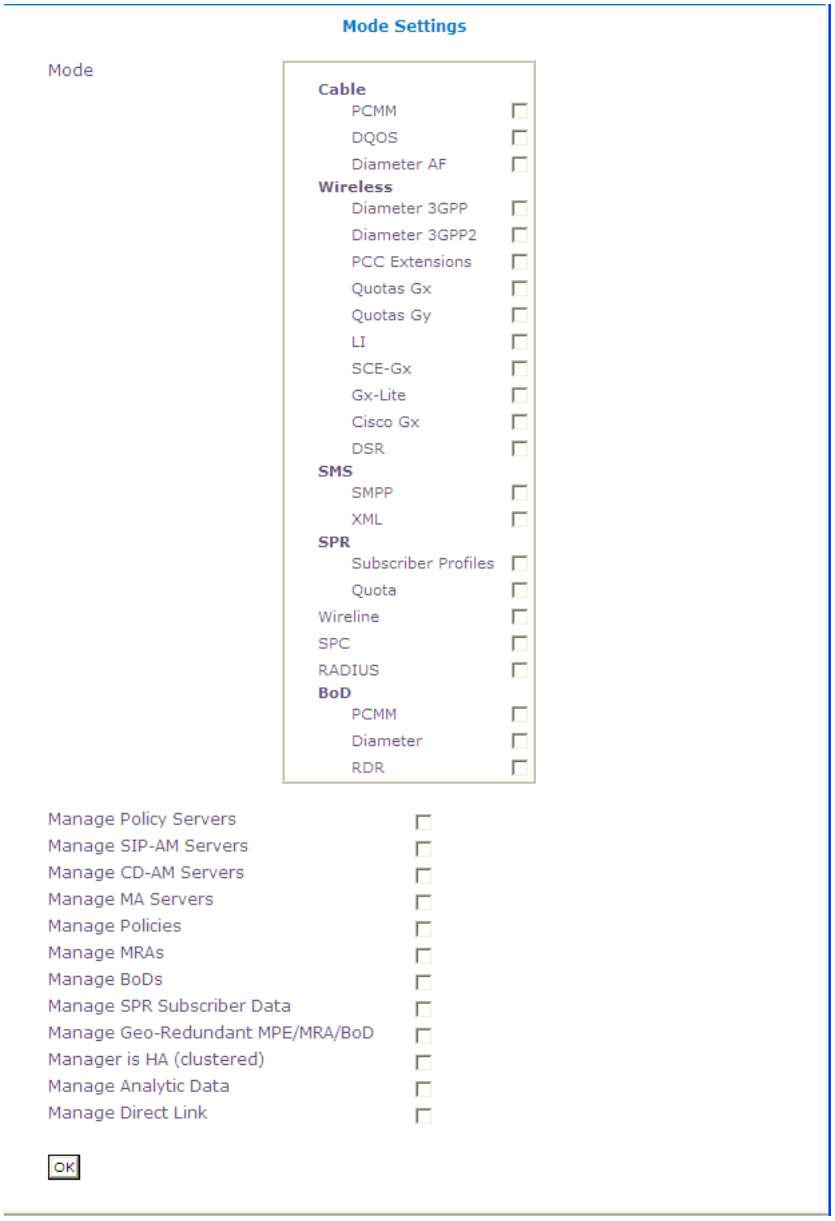

Figure 35: Mode Settings Page

Table 24: CMP Modes and Sub-Modes

Mode	Sub-Mode	Description
Cable Mode	Enables support of a cable carrier environment. Functions are described in the *Configuration Management Platform Cable User Guide*.	
	PCMM	Supports PacketCable MultiMedia functions.
	DQOS	Supports Dynamic Quality of Service functions.

Mode	Sub-Mode	Description
	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 Guide*.	
	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 *Configuring Lawful Intercept Application Note*.
	SCE-Gx	Supports the Cisco Service Control Engine Gx protocol. If this mode is selected, Diameter 3GPP and RADIUS must also be selected, and other Gx sub-modes must not be selected.
	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.
SMS Mode	Enables support of SMS servers. Functions are described in the *Configuration Management Platform Wireless User Guide*.	
	SMPP	Supports SMS using SMPP protocol.

Mode	Sub-Mode	Description
	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.	
	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 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.	
BoD Mode	Enables the Bandwidth on Demand Application Manager (BoD-AM), which support video on demand (VoD) servers. Functions are described in the *Bandwidth on Demand Application Manager User Guide*.	
	PCMM	Supports a network creating PacketCable Multimedia (PCMM) sessions.
	Diameter	Supports a network creating Diameter sessions.
	RDR	Supports a network containing Service Control Engine (SCE) devices transmitting Raw Data Records (RDRs).

Glossary

#

3GPP	3rd Generation Partnership Project. The standards body for wireless communications. 3rd Generation Partnership Project
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.

C

CMP	Configuration Management Platform

C

	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	Cable modem termination system An edge device connecting to subscribers' cable modems in a broadband network. A CMTS device can function as a PCEF device; see PCEF. Cable Modem Termination System: Equipment used by cable companies to provide high speed data services to cable subscribers.

D

Diameter	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. Diameter is the successor to the RADIUS protocol. The MPE device supports a range of Diameter interfaces, including Rx, Gx, Gy, and Ty.
DOCSIS	Data Over Cable Service Interface Specification - An international telecommunications standard for adding high-speed data transfer to an existing cable TV system.

D

Employed by many cable television operators to provide Internet access over their existing infrastructure.

E

event

In Policy Management, an expected incident that is logged. Events can be used for debugging purposes.

F

FQDN

Fully qualified domain name

The complete domain name for a specific computer on the Internet (for example, www.tekelec.com).

A domain name that specifies its exact location in the tree hierarchy of the DNS.

G

GUI

Graphical User Interface

The term given to that set of items and facilities which provide the user with a graphic means for manipulating screen data rather than being limited to character based commands.

I

IP

Intelligent Peripheral

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.

N

network device	A physical piece of equipment or a logical (software) entity connected to a network; for example, CMTS, video distribution router, gateway router, or a link. This may also include sub-components of network elements (such as an interface) or lower-level devices such as cable modems or CPEs.
network topology	A map of physical equipment or logical entities in a network.

O

OSS	Operations Support System Computer systems used by telecommunications service providers, supporting processes 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

PCC	Packet Call Center Policy and Charging Control
PCMM	PacketCable MultiMedia
policy group	An ordered group of policies, organized for ease of administration or deployment.

Q

QoS	Quality of Service Control mechanisms that guarantee a certain level of performance to a data flow.

S

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 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

V

VoIP	Voice Over Internet Protocol Voice communication based on the IP protocol competes with legacy voice networks, but also with Voice over Frame Relay and Voice and

V

Telephony over ATM. Realtime response, which is characterized by minimizing frame loss and latency, is vital to voice communication. Users are only prepared to accept minimal delays in voice transmissions.

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.